

EMC[®] ViPR[®] SRM

Version 4.2.1

Web Portal Guide

302-004-843

02

Copyright © 2013-2018 Dell Inc. or its subsidiaries All rights reserved.

Published April 2018

Dell 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." DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

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

CONTENTS

Figures		5
Tables		7
Chapter 1	Web Portal Guide	9
	Information you can find here.....	10
	Key terms.....	10
	Browsing reports.....	12
	Navigation styles	12
	Displaying and canceling reports.....	14
	User Interface menu options.....	15
	Search features.....	16
	Display options.....	20
	Exporting a report.....	25
	Using Tools	26
	Report headers.....	32
	Working with personal views.....	34
	Working with dashboards.....	35
	Working with table reports.....	35
	Working with graphical reports.....	38
	Reading baseline reports.....	40
	Working with topology reports.....	41
	Editing reports.....	41
	About the Edit mode interface.....	42
	Workflow for creating new reports.....	44
	Filtering & Expansion tab.....	45
	Report Configuration tab.....	58
	Descriptions of report types	61
	Report Details tabs.....	64
	Formula tab.....	109
	UI Interactivity and Table Interactivity tabs.....	112
	Pre-Generation tab	114
	Report URL syntax.....	114
	Defining custom colors.....	117
	Accessing reports from outside of the Console	118
	Getting the complete URL for the current report.....	119
	Frontend report lookup tool.....	119
	Report URL syntax.....	121
	Customizing user settings and custom reports.....	124
	Modifying your user data and password.....	125
	Setting your reporting preferences.....	125
	Setting the first report to appear after login.....	126
	Saving the report tree.....	126
	Import a report definition.....	126
	Find and fix broken links in reports.....	127
	View fixed links in Tomcat logs.....	129
	Using the Report Library.....	131
	Accessing SolutionPack reports	131
	Viewing the summary of a SolutionPack before installing it.....	131

Administering the interface.....	131
Modifying the default profile settings.....	131
Set first report after login.....	132
Global Portal Properties.....	132
Scheduled Reports.....	135
Deleting stored reports.....	136
User Sessions.....	136
Listing reports that are generating.....	137
Accessing modules.....	137
Deleting metrics.....	137
Configuring an SMTP server	138
Managing users and user rights.....	139
User Management.....	139
Roles management.....	142
Profiles.....	147
Importing users.....	149
Managing ReportPacks.....	149
Creating a ReportPack.....	149
Uploading a <ReportPack>.arp file.....	150
Exporting a ReportPack.....	150
Deleting a ReportPack.....	150
Adding a new report template to a ReportPack.....	150
Copying a report template.....	151
Deleting a report template from a ReportPack.....	151
Exporting a report as a template in XML format.....	151
Centralized Management.....	152
Server Management.....	152
Package management.....	157
License management.....	158
Logs and diagnostics.....	160
Discovery Center.....	161
Groups Management.....	167
Central Configuration Repository.....	179
SolutionPacks.....	180
Data Enrichment.....	184
Enabling Impact Analysis.....	198
Alerting.....	198
Change the refresh rate on alerting reports.....	199
Alerting infrastructure.....	199
Access the Alerting portal page.....	200
Alerting node tools.....	200
Alert definitions.....	201
Alert definition component reference.....	215
Creating adapters.....	226
Creating alert templates.....	228
Grouped boxes.....	231
Probing alerts.....	233
Create alerts from scheduled reports	236

FIGURES

1	Groups Management UI	170
2	Groups Management UI—Rules.....	171
3	Filter in an alert definition.....	203
4	Operator in an alert definition.....	204
5	Comparator with two outcome paths.....	205
6	Actions in an alert definition.....	206
7	Example alert definition.....	207

FIGURES

TABLES

1	Properties searched within each database.....	17
2	Time span examples.....	23
3	Report headers.....	32
4	Tools and information.....	32
5	Report tree tools.....	42
6	Edit mode tabs.....	43
7	Table report types.....	61
8	Graph report types.....	61
9	Aggregated report types.....	62
10	TopN report types.....	62
11	Misc report types.....	63
12	Mixed report types.....	63
13	Formatters for value columns.....	66
14	Summary of predefined group types.....	168
15	Field names for filtering on event sources.....	203

TABLES

CHAPTER 1

Web Portal Guide

- [Information you can find here](#)..... 10
- [Key terms](#)..... 10
- [Browsing reports](#)..... 12
- [Editing reports](#)..... 41
- [Accessing reports from outside of the Console](#) 118
- [Customizing user settings and custom reports](#)..... 124
- [Using the Report Library](#)..... 131
- [Administering the interface](#)..... 131
- [Managing users and user rights](#)..... 139
- [Managing ReportPacks](#)..... 149
- [Centralized Management](#)..... 152
- [Alerting](#)..... 198

Information you can find here

Find information on how to use and edit reports, manage users, and customize the interface to meet your needs.

- Browsing Reports — for users who want to manipulate how reports are displayed.
- Editing Reports — for users and administrators who want to edit and create reports.
- Customizing User Settings — for users who want to set personal preferences or manage custom reports.
- Report Library — for users and administrators who work with installed SolutionPacks.
- Administering the Interface — for administrators who want to customize the interface for all users.
- Managing Users — for administrators who want to add users and set up user restrictions.
- Managing ReportPacks — for users and administrators who want to create, import, and export ReportPacks.
- Centralized Management — for administrators who want to manage, configure, update, or expand system components.
- Alerting — for users and administrators who want to enable, disable, configure, or create alert definitions, define alerts from data in scheduled reports, probe alerts, or manage other alerting components.

For the latest related documentation on ViPR SRM, go to the community network at <https://community.emc.com/docs/DOC-64008>.

For the latest related documentation on Service Assurance Suite, go to the community network at <https://community.emc.com/docs/DOC-61185>.

For release notes, security guides, software updates, or information about products, go to <https://support.emc.com>.

Dell EMC M&R 6.8u4 Help version 10 last updated: 04/17/18

Key terms

Review these commonly-used terms before you begin using the interface.

aggregation

Two types of values are stored in the database: real-time values and aggregated values. Real-time values are the actual values that were collected by the collectors. Aggregated values are the compound of several real-time values. Aggregated values improve report performance and save disk space. Aggregated values are created and updated as real-time values are collected. This way, aggregated values are always current and precise. A default installation contains 1 hour, 1 day, and 1 week aggregates.

backend

The backend module receives, normalizes, and consolidates the collected data and pushes it to the database. The backend performs data aggregation.

browse mode

The 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. In

Browse mode, an expansion on a device property will appear as one node per device in the report tree.

collector

Collectors are lightweight modules that gather time series data from supported devices and configuration items and ingest it into the processing pipeline as raw values.

data point

A data point is one unit of data that contains a timestamp, metric, and value.

edit mode

In Edit mode, you can create new reports, edit existing reports, and customize the report tree. When the interface is in Edit mode, the report tree tools that are used to add, delete, cut, copy, paste, and link report nodes are displayed. In Edit mode, the nodes in the report tree are not expanded as they are in the Browse mode. The report pane displays the report configuration tabs.

expansion

Expansion dynamically creates sub nodes. Sub nodes add levels to a report so that 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.

filtering

A filter is an expression that limits the number of metrics displayed in a report. Filtering 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.

frontend

The frontend refers to the main user interface that displays time series and event-based data in reports. The frontend interfaces with the database.

metric

A metric is a quantity that can be collected and reported on, such as: CPU usage, throughput, and capacity.

node

Nodes appear in the report tree. You click a node to generate and see a report. Reports appear on the right side when you select a node. Nodes can have parent and child nodes (sub nodes).

property

A property is meta data that is associated with a metric. For example, an iflnOctets metric could also contain information about the device type supplying the value as well as the device name.

root node

A root node is the topmost node of a branch in the report tree.

Browsing reports

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

Navigation styles

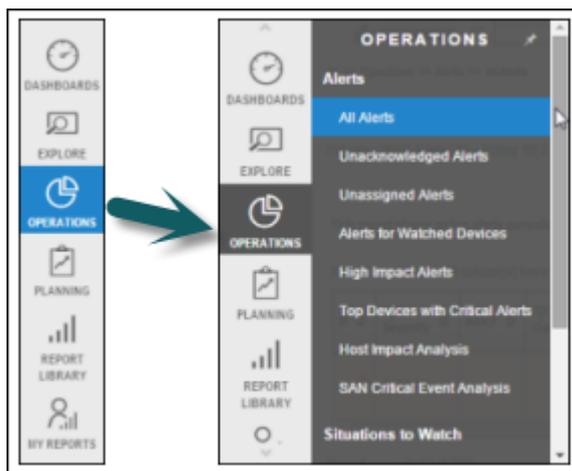
The navigation column on the User Interface provides access to all reports. You can choose between the Icon Navigation and Tree Navigation styles. The Administrator sets the default style for the installation.

Default style

The administrator sets a global default navigation style in the **Navigation Style** field at **Administration > Portal > Default Display**.

Each user account can override the global default style and set their own preferred style in the **Navigation Style** field at **Profile > View Profile > Preferences**.

Icon Navigation

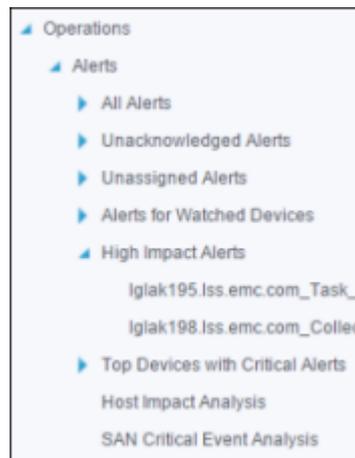


The Icon Navigation style helps preserve space on smaller-sized windows. Initially, it shows only a thin column of icons.

To navigate to a report, click an icon. A second column of categories appears. Click a category to expand it, and click again to select a report. Notice the scrollbar on the expanded second column.

The second column disappears when you click a report. To navigate further into the reporting structure, use the tabs, dashboards, and links in the reporting pane.

Tree Navigation



The Tree Navigation style can display the entire reporting structure, with multiple nodes expanded at once.

Use the arrows to expand or contract the nodes. Click any entry in the tree to display a report. You can also use the tabs, dashboards, and links in the reporting pane to navigate around the reporting structure.

Set the navigation style

You can change the navigation style for your user account.

The administrator sets a global default navigation style under **Administration > Portal > Default Display**. Your user account settings can override the global setting.

Procedure

1. In the banner at the top of the User Interface, click **Profile**  > **View Profile**.
2. Click **User Preferences**.
3. In **Navigation Style**, make a selection.

Option	Description
default	The administrator's global setting
tree	Hierarchical tree
icon	Thin icon column

4. Click **Save**.

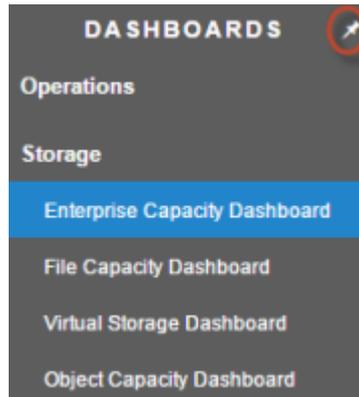
Make icon navigation column static

For the icon navigation style, use the pin to toggle the second column between temporary and static states.

In its default state, the second column disappears when you select a report. You can pin the second column so it remains visible and the report pane moves to the right to accommodate it.

Procedure

1. Click the pin at the top of the icon navigation bar.



2. Click the pin again to return the second column to a temporary display.

Displaying and canceling reports

When you navigate to a report on the Console, the system starts generating that report using the most recently collected data. If you navigate away from the page to another report, the system cancels the first report and starts generating the new report.

Here are some important points to remember about displaying reports.

Data collection occurs in scheduled intervals

There might be many data collectors in your system, each gathering data from different components in the infrastructure. Each collector is configured with a collection interval. Depending on how your administrators configured the various collectors, you might see some infrastructure activity reflected in reports within 15 minutes, and others not for a day or more.

One report per user session at a time

The Console actively generates one report per user session at a time. If you navigate to another report while the first one is still generating, the first report is cancelled, and the new one starts. If you are waiting for a long-running report to generate, do not click around examining other reports expecting the original report to finish running in the background.

Scheduled reports generate in the background

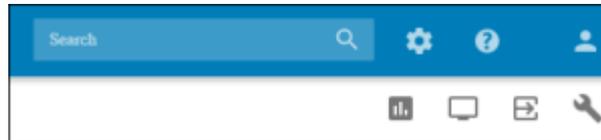
Scheduled reports do not affect the "one report per user session" rule. If you consistently depend on reports that take a long time to generate, you can put those reports on an automatic schedule, and view them in their completed state whenever you need them. Use **Tools > Schedule this Report**. You can request that the report be emailed, stored on a remote system, or stored for Console viewing under **My Reports > Stored Reports**.

Reports in Stored Reports are snapshots, and are not regenerated

When you navigate to a report in the **My Reports > Stored Reports** node, the snapshot of data from the stored date/time redisplay. The report is not regenerated. You can store a snapshot using **Tools > Store this Report**, or schedule the report as mentioned above.

User Interface menu options

The User Interface contains a set of menu options in the upper right corner of every report page.



Quick Search

Searches a predefined set of database properties for a requested value.

Administration

Contains links to administrative interfaces:

- **Administration**—Opens the **Administration** interface.
- **Centralized-Management**—Shortcut to **Administration > Centralized Management**.
- **Advanced Search**—Opens the Advanced Search dialog.
- **Modules**—Shortcuts to specific modules under **Administration > Modules**.

Help

Contains options for more information:

- **Documentation**—Opens the product Help topics.
- **About**—Information about the current software versions and licensing messages.

Profile

Contains Information and options for the current login session:

- Shows the user name of the logged in user.
- Shows the role of the user.
- **View Profile**—Displays the User Settings page, where you can set user-specific information, including a password change and user preferences, and manage custom reports.
- **Log Out**—Ends the current session.

Modifications

Contains options for changing or creating report definitions:

- **New Simple Report**—Use a wizard to quickly create a new report.
- **Edit Reports**—Enter Edit mode to change report definitions.

Display

Opens the dialog for choosing the report format and time span.



Export

Contains output formats for saving the currently displayed report. Click the Tool icon next to a format to configure the output.



Tools

Contains options for saving, sending, storing, and scheduling automatic runs of a report.

- **Show Report URL**
- **Print-Friendly page**
- **Store this Report**
- **Favorite this Report**
- **Schedule this Report**
- **E-mail me this Report**

Search features

Use the search features to retrieve information from the database and locate reports containing relevant information.

Quick Search

Searches a predefined set of database properties for a requested value. This search feature returns an organized set of links to the reports that contain the value.

To initiate a Quick Search, enter a search string in the **Search** text box in the banner area, and press **Enter**. The search string must be at least 3 characters and not contain wildcards.

Advanced Search

Provides a filter for defining detailed search conditions and an expansion field for defining what to return and how to organize the results.

To initiate an Advanced Search, click **Administration > Advanced Search**.

In either case, you can view the results in the report window and also in a **Search Results** node created in the report tree.

Quick search

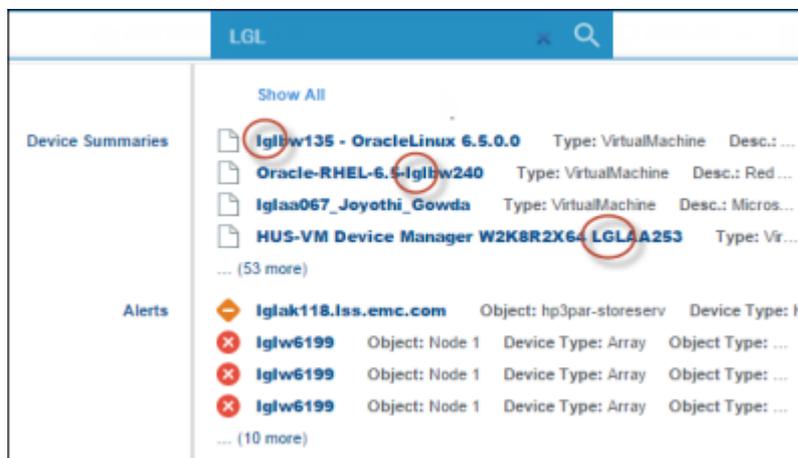
To perform a quick search, enter your search string in the **Search** text box in the banner, and press Enter. The search returns links to reports about the found values.

Note

The search string must contain at least 3 characters and no wildcards.

The search is case insensitive, and finds occurrences of the string anywhere in the property value (that is, at the beginning, middle, or end of the value). The result is a categorized list of links to reports that contain the matched value.

For example, the following search for LGL finds device names containing that value, and returns links to major reports about those devices. In this case, the links are to the device summary reports and to alerts on those devices.



The search looks at a limited set of predefined properties in each of the product databases. The following LGL table shows the database properties that are searched.

Note

The categories correspond to databases. Your installation might not use all of the databases.

Table 1 Properties searched within each database

Categories			
Device Summaries (APG)	Alerts	Compliance Breaches	VMware Events
device serialnb ip devdesc model vendor domain WWPNs for: <ul style="list-style-type: none"> • Server HBA ports • FC switch ports • Storage ports • VNX host ports 	severityAsString owner parttypedisplayname partdisplayname eventdisplayname category fullmsg eventSource devtype device	breachName device policy rule	datastore device Message VM Host User

The search results are links to reports. Results are presented in the following ways:

- A popup shows a preview of results. Click in the popup to go to the full result list or to a specific result.
- The full result list is a tabbed report by category.
- The navigation tree shows the search results in a node under **My Reports > Search Results**.

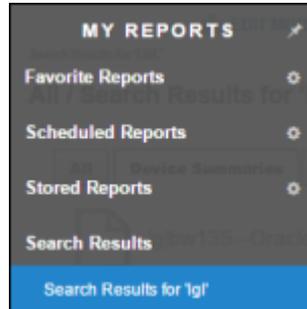
Procedure

1. In the **Search** box at the top of any report page, type a search string, and press **Enter**.

- The search is case insensitive. For example, 1g1 finds lgl and LGL.
 - The search string must contain at least 3 characters.
 - The search finds occurrences of the string in the beginning, middle, and end of values.
 - Do not use wildcards in the search string.
2. In the pop-up preview of search results, do any of the following:

Action	Result
Click Show All or one of the result categories on the left.	A tabbed Search Results report opens, showing all results, organized by category. Click a link in any of the lists.
Click a link in the list.	The selected report opens.

3. To return to the search results after navigating to a report, do either of the following:
- Use the browser's **Back** button.
 - Use the Search Results node in the report navigation tree, located under **My Reports > Search Results**.



Note

Only one search node is saved. If you perform another search, it overwrites the first search node.

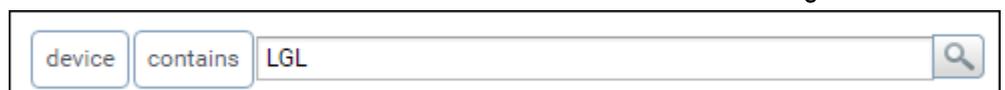
Advanced search

To perform an Advanced search, click **Administration > Advanced Search**. The dialog provides filtering and expansion fields.

Procedure

1. Click **Administration > Advanced Search**.
2. Right-click in the **Filter** box to define your search.

Here is a filter that searches for device names that contain the string LGL.



3. In the **Expansion** field, enter one or more properties to define nodes in the search results.

For example, if you enter `device`, the search results will contain nodes corresponding to each device that matches the filter. If you enter `bunit`

device, the search results will contain nodes for business unit from the records that match the filter, with sub nodes for the device values.

Click the **Property Selection Helper** for help in selecting properties.

4. In the **Search Base** field, select where you want the search to be applied and results to appear.

The search results node is subject to filters on nodes higher up in the hierarchy. Choose one of these options:

- Select **search from root** to add the search results node at the root.
- Select **search from the selected node** to add the search results under the node in the navigation tree selected in Step 1.

5. Click **Apply**.

The search results appear as a report. If multiple expansion fields were specified, click a row in the report to see the sub nodes.

To navigate to the results in the report tree, click **My Reports > Search Results**.

Note

Multiple expansion nodes are better viewed in the Tree Navigation style.

Example 1 Example

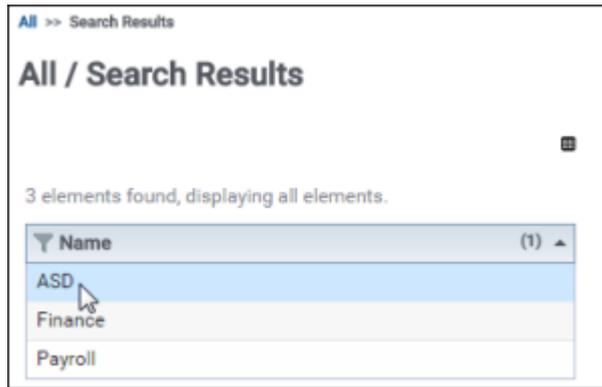
Here is a search definition that expands on `bunit` and then `device`.

The screenshot shows the 'Advanced Search' interface with the following settings:

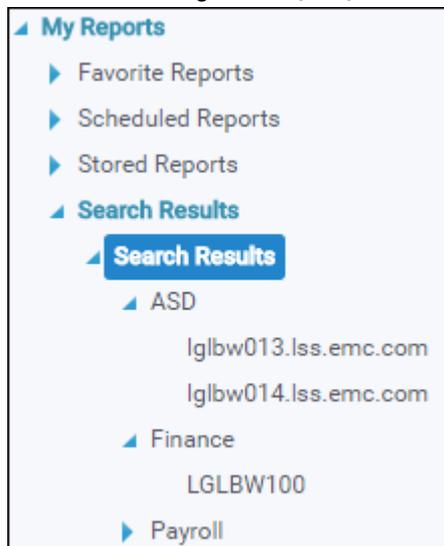
- Filter:** A text input field containing 'device contains LGL' with a help icon (question mark) to its right. Below it, the corresponding search syntax is shown as 'device=%LGL%'.
- Expansion:** A text input field containing 'bunit device' with a help icon (question mark) to its right.
- Search Base:** A dropdown menu with 'search from root' selected and a help icon (question mark) to its right.
- Apply:** A button located below the Search Base dropdown.

The **Search Results** report is a list of business units. Click a row to see the devices under a business unit.

Example 1 Example (continued)

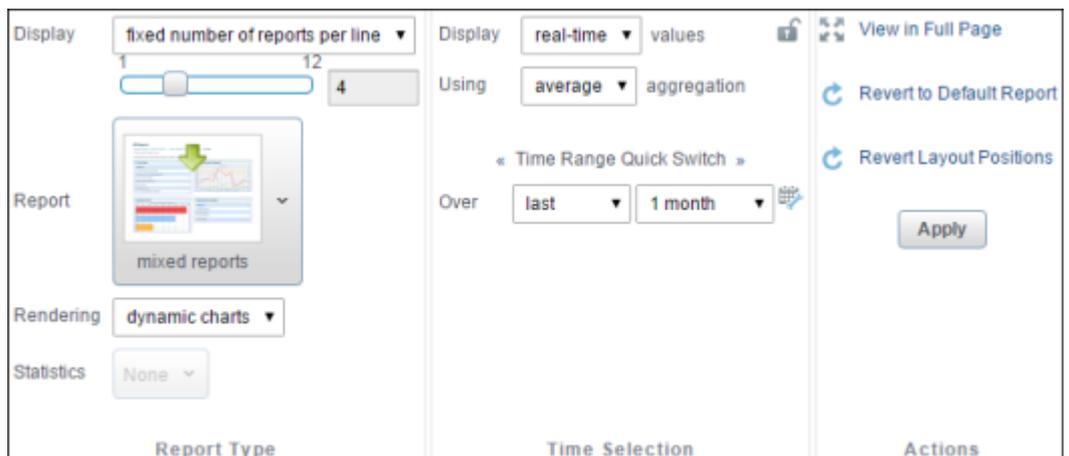


In the Tree Navigation style, you can expand multiple nodes:



Display options

Click the **Display**  icon to change how a report displays, including the report type, the time range, aggregation type of the metrics used, and report size.



Changing Report Type options

You can change some aspects of a report's presentation for your user account and current session.

Procedure

1. Click **Display** .
2. In the **Report Type** area, change one or more of the following parameters:

Option	Description
Display	Controls the width of reports. <ul style="list-style-type: none"> • Default display uses a default width. • Fixed number of reports per line controls the width of a single report based on how many reports you set per line. Use the slider that appears to set the number of reports per line. • Fixed width for reports sets the width of all reports on a line. Use the slider that appears to set the width. For example, you can set a line of gauge reports so they are all small or all large.
Report	Controls the report type. You can change this parameter only if other compatible report types are available. Click the down arrow in the report type icon. If other compatible report types are available, their icons appear. Click to select a different report type.
Rendering	Some report types offer different rendering options. Choose an option from the list.
Statistics	Some report types offer options for displaying statistics. Choose an option from the list.

3. Click **Apply** in the **Actions** area.

Changing Time Selection options

Options in the Time Selection area control the time period aggregates and the span of time covered in a report.

Procedure

1. Click **Display**.
2. In the **Time Selection** area, change one or more of the following parameters:

Display values

Choose the sampling period used to collect data from the database.

Sampling period is one of:

- **real -time**
- **1 hour**
- **1 day**
- **1 week**

Note

Selecting **real-time** for the display value gives the same output no matter which aggregation option is selected.

Using aggregation

Choose the type of sample collected from the database. Sample type can be:

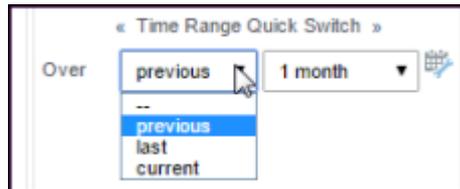
- **average**
- **min**
- **max**
- **sum**
- **last**
- **count**

Time Range Quick Switch

Time range defines the time that the report should cover. The range consists of a time span and a duration.

Time span

Choose the type of time span in the first drop-down list. Time span specifies when the reporting period starts and ends.

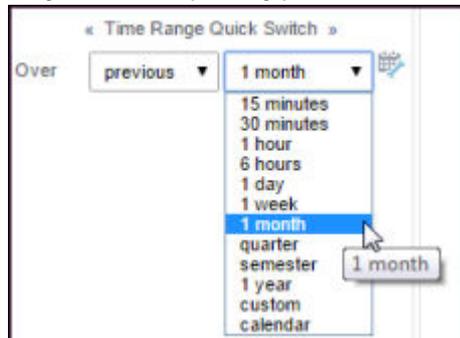


- **previous**—The reporting period starts and ends in the past.
- **last**—The reporting period starts at the appropriate interval counting back from the current time, and ends at the current time.
- **current**—The reporting period starts in the past and ends in the future, and includes the current point in time. (Metrics are for a partial reporting period.)

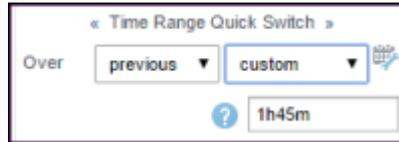
See the table below below for more information about time spans.

Duration

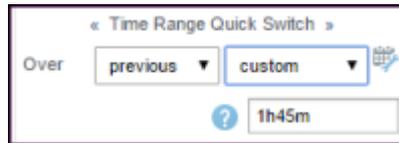
Choose a duration in the second drop-down list. Duration specifies the length of the reporting period.



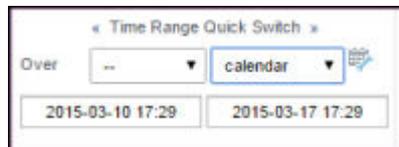
The **custom** option lets you specify an exact span of time, such as 1M2w (1 month and 2 weeks) or 1h45m (1 hour and 45 minutes).



The **custom** option lets you specify an exact span of time, such as 1M2w (1 month and 2 weeks) or 1h45m (1 hour and 45 minutes).



The **calendar** value lets you specify an exact date/time range, as shown here:



3. Click **Apply** in the **Actions** area.

Results

Table 2 Time span examples

Time Span Value	Explanation	
previous	The reporting period starts and ends in the past. Examples:	
	Previous hour	If the time is currently 10:15 AM, metrics are for 9:00 AM to 10:00 AM.
	Previous week	If today is Thursday, March 12, metrics are for the previous week, Monday March 2 to Sunday, March 8, inclusive.
	Previous month	If today is March 12, metrics are for the previous month, February 1 to February 28, inclusive.
last	The reporting period starts at the appropriate interval counting back from the current time, and ends at the current time. Examples:	
	Last hour	If the time is currently 10:15 AM, metrics are for 9:15 AM to 10:15 AM.
	Last week	If today is Thursday, March 12, metrics are for days Thursday to Thursday, March 5 to March 12.

Table 2 Time span examples (continued)

Time Span Value	Explanation	
	Last month	If today is March 12, metrics are for February 12 to March 12, inclusive.
current	The reporting period starts in the past and ends in the future, and includes the current point in time. (Metrics are for a partial reporting period.) Examples:	
	Current hour	If the time is currently 10:15 AM, metrics are for 10:00 AM to 11:00 AM.
	Current week	If today is Thursday, March 12, metrics are for Monday to Sunday, March 9 to March 15, inclusive.
	Current month	If today is March 12, metrics are for March 1 to March 31, inclusive.

Saving Display settings

Changes to Display settings normally apply to the current report and current session only. Use the following methods to make changes more permanent.

- Click the **Lock** icon. Your changes are applied to all reports that you browse during the current session. (Exception: On Dashboard reports of type **mixed using defaults**, the child reports always retain their default time settings.)
- Click **Tools > Save this Report**. All of your changes are stored as a report under **My Reports > Stored Reports**. This version of the report persists across sessions, but is visible only in your user account.
- To save your changes as the default display for all users, use **Modifications > Edit Reports**. On the **Report Configuration** tab, look for the following fields to change the reporting periods:
 - **Default Duration**
 - **Sampling Period**
 - **Sampling Type**

Locking Time Selection settings for your session

By default, changes to the Time Selection settings apply only to the currently displayed report. You might want to lock the settings so they apply to other reports.

If you apply new **Time Selection** settings and then click the lock icon, the new settings are applied to every report that you view during the current session, until you unlock them. Locking can be useful, for example, to have a snapshot in time for comparing equivalent graphs for a device.

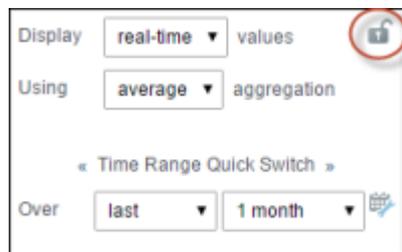
Note

In a report type of **mixed using defaults**, time setting changes, including locked values, do not apply to the child reports. The child reports in a **mixed using defaults** report always use the default time settings in the original report definition. Dashboards are often **mixed using defaults** reports.

To lock and unlock **Time Selection** settings:

Procedure

1. Click **Display**.
2. Make changes to the **Time Selection** settings.
3. Click **Apply** in the **Actions** area.
4. Click the **Lock** icon in the upper right corner of the **Time Selection** area.



5. To hide the **Display** menu, click outside of it.
You can now navigate to other reports and retain the locked time selections.
6. To unlock the settings, click **Display** and then click the **Lock** icon.
7. To return the current report to its default settings, click **Revert to default report** in the **Actions** area.

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**.

Exporting a report

You can export reports into PDF, Excel, CSV, XML, SVG, PNG, or JPEG formats.

The export feature sends the contents of the currently displayed report to the selected output format.

Note

On a tabbed report, the export applies only to the contents of the currently displayed tab.

Procedure

1. Click **Exports** .



Note

If a format name is grayed out, that format is not available for the currently displayed report.

2. To configure an output format:
 - a. Click the  icon next to the format name, if available.
 - b. Complete the format-specific dialog.

Format	Configuration fields
PDF	Configures page size, orientation, and layout.
Excel	Configures column layout and whether to show or hide titles, time ranges, and descriptions.
CSV	Configures column layout and whether to show or hide titles, time ranges, and descriptions. <div style="border: 1px solid black; padding: 5px;"> <p>Note</p> <p>Excel and CSV results are similar, but differences exist in column and title layouts.</p> </div>

- c. Click one of the following:
 - **Save and Export**—Save the settings on the dialog and export the report using the settings on the dialog. The changes pertain to all reports. They apply to your user account and persist across log-in sessions.
 - **Export**—Export the current report using the configuration settings but do not save the settings.
3. To export without formatting, click a format name in the list.
4. Respond to your browser's prompts to save the file in your desired location.

Using Tools

With utilities on the Tools  menu, you can use a wizard to quickly define a new report, get a report URL for referencing purposes, prepare a print-friendly page, save a snapshot of the currently displayed report, create report bookmarks, define a regular schedule for a report, and email a report.

Getting the complete URL for the current 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 currently displayed 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-friendly page**.
A page appears in a new tab without the navigation tree.
2. Use your browser's print feature to print the page, and then close the tab.

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**.

After having viewed a stored report, to navigate to the real-time data version, select the drop-down next to the report date and choose **Revert to real-time data**.

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. Optionally name the report.
3. Optionally click the **Use as Login Report** option if you want this report to display when you log in.
4. Click **Save**.

The report appears in the navigation tree under **My Reports > Favorite Reports**.

Displaying a favorite report at login

You can mark one of your favorite reports to automatically appear in the report pane after you log in.

Procedure

1. In the navigation tree, click **My Reports > Favorite Reports**.
2. Expand the blue bar for the report you want to display at login.
3. Click **Use as login report**, and then **Save**.

Creating a scheduled report

Scheduled reports are generated at fixed times and saved. The saved results can be stored under My Reports, emailed, transferred to a remote location, or analyzed by the alerting module.

Procedure

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

When the scheduled report is generated, it queries the database for the latest data.

4. Select **Active** to enable the schedule.

Deselect **Active** to disable report generation but save the scheduling and report disposition information. You can reactivate it later.

5. Use one or more of the other tabs to configure the disposition of the generated reports.

Tab	Description
Storage	Store the report under the My Reports node.
Email	Send the report to a list of recipients.
Remote Transfer	Send the report to a specified URL.
Alert	Send the report data to the alerting module.

6. Click one of the following:

Option	Description
Cancel	Exit the scheduled report dialog without saving the report.
Save	Save the new scheduled report (and use the report's default time settings as configured in the report definition).
Save modified report	Save the new scheduled report and use the time settings as configured in the Display menu

7. To make changes or add additional dispositions, edit the scheduled report.

Store a scheduled report in Stored Reports

You can store a generated report in the **Stored Reports** node on the Console. You can send email notifications every time a new version is stored.

A stored report is the default disposition of a scheduled report if you complete only the Scheduling tab and click **Save**. To view the stored report, navigate to **My Reports > Stored Reports > *scheduled_report_name***.

Procedure

1. Create a new scheduled report or edit an existing one.
2. Click the **Storage** tab.
3. Select **Store this report**.

The checkbox is selected by default when you create a new scheduled report.

4. For **Recipients**, type the email addresses to receive notifications each time the report is generated.

Leave this field blank if you do not want notifications.

5. Click **Save**.

Emailing a scheduled report

You can email the results of scheduled report runs to a list of recipients. You can choose one or more formats to send.

Before you begin

For successful emails, an administrator must configure an SMTP server as described [here](#).

Procedure

1. Create a new scheduled report or edit an existing one.
2. Click the **Email** tab.
3. Configure the email recipients, subject, and message of the email to send.
4. For **Formats**, select one or more formats for the report, such as PDF, CSV, or an image file.

Some format selections are greyed out if they are not appropriate for the report type.

A blue configuration bar appears for each format selected.

5. Optionally, click a blue bar to expand it and change the export settings for the format.
6. Click **Save**.

Results

After each scheduled run of the report, an email is sent to the list of recipients, with the report in each of the selected formats included as attachments.

Send a scheduled report to a remote location

You can transfer generated report files, in one or more formats, to specified URLs.

The file names of the transferred reports are the scheduled report name assigned on the **Scheduling** tab with an appended date, as follows:

reportName_date.extension

Procedure

1. Create a new scheduled report or edit an existing one.
2. Click the **Remote Transfer** tab.
3. Click **New Remote Location**.
4. For URL, type a well-formatted URL.

The ? icon next to this field contains more information and examples of acceptable URLs.

5. For **Formats**, select one or more formats to transfer, such as PDF, CSV, or an image file.

Some format selections are greyed out if they are not appropriate for the report type.

An expandable arrow appears for each selected format.

6. Optionally, click an arrow to expand a format type and change the export settings used by the format.
7. Click **Save**.

Send a scheduled report to the alerting module

You can send report data to the alerting module for analysis. The alerting module can then generate alerts or other actions based on the report data.

Before you begin

To generate alerts or notifications based on the data, an alert definition, defined in the alerting module, is required with these components:

- **Filter** — Identifies the report name to process. The <name> element in the XML file identifies the report name.
- **Conditions, comparators, etc.** — Identifies the metrics to analyze and the conditions that cause actions. The <th> elements in the XML file identify metric names.
- **Actions** — Defines the actions, such as writing to a log, sending an email, or generating an SNMP alert message.

The alert definition must be enabled to process the data.

Use this tab to send the report data, in the form of an XML file, to the alerting module.

Procedure

1. Create a new scheduled report or edit an existing one.
2. Click the **Alert** tab.
3. Select the alerting module to start sending data to it.

Deselect the alerting module if you no longer need to send data to it.

4. Click **Save**.

Results

Every time the report is generated, an XML file containing the data is forwarded to the alerting module. Nothing else happens on the alerting side if there is no enabled alert definition to process the data.

Using advanced settings

In typical situations, you do not need to use this tab.

The system uses default settings to manage memory and resources for report generation. In a few unusual circumstances, those default settings might not be adequate. This tab provides a way to override the system defaults for reports that need additional resources.

If a scheduled report is not completing as expected or is consuming excessive resources that you want to limit, you can edit the scheduled report and adjust the advanced settings.

Procedure

1. Click **My Reports > Scheduled Reports > *scheduled_report_name***.
2. Click the **Advanced Settings** tab.
3. Click the ? icon next to each field to determine whether to change the setting.
4. Click **Save**.

Edit scheduled report settings

You can change the settings on a saved scheduled report,

Procedure

1. Click **My Reports > Scheduled Reports > *scheduled_report_name***.
2. Change any setting on any tab.
3. Click **Save** on any tab.

Launching a scheduled report now

You can launch a scheduled report at any time.

Procedure

1. In the report tree:

Navigation Style	Description
Tree	Go to My Reports > Scheduled Reports .
Icon	Go to My Reports > Scheduled Reports >  . The  appears in the menu to the right of Scheduled Reports .

2. In the right pane, right-click the report you want to run, and select **Launch now**.
Watch the **Status** column for notification that the report is finished running.
3. The finished report appears in the location configured on the Stored Report tab for the scheduled report. The default location is in the report tree, under **My Reports > Stored Reports**.

Emailing a report to yourself

You can send a PDF of the currently displayed report to your email address.

Before you begin

1. To add or verify your email address, click **Settings** in the banner. On the **User** tab, complete the **Email Address** field.

- For successful emails, an administrator must configure an SMTP server as described [here](#).

Procedure

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

A PDF of the current report is sent to the email address specified in your user settings.

Report headers

The title area of a report contains useful tools, links, and information.

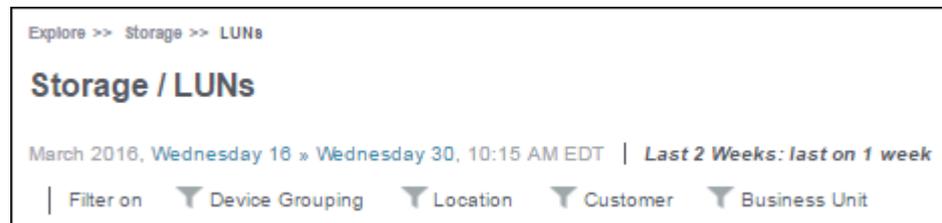


Table 3 Report headers

Line	Short Description	Explanation
1	Breadcrumb line	Shows the report path of the current report. Some components in the path are active links.
2	Report title	Shows the last two components of the report path.
3	Date/time span	Shows the date and time span being reported on, followed by the date/time settings that control the time span. Default date/time settings are specified in the report definition. You can change the settings for your user session on the Display  menu.
4	Group filters	(not on all reports) The Filter on line shows optional filters defined for the report. If available, these buttons provide an easy way to limit the report content. For example, use the Location filter to report on only selected location values.

The column headers in table reports also provide useful tools and information.

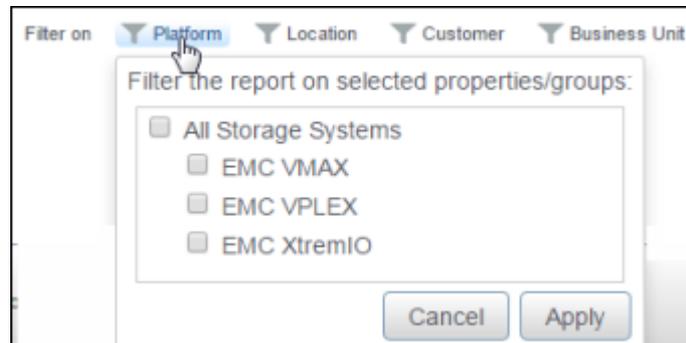
Table 4 Tools and information

Tool	Explanation
	Filters the report based on values in the column.
	Indicates that the report is sorted based on the column. Indicates whether the sort order is ascending or descending. To perform a sort based on a column contents, click the header. Click the header again to reverse the sort.
Tooltip	Hover the cursor over a column heading to show a column description. You can customize this description in Modifications > Edit Reports .

Filtering a report using group filters

Group filters appear across the top of the report. They present a dialog of checkboxes that let you select multiple values to filter on.

The dialog lists only the values that are relevant to the current report. The list reflects values found in your database and is further limited by the report filter in the report definition. For example, the following filter on the Platform property lists only the platforms that are being monitored at your installation and are relevant to the report definition:



Note the following information about group filters:

- If there are no values in a filter dialog box, your installation is not populating that property or no values apply to the current report.
- In a table report, the values that appear in a group filter dialog are not limited by column filters that might be in effect on the table. For example, a **Situations to Watch** report might contain a column filter that eliminates rows based on values in that column. In that way, a **Situations to Watch** report typically contains just a few rows. The group filter dialog, however, continues to list all values relevant to the entire report, not just to the displayed rows.
- A group filter can be defined for any database property using features in EDIT MODE.
 - When a group filter is defined on a data enrichment property maintained in the **Administration > Centralized Management > Groups Management** module, the values in the dialog are the group names as defined under Groups Management. Some groups have predefined names. Other groups define a Default group, and expect each installation to create additional groups.
 - When a group filter is defined on a collected property, the values in the dialog are the property values from the database.

Use the following procedure to set group filter values.

Procedure

1. Click the icon for a group filter.

A dialog of values appears.

2. Select one or more values to include in the report, and click **Apply**.

The icon for the applied filter changes color to indicate that the filter is active. The report redisplay, including only the data for the values you specified in the filter. For example:

- In a table report, the filter eliminates rows from the table.
- In a consolidated bar chart, the filter eliminates data from the metric calculations.

3. If multiple filters are available, optionally apply additional filters to further limit the displayed results.
4. To cancel a filter, click the filter icon and select **Clear**.

Working with personal views

You can create your own personal dashboards of favorite reports.

Creating a personal view of selected reports

A personal view is a dashboard of several of your favorite reports.

When you create your first personal view, a new node is created under My Reports. The new node is also named My Reports. You can define more than one personal view.

Procedure

1. Mouse over the upper right corner of a report.
A set of icons appears.
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.
4. To view your personal view, click **My Reports > My Reports > *personal_view_name*** in the navigation tree.

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**.

Working with dashboards

A dashboard is a collection of reports that all display on the same page. You can rearrange and resize the reports in a dashboard.

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.

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**.

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**.

Working with table reports

Users can customize table reports while in Browse Mode..

Applying filters to table columns

Filtering on a table column redisplay the report, showing only the rows that match the filter condition.

The filter icon () in a column header indicates that the column can be filtered. Columns with metrics, dates, and times cannot be filtered.

Procedure

1. Click the  icon in the column header.
2. In the text box that appears, enter the filter value using any of these methods:

Note

The values are case-sensitive.

- Type a valid value for the column
- Type a space, wait for the system to show suggestions, and select from the list
- Start typing a value, wait for the system to show suggestions that start with your entry, and select from the list
- Type a value containing wildcards. Supported wildcards are:

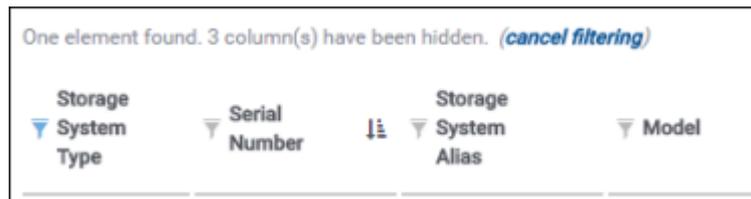
%	Matches any character any number of times. For example: <ul style="list-style-type: none"> ▪ <code>VPLEX%</code> matches any value that starts with VPLEX ▪ <code>%04</code> matches any value that ends with 04 ▪ <code>%Unified%</code> matches any value that contains the characters Unified in the beginning, middle, or end
*	Same as %.
_	Matches any one character. For example: <code>__%.%.%</code> matches any IP address whose first component is 2 digits.

3. Press **Enter**.

The report redisplay, showing only those rows with values that match the filter.

In addition, the following visual cues remind you that a filter is in effect:

- The filter icon for the filtered column is blue.
- The phrase **cancel filtering** appears in the sentence above the table.



4. To revert to the original report, click **cancel filtering**.

Sorting on a table column

You can change the sort order of a table report. The column headers indicate the sort columns.

There are two ways to change the sort order of a table report.

Temporary change

Click the column header to sort the table by that column.

Customize table columns

Use the **Customize Table Columns**  icon above the upper right corner of the table to set a more permanent sort, and to specify more than one sort column.

For a temporary sort change, use this procedure:

Procedure

1. Click a column header.

The table redisplay in resorted order, based on the column you requested. The



icon appears in the column header, indicating that the report is sorted by this column.

2. To reverse the sort order in that column, click the icon.
3. To sort by a different column, click the new column header.

The report is resorted based on the new sort request.

Show, hide, rearrange, and sort table columns

The **Customize Table Columns** icon lets you make custom changes to a table report while still in Browse mode.

Show or hide columns

You can customize a report by hiding unwanted columns.

Some reports define hidden columns. For example, many reports define columns for the business unit, customer, and location data enrichment fields, but hide the columns by default. If your installation defines values and data enrichment rules for one or all of those optional fields, those columns become meaningful.

Rearrange columns

You can move columns to appear in your preferred order.

Sort

You can override the default sort order defined for a table. You can specify one or multiple columns to use for sorting.

Procedure

1. Click **Customize Table Columns** , located in the menu that appears when you hover the cursor over the upper right corner of the table.

The **Table Customization** dialog appears.

2. To change the order of columns in the table, drag column names up or down in the **Displayed Columns** list.
3. To show or hide columns in the table, choose one of the following:

Option	Description
Display the column	The column is always displayed.
Hide the column if it is empty	The column is hidden if its cells on the current report page are all empty. With this option, a sparsely populated column might appear on one page and be hidden on the next page.
Never display the column	The column is always hidden.

Note

These options might not always be available. The underlying report definition controls whether users can show or hide a column.

4. To sort the table based on the values in one or more columns:
 - a. Drag the columns to use for sorting to the **Sorted Columns** list.
Drag a column into the sorted columns box, not the white space below it.
 - b. For each sorted column, select **Ascending** or **Descending**.
 - c. Drag to rearrange the sorted columns according to the sort order needed.
5. Click one of the following:

Button	Description
Apply	Apply changes on the report until it expires
Save and Apply	Apply changes and save permanently (until you save other changes)
Revert to default report	Cancel changes and revert to the report as defined in Edit mode
Revert to saved customization	Return to values saved with Save and Apply

Using the Actions Menu to take action from a report

Right-click on a table cell or icon to display an Actions Menu. The selections on the menu depend on the current report.

Procedure

1. Right-click a report table cell or icon.

Depending on the report, the Actions Menu appears containing one or more actions that you can take on the data in the row.

Note

Not all tables have Action Menus associated with them.

2. Select an action from the menu.

Working with graphical reports

Users can customize graphical reports while still in Browse Mode.

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**.

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.

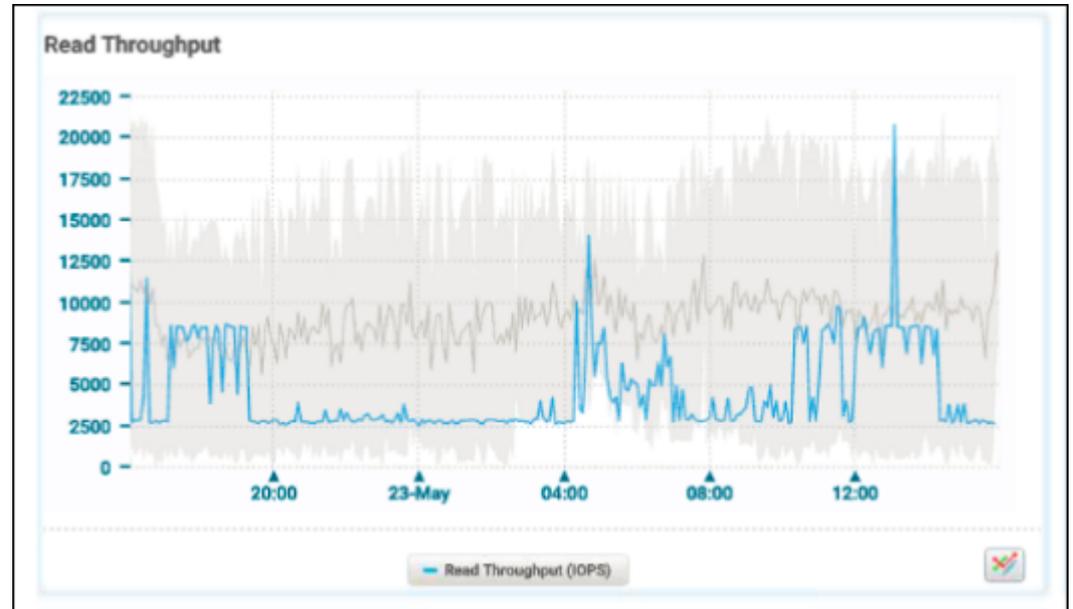
Procedure

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.

Reading baseline reports

Baseline reports show performance information about a single metric over a 24-hour period. The baseline helps you understand whether the current value of the metric is normal or a deviation that indicates a problem.

The following example baseline shows current and baseline values over time for a read throughput metric.



A baseline report contains the following information:

- Blue line—Shows the value of the metric over the last 24 hours.
- Gray line—Shows the average of the metric over the past 4 weeks, for the same day of the week and hours represented by the blue line.
- Shaded gray area—Shows the range of all values collected over the last 4 weeks for the same day of the week and hours represented by the blue line.

A baseline report needs more than 1 week of data before it can display the gray line representing the baseline average. When less than 1 week's values are available, only

the blue line showing values for the last 24 hours appears on the report. A baseline is fully functional when 4 or more weeks of data are available.

Baselines are particularly useful for interpreting performance metrics. The baseline lets you compare the current value of a metric to the average value over time. A deviation from the average stands out visually, drawing the attention of the user. You can also notice trends in the performance over time.

Working with topology reports

Topology reports show relationships between components in the infrastructure. You can drill down from a topology map to detailed reports about the components.

Topology reports

A topology report displays a graphical representation of the configuration items in the infrastructure and how they are linked together.

Topology reports depend on successful discovery of configuration items. Only discovered configuration items appear in the report.

A topology report consists of connected nodes. A node represents a configuration item.

- Click a node icon to drill down to the home report or other available details for that node below the topology map.
- Click the + symbol on a node icon to expand the topology to show more detail.
- Click the - symbol in the lower left corner of an expanded node to contract it.
- Hover the cursor over any node or port symbol to display a tooltip. Tooltips show details such as model names, utilization metrics, or port numbers.

Finding and filtering nodes on topology reports

You can find and filter specific nodes on a topology report.

Procedure

1. On a storage system mixed report, hover the cursor over the **Topology Map** title.

A text box and several buttons appear.

2. To find a node on the topology report, type the beginning of the node name in the text box and click **Find a node** (magnifying glass).
3. To restrict the map to show one node only, type the beginning of the node name in the text box and click **Filter nodes** (funnel).

Editing reports

Use Edit mode to create new reports, edit existing reports, and customize the report tree.

Your user account must have appropriate permissions to enter Edit mode.

To enter Edit mode, click **Modifications**  > **Edit Reports**. Although you can navigate to different reports from Edit mode, it is easier to navigate to the report you want to edit first.

Changes that you save in Edit mode are visible to all users, unless you make the changes under the **My Reports** node.

About the Edit mode interface

The interface changes in Edit mode, providing ways to change both the current displayed report and the report navigation tree.

Report Navigation

The report navigation column on the left changes into a tree of nodes and displays a set of icons across the top for editing the tree. With the icons, you can add a new report to the tree, copy and paste a report to a different location in the tree, or delete a report.

Report window

The reporting window on the right changes and displays a set of tabs that define all attributes of the current report. You can make changes to the report definition by changing settings in the tabs.

Report tree tools

The report tree tools are icons that appear above the report navigation tree in Edit mode. The icons provide report tree management features, such as creating a new report, copying and pasting reports to another location in the tree, and deleting reports.

Table 5 Report tree tools

Icon	Description
 Create a new report	Creates a new report node as a child of the selected one. The node is called New Report and can be customized.
 Cut the selected report	Cuts a report node to be pasted elsewhere in the tree. After you paste it, the report node is deleted from its original location.
 Copy the selected report	Copies a report node and its children to be pasted elsewhere in the tree. Although the report node and its children appear in its original and new location, copied report nodes are not linked in any way. Modifications made to a copy are not reflected in the original report. Use the Link tool to create shortcuts.
 Paste as child report	Pastes a copied or cut report node.
 Paste as a link	Creates a link to a previously copied report node. A linked node is not a copy of the original report node, but a reference to it. Any modifications made to the original report node are applied to the link node. Be careful because the link will be broken if you remove, move, or alter the tree structure around the linked node. A broken link has a warning icon.

Table 5 Report tree tools (continued)

Icon	Description
 Delete the selected report	Deletes the selected report node and all of its children. This action can not be reverted.

Report editing tabs

The report editing tabs appear in the reporting window in Edit mode.

Each tab contains parameters used to generate the report. These parameters define the report presentation, the report contents, and how the report interacts with other reports.

Table 6 Edit mode tabs

Tab	Description
Filtering & Expansion	Defines the basic properties of the report node, what will be reported on, and how the data fields are expanded within the report.
Report Configuration	Defines the core settings for the report, the report type, the time period of the report, and level of data aggregation. The report type (such as table, graph, heat map, mixed, and so on) determines the parameters that appear on the Report Details tab.
Report Details	Defines the detailed appearance and contents of the report. The parameters on this tab are specific to the report type. For a table, this is where you define all of the attributes of columns in the table. For a graph, this is where you define the attributes of the axes.
Formula	Manages formulas applied to the node.
Interactivity	Controls the interaction between reports.
Pre-generation	Generates reports on a predetermined schedule.

Using My Reports as a workspace

The **My Reports** node provides a workspace for creating new reports, editing reports, and testing changes.

Changes that you save while in Edit Mode are visible to all users, unless you make the changes in the **My Reports** node.

Note

It is not recommended to directly update the global reports.

Use any of the following methods to move a report into the **My Reports** node.

1. Pin the report to a personal view.
2. Copy and paste a report from the main tree into the **My Reports** node. Do not use a link if you want to keep your copy separate from the copy in the main report tree.

To make your changes visible to all users, coordinate with an administrator to prevent multiple users from simultaneously editing the same template. Administrators might have you:

- Edit directly.
- Copy and paste.
- Extract the report as a template, add the template to a ReportPack, and hook the ReportPack to nodes in the main report tree. This method provides the best way to control the hooks and manage the tree hierarchy.

Simultaneous edits by multiple users

If more than one user has write access to reports outside of **My Reports**, then it is possible for two users to edit report templates at the same time, affecting each other's modifications.

Templates or ReportPacks can become corrupted as a result. Only when the first user is finished editing, saves the changes, and logs out of the session can another user safely edit a report affecting the same template.

The following permission controls and change management procedures are recommended to avoid potential problems with multiple users affecting the same template:

1. Write access to global report templates should be assigned to a limited number of Administrator users. Assign this limited permission on **Administration > Roles > Template Access**. Most users should have read-only access to templates.
2. Instruct users to edit reports only under **My Reports**. This is always safe because the changes are visible to only the current user and templates are not affected. No special permission is required.
3. Implement an internal change management process for users to submit requests to an admin requesting changes to the global reports.
4. Always make a backup of the report tree before you begin making major changes to reports or the report tree outside of the **My Reports** node. You can create a report tree backup on **Settings > Custom Reports > download a backup**.

Workflow for creating new reports

When creating a new report, decide first on the metrics to include, and then set the report type and the report details.

The workflow for creating new reports is as follows:

1. Enter Edit mode.
2. Using the tools at the top of the report navigation tree, create the node for the new report in the report tree. There are two basic approaches:
 - Copy an existing report to use as the basis for the new report
 - Create a new report from scratch

Either way, the remaining part of the workflow is the same.

3. 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.
4. 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.

5. Depending on the report type selected, the **Report Details: <Report Type>** tab appears where you set the parameters for presenting information in the report.
6. Using the **Formula** tab, you can apply formulas to any node of the tree to perform advanced computations on any number of values.

Workflow for creating a mixed report

Mixed reports let you combine several reports in a single report page or apply an overlay that displays events on top of a regular graph.

The following types of mixed reports are supported.

mixed report	Uses the report type from any child node but retains the selected time period and aggregation from the selected node. This report is compatible with the mixed defaults report.
mixed using defaults	Uses the default report settings from any child node. This report is compatible with the mixed report.
overlay	Displays events superimposed on top of a graph report and a child table report listing a superset of events plotted in the overlays.
tab report	Shows a tabbed interface for navigating to reports.

A mixed report uses the child nodes in the report tree as the source of the reports to include. To create nodes, use the report tree tools to add a new report or copy and paste reports.

Procedure

1. Create a new node in the report tree for the mixed report.
2. Under the new node, create child nodes for the reports in the mixed report.
3. In the report tree, select the parent node.
The configuration tabs for that report appear in the report window.
4. On the **Report Configuration** tab, select the type of mixed report.
Select **mixed**, **mixed using defaults**, **overlay**, or **tab**.
5. On the **Filtering & Expansion** tab, expand the **Child Node Ordering** section.
The child nodes from the report tree are automatically listed there.
6. Edit the ordering of the nodes, if needed.
7. Click **Save**.
8. Click **Browse mode** to see the shell of the mixed report.
For example, if you selected a tab report, you can see the tabs across the mixed report page.
9. Click **Edit mode** and complete the configuration for the mixed report and all of its child reports.

Filtering & Expansion **tab**

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.

Predicate	Description
#<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
*	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 the following expansion types:

- sql
- regex
- dual-regex
- split
- hide
- in-list

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.

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>`

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>`

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>`.

Hide expansion (type=hide)

This expansion is used to prevent a property from being displayed when expanding on multiple properties.

For example, suppose you want to expand a node on `device,deviceid` because you have some hosts sharing the same device property. `Deviceid` is used to prevent hosts with the same device property from being merged into a single node, but the `deviceid` property is a long text string that you don't want to have displayed. By using the hide expansion (`device,deviceid<type=hide>`), the expansion works as it normally would only the `deviceid` is not displayed.

Example of hosts without using the hide expansion:

- example-host1, afekru23417
- example-host2, fdg98713424
- example-host3, dsfhng32442
- example-host3,dfsanl123412
- example-host4,134wf343244

Example of hosts using the hide expansion:

- example-host1
- example-host2
- example-host3
- example-host3
- example-host4

In-list expansion (type=in-list)

This expansion checks if the value of a property is inside a list of another property value. Expansion occurs only if the value of a property in the parent is inside the property list value of a child.

This expansion uses these parameters:

Values separator	The value separator.
List property	Property containing the list.
Name override	Can be used to override the generated node names.
Level up	Tells how many nodes should be skipped in the filter computation.

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**.

Add or remove group filters

A group filter appears in the header of a report. It provides a convenient way for users to filter a report on a predefined property or set of properties.

A group filter presents values for one or more properties in a checkbox format. Users can select multiple property values under a property, and the report is filtered to show data for the selected values. Any property that appears in the report can be used in a group filter, including the data enrichment properties.

Use this procedure to add or delete a group filter. For group filter syntax information, see Group Filter in the Help topic named **Filtering and expansion parameters**.

Procedure

1. In Edit mode, click the **Filtering & Expansion** tab.
2. Expand the **Group filter** section.
3. To remove a filter:
 - a. Click **Switch to input field**.
 - b. Remove a property from the expression.
4. To add a new filter:
 - a. Click **Add Property**.
 - b. Click **Simple** at the top of the dialog.
 - c. Select the property you want to create a filter for.

To be meaningful, the property must be used in the report. Use the **Report Details** tab to research the property names used in the report.

- d. To combine properties in the same filter, select multiple properties and then click **Selected properties must be grouped** at the bottom right of the dialog.
 - e. Click **OK**.
 - f. Click **Switch to list** and edit the expression if needed.
5. 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. The filter is used to craft the expansion, but the filter will not be used when you drill down. This is useful when you are doing a filter on a property which is not shared by every raw data (like devdesc).
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. The system expands on everything and then applies the filter for underlying data. This is useful when you have inventory reports, where some expansions can have zero children.
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. The system expands on matching raw data. The filter is kept when drilling down.

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.

Group filter

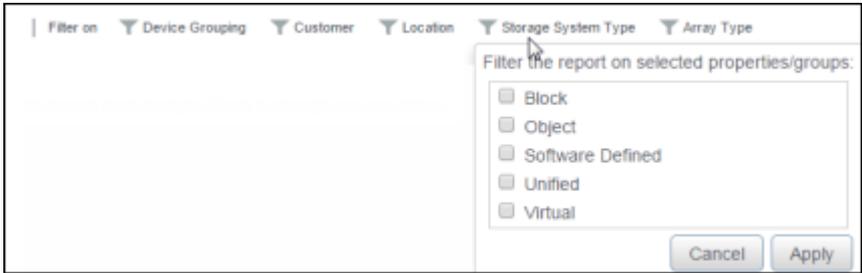
This setting creates property-specific filters in Browse mode. The result is a filter icon with the property name at the top of the report, and an associated dialog box that lists values in a checkbox format. This feature works with any property, including the property tagging filter (PTF) groups that are managed in the **Centralized Management > Group Management** interface.

Add Property adds a new property to the group filter expression.

Simple	<p>Adds a simple filter for a database property, including flat groups managed in Centralized Management > Groups Management. To add a simple filter:</p> <ol style="list-style-type: none"> 1. Click Add Property > Simple. 2. Select the property. Use the Report Details:Table tab to research properties that appear in the table. 3. To combine multiple properties into the same filter, select the properties and then click Properties must be grouped.
---------------	--

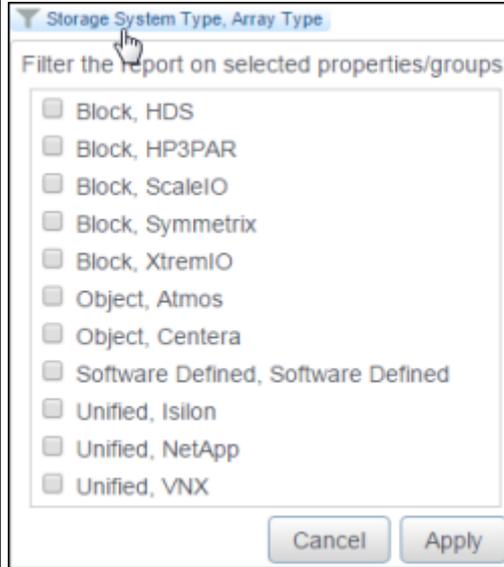
<p>Complex</p>	<p>Adds a filter with a dialog of choices that matches the hierarchy of group names in a hierarchical group type managed in Centralized Management > Groups Management. For example, in Groups Management, Device Grouping is a hierarchical group that expands on the <code>nodegrp</code> property. To add a filter for a hierarchical group type:</p> <ol style="list-style-type: none"> 1. Click Add Property > Complex. 2. Choose Expand on grouping ui formatted data enrichment properties. 3. Follow the steps in the expansion wizard. 4. To match the group management interface, use the / character as the hierarchical (level) separator and the character as the expansion-separator. 5. Click OK. <p>Here is the complex expansion for the <code>nodegrp</code> property. It creates a hierarchical dialog box of filter values that matches your hierarchy of device group names defined in Groups Management.</p> <pre>nodegrp<type=grouping-ui;level-separator=/;expansion-separator= ></pre>
-----------------------	--

All selected properties are combined into a single group filter expression. After adding properties, click **Switch to List** to see the entire group filter specification. You can make adjustments to the syntax. The syntax for multiple filters is described below:

<pre>prop1 prop2</pre>	<p>Creates separate filters and dialogs for each property. For example, the following expression creates a filter for a complex expansion on the <code>nodegrp</code> property and several additional filters for simple properties.</p> <pre>nodegrp<type=grouping-ui;level-separator=/;expansion-separator= > customer location sstype arraytyp</pre> <p>Below is the result in Browse mode, with the dialog for the Storage System Type filter open.</p>
	
<pre>prop1,prop2</pre>	<p>Combines two properties into the same filter and generates a dialog with all value combinations in a flat list.</p>

Here is an example with two properties, and the resulting dialog in Browse mode.

(sstype, arraytyp)

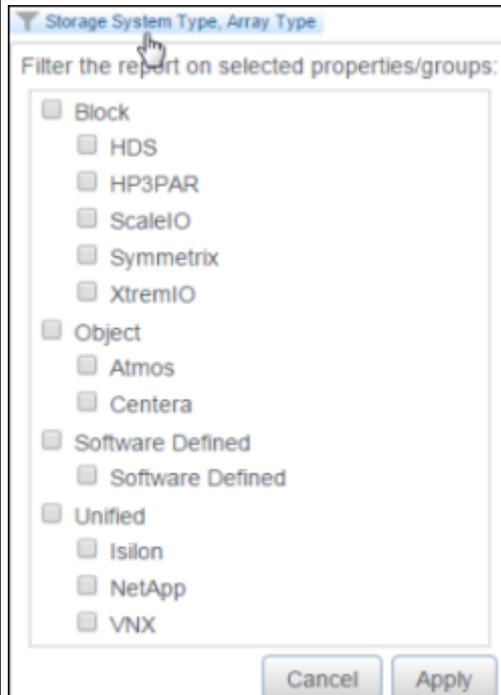


(prop1), (prop2)

Combines two properties into the same filter and generates a dialog with value combinations in a hierarchical list.

Here is an example with two properties, and the resulting dialog in Browse mode.

(sstype), (arraytyp)



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 tab

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.

The report types are divided into the following categories:

- Tables
- Graphs
- Aggregated
- TopN
- Misc
- Mixed

Within a report category, all of the report details are the same. This means that you can change the **Report Type** to another one in 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.

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.

Outages type

Set the type to the category set in the Outage Manager.

Outages property

Set the value to the property you would like to use for matching. If the value of the property matches one of the objects in the Outage Manager, that outage will be applied.

Descriptions of report types

In **Edit** mode, on the **Report Configuration** tab, choose the **Report Type**.

Tables

The following report types appear in the **Tables** category.

Table 7 Table report types

Report 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.
item	Displays a limited amount of data as an icon. You can define up to 5 data values (columns) and place them in desired locations within the icon. There can be only one sub node (that is, only one row when converted to a standard table).

Graphs

Table 8 Graph report types

Type	Description
simple chart	Gathers variables from the selected node. Draws a single graph showing variables as points over a time series.

Table 8 Graph report types (continued)

Type	Description
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.
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.

Aggregated**Table 9** Aggregated report types

Type	Description
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.
status	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.
heat map	Displays metrics in a two-dimensional grid. The cells are color-coded according to defined thresholds.
tree map	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.
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.

TopN**Table 10** TopN report types

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

Table 10 TopN report types (continued)

Type	Description
	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
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.

Misc**Table 11** Misc report types

Type	Description
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.
icons	Displays icons and associated descriptions.
external	Displays a referenced web page.
map	Displays child nodes as a point on a geographical map.
topology	Displays a graphical representation of discovered configuration items in an infrastructure and how they are linked together.

Mixed**Table 12** Mixed report types

Type	Description
mixed	Displays the set of child reports defined in the report tree on a single report page. All child reports use the selected time range and aggregation of the parent mixed report.
mixed using defaults	A mixed report that always uses the default time range and aggregation defined for each child report.
overlay	Plots events defined in a table report onto a graph.
tab	Displays a parent report with child reports as interactive tabs.

Report Details tabs

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

The **Report Details** tab name and the parameters appearing on the tab change 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

You can add columns to a table, change the configuration of sub nodes that populate the rows, or change the formatting of cells in the table.

Table columns are defined as attributes, properties, values, or aggregations or combinations of other columns. You can add new columns and change the order of the columns.

A table row is generated for each sub node. Usually, the sub node is expanded on a property.

Setting threshold definitions in tables

Set thresholds to visually alert you when a metric has exceeded a threshold.

Procedure

1. Click the **Report Details: Table** tab.
2. Expand a value.
3. Expand **Value Settings**.
4. In the **Thresholds Definition** section, select the name of the threshold, enter a value, and select a color for each threshold you want to define.
5. In **Value Formatter**, select how exceeded thresholds will be displayed in the report.
6. 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**.

Change a column header tooltip

You can customize the tooltip that describes the column when a user hovers the cursor over the column header.

Procedure

1. Click the **Report Details: Table** tab.
2. Expand the blue bar for the column whose description you want to change.
3. Change the contents of the **Description** field.
4. Click **Save**.

Change refresh rate

For table reports, you can turn the automatic refresh feature on or off, or change the refresh rate.

Procedure

1. Click the **Report Details: Table** tab.
2. Click to expand the **Display Options** section.
3. For **Refresh Interval (secs)**, enter an integer to set a new rate, or enter 0 to turn off automatic refresh.

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. Add a column by clicking **+Attribute**, **+Property**, **+Value**, **+Aggregated**, or **+Combined**.

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.
Aggregated	An aggregated value (average, minimum, maximum, or sum) of previously specified value columns.
Combined	Displays the contents of multiple previously specified columns in a single column.

3. In **Column Name**, enter a name for the column.
4. Select the attribute, property, or value to display in the column.
5. In **Column Description**, enter a description of the column that appears as a tooltip when the user hovers the cursor over the column header.

6. Click **Save**.
7. Click **BROWSE MODE**.
8. Verify that the new column appears in the table.
 If not, do the following:
 - a. Hover your cursor over the upper right corner of the table, and click the **Customize Table Columns**  icon.
 - b. Find the new column name in the list and make sure that **Display the column** is selected.
 - c. Click **Save and Apply**.
 This adds the new column's display preferences to other previously saved custom settings.
 - d. If the new column still does not appear, return to Edit Mode, click the **Report Details: Table** tab, and check the value in the **Advanced Properties > Display Customization** field.
9. Instruct other users who saved customizations on that report to enter the Table Customization dialog and click **Save and Apply** again.

This action is required to add the new column to the saved settings.

Formatting metrics in a value column

You can assign formatters to define the presentation of the data in the cell.

The following formatters are available for table columns that are of type Value. Many of the formatters depend on threshold settings. A column can have multiple formatters. For configuration details about a formatter, see [Table report parameters](#) on page 70.

Table 13 Formatters for value columns

Formatter	Description
Number	Displays the value in numeric form, with specified rounding and prefix or suffix symbols, such as currency symbols.
Value to string	Converts a value to a specified string. You define value-string pairs in the formatter configuration.
Date and time	Displays the date and time when the value was collected, converted from the UNIX timestamp. You specify the format in the formatter configuration. Options are: date and time, date, time, using several format variations (full, long, medium, short) as well as a custom format based on Java's SimpleDateFormat.
Duration	Formats a number of seconds into larger units, such as weeks, days, hours, and minutes.
Bar inline chart	Displays an inline bar graph whose number of bars depends on the selected sampling period and report time range.
Performance inline chart	Displays a horizontal bar showing the value and where it is located compared to the column's defined thresholds.
Time value inline chart	Displays an inline line graph that plots the metric values over the report's time period.

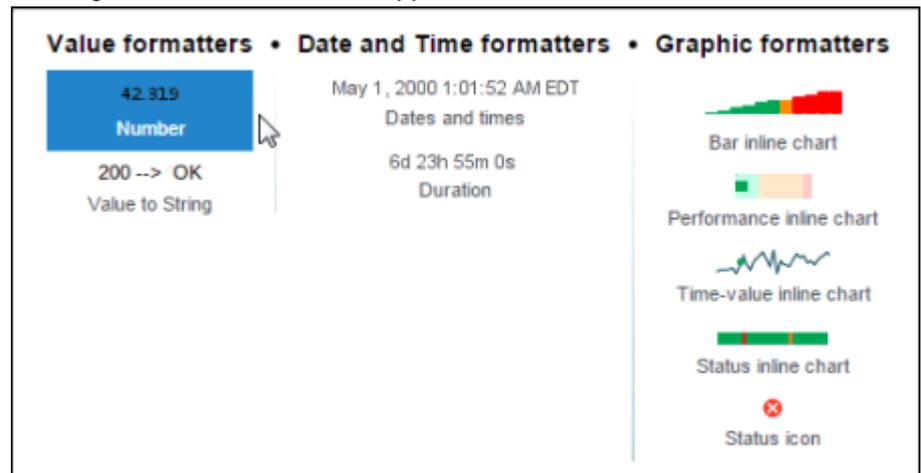
Table 13 Formatters for value columns (continued)

Formatter	Description
Status inline chart	Displays an inline status bar showing the status value over the time period. The line shows the status of each point of a value. The status colors of the band are set according to the threshold severities set on the column.
Status icons	Displays the status icon for the value according to the column's defined thresholds.
Additional Formatters located in the Decorations section	
Row and cell background formatters	Changes the background color for a row or a column cell.
graphic tooltip	Adds or removes a contextual graph that appears like a tooltip when the user clicks a cell.

Procedure

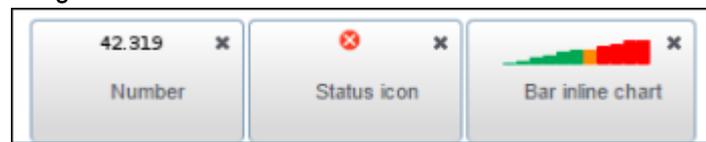
1. Click the **Report Details: Table** tab.
2. To add a formatter to a Value column:
 - a. Click to expand the blue bar for a Value column.
 - b. Click **Value Settings**.
 - c. In the Value Formatters section, click **+Formatter**.

A dialog of available formatters appears.



- d. Click a formatter to select it.

The formatter is added to the set of formatters for the column. A rectangle represents each assigned formatter. Here is an example of three formatters assigned to a column.



Note

For new columns, the **Number** and **Status icon** formatters appear by default.

Note

To delete an unwanted formatter, click the **X** in the upper right corner.

- e. Repeat these steps to add another formatter.

A column can have multiple formatters. A scroll bar might appear after adding several formatters.

3. Click a formatter box to expose and configure its optional settings.

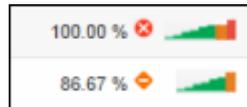
For example, the **Number** formatter lets you define the number of decimal positions, and prefix or suffix strings (for currency symbols, among other uses).

4. (Optional) Rearrange the order of the formatters by dragging and repositioning them.

The order of the formatters controls the order of appearance in the table cells.

5. (Optional) Under **Decorations**, configure backgrounds, and add or remove the graphic tooltip.
6. Verify the look of a cell in the **Column Preview** section.
7. Click **Save**.
8. Click **BROWSE MODE** to return to the report and view the new formats.

Here is an example of two cells from a column using the Number, Status icon, and Bar inline chart formatters. The number formatter is configured using a % symbol as a suffix.



Formatting text and backgrounds in columns

You can use text decorations for any column type to define the formatting of the text in a cell. For value columns, you can also apply colored backgrounds to cells.

Procedure

1. In Edit Mode, click the **Report Details: Table** tab.
2. Expand **Value Settings**.
3. In the **Decorations** section, configure the available formatting fields.
 - **Cell background** applies the threshold severity color to the cell. This option is available only for value columns.
 - **Row background** applies the threshold severity color to the entire row. This option is available only for value columns.
 - **Text Mode**
 - **Default Style** has no interactive action.
 - **Link Style** adds an underline when the cursor hovers over the text, indicating a link. To add the link target, see the **Table Interactivity** tab.

- **Text Color** applies colors to the column text. Choose a predefined Default or Severity color, or choose Custom Color and then define the color.
 - **Text Style** applies font styles to the column values.
 - **Graph tooltip** shows a contextual graph above the cell when the cell is clicked. This option is available only for value columns.
4. Verify the look of the cell in the **Column Preview** section.
 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.

Adding a total row at the bottom of a column

You can add a row to display the total or average value of all of the visible rows for a particular column at the end of that column. For example, suppose you have a column that lists individual storage capacities. You can add a row at the end of that column that displays the sum of all of the individual capacities to give you the total capacity.

Procedure

1. Click the **Report Details: Table** tab.
2. Click the **Value** icon of the column in which you want to add a summary row.
3. Expand **Summary settings**.
4. In **Label**, enter the name for the sum or average that will be calculated and displayed, such as Total or Average.
5. In **Generate column**, select either **sum** or **average**.

6. Click **Save**.

Table report parameters

You define the columns in a table report on the **Report Details: Table** tab. Parameters are based on whether the column is an Attribute, a Property, a Value, an Aggregate of the value from other columns, or a Combined cell showing data from other selected columns.

Attribute

An attribute column displays information about the direct children of the currently selected parent node.

Column Name	A unique name that appears in the table header.
Description	The description that appears when the user hovers the cursor over the column header. You can edit the description. See the tooltip for information about supported HTML tags, attributes, and classes.
Attribute	<p>Defines the property to retrieve from the requested node.</p> <ul style="list-style-type: none"> • 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.
<input type="checkbox"/> Advanced Properties:	
Display Customization	<p>Selects whether the column is displayed in Browse mode. The Hide if empty choice means that the column is hidden if its cells on the current report page are all empty. A sparsely populated column might appear on one page and be hidden on the next page.</p> <ul style="list-style-type: none"> • Locked Mode selections make a permanent selection for all users. Choices are Always display, Hide if empty, or Always hide. • User Customizable selections make the initial selection for all users but allow users to change that selection in Browse mode, using the Table Customization icon.

	<hr/> <p>Note</p> <p>If a new column is added after users have saved customizations, those users will not see the new column unless they return to the Table Customization dialog and click Save and Apply again. This action is required to add the new column to their other saved preferences and make the new column visible in the report.</p> <hr/> <p>Also see the note in the next field.</p>
Display Condition	<ul style="list-style-type: none"> • Show displays the column. • Hide hides the column. • condition lets you specify a regular expression to match against the node name. <ul style="list-style-type: none"> ▪ If condition is TRUE, the column shows. ▪ If condition is FALSE, the column is hidden. <hr/> <p>Note</p> <p>This hide condition overrides any choices made in the Display Customization field above.</p> <hr/>
Decorations	<p>Change the look of the text in the column.</p> <ul style="list-style-type: none"> • Text Mode <ul style="list-style-type: none"> ▪ Default Style has no interactive action. ▪ Link Style adds an underline when the cursor hovers over the text, indicating a link. To add the link target, see the Table Interactivity tab. • Text Color applies colors to the column text. Choose a predefined Default or Severity color, or choose Custom Color and then define the color. • Text Size sets the font size for column values. Choose S for small, M for medium, L for large, or XL for Extra Large. • Text Style applies font styles to the column values. Choose B for bold style, I for italics style, and U for underlined style.
Column Preview	Shows how a cell in the column will appear, based on the currently chosen decorations.

Property

A property column displays information 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.
Description	The description that appears when the user hovers the cursor over the column header. You can edit the description. See the tooltip for information about supported HTML tags, attributes, and classes.
Property	Type the name of a property or click the Property Selection Helper icon and select one or more properties.

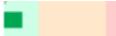
<p>Selected Value</p>	<p>Specifies the source of the value to retrieve. Choose one of the following:</p> <ul style="list-style-type: none"> • use filter below allows the configuration of a specific filter using the Filter Wizard. The wizard selects the value to display from the child. • use default formula result displays the result from the default formula configured on the child node. Select this option only when a default formula is configured. To check or set a default formula, click the child node in the navigation tree and go to the Formula tab on the child. When multiple formulas are defined, one of them can be identified as the default formula in the formula's Results Returned section. • use formula result displays the value retrieved from the result of a formula applied on the children. Select the formula from the drop-down list.
<p>Filter to Apply</p>	<p>Refines the properties to display in this column. This filter applies to each line of the table, which represents each child node.</p>
<p><input type="checkbox"/> Advanced Properties</p>	
<p>Display Customization</p>	<p>Selects whether the column is displayed in Browse mode.</p> <ul style="list-style-type: none"> • Locked Mode selections make a permanent selection for all users. • User Customizable selections make the initial selection for all users but allow users to change that selection in Browse mode, using the Table Customization icon.
<p>Display Condition</p>	<ul style="list-style-type: none"> • 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.
<p>Decorations</p>	<ul style="list-style-type: none"> • Text Mode <ul style="list-style-type: none"> ▪ Default Style has no interactive action. ▪ Link Style adds an underline when the cursor hovers over the text, indicating a link. To add the link target, see the Table Interactivity tab. • Text Color applies colors to the column text. Choose a predefined Default or Severity color, or choose Custom Color and then define the color. • Text Size sets the font size for column values. Choose S for small, M for medium, L for large, or XL for Extra Large. • Text Style applies font styles to the column values. Choose B for bold style, I for italics style, and U for underlined style.
<p>Column Preview</p>	<p>Shows how a cell in the column will appear, based on the currently chosen decorations.</p>

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	<p>Specifies the source of the value to retrieve. Choose one of the following:</p> <ul style="list-style-type: none"> • use filter below allows the configuration of a specific filter using the Filter Wizard. The wizard selects the value to display from the child. • use default formula result displays the result from the default formula configured on the child node. Select this option only when a default formula is configured. To check or set a default formula, click the child node in the navigation tree and go to the Formula tab on the child. When multiple formulas are defined, one of them can be identified as the default formula in the formula's Results Returned section. • use formula result displays the value retrieved from the result of a formula applied on the children. Select the formula from the drop-down list.
Filter to Apply	Appears only when filter is selected in the previous field. Refines the values to display in this column. This filter applies to each row in the table. A row represents a child node.
 Time Management	
Sampling Period	<p>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.</p> <p>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.</p> <hr/> <p>Note</p> <p>When you select a small sampling period, such as real-time, allow for the selection of higher periods for better performance.</p> <hr/>
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.

<p>Recover</p>	<p>Determines whether to retrieve all values over the report time range or only the last value.</p> <ul style="list-style-type: none"> • If all values are retrieved, then they are aggregated into a single value using the function selected from the Temporal Aggregation drop-down list. • If only the last value is retrieved, the time interval used for that retrieval changes based on the report's sampling period setting and the period of time set in the Time Threshold parameter. See the tooltip for examples.
<p>Temporal Aggregation or Time Threshold</p>	<p>Related to the Recover parameter.</p>
<p><input type="checkbox"/> Value Settings</p>	
<p>Scaling Mode</p>	<p>If the unit of the values does not correspond to what you want to display, this field provides a way to convert the values.</p> <ul style="list-style-type: none"> • none does not convert values. • multiply multiplies all values by the decimal value you supply in the by field that appears when you make this selection. • divide divides all values by the decimal value you supply in the by field that appears when you make this selection. • by unit converts the existing unit values into new unit values without requiring you to provide the formula. For example, using selections from the drop-down lists that appear, you can convert Celsius to Fahrenheit, or Packets per second into KPkts per second. Explore the drop-down lists to see all of the supported conversions.
<p>Value Display</p>	<p>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, only critical values, and so on. This parameter lets you to generate exception reports that present only relevant data.</p>
<p>Thresholds definition</p>	<p>Defines threshold values for this column.</p>
<p>Value Formatters</p>	<p>Controls the presentation of the metric data in each cell. Click +Formatter to view and select a formatter. After adding a formatter, click it to expose its configuration options. The available formatters and their configuration options are listed here. Many depend on threshold settings.</p>
<p>Number</p> <p>4.234</p> <p>\$4.23</p> <p>4.23£</p> <p>4%</p>	<p>Displays the value in numeric form.</p> <ul style="list-style-type: none"> • Decimal Rounding — Number of decimal places to display; default is 2. • Value Prefix — string to display before the metric, such as a currency symbol. • Value Suffix — string to display after the metric, such as a currency symbol, % symbol.

	<ul style="list-style-type: none"> • Apply threshold severity color? — Select to display the value using the color of the matching threshold category.
<p>Value to string</p> <p>200 --> OK</p>	<p>Converts a value to a specified string.</p> <ul style="list-style-type: none"> • Decimal Rounding — Number of decimal places to display; default is 2. • No-match fallback string — Type the string to use if a value does not match any category added under + Add value-strings • Apply threshold severity color? — Select to display the value using the color of the matching threshold category. • + Add value-strings — Click to supply the value-to-string conversions you want.
<p>Date and time</p> <p>May 1, 2000 1:01:52 AM EDT</p>	<p>Displays the date and time when the value was collected, converted from the UNIX timestamp.</p> <ul style="list-style-type: none"> • Format — Select a date/time format from the list. • Apply threshold severity color? — Select to display the value using the color of the matching threshold category.
<p>Duration</p> <p>6d 23h 55m 0s</p>	<p>Formats a number of seconds into a duration phrase consisting of larger units (combinations of days, hours, minutes, and seconds).</p> <ul style="list-style-type: none"> • Format — Select a duration format from the list. • Apply threshold severity color? — Select to display the value using the color of the matching threshold category.
<p>Bar inline chart</p> 	<p>Displays a horizontal bar graph whose number of bars depends on the selected sampling period and report time range.</p> <ul style="list-style-type: none"> • Bar width — Specify the size of each bar in pixels. • Bar spacing — Specify the spacing between bars in pixels.
<p>Performance inline chart</p> 	<p>Displays a vertical 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. There are no configuration options.</p>
<p>Time value inline chart</p> 	<p>Displays a line graph that plots the metric values over the report's time period. There are no configuration options.</p>
<p>Status inline chart</p> 	<p>Displays a horizontal status bar showing the status value over the time period. The line shows the status of each point of a value. The status colors of the band are set according to the threshold severities set on the column.</p> <ul style="list-style-type: none"> • Bar width — Specify the size of each bar in pixels. • Bar spacing — Specify the spacing between bars in pixels.
<p>Status icon</p> 	<p>Displays the status icon for the value according to the column's defined thresholds.</p>

	There are no configuration options. The status icons are fixed in the system.
Decorations	<ul style="list-style-type: none"> • Cell background applies the threshold severity color to the cell. Choose one of the following: <ul style="list-style-type: none"> ▪ none—no color applied ▪ lighter color—applies a lighter shade of the appropriate threshold color ▪ normal color—applies the threshold color • Row background applies the threshold severity color to the entire row. The same options are available as for cell background, described above. • Text Mode <ul style="list-style-type: none"> ▪ Default Style has no interactive action. ▪ Link Style adds an underline when the cursor hovers over the text, indicating a link. To add the link target, see the Table Interactivity tab. • Text Color applies colors to the column text. Choose a predefined Default or Severity color, or choose Custom Color and then define the color. • Text Size sets the font size for column values. Choose S for small, M for medium, L for large, or XL for Extra Large. • Text Style applies font styles to the column values. Choose B for bold style, I for italics style, and U for underlined style. • Graph tooltip shows a contextual graph above the cell when the user clicks the cell.
Column Preview	Shows how a cell in the column will appear, based on the currently chosen decorations, formatters, and thresholds.
<input type="checkbox"/> Advanced Properties	
Display Customization	<p>Selects whether the column is displayed in Browse mode.</p> <ul style="list-style-type: none"> • Locked Mode selections make a permanent selection for all users. • User Customizable selections make the initial selection for all users but allow users to change that selection in Browse mode, using the Table Customization icon.
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 - 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: <ul style="list-style-type: none"> <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 <code>v > 10.00 or v < 5.00</code>: the value is greater than 10.00 or less than 5.00 <code>v != 0.00 and not empty(v)</code>: the value is not equal to 0.00 and not empty <code>v == 0.00 && not empty(v)</code>: the value is equal to 0.00 and not empty
<input type="checkbox"/> Summary Settings	
Label	Label for the summary cell for this column. It will be put in front of the aggregated value. If left blank, the cell will only contain the value. If all the summary cells have the same label, they will all be merged into one.
Generate column	Choose the aggregate you wish to apply to the summary column, either sum or average .

Aggregate

This column type displays an aggregated value obtained from other selected columns. The **Aggregate** parameter defines the type of aggregation applied to the selected columns.

Column Name	A unique name that appears in the column header.
Description	The description that appears when the user hovers the cursor over the column header. You can edit the description. See the tooltip for information about supported HTML tags, attributes, and classes.
Aggregate	Select one of the following: <ul style="list-style-type: none"> Sum—In the current row, adds the values in all of the selected columns and displays the total. Min—In the current row, compares the values in all of the selected columns and displays the lowest value. Max—In the current row, compares the values in all of the selected columns and displays the highest one. Average—In the current row, computes the average from all of the selected columns and displays the result.
Aggregated Columns	Choose the columns to include in the aggregated value.

<input type="checkbox"/> Value Settings:	
Scaling Mode	<p>If the unit of the values does not correspond to what you want to display, this field provides a way to convert the values.</p> <ul style="list-style-type: none"> • none does not convert values. • multiply multiplies all values by the decimal value you supply in the by field that appears when you make this selection. • divide divides all values by the decimal value you supply in the by field that appears when you make this selection. • by unit converts the existing unit values into new unit values without requiring you to provide the formula. For example, using selections from the drop-down lists that appear, you can convert Celsius to Fahrenheit, or Packets per second into KPkts per second. Explore the drop-down lists to see all of the supported conversions.
Value Display	<p>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, only critical values, and so on. This parameter lets you to generate exception reports that present only relevant data.</p>
Thresholds definition	<p>Defines threshold values for this column.</p>
Aggregated Formatters	<p>Controls the presentation of the data in the cell. Click +Formatter to view formatter names and select one. After adding a formatter, click it to expose its configuration options.</p> <p>The following aggregate formatters are available. Choose either or both.</p>
<p>Number</p> <p>4.234</p> <p>\$4.23</p> <p>4.23£</p> <p>4%</p>	<p>Displays the value in numeric form.</p> <ul style="list-style-type: none"> • Decimal Rounding — Number of decimal places to display; default is 2. • Value Prefix — string to display before the metric, such as a currency symbol. • Value Suffix — string to display after the metric, such as a currency symbol, % symbol. • Apply threshold severity color? — Select to display the value using the color of the matching threshold category.
<p>pie chart</p> 	<p>Displays a small pie chart showing the proportion of values falling into the various defined threshold levels. The colors reflect the threshold levels.</p>
Decorations	<p>Change the look of the background and the text in the column.</p> <ul style="list-style-type: none"> • Cell background applies the threshold severity color to the cell. Choose one of the following: <ul style="list-style-type: none"> ▪ none—no color applied ▪ lighter color—applies a lighter shade of the appropriate threshold color

	<ul style="list-style-type: none"> ▪ normal color—applies the threshold color • Row background applies the threshold severity color to the entire row. The same options are available as for cell background, described above. • Text Mode <ul style="list-style-type: none"> ▪ Default Style has no interactive action. ▪ Link Style adds an underline when the cursor hovers over the text, indicating a link. To add the link target, see the Table Interactivity tab. • Text Color applies colors to the column text. Choose a predefined Default or Severity color, or choose Custom Color and then define the color. • Text Size sets the font size for column values. Choose S for small, M for medium, L for large, or XL for Extra Large. • Text Style applies font styles to the column values. Choose B for bold style, I for italics style, and U for underlined style. • Graph tooltip adds or removes a contextual graph that appears like a tooltip when the user clicks a cell.
Column Preview	Shows how a cell in the column will appear, based on the currently chosen decorations.
<input type="checkbox"/> Advanced Properties	
Display Customization	<p>Selects whether the column is displayed in Browse mode.</p> <ul style="list-style-type: none"> • Locked Mode selections make a permanent selection for all users. • User Customizable selections make the initial selection for all users but allow users to change that selection in Browse mode, using the Table Customization icon.
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- 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?	<p>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.</p> <p>In the following examples, rows are only shown when:</p> <ul style="list-style-type: none"> • <code>empty (\$v)</code>: the value is empty

	<ul style="list-style-type: none"> • <code>(\$v >= 0 or \$v < 0) and not empty(\$v)</code> : the value is (greater or equal to 0.00 or less than 0.00) and not empty • <code>\$v > 10 or \$v < 5</code>: the value is greater than 10 or less than 5 • <code>\$v != 0 and not empty (\$v)</code>: the value is not equal to 0 and not empty • <code>\$v == 0 && not empty(\$v)</code>: the value is equal to 0 and not empty
<input type="checkbox"/> Summary Settings	
Label	Label for the summary cell for this column. It will be put in front of the aggregated value. If left blank, the cell will only contain the value. If all the summary cells have the same label, they will all be merged into one.
Generate column	Choose the aggregate you wish to apply to the summary column, either sum or average .

Combined

This column type renders the data from other selected columns in a single cell.

Column Name	A unique name that appears in the column header.
Description	The description that appears when the user hovers the cursor over the column header. You can edit the description. See the tooltip for information about supported HTML tags, attributes, and classes.
Combined Formatters	<p>Specifies the columns to copy.</p> <ol style="list-style-type: none"> 1. Click +Formatter and select Copy Column Content. 2. Click the formatter and select a column name from the drop-down list. 3. Repeat the previous steps one or more times to add the second and additional columns to the combined column. 4. Drag the formatters to rearrange their order of appearance. <p>The Copy Column Content formatter is the only formatter available in a combined column. You can format the data using formatters in the original columns.</p>
Column Preview	Shows how a cell in the column will appear, based on the currently chosen columns.

Display Options

The **Display Options** section applies to the entire report, rather than to a specific column.

Paging mode	<p>Indicates the number of lines to display on one page of a table.</p> <ul style="list-style-type: none"> • 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.
Refresh interval (secs)	Controls the rate of automatic refreshes. Leave blank or enter 0 to turn off auto refresh.

Item reports (hero dashboards)

An item report is a type of table report. Each row is displayed as a self-contained item, with the column information mapped to a position in the item.

A series of item reports can create an easy-to-read dashboard, sometimes called a hero dashboard, that makes important information easily visible.

Here are example item reports on a hero dashboard:

Storage Groups not Remotely Replicated

42 Out of 942 Storage Groups

LUNs with Service Level Violations

Critical Marginal

Host File Systems with < 30 days till full

19 Out of 599 Host File Systems

Changing a table report to an item report

You can start with a table report and redefine it as an item report, mapping selected columns to the item.

For example, here is a simple table report from the SolutionPack for EMC M&R Health.

All >> Content Library >> Dell EMC M&R Health >> Misc. Reports >> Daily Dashboard >> File Systems

Daily Dashboard / File Systems

November 2016, Monday 28 > Tuesday 29, 3:34 PM EST | *Last 1 Day*

2 elements found, displaying all elements.

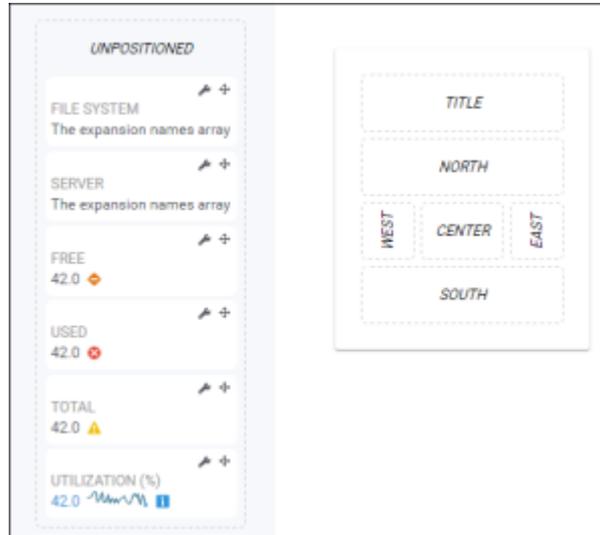
Server	File System	Total	Free	Used	Utilization (%)
ljsa128.lss.emc.com	/boot	159.8 MB	128.2 MB	31.6 MB	19.8
ljsa128.lss.emc.com	/	119.8 GB	111.8 GB	8.0 GB	6.7

The following procedure changes this report to an item report.

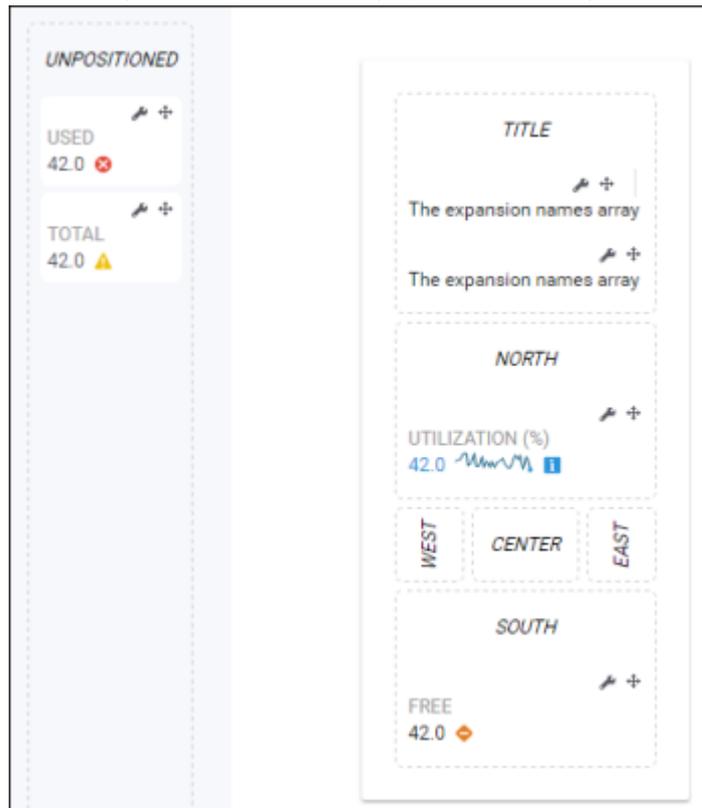
Procedure

1. Go to **Report Library > EMC M&R Health > Miscellaneous Reports > Daily Dashboard**.
2. Scroll down and click the **File Systems** report title.
This isolates the File Systems report to a page by itself.

3. Click **Modifications > Edit Reports**.
4. Click the **Report Configuration** tab.
5. For **Report Type**, in the **Tables** category, select **Item**.
6. Click the **Report Details: Item** tab.
7. Expand **Columns Positioning**.
8. For **Container Display Mode**, select **Cards**.
9. For **Item Alignment**, select **left**.
10. Drag columns from the **Unpositioned** section onto the design map at the right.
For example, here is the starting point, with all columns unpositioned.

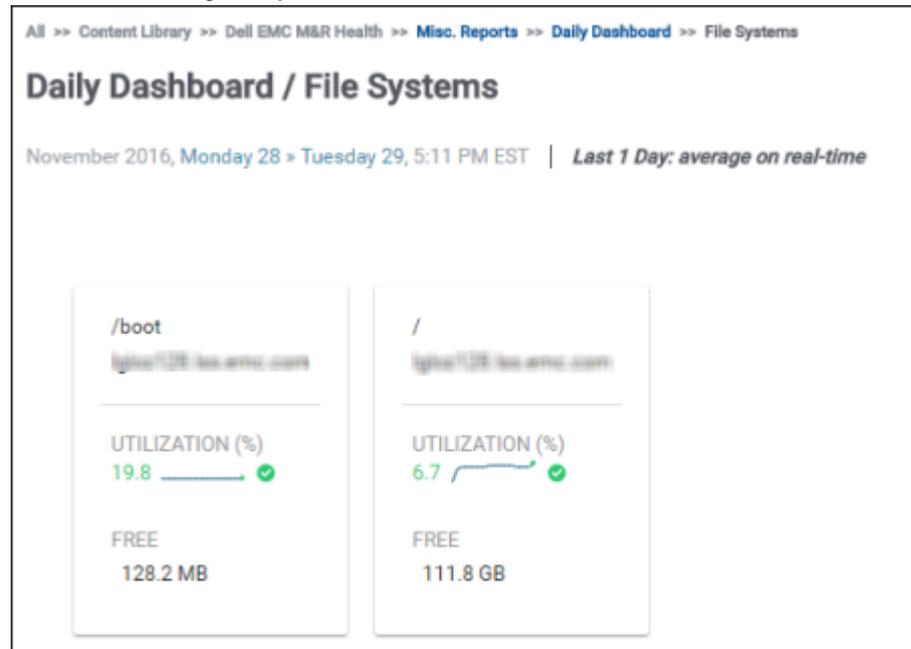


Here is the designed item, with 4 of the 6 columns mapped. The columns that remain unpositioned will not display in the item report.



11. Click **Save**.
12. Click **BROWSE MODE**.

Here is the changed report.



Item report parameters

You define information to include in an item report and design the look of the report on the **Report Details: Item** tab.

An item report is a type of table report. To specify the information to include in the report, you define columns, the same as for other types of table reports. A column is defined based on an Attribute, a Property, a Value, an Aggregate of other columns, or a Combined cell of other columns. For more information about adding and configuring columns, see [Table report parameters](#) on page 70.

The following parameters control the look of the item report, after you have defined the columns.

<p>Container Display Mode</p>	<p>Defines the overall style of each item (each converted row).</p> <ul style="list-style-type: none"> • Basic—Borders are not used. Column labels (subtitles) are not shown. • Cards—Each item is enclosed in a border. Column names appear as subtitles for each included column of information.
<p>Item Alignment</p>	<p>Defines whether the item report (all items in the report) is aligned to the left margin, centered, or right margin of the page.</p>
<p>Design Map</p>	<p>Defines which columns to include in the item report, and where to position each column.</p> <p>Drag a column from the unpositioned section to the desired location in the design map. Sections in the design map can hold multiple columns.</p> <p>Click the Tools icon associated with a column to access the column formatting parameters.</p>

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.

Setting threshold definitions in graphs

Set thresholds to display a dotted line in the graph that will let you see when a metric has exceeded the threshold.

Procedure

1. Click the **Report Details: Graph** tab.
2. Under **Thresholds Definition**, select the name of the threshold, enter a value, and select a color for each threshold you want to define.
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.

Note

Summary statistics do not exclude values of defined maintenance periods from the computation.

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.

Report types include horizontal bars, single horizontal bar, time ratio bars, and pie chart.

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.

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 **Value display**, select **absolute** to show the actual metric value, or **relative** to show metrics as percentages of the total along with the absolute values.
5. In **Number of sections**, enter the maximum number of sections to display.
6. In **Others section color**, select the color for the Other section.
7. Click **Expand Others section** to expand the Other section into subsections.
8. 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 Graph** tab.

Scaling Mode

If the unit of the values does not correspond to what you want to display, this field provides a way to convert the values.

- **none** does not convert values.
- **multiply** multiplies all values by the decimal value you supply in the **by** field that appears when you make this selection.
- **divide** divides all values by the decimal value you supply in the **by** field that appears when you make this selection.
- **by unit** converts the existing unit values into new unit values without requiring you to provide the formula. For example, using selections from the drop-down lists that appear, you can convert Celsius to Fahrenheit, or Packets per second into KPkts per second. Explore the drop-down lists to see all of the supported conversions.

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 above-selected metric.

Properties in Legend

These properties display in the graph's legend to identify the variables. To select the properties to appear in the legend, click **Edit Properties**. To reorder their appearance in the legend, hover the cursor over a row and drag it to the desired position.

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. Select **alphabetical** to order the bars using the metric names. Select **descending** to order the bars according to the metric values.

Section Grouping Mode

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

Value display

Determines whether metric values are shown as actual metric values (absolute) or as percentages of the total in addition to the absolute values.

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

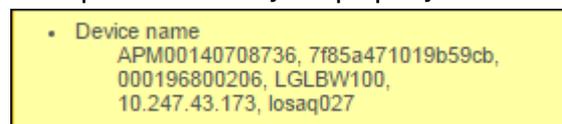
Sets the color of the Others section if defined.

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.

Tooltip Information

Choose properties to add to the tooltips that appear when a user hovers the cursor over a location in the report. The additional information is added in the form of a description followed by the property values. For example:



To select properties, click **Add Property**.

- Type the database property name, such as `device`, or click the property selector icon for help.
- Choose the description to appear on the first line.
 - **default** uses the system-provided description for the property. For example, `Device name` is supplied for the device property.
 - **none** does not use a description.
 - **custom** lets you supply a phrase.

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 marker 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 marker. The size of the marker is based on a given metric value, such as the number of devices at the site.
- **Color** sets the color of the marker. The color relies on the threshold definition where green represents no thresholds, yellow represents a major threshold, and red represents a critical threshold.
- **Marker** sets the shape of the marker, such as circle, square, triangle, inverted triangle, pentagon, and sector.

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." The following example overrides the location of device1 with latitude and longitude values: `device1=42.2251051 -71.5315623`

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 **Marker** tab.
3. Enter the minimum and maximum values for marker size.
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.
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.

Marker	<p>Defines the icon or area marker for each point. Current choices for icons are: circle, square, triangle, inverted triangle, and pentagon.</p> <ul style="list-style-type: none"> • 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. • 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. <p>The current choice for the area marker is sector.</p> <ul style="list-style-type: none"> • Azimuth - Indicates the orientation of the area marker (N, S, E, or W). This value is mandatory. The permissible values are 0-359 degrees. There is no default value. • Radius - Radius value for the area marker in meters. This is an integer value and it is optional. The default value of 500 meters is used if no column is configured for the setting. • Beam width - Width of the area marker in degrees. This value is optional. The permissible values are 0-359 degrees. The default value of 120 degrees is used if no column is configured for the setting. <p>Sector markers become larger when zooming in and smaller when zooming out, but icon size does not change with zoom level.</p> <p>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.</p>
---------------	---

Displayed Information

For information on these settings, see [Table report parameters](#) on page 70.

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.

Axes Configuration

Select one of the following expansion types for each axis.

- Select **use report expansion** to expand the axis using the parameters specified for the report on the Report Configuration tab. The ordering on the axis is defined by the **Order** field, below.
- Select **use time expansion** to build a heat map that displays the evolution of variables based on a unit of time (e.g. a day, a week, a month, and so on.) If you select this value, the **Sampling Period** and **Over** fields appear. For example, you can expand an axis to show a Sampling Period of every day over some time period, such as a week or a year.

If the report duration is 1 year and the X axis time range is for each 1 day over 1 week, each heatmap value is an aggregate (such as an average, minimum, or maximum) of all values of each time period included in the report duration.

Scaling Mode

If the unit of the values does not correspond to what you want to display, this field provides a way to convert the values. Select one of the following:

- **none** does not convert values.
- **multiply** multiplies all values by the decimal value you supply in the **by** field that appears when you make this selection.
- **divide** divides all values by the decimal value you supply in the **by** field that appears when you make this selection.
- **by unit** converts the existing unit values into new unit values without requiring you to provide the formula. For example, using selections from the drop-down lists that appear, you can convert Celsius to Fahrenheit, or Packets per second into KPkts per second (or MPkts, GPkts, or TPkts per second). Explore the drop-down lists to see all of the supported conversions.

Order

Defines how to order rows and columns in the heat map when the axes configurations are **use report expansion**. See the tooltip next to the field for explanations of each option.

Show Grid

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

Use Gradient Color

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

Thresholds Definition

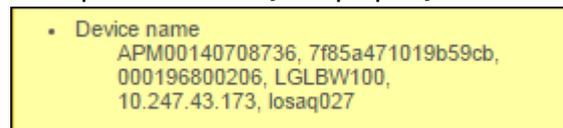
Thresholds do not appear in the report unless you give them values. You can add as many thresholds as necessary, giving each a different value and color. If there is a second y-axis in the graph, the thresholds you set here are applicable only to the main y-axis.

Properties in Legend

These properties display in the graph's legend to identify the variables. To select the properties to appear in the legend, click **Edit Properties**. To reorder their appearance in the legend, hover the cursor over a row and drag it to the desired position.

Tooltip Information

Choose properties to add to the tooltips that appear when a user hovers the cursor over a location in the report. The additional information is added in the form of a description followed by the property values. For example:



To select properties, click **Add Property**.

- Type the database property name, such as `device`, or click the property selector icon for help.
- Choose the description to appear on the first line.
 - **default** uses the system-provided description for the property. For example, `Device name` is supplied for the device property.
 - **none** does not use a description.
 - **custom** lets you supply a phrase.

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: <ul style="list-style-type: none"> • 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. <hr/> 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. <hr/>
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.

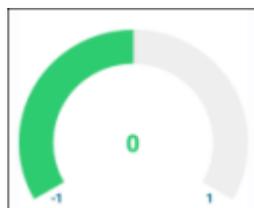
Column Name	A unique name that appears in the table header.
Selected Value	Retrieves a value from the following: <ul style="list-style-type: none"> • filter allows the configuration of a specific filter through the Filter Wizard, which selects the metrics to display from the child. • formula result displays 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	<p>The time interval relative to the report time value, which acts like a tolerance interval for data retrieval.</p> <p>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.</p> <hr/> <p>Note</p> <p>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.</p> <hr/>
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.

Gauges

A gauge shows one value in a meter. The report definition defines the value being measured, the upper and lower bounds of the gauge, and optional thresholds for identifying severity.

In the following gauge, the boundaries are defined as -1 and 1. The current measurement is 0.



Thresholds defined on the measured value can change the color of the gauge. For example, a green gauge indicates that the current value falls within the defined range for Normal, whereas a red gauge indicates a Critical value, according to the defined thresholds.

Gauge report parameters

The **Report Details: Gauge** tab defines the metric to measure, its lower and upper bounds, and threshold ranges.

Value

This section defines the metric value that the gauge is reporting.

Column Name	A unique, descriptive name. This name does not appear in the report.
Selected Value	<p>Specifies the source of the value to retrieve. Choose one of the following:</p> <ul style="list-style-type: none"> • use filter below allows the configuration of a specific filter using the Filter Wizard. The wizard selects the value to display from the child. • use default formula result displays the result from the default formula configured on the child node. Select this option only when a default formula is configured. To check or set a default formula, click the child node in the navigation tree and go to the Formula tab on the child. When multiple formulas are defined, one of them can be identified as the default formula in the formula's Results Returned section. • use formula result displays the value retrieved from the result of a formula applied on the children. Select the formula from the drop-down list.
Filter to Apply or Use time settings from	<p>Filter to Apply appears when use filter below is selected in the previous field. The filter defines the value to display in the gauge. This filter applies to each row in the table. A row represents a child node.</p> <p>Use time settings from defines which time settings to use for the metric retrieval.</p>
<input checked="" type="checkbox"/> Time Management (appears only when Selected Value is a filter)	
Sampling Period	<p>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.</p> <p>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.</p>

	<p>Note</p> <p>When you select a small sampling period, such as real-time, allow for the selection of higher periods for better performance.</p>
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...	<p>Determines if all values are retrieved over the report time range, or only the last value.</p> <ul style="list-style-type: none"> • If all values are retrieved, then they are aggregated into a single value using the function selected from the Temporal Aggregation drop-down list. • If only the last value is retrieved, the time interval used for that retrieval changes based on the report's sampling period setting and the period of time set in the Time Threshold parameter. See the tooltip for examples.
Temporal Aggregation or Time Threshold	Specify settings related to the Recover parameter, above.
<input type="checkbox"/> Value Settings	
Scaling Mode	<p>If the unit of the values does not correspond to what you want to display, this field provides a way to convert the values.</p> <ul style="list-style-type: none"> • none does not convert values. • multiply multiplies all values by the decimal value you supply in the by field that appears when you make this selection. • divide divides all values by the decimal value you supply in the by field that appears when you make this selection. • by unit converts the existing unit values into new unit values without requiring you to provide the formula. For example, using selections from the drop-down lists that appear, you can convert Celsius to Fahrenheit, or Packets per second into KPkts per second. Explore the drop-down lists to see all of the supported conversions.
Thresholds definition	Defines the thresholds that determine the color of the gauge display.

Lower Bound and Upper Bound

This section defines the values for the lower and upper bounds of the gauge. The parameter names and choices are the same for both lower and upper bounds.

Use value from	<p>Defines how to obtain the boundary value.</p> <ul style="list-style-type: none"> Choose static value to specify a fixed value for the bound. Choose one of the other selections to define a value in the same way that you would define values for table columns. You can choose to define an attribute, property, or value column.
additional parameters	<p>The remaining parameters change depending on the selection above.</p> <ul style="list-style-type: none"> For static value, specify the fixed value to appear for the bound. For the other selections, see the descriptions for attribute, property, and value columns in Table report parameters on page 70.

Legend

Properties in Legend	Click Edit Properties to configure the properties to show in the legend.
Legend Visibility	Choose whether to show or hide the legend.

Editing external reports

An external report is a web page displayed in the report pane.

The **Report Details: External** tab contains the URL of the page to display.

Other report definition tabs, such as **Filtering & Expansion** and **Report Configuration** define the report.

Procedure

1. Click the **Report Details: External** tab.
2. Edit the URL in the **External report address** field.
3. Edit other tabs if needed.
4. Click **Save**.

External report address

The **external report address** field on the **Report Details: External** tab contains the URL of the page to display in the report.

With an external report, you can access and display data from external software and display the retrieved data in the report pane. The request for data is included in the URL.

The URL can include the following parameters.

Parameter	Example
<p><code>#{property}</code></p> <p>where:</p> <p><i>property</i> is an APG property name from the APG databases.</p>	<p>To filter results from an external database by device, use <code>#{device}</code> in the URL</p>

Parameter	Example
<p>@{attribute}</p> <p>where:</p> <p>attribute is an attribute of the report nodes.</p>	<p>If the Filtering & Expansion tab defines device filters, you might want to display the filter values in the report. To do so, use @{nodeFilter} in the URL.</p>

Example URLs

To query Bing:

```
http://www.bing.com/search?q=${parttype}
```

To query an external defect database and list entries associated with a specific device:

```
http://jira.yourcompany.com:8443/issues/?jql=project%3DEMP AND text~"${device}"
```

To open the vsphere client software:

```
https://${device}/vsphere-client/?csp
```

@{attribute} values

Use @{attribute} to display an attribute of the report node in the report. These attribute parameters represent the same parameters that you can select for an attribute column in a table report.

Note

These attributes are the same attributes that you can select in a table report for an attribute column.

@{attribute}	Description
@{singleNodeid}	Displays the node short ID relative to its parent.
@{id}	Displays the complete node ID starting from the root of the tree.
@{name}	Displays the name of the node.
@{fullName}	Displays the complete name of the node, starting from the root of the tree.
@{expandedNames[0]}	Is only applicable if a multi-expansion is applied to the node that appears in the table. The index references a zero-based array that contains each part of the multi-expansion.
@{nodeFilter}	Displays the node short filter, relative to its parent.
@{filter}	Displays the complete node filter, starting from the root of the tree.
@{childCount}	Displays the number of children of that node.

@{ <i>attribute</i> }	Description
@{ <i>other</i> } where: <i>other</i> is any other node attribute. For example, @{depth}.	Displays any other attribute of the node. Useful for debugging purposes.

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.

Layers

Layers and map types are collections of nodes and connections between nodes that present a specific view, such as a physical connectivity view or a logical connectivity view. The difference is that you can select multiple layers to be simultaneously displayed for a report, but you can display only one map type for a report at any time.

Display layers through the Layers panel and display a map type with the **Map Type** drop-down menu. This menu is only visible in topology maps that have been defined to use layers.

Overlays

An overlay is a defined subset of visual indicators, such as bullets, images, spark lines, and color, that provide additional information about nodes and connections in a network. For example, overlays might be defined to inform users about availability, performance, and health. A user can manually display or hide an overlay by selecting or unselecting it in the **Layers and Overlays** menu on the topology map.

Hop Count

The hop count is the number of devices that data must pass through from a source to a destination. Topology maps can be defined to display the **Increase Hop Count**, **Decrease Hop Count**, and **Reset Hop Count** icons that enable the number of elements being displayed to increase or decrease according to the number hops selected. The hop icons are not displayed in topology maps that have not been defined to support the feature.

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 <code>devtype=="Router" & device=="R1"</code> ; 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 <code>devtype=="Switch"</code> , 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 WebApps-Resources module. For example, <code><APG>/Custom/WebApps-Resources/Default/topomap</code> .
Topology Service	Uses the Topology service to generate the nodes and relationships.

Topology

Select a topology service that is configured for the frontend application. These are set in the web application's dedicated XML configuration (for example, in `/opt/APG/Web-Servers/Tomcat/Default/conf/Catalina/localhost/APG.xml` on UNIX or in `C:\Program Files\APG\Web-Servers\Tomcat\Default\conf\Catalina\localhost\APG.xml` on Windows.)

XML data retrieval

This defines how the XML describing the report will be retrieved. In each case (file path or URL), dynamic values can be injected at runtime through the use of placeholders. To inject a property, use `${property}`. For example: `http://www.company.com/query?mydevice=${device}`. For topology map services that support it, "hop-count" and "map-type" can also be specified with their default values, using `@{hop-count[:default-value]}` or `@{map-type[:default-value]}`. For example:

<http://www.company.com/query?hopcount=@{hop-count:1}&maplayer=@{map-type}>.

Layout

Use this setting to determine the layout of the report.

Organic	Positions the nodes based on a physics simulation of the interacting forces.
Balloon	Positions the child nodes radially around their parents.
Circular	Positions the nodes in a circle.
Hierarchical	Positions the nodes in a hierarchical structure.
Tree	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`.

Overlays

Name

Assign a name for the overlay. The name will appear on a menu of layers and overlays that is displayed on each Topology report.

Display

The value `visible` causes the overlay to be displayed by default on the Topology report. The value `hidden` causes the overlay to be hidden by default.

Templates

Use the pencil icon to display a list of templates. Select the template in which overlays are defined.

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

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.

Tooltip Information

Tooltips appear when the user hovers the cursor over a status icon. A tooltip displays one or more property values and, optionally, property names. You specify the properties that appear.

Click **Add** to add a new property to the tooltips. Two fields appear.

- In the first field, type a database property name or click the **property selection helper** and select a property from the list.
- In the second field, specify the label to appear with the value in the tooltip.

Default	Uses the system default name, as shown in the property selection helper .
None	Does not use a label.
Custom	Lets you supply the label.

Using threshold definitions

Setting a threshold enables you to visually see when a value has exceeded a threshold in a report.

Threshold definitions can be used in almost all report types. With thresholds, you can set a threshold value and assign it a color to represent that threshold. You can define as many thresholds as you need.

In regular graphs, thresholds appear as dotted lines at the defined values and with the corresponding colors, along with an explanation of each threshold in the legend so you can gauge values that have crossed your defined thresholds. In graphs where there is a second y-axis, the thresholds only apply to the main y-axis.

In tables, status tree map reports, and heat map reports, values that fall within a severity range appear in the selected color or display the related severity icon.

Thresholds define the value of the lower boundary of a severity range. The higher boundary of the range is defined by the next highest threshold value. Inclusion of the boundary values in the range is based on severity of the range. The highest severity always wins the boundary value.

You can define ascending (higher value is more critical) severity, descending (lower value is more critical) severity, or a combination of both ascending and descending severities, as shown in the examples below.

Example 1: As a value increases, the severity increases

Suppose you want to set the following thresholds for a CPU utilization value:

- **Normal** for any value less than 20% utilization
- **Minor** for any value equal to or greater than 20% and less than 50% utilization
- **Major** for any value equal to or greater than 50% and less than 80% utilization
- **Critical** for any value equal to or greater than 80% utilization

The thresholds would be set as follows:

- **Normal:** -Infinity
- **Minor:** 20
- **Major:** 50
- **Critical:** 80

Example 2: As a value decreases, the severity increases

Suppose you want to set the following thresholds for a free space on disk value:

- **Critical** for any value equal to or less than 20 GB
- **Major** for any value greater than 20 GB and less than or equal to 30 GB
- **Minor** for any value greater than 30 GB and less than or equal to 80 GB
- **Normal** for any value equal to or greater than 80 GB

The thresholds would be set as follows:

- **Critical:** -Infinity
- **Major:** 20
- **Minor:** 30
- **Normal:** 80

Example 3: As a value increases or decreases too much, the severity increases

Suppose you want to set the following thresholds for a packet jitter value (which can be positive or negative):

- **Unknown** for any value less than -30 ms
- **Critical** for any value equal to or greater than -30 ms and less than -15 ms
- **Major** for any value equal to or greater than -15 ms and less than -10 ms
- **Minor** for any value equal to or greater than -10 ms and less than -5 ms
- **Normal** for any value equal to or greater than -5 ms and less than 5 ms
- **Minor** for any value equal to or greater than 5 ms and less than 10 ms
- **Major** for any value equal to or greater than 10 ms and less than 15 ms
- **Critical** for any value equal to or greater than 15 ms and less than 30 ms
- **Unknown** for any value greater than 30 ms

The thresholds would be set as follows:

- **Unknown:** -Infinity
- **Critical:** -30
- **Major:** -15
- **Minor:** -10
- **Normal:** -5

- **Minor:** 5
- **Major:** 10
- **Critical:** 15
- **Unknown:** 30

Example 4: Using a Custom severity

A standard severity has a fixed level. In increasing order (from lower to higher), the highest severity always gains the range border values:

- **Undefined**
- **Normal**
- **Informational**
- **Unknown**
- **Minor**
- **Major**
- **Critical**

For example, if you have the following threshold definitions:

- **Major:** -Infinity
- **Minor:** 0
- **Critical:** 10

Then the following ranges are defined:

- **Major:** -Infinity <= value <=0
- **Minor:** 0 < value < 10
- **Critical:** 10 <= value

The range that has the Minor severity does not own the values 0 and 10 because the Minor severity is less than Major (which owns 0) and Critical (which owns 10).

Unlike a standard severity, a Custom severity is not fixed; it is the same as the severity of the range defined before it (in terms of value range). For example, if you have the following thresholds set in a table:

- **Major:** -Infinity
- **Minor:** 0
- **Custom:** 5
- **Major:** 10

Then the following ranges are defined:

- **Major:** -Infinity <= value <= 0
- **Minor:** 0 < value < 5
- **Custom:** 5 <= value < 10
- **Major:** 10 <= value

The Custom range has the same severity level as the one set in the range defined before it (Minor). This is why the range using the Custom severity "owns" value 5 but not value 10 (which is owned by a range which has a higher severity: Major).

Mixing standard and Custom thresholds is not recommended.

If you use only Custom severities in all ranges, the behavior will be constant because the actual severities will all be the same. For example, suppose the following thresholds are defined:

- **Custom 1:** -Infinity
- **Custom 2:** 0
- **Custom 3:** 100

Then the following ranges are defined:

- **Custom 1:** $-\text{Infinity} \leq \text{value} < 0$
- **Custom 2:** $0 \leq \text{value} < 100$
- **Custom 3:** $100 \leq \text{value}$

If you define ranges with increasing severity:

- **Normal:** -Infinity
- **Minor:** 0
- **Major:** 100
- **Critical:** 1000

Then the following ranges are defined:

- **Normal:** $-\text{Infinity} \leq \text{value} < 0$
- **Minor:** $0 \leq \text{value} < 100$
- **Major:** $100 \leq \text{value} < 1000$
- **Critical:** $1000 \leq \text{value}$

If you define ranges with decreasing severity:

- **Critical:** -Infinity
- **Major:** 0
- **Minor:** 100
- **Normal:** 1000

Then the following ranges are defined:

- **Critical:** $-\text{Infinity} < \text{value} \leq 0$
- **Major:** $0 < \text{value} \leq 100$
- **Minor:** $100 < \text{value} \leq 1000$
- **Normal:** $1000 < \text{value}$

The highest severity always wins the border value.

Formula tab

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 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. <ul style="list-style-type: none"> • 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	<p>Enables you to combine parameter instances and types of Filter on this Node, Formula Result, Constant Value or Property Value.</p> <p>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.</p> <p>Some formulas treat combined parameters in other ways. Refer to the formula and parameter descriptions for more information.</p>
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.

Copying a formula

You can copy and paste formula definitions from one report definition to another report definition.

Procedure

1. Navigate to the report that contains the formula you want to copy.
2. Click the **Report Details: Formula** tab.
3. Check the formulas that you want to copy.
4. Click **Copy**.
5. Navigate to the report where you want to paste the formula.
6. Click the **Report Details: Formula** tab.
7. Click **Paste**.
Paste appears next to **Add Formula**.
8. Click **Save**.

UI Interactivity and Table Interactivity tabs

The interactivity parameters control the results when the user clicks on elements in the report.

Three levels of interactivity are provided.

Default behavior

Check or uncheck these parameters to switch them on or off. The underlying actions are preconfigured and not editable.

- **Default pass through** sends events performed on a child report to the mixed report.
- **Drilldown on row click** enables or disables a jump to a more detailed report when the user clicks a row in a table report.
- **Drilldown on title click** enables or disables the jump to the full page report when the user clicks the title of a report. For example, a mixed report or list report contains titles to other reports.
- **Action Menu** enables or disables the context popup menus that exist on some reports. For example, the alerting reports have a context menu for managing alerts. When there is no context menu, the installed default is unchecked.
- **Allow multiple row selection** enables or disables the appearance of the checkbox column as the first column in a table report.

Note

This field controls whether the multi-row selection column appears. Customization activity is required to associate actions with the checkbox column.

Custom behavior

Click buttons in this section to configure custom behavior.

Jump to report links an element in a report to another report. The Destination Report Path includes the UID, the node path of the report you are linking to, and optionally, expressions that filter the initial display of the report.

For an example report path, click the information icon next to the **Destination Report Path** field.

Custom triggers and adapters

Click **Switch to Advanced Mode** to expose additional customization features. You can create custom triggers for interactive actions and custom adapters to handle additional actions.

Enable multi-row selection and action menu

You can configure an existing action menu to respond to multiple row selections.

Before you begin

This procedure assumes that the action menu is already defined and associated to the report.

By default, a right-click action menu applies to a single row. For example, in any alerting report, right-click text in a row to view and use the action menu.

This procedure enables the action on multiple rows.

Procedure

1. Navigate to the table report that has an action menu associated with rows.
2. In **Edit Mode**, click **Table Interactivity**.
3. Select the following fields:
 - **Action Menu** (should already be enabled)
 - **Allow Multiple Row Selection**
4. Click **Save**.
5. Click **Browse Mode**.

Notice that the column of checkboxes is now visible.
6. Test whether the action menu operates on multiple rows:
 - a. Select the checkboxes on several rows.
 - b. Right-click the text in one of the selected rows.
 - If the action menu appears, the feature is operational. You can exit the menu and deselect rows.
 - If the message `No actions available` appears, continue with the next step.
7. On the command line for the Frontend server, navigate to the following folder:


```
<APG>/Custom/WebApps-Resources/Default/actions/
```
8. Open the XML file for the appropriate action menu. For example, the default alerting context menu is defined in `event-mgmt.xml`.
9. Add the `multiple` attribute to the `<script>` element as shown here:

```
<script file="ack-event" result="notification"
  timeout="10000" multiple="multiple-execution">
```

10. Save the file and retest.

Pre-Generation tab

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. This tab is not available for reports those are under **My 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>	<p>The code numbers for the display modes:</p> <ul style="list-style-type: none"> • 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>mode=<code></code>	<p>The codes for the report modes:</p> <ul style="list-style-type: none"> • 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>	<p>The period, in seconds, of the selected aggregate. Zero indicates real time. This period must exist in the database.</p>
<code>type=<code></code>	<p>Not applicable when the period is zero. The code numbers for the aggregate types are:</p> <ul style="list-style-type: none"> • 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>	<p>A list of indexes separated by underscores that restricts the displayed variables.</p>

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>	<p>The code to use as the default for the report time range description:</p> <ul style="list-style-type: none"> • 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>	<p>The code to use as the default for the report time range:</p> <ul style="list-style-type: none"> • 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 <p>You can use as many duration codes as needed for the report time range but you cannot duplicate a code.</p> <p>You can also use calendar in the duration code followed by <code>start</code> and <code>end</code> or <code>start_ts</code> and <code>end_ts</code>, as explained in the following rows.</p>
<code>start=<date>end=<date></code>	<p>The time range of the requested report. The <code><date></code> field has the YYYY-MM_DD HH:MM format.</p> <hr/> <p>Note</p> <p>URL encoding may transform spaces into plus signs and colons and periods into %3A.</p> <hr/>
<code>start_ts=<timestamp>end_ts=<timestamp></code>	<p>The time range using timestamps of the requested report. These are UNIX timestamps.</p>

Advanced time management

Syntax	Description
--------	-------------

<code>itz=<timezone code></code>	The time zone for the report, which uses standard zone names such as America/Montreal.
<code>tf=<time filter.expression></code>	The maintenance period for the report. The <code><time filter.expression></code> field consists of <code><dom>;<dow>;<hod></code> to represent the following: <ul style="list-style-type: none"> • <code>dom</code>=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. • <code>dow</code>=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 <code><base></code> 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.

Defining custom colors

Custom colors are supported for custom threshold definitions and for text decorations.

Custom colors are supported for the following definitions in Edit mode.

- The **custom** threshold type lets you define a custom color.
- Text decorations and backgrounds in table columns let you define custom colors.

In the above field definitions, when you click the color patch next to the field, the custom color map appears. In other situations where the color is fixed, clicking the color patch does not open the color map.

The colors associated with threshold severity levels are fixed as follows:

- Normal — Green
- Critical — Red
- Major — Orange
- Minor — Yellow
- Informational — Blue
- Unknown — Grey

Use the following procedure to use the color map to define a custom color.

Procedure

1. Click the color patch in a location where a custom color is supported.

The color map appears.



2. Slide the arrows on the color bar to choose the color family.
3. Drag the white circle to a location on the color map to fine-tune the color within the family.

You can also use the code designations to define the color if you know them.

4. Click **Apply**.
5. To reuse a defined custom color in another definition, copy the HEX color code in the # field and paste the value into the same field in the new definition.

Accessing reports from outside of the Console

You might want to bookmark a report, include the report or its URL in another document or email, or have a third-party application access and display the most current version of a report.

The following methods for accessing a report from outside of the Console are supported:

- Copy the URL of a displayed report in the Console, and save it for use elsewhere

- Use the Frontend Report Lookup tool to access a report using unique IDs and report paths
- Construct a URL using supported parameters

Getting the complete URL for the current 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 currently displayed report appears in a popup.

2. Copy the URL.

Frontend report lookup tool

The Frontend report lookup tool finds a report using report names and unique identifiers (UIDs), rather than the URL syntax.

You can use the lookup tool in a web browser URL field. A third-party application can use the lookup tool to request a report and display the report in its own context.

Syntax

```
http://server:port/APG/lookup/{uid}/report_name/report_name/.../*
```

where:

server

Is the Frontend server name.

port

Is the Frontend port access; default is 58080.

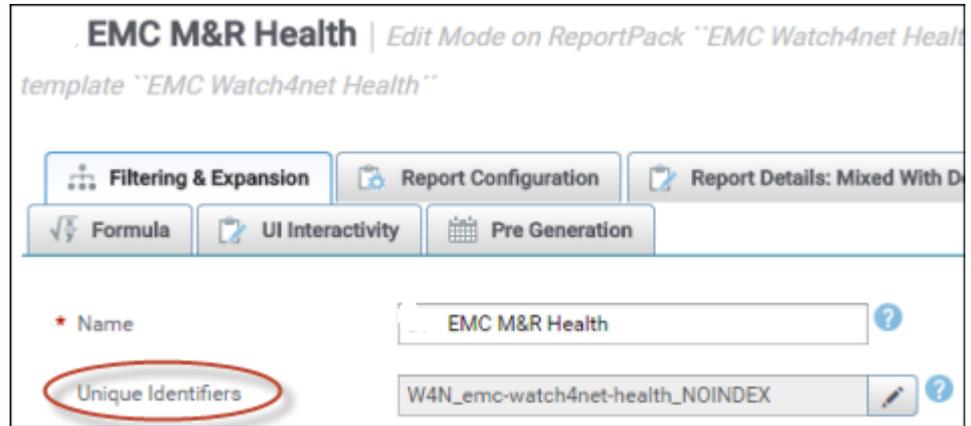
uid

Is the unique identifier assigned to a top-level branch in the reporting tree.

To find a UID:

1. Navigate to the report.
2. Click **Modifications**  > **Edit Reports**.
3. If a message appears stating that the report is linked, click the link in the **Linked to** field.
4. On the **Filtering and Expansion** tab, look for the **Unique Identifier** field.
5. If the **Unique Identifier** field does not exist, click a node that is higher in the reporting structure.

The following image shows the UID for the **EMC M&R Health** node in the reporting tree.



report_name

Is the portion of the report path after the report represented by the *uid*. The report path is a combination of the breadcrumb at the top of the report and tab names if they exist.

For example, consider the following report path:



The lookup path for the Arbiter Load Balancer report is:

```
myserver.emc.com:58080/APG/lookup/{W4N_emc-watch4net-health_NOINDEX}/Collecting Level Performance / Arbiter Load Balancer
```

*** (as last component of report_names)**

The asterisk is required if there are more reports under the requested report. In most cases, it is required because drill down reports are children of a report. In the Alerting examples below, the asterisk is required because of drill downs to individual alerting reports.

Examples

To use the following examples, replace the *server_name:port*.

Example

Report Path:

Dashboards > Operations > Alerts Summary > Alerts by Severity > MAJOR

Lookup string:

```
http://myserver.emc.com:58080/APG/lookup/{DASHBOARDS-BRANCH}/Operations/Alerts Summary/Alerts by Severity/MAJOR/*
```

Example

Report Path:

Report Library > Oracle Database > Inventory > Tablespaces

Lookup string:

```
http://myserver.emc.com:58080/APG/lookup/{ORACLE_INV_DATABASE-
INVENTORY}/Tablespaces/*
```

Example

Report Path:

Report Library > EMC M&R Health > Collecting Level Performance > Arbiter Load Balancer

Lookup string:

```
http://myserver.emc.com:58080/APG/lookup/{W4N_emc-watch4net-
health_NOINDEX}/Collecting Level Performance/Arbiter Load
Balancer/*
```

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>collapse=<nodeid></code>	Collapses the selected node and any child nodes.
<code>collapseALL</code>	Collapses the whole tree.

Report settings

Syntax	Description
<code>display=<code></code>	<p>The code numbers for the display modes:</p> <ul style="list-style-type: none"> 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

mode=<code>	<p>The codes for the report modes:</p> <ul style="list-style-type: none"> • 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>	<p>Not applicable when the period is zero. The code numbers for the aggregate types are:</p> <ul style="list-style-type: none"> • 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:

	<ul style="list-style-type: none"> • 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>	<p>The code to use as the default for the report time range:</p> <ul style="list-style-type: none"> • 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 <p>You can use as many duration codes as needed for the report time range but you cannot duplicate a code.</p> <p>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.</p>
start=<date>end=<date>	<p>The time range of the requested report. The <date> field has the YYYY-MM_DD HH:MM format.</p> <hr/> <p>Note</p> <p>URL encoding may transform spaces into plus signs and colons and periods into %3A.</p> <hr/>
start_ts=<timestamp>end_ts=<timestamp>	<p>The time range using timestamps of the requested report. These are UNIX timestamps.</p>

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>	<p>The maintenance period for the report. The <time filter.expression> field consists of <dom>;<dow>;<hod> to represent the following:</p> <ul style="list-style-type: none"> • 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.

	<ul style="list-style-type: none"> • <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.

Customizing user settings and custom reports

You can customize user account information, such as the password and associated email address, general UI display preferences, and individual table report display preferences. You can manage custom reports and run the Broken Links Detection Tool.

To access the **User Settings** dialog, click **Profile**  > **View Profile** in the banner area of the User Interface.



Modifying your user data and password

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

Procedure

1. Click **Profile**  > **View Profile**.
2. On the **User Data** tab, make changes to name and email address.
An email address is required for certain features, such as emailing reports or requesting notifications.
3. On the **Change Password** tab, change your password.
4. Click **Save**.

Setting your reporting preferences

You can set the behavior for displaying reports for your user account. The choices you make here override the default settings made by an administrator for the portal or in the user profiles.

Procedure

1. Click **Profile**  > **View Profile**.
2. Click the **Preferences** tab.
3. Select the **Language** to use.
4. Select the **Navigation Style** to use.

Option	Description
default	The administrator's setting for the portal
tree	Hierarchical tree
icon	Thin icon column

5. In **Report Auto Refresh Rate**, leave blank to turn off auto refresh, or type the interval between redispays. The default setting comes from the **Administration > Profile** configuration.
6. In **Background Reports**, select one of the following:

Option	Description
Show the Question	<p>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.</p> <hr/> <p>Note</p> <p>For any user session, only one background report can run at a time. If you answer yes when another report is still running, the running report is aborted and will not be generated.</p> <hr/>

Option	Description
No Background Reports	No question appears because generating a report in the background is not an option.

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

Setting the first report to appear after login

A user account can set a favorite report that always appears first after login.

Procedure

- Go to the report that you want to set as the login report.
- Click **Tools > Favorite this report**.

If the report is already favorited, an error states that the link requires a unique name. In this case, you can either:

- Rename the link to make it unique.
 - Click the **Trash** icon to delete this entry, and expand the existing entry of the same name.
- Click the **Use as Login Report** button.
 - 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

- Click **Profile**  > **View Profile**.
- Click the **Custom Reports** tab.
- Click **download a backup** in the text.
- Click **Save** and navigate to the place to store the report tree.

Import a report definition

You can import the XML definition of a node or report into your report tree.

Note

To export the XML definition of a report:

- Navigate to the report whose definition you want to export.
- Click **Exports > XML format**.

Follow these steps to import the definition into another installation.

Procedure

1. Click **Modifications > Edit Reports**.

You can only perform this task in edit mode.

BROWSE MODE displays when you are in edit mode.

2. In the left navigation, 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.
Selecting **My Reports** or a node under **My Reports** is recommended.
3. Click **Profile**  > **View ProfileCustom Reports**.
4. Next to **Upload a Branch**, click **Browse** to locate the XML file to upload.
5. In **Conflict Management**, select the preferred action in case the node or report that you are uploading is using the same ID as an existing node or report:

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

6. Click **Save**.
7. Click **BROWSE MODE** to view the imported node or report.

Find and fix broken links in reports

The Broken Links Detection Tool scans your entire report tree and identifies all report links that cannot be resolved. It fixes links if possible and provides best guess suggestions for resolving others.

Reasons for broken links

As reports are moved, removed, updated, or disabled in the report tree, links to those reports from other reports must be changed. The old links no longer work. In addition, pre-generated reports and reports in the **My Reports** node, such as pinned reports, scheduled reports, and favorited reports are based on links that might be broken when reports are moved, removed, or disabled.

Changed UIDs also result in breaking links to reports that were linked or hooked to the original UIDs.

SolutionPack upgrades that include moved or updated reports can impact links. For this reason, whenever a SolutionPack upgrade occurs, the upgrade process schedules the Broken Links Detection Tool to run after a timed waiting period. If you are sequencing multiple SolutionPack upgrades closely together, the waiting period is moved out with each upgrade, so that the Broken Links Detection Tool runs only once after all of the upgrades seem to be finished.

You can run the Broken Links Detection Tool on demand at any time.

Fixing broken links

Many broken links are fixed automatically by the tool during its execution. These are the links that are known with 100% certainty to be remapped to other locations. The tool does not show the automatically fixed links. If you are

interested in viewing them, you can change the logging level of the daily Tomcat log file. See [View fixed links in Tomcat logs](#) on page 129.

For suggested fixes that do not rate a 100% confidence, the tool presents you with the path name of the broken link, a suggested path for fixing the link based on certain assumptions, and a confidence percentage for how accurate the suggestion might be. You can select whether you want to apply the suggested fix, and the tool will apply the fix.

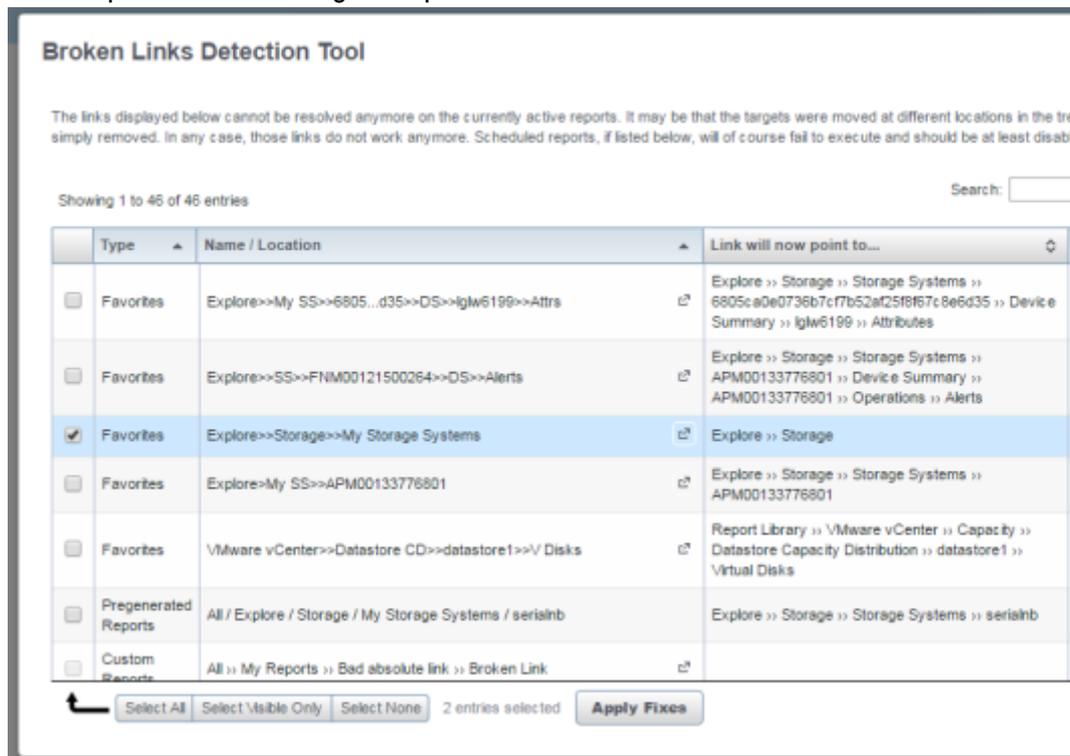
If you believe the suggestion is incorrect, or if there is no suggestion provided, you must manually fix the link.

The following procedure describes how to run the tool and how to resolve the detected broken links that were not fixed automatically by the tool.

Procedure

1. Click **Profile** > **View Profile** > **Custom Reports**.
2. In the **Broken Links Detection** section, click **Open Tool**.

This button launches the Broken Links Detection task. The task runs on your entire report tree. The dialog that opens shows the results of the run.



The dialog shows the following information for each broken link detected:

Column	Description
Type	The type of link that is broken. Examples are: Custom reports, Favorites, Pinned, Scheduled, Pre-generated.
Name/Location	The report containing the broken link.
Link will now point to...	The report path of the proposed new link.
Confidence	Percentage of confidence that the new link is correct.

Column	Description
	<ul style="list-style-type: none"> • When Confidence is 100%, the tool fixes the link automatically and it is not listed here. The 100% confidence rating occurs when changes match those that are recorded in the mapping files installed with the tool. • When Confidence is not 100%, the value is based on how many components in the broken URL were mapped to known new values or new values that are similar to the original. • The lowest level confidence suggestions are based on similar node names to the original path, and similar depths of levels in the hierarchy. You should check the suggestion carefully to ensure it is correct, and if not, fix the link manually. • A 0 level confidence rating indicates that the tool could not find any similar report path to suggest. The old and new report path names are too different from each other to be matched. A manual fix is required.

3. Analyze the suggestion in each row.
4. To accept a suggestion, click the box in the first column.
5. Click **Apply Fixes**.
6. When a suggestion is not correct or when there is no suggestion at all, manually fix the link as follows:
 - a. Click the **Go To** icon in the **Name/Location** column.
 - b. Use any of the following suggestions to manually correct the link path:
 - For Scheduled Reports, Pinned Reports, and Favorites, it is easiest to recreate the link using the User Interface, and delete the outdated report.
 - Correct a link manually: Click **Modifications > Edit Reports**, and change either the incorrect UID or the incorrect report path in the report definition. There are various places in a report definition where a link to another report might occur. Save the change, and return to Browse Mode.
 - If there are many reports linking to the same report path, and that report path is not being found, causing many failures for the same reason, consider the following approach: revert to the previous version, add a UID to the report that is not being found, then upgrade, add the same UID into the report definitions in the new location, and run the tool again. You should analyze whether this advanced approach is better than correcting each link manually.
7. Rerun the Link Detection Tool.
8. Repeat these steps until all broken links are fixed.

View fixed links in Tomcat logs

Fixed links are logged in the daily Tomcat log file if the logging level is set to FINE or FINEST.

By default, the logging level does not produce much information about fixed links. You need to change the logging level.

Note

Additional log entries will increase IO activity and can significantly impact performance.

The log filename is `catalina.<date stamp>.log` located on the Frontend server here:

```
/opt/APG/Web-Servers/Tomcat/Default/logs
```

The configuration file for changing the log level is:

```
/opt/APG/Web-Servers/Tomcat/Default/conf/logging.properties
```

You can change the logging level and access the log files on the web portal.

Procedure

1. Go to **Centralized Management > Logical Overview > Miscellaneous > Web-Servers > Tomcat:: *instance name* - *server name***.
2. To change the logging level:
 - a. In the right pane, expand the **Configuration Files** blue bar.
 - b. Locate the `conf/logging.properties` file and click the **Edit** icon on the row.
 - c. Add the following line to the end of the file:

```
com.emc.mnr.links.level=ALL
```

- d. To enable FINEST logs, locate this existing line:

```
1catalina.com.watch4net.apg.logging.jul.handler.RotateFileHa  
ndler.level  
= FINE
```

- e. Change `FINE` to `FINEST`.
 - f. Save the file.
 - g. Restart Tomcat.
 - h. Rerun the Detect and Fix Links tool to start capturing the additional log entries.
3. To view the log entries:
 - a. Expand the **Log Files** blue bar.
 - b. Download or view a `catalina.<date stamp>.log` file.

Using the Report Library

The Report Library node in the navigation tree holds all of the reports for the installed SolutionPacks.

Each installed SolutionPack has a corresponding node under Report Library.

Accessing SolutionPack reports

You can access the reports for installed SolutionPacks in Browse Mode.

Procedure

1. Click **Report Library**.
2. Click a SolutionPack name.

The selected SolutionPack node expands, showing a hierarchical tree of categories and reports offered by the SolutionPack.

3. Click a report.

Results

The selected report appears in the report pane.

Viewing the summary of a SolutionPack before installing it

You can view the details of a SolutionPack in the **SolutionPack Center** before installing the SolutionPack.

Procedure

1. Click **Administration**  > **Centralized Management** > **SolutionPacks**.
2. Click **SOLUTIONPACK CENTER**.
3. Select the SolutionPack in the **Browse and Install SolutionPacks** page.

Administering the interface

Administrators can customize the user interface, set global preferences, manage user logins and sessions, and edit scheduled reports directly from the portal.

Modifying the default profile settings

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

Both UI pages affect the exact same configuration settings.

Procedure

1. Navigate to either of the following:
 - **Administration**  > **Portal** > **Default Display**.
 - **Administration**  > **Profile** > **Default Profile**.
2. To upload and use a new logo image:
 - a. Click **Upload New Logo**.

- b. Click **Choose file**, browse to the new image to upload, and click **OK**.
- c. Select the new logo file name in the **Logo** drop-down list.
- 3. Select the **Locale**.
- 4. Select the **Time Zone**.
- 5. Click **Save**.

Set first report after login

Administrators can control which report appears first after login for each user profile.

Each user account can override the Profile setting by clicking **Tools > Favorite this report** and setting the associated **first report after login** button.

Procedure

1. Click **Administration > Profiles > *profile_name* > Customizable Settings**.
2. For **Login Report**, provide the report identifier, report URL, or report lookup value.
Click the tooltip for formatting information.
3. Click **Save**.

Global Portal Properties

You can set global properties for the frontend interface.

Click **Administration > Portal** to access the **Global Portal Properties** page.

Tab	Settings
Frontend Status	<ul style="list-style-type: none"> • Set autocomplete behavior for logging in. • Enable/disable maintenance mode for the portal. • Add a message to the login screen.
Default Display	Set the portal-wide logo, banner, language, time zone, and navigation styles.
File Customization	Customize the look and actions in the portal by uploading custom CSS and JavaScript files. Enable keyboard shortcuts and upload hotkey definitions.
Duration Aliases	Define the aliases to use on the portal for time range duration values, such as quarter (for 3 months) and semester (for 6 months).
External Users	<ul style="list-style-type: none"> • Link the Admin Role or Custom Reports Role to external users. • Allow or disallow external users to schedule reports.

Setting login properties and welcome message

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

Procedure

1. Click **Administration  > Portal**.
2. Click the **Frontend Status** tab.

3. For **Login Page Autocomplete**, choose from the following options:
 - Retain the default browser behavior
 - Disable browser autocomplete for both the username and password.
 - Disable browser autocomplete for the password only.
4. Type a customized splash screen message in **Welcome Message**.
HTML code is accepted.
5. Click **Save**.

Disabling user logins

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

Procedure

1. Click **Administration**  > **Portal**.
2. Click the **Frontend Status** tab.
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**.

Set default navigation style

The navigation style controls the look and feel of the left navigation column in the reporting interface.

This procedure sets the global default navigation style. Users can override the global default style for their user account on **Profile**  > **View Profile** > **Preferences**.

Procedure

1. Click **Administration**  > **Portal**.
2. Click the **Default Display** tab.
3. In **Navigation Style**, make a selection.

Option	Description
Use icon navigation	The navigation column contains icons that expand into reporting categories
Use tree navigation	The navigation column is a hierarchical tree of report names

4. Click **Save**.

Customizing the look of the interface

You can change the look of the interface by applying custom CSS and JavaScript files.

The selected files apply to the reporting interface, the alerting interface, and the Centralized Management interface. The CSS files can customize elements in the interface, such as buttons and menus. The JavaScript files run when a page is loaded.

Use the **File Customization** tab to:

- Upload custom CSS and JavaScript files

- Select the CSS and JavaScript files to apply
- Download, edit, and reload any CSS or JavaScript file

Procedure

1. Click **Administration**  > **Portal**.
2. Click the **File Customization** tab.
3. If needed, click an **Upload** button to upload a file into the system.
4. Click to select one or more files to apply to the interface.
You can select multiple .css and multiple .js files.
5. To edit files:
 - a. Select the files to edit.
 - b. Click the appropriate **Download** button.
 - c. Edit the files offline, and then upload them.
6. 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**  > **Portal**.
2. Click the **File Customization** tab.
3. Check **hotkeys.js**.
4. Click **Save**.
5. Press **Alt + h** to view the keyboard shortcuts.

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**  > **Portal**.
2. Click the **Duration Aliases** tab.
3. Edit the default values for the aliases.
4. To add another alias, click **Add period alias** and provide the values.
5. Click **Save**.

Assigning system roles to external users

You can assign the Admin Role and the Custom Reports Role to a user who is authenticated in an external system, such as LDAP. You can enable or disable permissions for external users to schedule reports.

Procedure

1. Click **Administration** > **Portal**.

2. Click the **External Users** tab.
3. For **Admin Role** and **Custom Reports Role**, identify an external user that is already defined under **Administration > Roles > External Users**.

Note

Only one external user is accepted in each field.

- For LDAP, copy a Distinguished Name (DN) from **Roles > External Users** .
For example:
`CN=test,OU=my_ldap,DC=1ss,DC=emc,DC=com`
 - Otherwise, copy a User Principal from **Roles > External Users**.
4. Click the **Schedule Reports** checkbox to enable external users to schedule reports.
 5. Click **Save**.

Scheduled Reports

You can access all scheduled reports from one portal page.

Click **Administration > Portal > Scheduled Reports**. The following sub-reports appear on the page: All scheduled reports appear, including those that are pending, currently running, pre-generated, such as the daily health report, and any invalid reports.

- **Users Scheduled Reports**—Lists users that have defined scheduled reports. Click a user to list and manage the reports scheduled by that user.
- **Pre-Generated Reports**—Lists reports that have a defined pre-generation schedule in their report definition.
- **Running Reports**—Lists reports that are currently running.
- **Invalid Scheduled Reports**—Lists reports that are scheduled but cannot run successfully.

Managing scheduled reports

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 > Portal > Scheduled Reports**.
2. In the **Users Scheduled Reports** table, click a user name.

The **Scheduled Reports** table appears, listing all of the reports that the selected user has scheduled.

3. Right-click 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.

4. Choose an option from the right-click menu.

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**.

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**.

User Sessions

You can view information about all current user sessions.

Click **Administration > Portal > User Sessions** to access the User Sessions page. The page shows the following information for each current user session:

- Login name—Multiple users can log in using the same login name.
- Login date and time
- Last access date and time
- Client IP address
- Client host name
- Status

Canceling a user session

You can cancel a user's login session, which logs the user out of the interface.

Procedure

1. Click **Administration**  > **Portal** > **User Sessions**.
2. Right-click the user name, and choose **Kick Out**.

Note

You cannot cancel the session of an administrator.

3. Click **OK** to confirm the action.

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**.

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.

Configuring an SMTP server

Configure an SMTP server to enable the email features in the product.

Procedure

1. Click **Administration**  > **Modules** > **Alerting**.
2. In the banner, click **Global Settings** .
3. Configure the SMTP fields.
4. Click **Save**.

You have successfully set the SMTP variables on the Backend (alerting) server. In a 1-server setup, these settings also apply to the Frontend server.

5. In a setup with more than one server, set the SMTP variables on each Frontend server.

Note

This step is required in a 4-VM vApp, or if the installation includes more than one Frontend.

- a. On the Backend server, copy the SMTP variables in `/opt/APG/bin/apg.properties`.
- b. On each Frontend server, paste the variables into `/opt/APG/bin/apg.properties`.
6. Restart the Tomcat server.
 - a. Go to **Administration** > **Centralized Management** > **Logical Overview** > **Miscellaneous** > **Web Servers**.
 - b. Click a Tomcat server.
 - c. Click **Restart**.

Managing users and 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 filter 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.

User Management

A user account and associated password is required to log onto the system. A user account can be shared, and multiple online sessions can log in with the same account simultaneously.

The **User Management** page lists all user accounts, their status, their profile, whether the profile is enabled, and some optional information about the user, such as name and email address. From this page, you can add new user accounts and edit existing user accounts.

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. 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. Click a user.
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. Click a user.
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**.

Test user settings

You can test user settings to check how the settings are applied. Test mode applies settings as if you were logging in as this user.

Procedure

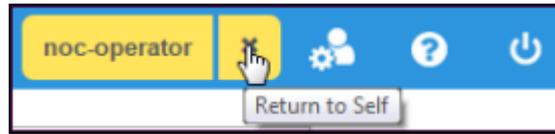
1. Click **Administration** .
2. Under **User Management**, click **Users**.
3. Right-click the user you want to test, and choose **Test User**.

You are launched into the User Interface, interacting with the Console as if you had logged in as the user being tested. The user name that you are testing appears in the banner.



4. Browse reports and edit reports, and perform any other regular functions.
Any changes you make, such as editing report settings, are saved with their new settings.

5. To exit test mode, click the **X** next to the test user's name in the banner.



Copying users

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

Procedure

1. Click **Administration** .
2. Click **Users**.
3. Type the settings for this new user.
4. Click **Save**.

Editing users

You can modify the settings of existing users.

Procedure

1. Click **Administration** .
2. Click **Users**.
3. Right-click 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. Click **Users**.
3. Right-click 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**.

Roles management

A role defines rights and restrictions that apply to users assigned to the role.

The **Roles Management** page provides a view of the defined roles in your system, their descriptions, and how many users currently use each role. From this page, you can create new roles or edit existing roles. You can also enter Advanced Mode to

establish inheritance relationships, assign administration tasks to a role, and reset default roles to their original settings.

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**.

Profiles

A profile defines user account attributes that apply to a group of users.

A profile defines the following attributes:

- Locale (language)
- Time zone
- Logo to display in the banner
- Login report to display when the users log in
- Report auto refresh rate

The **Profiles Management** page shows the profiles that are defined in your system and the number of users assigned to each profile. From this page, you can create new profiles, edit existing ones, and assign users to profiles.

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.

To have the default login report display an icons report, enter: ?
`report&select=0&mode=ico&display=2`
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**.

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`.

Managing ReportPacks

ReportPacks are pre-defined reports that are dedicated to a specific technology or vendor. Using a ReportPack saves you customization time and effort.

ReportPacks include built-in metrics, formulas and analytics. They also include report templates.

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. Click **Administration > ReportPacks**.

4. Click **New ReportPack**.
5. On the **ReportPack Information** tab, type the **Name**, **Version**, and **Description** of the ReportPack.
6. Click **Save**.

Uploading a *<ReportPack>.arp* file

You can add a *<ReportPack>.arp* file that contains required information and report templates. The ReportPack displays 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.

Only .arp file types may be uploaded.

4. Click **OK**.

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 > ReportPacks**.

2. Click the ReportPack name.
3. Scroll down, and click **+New Template**.
4. On the **Template Information** tab, type the **Name**, **Version**, and **Description** of the report.
5. 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.

6. Click **Choose File**, navigate to the template file to add, and follow your browser prompts to select or open the file.
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.

Exporting a report as a template in XML format

You can export a report in XML format which can then be used as a report template.

Procedure

1. Navigate to the report.
2. Click **Edit Mode**.
3. Click **EXPORT** in the banner.
4. Choose **Report Template**.
5. Follow your browser prompts to save the file.

You can add the extracted template to a ReportPack.

Centralized Management

Centralized Management provides access to servers, system modules, tasks, SolutionPacks, and ESRS configuration. It also provides access to configuration files, system logs, data enrichment modules, and the discovery center.

To access Centralized Management:

Procedure

1. In the banner area of the User Interface, click **Administration**  > **Centralized Management**.
2. Use the navigation column on the left to access a server or open a management node.

Server Management

The Physical and Logical server overviews enable you to configure, monitor, and manage the servers in your deployment.

After you register the first server, Centralized Management connects to it and retrieves information such as the server status, the modules installed on it, and provides access to it through the navigation tree.

Physical Overview

The Physical Overview displays the registered servers and the modules installed on each server. To get more details about a server, click the server's name. The charts at the top of the page can be configured to display a combination of historical data and real-time data.

Logical Overview

The Logical Overview shows the distribution of the modules based on their categories. To get more details about a module, click the module's name.

Each module page provides:

- Module properties - category, name, version, instance name, install path, and description
- Service status - displays the service status and enables you to stop, start, and restart a service
- Configuration files - enables you to view and edit the module configuration files
- Log files - enables you to view and tail module log files

Registering additional servers

This topic describes how to register remote servers using Centralized Management.

Before you begin

The additional server must be installed.

Procedure

1. Click **Administration** > **Centralized Management**.
2. Click **CONFIGURATION**.
3. Click **Register a Server**.
4. Enter the following server details:

- Server HostName - links the server to a device in the APG database. So in order to get historical data for the device, we recommend you use the same name as the one used by APG, which is the server hostname. You can confirm this by looking at APG Health reports in the frontend. Server HostName is also used as the server display name if you do not create an alias.
- Server Alias - customizes the server display name in Centralized Management without any other impact. If you do not create a Server Alias, the hostname is used.
- Server Description - is displayed on the Centralized Management home page in the Physical Overview under each server name. It can be left blank or used to describe a server role.
- Gateway URL - contains the URL that points to the remote WebService Gateway. The default configuration is: `https://server.name.or.ip:48443/`. The WebService Gateway service must be running on the remote server for this to be working.
- Username - is the one from the Gateway credentials. It's configured at its level, and described in its own documentation. Default value is admin.
- Password - is the one from the Gateway credentials. It's configured at its level, and described in its own documentation. Default value is changeme. On an already configured server, Password field is never displayed again. There is no need to fill it before saving: if left blank, it won't be removed, but rather left untouched. The only means to empty a saved password is to empty the username as well.
- SSL Validation - can be checked in order to validate the SSL certification before instantiating communication. When using a self-signed certificate, it should be left unchecked. When using a certificate recognized by a certificate authority, it should be checked.
- Operating System - will be used to select the packages to send and use on this server. Be sure to set it to the correct value, as setting it to a wrong value causes invalid or corrupted installations.
- Health Values - optional, configures the type of data that will be recovered and displayed for this server. Real-time data is collected through the APG Self-Monitoring Collector which provides real-time access to the server metrics. Historical data is retrieved from the usual APG database(s) and is also based on the metrics collected by that collector. Databases are defined in the web application context file. See the section on page 61. The physical overview displays historical data if it is enabled. The per-server view displays real-time data if it is enabled, and if not it falls back to historical data. If it is disabled, charts are hidden.
- Report in APG - optional, can be used as a shortcut to access a report in the APG frontend. Enter the targeted URL and this creates a link to it in the global physical overview. For example, `http://apg.server.name.or.ip:58080/APG/#0-4-52664238`.
- Server Tags - adds the ability to specify string tags for servers. Tags provide useful information about a server that help to select the recommended server for installing a SolutionPack component. For example, a server tag could be the last part of the ID of a SolutionPack component, such as "collect." A SolutionPack component with the ID `emc-vnx-collect` would be installed on a server with the tag "collect."

Starting and stopping services

The services panel displays the modules that have a service configuration, including those modules whose service entry has been removed.

The services table displays the module information, its current service status, and an estimate of its last status change. The status column can have multiple values depending on both the service properties and its current state:

- **Started** - the service is currently running.
- **Stopped** - the service has been stopped.
- **Not installed** - the service has been removed and is no longer available.
- **Unknown** - the current status cannot be determined. This is usually caused by a service that requires administrative privileges to run, such as the Collector Manager when you use the ICMP Collector on UNIX.

Services that require administrative privileges cannot be managed through Centralized Management.

If the **Start**, **Stop**, and **Restart** buttons are not available, the service cannot be managed through Centralized Management.

- **No button visible** - occurs on a service that needs to be running for Centralized Management to access the server; as a consequence, it cannot be started, stopped, or restarted. You must manually connect to the server using SSH or a remote desktop to perform these operations, for example the Webservices Gateway module.
- **All buttons are grayed** - Centralized Management does not have enough rights to perform operations on this service. This can apply to any service that requires administration privileges, such as the Collector Manager when running the ICMP Collector.
- **Only a restart button** - this occurs when looking at the Tomcat server used to run Centralized Management. The restart command is available, but if it were to fail, you must manually connect to the server to fix the issue.

Procedure

1. Click **Administration > Centralized Management**.
2. Click a server under **Physical Overview**.
3. Click the service you want to stop or start.
4. Click **Start**, **Stop**, or **Restart**.

Distribute startup of UNIX services

During a system startup on UNIX servers, problems can occur if all services attempt to start at the same time. You can distribute service startup by editing the `conf/unix-service.properties` files for each service.

Before you begin

This procedure only applies to UNIX configurations.

Procedure

1. Click **Administration > Centralized Management**.
2. Under **Physical Overview** or **Logical Overview**, navigate to a service name.
For example, collector-managers are services.

3. In the right pane, expand the blue bar for **Configuration files**.
4. Locate the `conf/unix-service.properties` file, and click the **Edit** (pencil) icon on the row.

In a long list of files, it might be on the second page.

5. Locate lines similar to the following in the middle of the file:

```
start.mode=java
start.target=com.watch4net.apg.v2.collector.Bootstrap
start.param.1=main
start.param.2=start
stop.mode=kill
# timeout is in seconds
stop.timeout=60
```

6. After `start.param.2`, add the following parameter

```
start.delay=n
```

where *n* is the number of seconds to delay the start of this service on system startup.

For example:

```
start.mode=java
start.target=com.watch4net.apg.v2.collector.Bootstrap
start.param.1=main
start.param.2=start
start.delay=40
stop.mode=kill
# timeout is in seconds
stop.timeout=60
```

7. Repeat these steps for each service, altering the `start.delay` value for each service to create a smooth startup.
8. Click **Save**.

Installing modules

Additional modules can be installed on a server.

Procedure

1. Click **Administration > Centralized Management**.
2. Use the **Physical Overview** or **Logical Overview** node to navigate to a server.
3. Click the **SolutionPacks & Modules** tab.
4. Click **Install**.
5. Select the packages you want to install.
6. Click **Launch**.

The Module Installation dialog box appears.

7. Enter an instance name.
8. Click **Launch**.

During the installation process, questions will appear in the display. You must answer each question within five minutes or the module installation process will

be interrupted, which can cause an inconsistent state on the remote server for the module.

You can remove a module that has been installed by mistake or been moved to another server by clicking **Uninstall**.

Updating modules

When new versions of a module are available, either by running the setup of a new release or by uploading new packages to the central repository, an update button is displayed on the module page.

If different versions are available, the update button will have the option to select the version you want.

Procedure

1. Click **Administration > Centralized Management**.
2. Use the **Logical Overview** node to navigate to the module to be updated.
3. Click the **SolutionPacks & Modules** tab.

If the module has an update available, the **Update to Latest Version** button is displayed at the top right of the page.

4. Click **Update to Latest Version**.

Accessing and editing configuration files

Using Centralized Management, you can edit, upload, and remove module configuration files.

Each module manages its own configuration files, even if the module is used within one another. For example, a collector or filter configuration that is used in a collector manager, is accessible only on the corresponding collector or filter page. The configuration files table only displays the files that are identified by the module as configuration files.

Procedure

1. Click **Administration > Centralized Management**.
2. Use the **Physical Overview** or **Logical Overview** branch to navigate to the module, service, or task whose configuration file you want to access.
3. Expand the **Configuration Files** block.
4. Click the pencil icon to open the configuration file.
5. Edit the file.

Be careful when making changes. The configuration file will not be checked to ensure that the changes you make are valid.

6. Click **Save**.

The file will be updated on the server.

Viewing log files

Log files are useful to check on a module that has recently been started or to track down a specific error that has occurred.

Log files are listed in the order of the last modification, so you usually want to take a look at the first file in the table. The number of files is defined in the logging configuration file, which is called logging.properties.

Procedure

1. Click **Administration > Centralized Management**.
2. Use the **Physical Overview** or **Logical Overview** branch to navigate to the module for which you want to view the log.
3. Expand the **Log Files** block.
4. Click the tail file or view file icon to open the log file.
 - View File - the view is not updated if the log file contents change on the disk.
 - Tail File - new lines are appended to the end of the view as they are written to the log file.
5. Click **Close**.

Viewing module tasks

The Task Scheduler is a service that runs scheduled tasks. The Task Scheduler replaces cron on UNIX-based systems and Schedule Tasks on Windows servers. The Task Scheduler runs important tasks, such as refreshing the Frontend-Search index periodically, or performing database maintenance tasks to optimize performance. Scheduled tasks can be viewed for each module.

A task has four parts:

- a schedule indicating when it will be executed
- the maximum time the task can take to run
- an action to take upon execution
- locks that prevent tasks from running simultaneously

Procedure

1. Click **Administration > Centralized Management**.
2. Under the **Physical Overview** node, click a server name.
3. Click the **Tasks** tab.

A table of tasks shows run statistics and current task status.

4. Click a task to access the following:
 - Configuration files for the task. You can view or edit the files.
 - Log files for each execution of the task. These can help troubleshoot tasks that did not complete successfully.
 - **Disable** button. You can cancel a pending task execution or interrupt it once it has started. There is no confirmation and the task is immediately cancelled.
 - **Run now** button. You can start execution immediately, rather than wait for the next scheduled run.

Package management

Using Centralized Management, you can manage how packages are used when installing or updating modules.

A package is a module that is frozen at a certain code version. Packages are used to determine when new versions of modules are available and to install and update

module instances. Modules are components of the core software, such as: backends, web servers, collectors (XML Collector, SNMP Collector, etc.), filters (Property Tagging Filter, Cross-Referencing filter, etc) and so on. Modules can be installed and managed separately. When a module is installed, that instance of the module has a unique instance name so that it can operate separately from other instances of the same module. By default, the name of a module instance is "Default."

The Packages Listing displays the package names, versions, and operating systems. Operating system is important as it determines whether the package can be installed on a particular server or not. You can mouse over the Status column to see where packages are installed. You can install a package multiple times on the same server.

Package synchronization

Packages can be retrieved, cleaned up remotely, or discarded from the Centralized Management package repository, which by default is located at Tools/Module-Repository.

Synchronization has the following options:

- Retrieve the last packages - connects to each server and fetches the latest version of their packages. This is the easiest way to install new packages as it runs the installer and then runs the synchronization.
- Retrieve all the packages - retrieves every available package from the servers, not just the latest ones. This is useful if you need a version that is not the most recent one.
- Remove remote packages - performs a cleanup of the remote servers. For each server, every package that is not installed on it is removed from its repository to free the disk space on the server.
- Remove old versions from repository - performs a cleanup on the Centralized Management package repository and keeps the most recent version of each package, as well as the versions that are currently installed on the remote servers.

A summary is displayed for each server with the count of globally added or removed packages and the file size. For example, on the central repository, adding two packages and removing five packages results in a count of negative three packages.

License management

Using Centralized Management you can add, delete, and check for expired licenses. Licenses are linked to a host.

Obtaining a permanent license

Obtain a permanent license by creating a Service Request (SR).

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	<pre>/opt/APG/bin/manage-licenses.sh host-id</pre>
Windows	<pre>C:\Program Files\APG\bin\manage-licenses.cmd host-id</pre>

2. Copy the host ID information to a clipboard or to a text file.
3. For product licenses, go to EMC Online Support at <https://support.emc.com> .
4. Browse to **Service Center > Product Registration & Licenses**.
5. Paste the host ID information into the Service Request.

Obtaining a temporary license

Obtain a 30-day temporary license file by creating a Service Request (SR).

Procedure

1. For product licenses, go to EMC Online Support at <https://support.emc.com>.
2. Browse to **Service Center > Product Registration & Licenses**.
3. Create the Service Request.

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. Click **Administration > Centralized Management**.
2. Click **Licenses Management**.
3. In the right pane, click **Upload**.
4. In the **License Upload** dialog box, browse to your license file. Click **OK**.
Licenses in a .zip file are extracted and will appear in the **Licenses Listing** table.
5. When **License Upload Complete** appears, click **OK**, and then click **Continue**.
6. If you have remote servers in a distributed environment, click **Synchronize** to add the licenses to the remote servers.

Synchronizing licenses

Use license synchronization when licenses have been installed or removed while the server was not reachable or unregistered. Synchronization uses the licenses that are in the Centralized Management license repository.

Procedure

1. Click **Administration > Centralized Management**.

2. Click **Licenses Management**.
3. In the right pane, click **Synchronize**.
4. In the **License Synchronization** dialog box, select one of the following:

Option	Description
Uninstall licenses	Uninstalls and removes licenses that are not available in the Centralized Management license repository. You cannot recover these licenses except if you have stored the license files elsewhere.
Install new licenses	Searches the licenses in the Centralized Management repository and makes sure they are properly installed on every remote server that is reachable. This is the most commonly used synchronization option.

5. Click **OK**.

Logs and diagnostics

The system modules write to log files that are accessible from the Console. The Diagnostic files are a collection of configuration and log files that can help analyze problems.

For information about log entries or the contents of configuration files, see the *EMC M&R Advanced Administration Guide*.

Accessing log files

You can download or view the log files generated by the system modules. You can watch the end of any log file.

Procedure

1. Go to **Administration > Centralized Management > Logs and Diagnostics > Log Files**.

The table shows all of the available log files, when they were last modified, and their size.

2. To view the end of any file, click the **Tail** icon in the **Filename** column.

A window opens showing the last set of entries in the file.

3. To watch as new entries are added, click the **Tail** icon, and then scroll to the end of the window.

As new entries are written to the log, they are visible at the end of the list and the display rolls up.

4. To view the entire contents of a file that is less than 1 MB, click the **View File** icon in the **Filename** column.

5. To download files, click the checkboxes to select the rows, and click **Download**.

Generating and downloading diagnostic files

Diagnostic files are a collection of system configuration files and log files that can help in researching problems.

Procedure

1. Go to **Administration > Centralized Management > Logs and Diagnostics > Diagnostic Files**.

2. Click **Generate diagnostic files**.

This action generates a zip file of XML, CSV, properties, script, and java files that define your current configuration and implementations. Some log files are also included.

When generation is completed, a **Download** button appears.

3. Click **Download** to access the .zip file.

Discovery Center

Discovery Center provides a central location to view and manage all devices that are being monitored by the SolutionPacks. In Discovery Center, you can add new devices to be monitored, change device connection credentials and parameters, and test connectivity to devices.

There are several ways to add new devices:

- Add a single device manually
- Add devices in bulk by importing a CSV file
- Add devices using discovery

Discovery Center contains two sections:

- The **Inventory Management** section is where you perform all add device and manage device activities.
- The **Discovery Center Backends** section contains collector information required by the automatic discovery operations.

Manage devices

The Inventory Management node is organized by device type. You perform all device management activities on the device type pages.

The list of device types on the **Inventory Management** page is based on the SolutionPacks that are installed.

Click a device type row to manage devices of that type. Management activities include:

- Review the list of devices that are actively being monitored.
- View and change connection credentials and other device-specific parameters
- Test connectivity to each device
- Add or delete devices

Viewing all known devices and testing connectivity

You can view a list of the devices that are known to the system. You can verify the connection parameters and availability for a device.

Procedure

1. Navigate to **Centralized Management > Discovery Center > Inventory Management**.

The table lists all of the device types that are being monitored and the number of devices in each type.

2. Click the device type of the device you want to test.

The **Collected Devices** tab lists all of the devices currently being monitored.

3. Select a device to test connectivity to it.
4. Click **Test** to verify that the device can be reached.

Changing device configuration

You can change the connection parameters and other configurations for a device.

Procedure

1. Navigate to **Centralized Management > Discovery Center > Inventory Management**.
2. Click the device type of the device you want to test.
3. On the **Collected Devices** tab, click the row of the device to test.
The configuration dialog for that device appears.
4. Change the parameters as needed.
See the product SolutionPack documentation for information.
5. Navigate to **Centralized Management > Discovery Center > Inventory Management**.
6. Click **Test** to verify connectivity.
7. Click **OK** to save the changes and exit the dialog.
The device row now appears blue and in italics, indicating that the changes must be distributed to the collector servers.
8. Click **Save**, and then **OK** to confirm the save.
9. Click **OK**.

Adding a new device manually

Using **Discovery Center** you can manually add a new device to be monitored.

Before you begin

To add a new device, you must already have a SolutionPack installed that supports that device type.

Procedure

1. Navigate to **Centralized Management > Discovery Center > Inventory Management > *device_type***.
2. On the **Collected Devices** tab, click **Add new device**.
3. In the device configuration dialog, enter the parameters for the new device.
The configuration dialog is device-specific. See the product SolutionPack documentation for information about each field.
4. Click **Validate and add**.
The validation tests connectivity to the device using the provided information. If an error indicator appears, correct the information and click **Test** to try again.
5. Click **OK** to confirm the add.
The new device appears in the **Collected Devices** table in blue and italicized, indicating that it is not yet saved in the system.
6. Click **Save**.
7. Click **Ok** to confirm the save.

8. Click **Ok**.

The **Status** column represents the discovery results. You can click on the status icon to view the discovery results.

If the connectivity status is a green check, the new device is now being monitored.

Adding devices using CSV files

You can import a properly formatted CSV file to add devices.

Each device type provides a template that describes the required format of the CSV file for the device type. You can also export existing devices into a CSV file.

Importing a CSV file

You can import a CSV file containing information about new devices to be monitored.

Before you begin

To import new devices, you must already have a SolutionPack installed that supports the device type to be imported.

To get a template of the CSV file for a specific device type, use the **Export Template** button.

Procedure

1. Navigate to **Centralized Management > Discovery Center > Inventory Management**.
2. Click the device type you would like to add.
3. Click **Import**.

The **Import new devices** popup appears.

4. For **Merge the devices to the existing ones?**

Option	Description
Do not check the option	Overwrite the current list of devices with the devices from the CSV file
Check the option	Keeps the current list of devices and add (merge) devices contained in the CSV file to the current list

5. Click **Choose File**.
6. Navigate to your CSV file.
7. Click **Ok**.
8. Click **Continue**.
The new devices appear in the Collected Devices table in blue and italicized.
9. Click **Save**.
The **Save Devices** popup is displayed.
10. Click **Ok** to overwrite or merge the devices.
11. Click **Ok**.

Exporting devices

You can export the list of devices currently being monitored.

Procedure

1. Navigate to **Centralized Management > Discovery Center > Inventory Management**.
2. Click the device type of the devices that you want to export.
3. Click **Export**.

Follow your browser's prompts to save the file.

Exporting a CSV file template

A template shows the expected format of the CSV file for the bulk import of devices. The template includes headers.

Before you begin

To export a CSV template file, there must already be one device of that device type available.

Procedure

1. Navigate to **Centralized Management > Discovery Center > Inventory Management > *device_type***.
2. Click **Export Template**.
3. Follow your browser's prompts to save the file.

Add devices using discovery

The discovery feature uses saved information in discovery groups to find new devices. A discovery group is specific to a device type.

The following procedures are required to implement discovery for a device type:

1. Register a collector server that supports discovery for the device type.

Note

The discovery method is supported by many, but not all, device types.

2. Create one or multiple discovery groups for the device type.
3. Trigger discovery for a discovery group.
4. Distribute the discovery results to the collector.

Registering a new collector server

To use the automatic discovery features, you must first register the collector server that supports the device type you want to discover.

Procedure

1. Go to **Centralized Management > Discovery Center > Discovery Center Backends**.

The table lists your system's Backend servers.

2. Click the row for a Backend server.

The collector servers that are registered to perform discovery are listed. The table also shows the discoverable device types supported by each collector server.

Note

If the table is empty, no collectors are registered.

3. To see a list of unregistered collector servers, click **Register**.
4. Select one or more servers from the list, and click **Register**.

If you are collecting VMware vCenter events, select the primary backend you want to use for discovering the events.

When registration finishes, all of the currently registered collectors are shown, with their supported discoverable device types.

Create a discovery group

A discovery group stores the connection information, credentials, and other configuration information required to discover a group of devices. For example, you might set up discovery groups to store IP address ranges or subnets and appropriate connection credentials.

Before you begin

Before you can create discovery groups and use the discover feature, the collection server associated with the device type must be registered.

- If the **Discovery Groups** and **Discovery Results** tabs are grayed out, the collection server is not registered.
- If the **Discovery Groups** and **Discovery Results** tabs are not shown, discovery is not supported for the device type.

Procedure

1. Go to **Centralized Management > Discovery Center > Inventory Management > *device_type***.
2. Click the **Discovery Group** tab.
3. Click **Add New Discovery Group**.
4. Type a name for the discovery group and click **OK**.
5. Click the discovery group name.

Tables for different types of discovery information appear.

6. Add entries into the tables to provide the discovery information for the group.
-

Note

The discovery group fields are different for each device type. For information about the fields, see the SolutionPack installation documentation.

Use the following buttons to add information into the tables:

Button	Description
Add new entry...	Add discovery information manually.
Import...	Import a file containing discovery information.
Export	Export discovery information to back it up or reuse it for other discovery groups.

Button	Description
Export template	Export a template to your Downloads folder. You can complete the template with discovery information, and then import the file with the Import button.

7. Click **Save**.

You can now choose the group in a discover operation.

Discover devices

Discovery finds devices based on the seed information in a discovery group.

Before you begin

Use the **Discovery Group** tab to create a discovery group or research the settings in discovery groups.

Procedure

1. Go to **Centralized Management > Discovery Center > Inventory Management > *device_type***.
2. On the **Collected Devices** tab, click **Discovery**.
If the **Discover** button is not available, make sure the appropriate collector is registered and that at least one discovery group is defined.
3. Select a **Discovery Group** name.
4. Select a **Discovery Mode**.
 - Use **Full Discovery** the first time you discover a group.
 - Use **Incremental Discovery** to discover a newly added device.
5. For **Automatically distribute results?:**
 - Select this option to distribute the results of discovery to the collector.
 - Do not select this option if you want to review and approve the discovery results before incorporating them into the system.
6. Click **OK**.
If you requested automatic distribution, the results of the discovery are visible on the **Collected Devices** tab, and also on the **Discovery Results** tab.
7. If you did not request automatic distribution, go to the **Discovery Results** tab to review the results and distribute them.

Distribute (import) discovery results

If you did not request the discovery process to distribute results, you can examine and distribute the results on the **Discovery Results** tab.

Procedure

1. Go to **Centralized Management > Discovery Center > Inventory Management > *device_type***.
2. Click the **Discovery Results** tab.
The table shows an overview of discovery requests, by discovery group.
3. To see details about the discovered devices, click a discovery group name.
4. To import the discovered devices into the system (so that monitoring activities can start on them), click **Import to Collected Devices**.

5. Complete the pop-up dialog as follows:

Option	Description
To delete all existing devices of this device type, and add the results of this discover group discovery	Click OK.
To retain existing devices, add the newly discovered devices, and update any existing devices if configuration changes were discovered	Click the Merge checkbox and then click OK.

6. To view the new set of devices being monitored, click the **Collected Devices** tab.

Review your devices and credentials to avoid lockout of devices due to multiple attempts of incorrect credentials. We recommend creating groups in such a way that devices have a minimal set of credentials to be tried against. We recommend using common public-private key pairs for multiple devices.

Groups Management

Groups Management is a data enrichment interface. It provides a quick way to add additional properties to collected records.

The Groups Management UI lets you create new values for a preconfigured property. You also create the match rules (called membership rules in the UI) that filter collected records for the purpose of assigning the new value to the matching records. For example, in the Customer group, you can create a new Customer value and then create the rules defining which collected records should be enriched with the new value.

Predefined groups are installed with SolutionPacks that use them. For example, any of the SolutionPacks that support LUN service level tagging will install the Service Level by LUNs group. If none of the corresponding SolutionPacks are installed, you will not see that group in your Groups Management node.

Each group is defined in a configuration file. These configuration files are used to specify the key-properties and new-properties for the Property-Tagging-Filter, the type of group (flat or hierarchical), and the properties to display in the Groups Management preview table. An advanced user can create a new configuration file to create a new group under the Groups Management node.

Flat and hierarchical groups

A group type is either flat or hierarchical.

Flat groups

In flat groups, collected data can match only one rule and the property being added can have only one value per matched database entry. The Service Levels by LUNs is an example of a flat group. The Service Level for a given LUN can be gold or platinum, but not both. The Service Level property (`psvc1vl1`) can store only one value.

In a flat group type, the group names appear in a simple list. The order of the group names is important. The matching is performed in the order of the groups listed, and the first match wins.

Hierarchical groups

In hierarchical groups, collected data can match one or more rules. As a consequence, the new property can store one or more values. These values are separated by pipes (|).

Platform is an example of a hierarchical group type. A given device can match the rules for All Hosts and also for Unix Hosts. Therefore, the stored property value could be "All Hosts|unix".

Note

A subgroup does not inherit the rules from its parent. The underlying rules must be constructed correctly to support the hierarchy.

The ordering of the rules does not matter in a hierarchical group.

Summary of groups

The Groups Management UI supports a predefined set of group types, each with a predefined purpose. It is important to use and populate each group type according to its intended purpose.

The following table contains a high-level description and intended purpose for each of the group types supported by the Groups Management UI. Each group type tags collected data with one or more additional properties.

Note

Your installation might not contain all of these group types. They are installed by the SolutionPacks that use them.

Table 14 Summary of predefined group types

Group type	Purpose
Business Unit Location Customers	<p>These three properties are used for filtering and reporting based on customer-defined attributes. Many reports contain predefined (but hidden) columns for these properties. Those same reports include predefined group filters for these properties, for easy filtering by selected customer, location, or business unit values.</p> <p>Installations can use these three properties for any desired purpose. For example, to track storage usage by applications, you might populate the Business Unit property with application names, and set up the rules to filter on storage assets used by specific applications.</p>
Platform	<p>This property is used for filtering and reporting based on platform type. Each installed SolutionPack adds to the list of predefined Platform values.</p> <p>Many reports contain the Platform group filter, for easy reporting by selected platforms.</p>
Service Level by LUNs Service Level by File Share Service Level by Bucket	<p>These properties are used for Service Level Agreement (SLA) classifications for global reporting by service levels. In addition, a cost per GB property, required for chargeback reporting, is associated with each service level. Predefined service levels are installed; users can create customized service levels as well, with restrictions.</p>

Table 14 Summary of predefined group types (continued)

Group type	Purpose
Block Chargeback Grouping File Chargeback Grouping	<p>These properties define rules that group storage assets for the purpose of global chargeback reporting. The chargeback reports show costs by service levels within each grouping. Predefined groupings are installed; users can create customized groupings for reporting purposes.</p> <hr/> <p>Note</p> <p>Object chargeback grouping uses the object namespaces and buckets defined on ECS assets.</p> <hr/>
ECS Capacity Rates	<p>This group defines several properties used for billing rates for AWS S3 object storage managed on ECS assets. The billing reports appear under Report Library > Amazon AWS > Operations.</p>

Setting cost basis for chargeback reports

Chargeback reports use a cost per GB configured on the service level.

Procedure

1. Click **Administration > Centralized Management > Groups Management > *service_level_group_type***.
 - For block chargeback, select **Service Level by LUNs**.
 - For file chargeback, select **Service Level by File Share**.
 - For object chargeback, select **Object Service Level**.

2. Click to select a service level name, and then click **Edit**.

3. In **cost per GB**, type the cost basis for this service level.

Cost is numerical with no assumption of any specific currency. Enter an integer or a decimal value. If you enter 3, the cost for 100GB is 300. If you enter .3, the cost for 100GB is 30.

4. Click **Save**.

Overview of the Groups Management UI

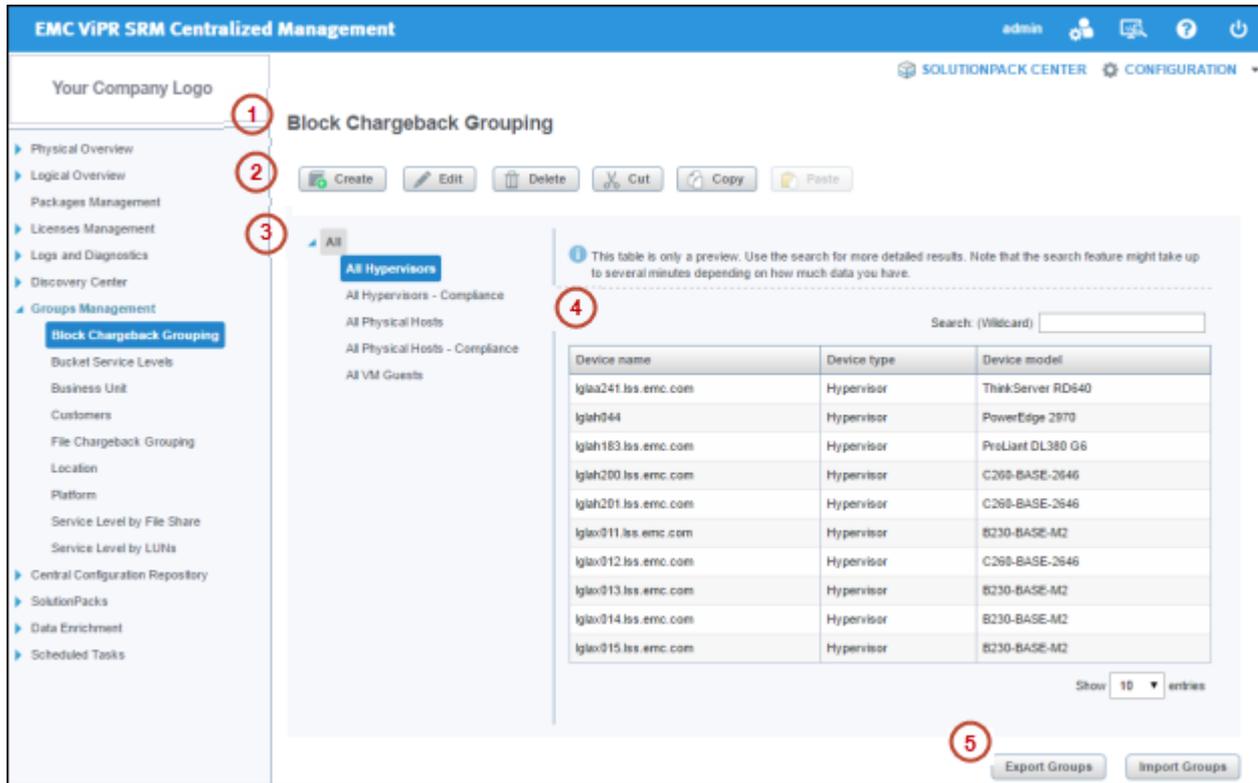
For each group type, you can create new group name values and edit the rules of membership in the groups.

The following image shows the Groups Management UI for the Block Chargeback group type.

Note

The interface is slightly different for hierarchical and flat groups but the concepts are the same.

Figure 1 Groups Management UI



1	Group Type	The group type represents predefined properties in the database. In this case, Block Chargeback Grouping represents the <code>nodegrp</code> property already defined in the database.
2	Action Buttons	The action buttons are used to manage the group of property values. For example: <ul style="list-style-type: none"> • Create creates a new possible value for the <code>nodegrp</code> property. • Edit edits the rules of membership associated with a property value. <p>The set of actions is different for flat and hierarchical group types.</p>
3	Group names	The group names are the possible values of the data enrichment property. In this example, the <code>nodegrp</code> property can have the value All Hypervisors, All Hypervisors - Compliance, or any other value in the hierarchy.
4	Preview table	The preview table lists a partial view of objects in the database that match the rules for a selected value. For example, with All Hypervisors selected, the preview table lists hypervisors but not physical hosts or VM guests.
5	Export/Import actions	These actions manage the underlying CSV file associated with this group type.

Select a group name and then click **Edit** to see the match rules for that group name. For example, select **All Hypervisors** and click **Edit** to see the rules that control when the value All Hypervisors is assigned to the `nodegrp` property in a collected entry.

Figure 2 Groups Management UI—Rules

Use the preview table

The preview table shows the members of a selected group and some additional information about the members.

Procedure

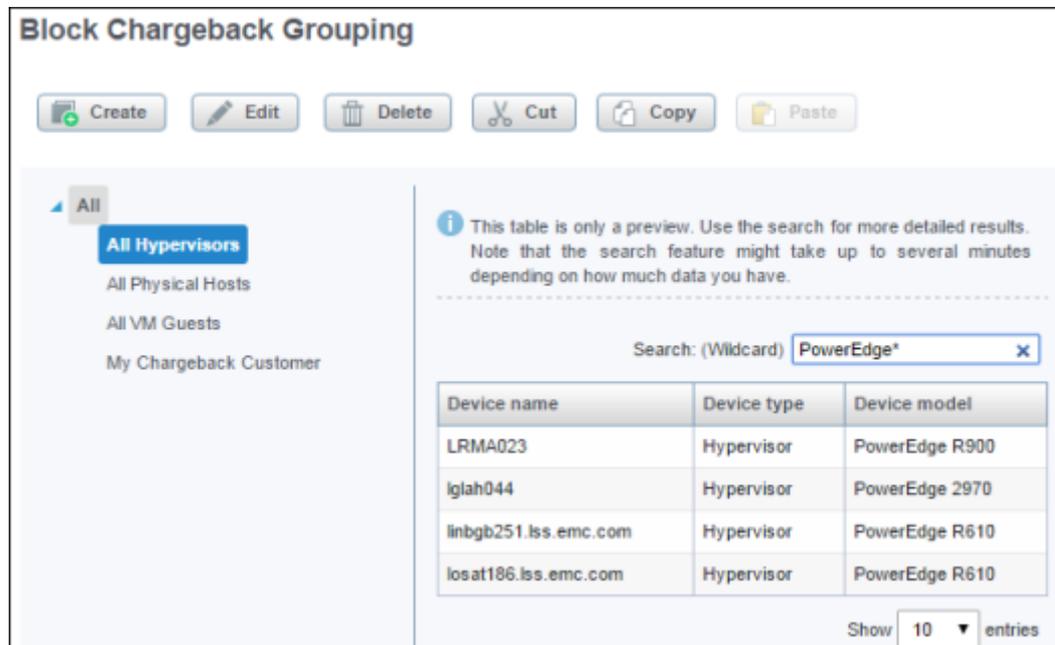
1. Go to **Administration > Centralized Management > Groups Management > *group_type***.

The page opens with the preview table showing the first 10 members of the first group on the page.

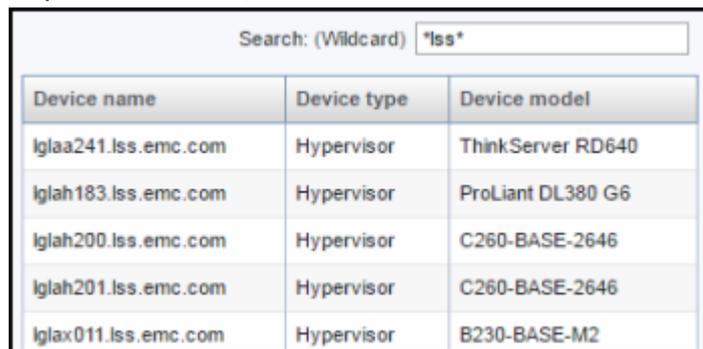
2. To see more members, use the **Show entries** drop-down box to adjust the number of entries displayed.
3. In cases where there are potentially many members, search for specific members using the **Search** text box.

Type an exact value for any of the columns in the table, or use the * wildcard for a partial value search. The search operates on any of the columns in the preview table.

For example, the search in the following figure (`PowerEdge*`) finds the members of the All Hypervisors group that match a device model starting with PowerEdge.



The wildcard can appear anywhere in the search phrase. For example, the search in the following figure (*lss*) finds all entries with the character sequence of lss anywhere in a value.



4. To preview another group, click to select the desired group name.

Create a group

The **Create** action creates a new property value and specifies the match rules for when that value is tagged to collected data.

Following is a general procedure for creating a group. Some group types, because of their intended purpose, have logical restrictions on the membership rules. For these and other limitations, see the individual group descriptions.

Procedure

1. Go to **Administration > Centralized Management > Group Management > group_type**.
2. In a hierarchical group type, click the parent under which you want the new group to appear.

Note

This step does not apply to flat groups.

3. Click **Create**.
4. In the first text field, enter the new group property value.
For example, to create a new Customer value, enter the Customer name. This is the value that will be added to tagged records.
5. If additional text boxes are displayed for additional properties, enter those values.
For example, the service level groups display the Cost per GB field.
6. For flat or hierarchical groups, add dynamic rules to define the records that will be tagged with the new property values.

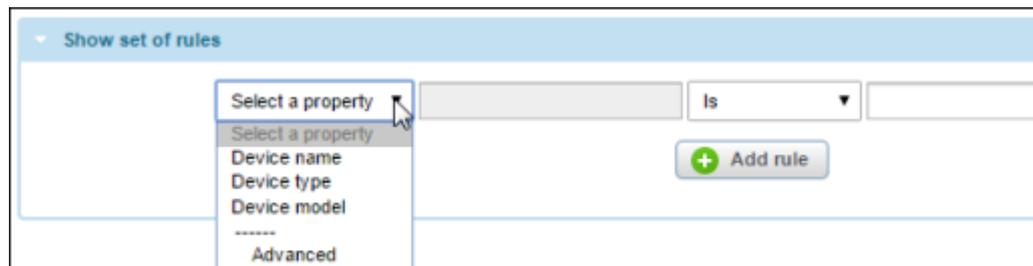
Note

For flat groups, dynamic rules are the only way to add members. For hierarchical groups, the UI shows tabs for **Dynamic Members** and **Static Members**. Click **Dynamic Members**.

A dynamic member definition consists of one rule or multiple rules ANDed together to create a rule set. You can OR multiple rule sets to create a detailed filter.

The following steps provide guidelines for creating dynamic rules.

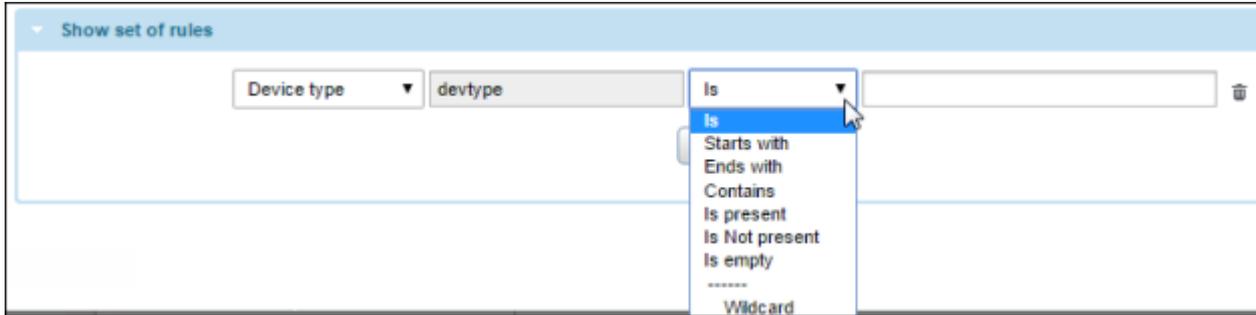
- a. From the drop-down list, choose **Select a property**.



Then do either of the following:

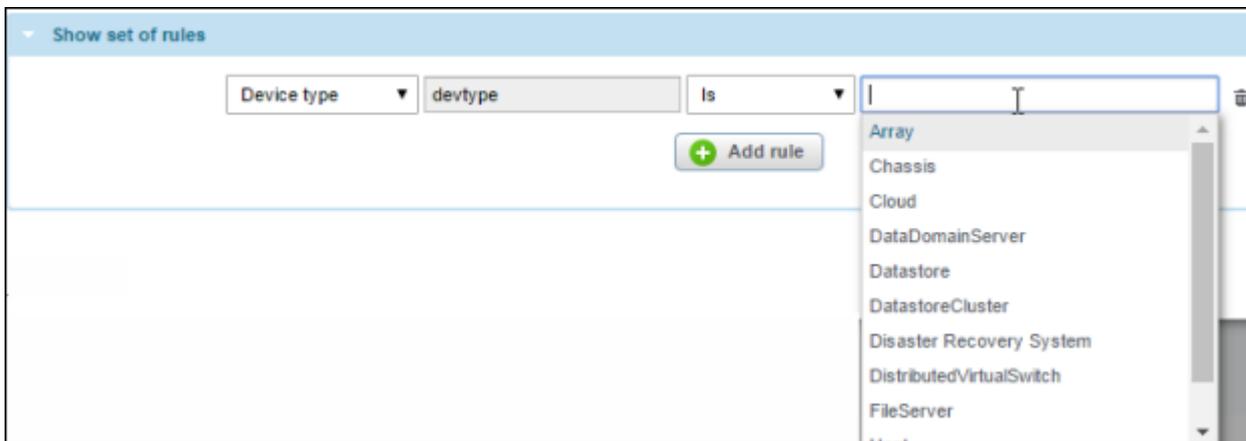
- Select one of the suggested property descriptions. The actual property name is automatically provided in the next text box. For example, if you select **Device type**, the property name `devtype` appears in the text box.
- Select **Advanced**, and then provide a database property name in the next text box.

- b. In the second drop-down list, select an operation. Select **Wildcard** to use wildcards in the value.

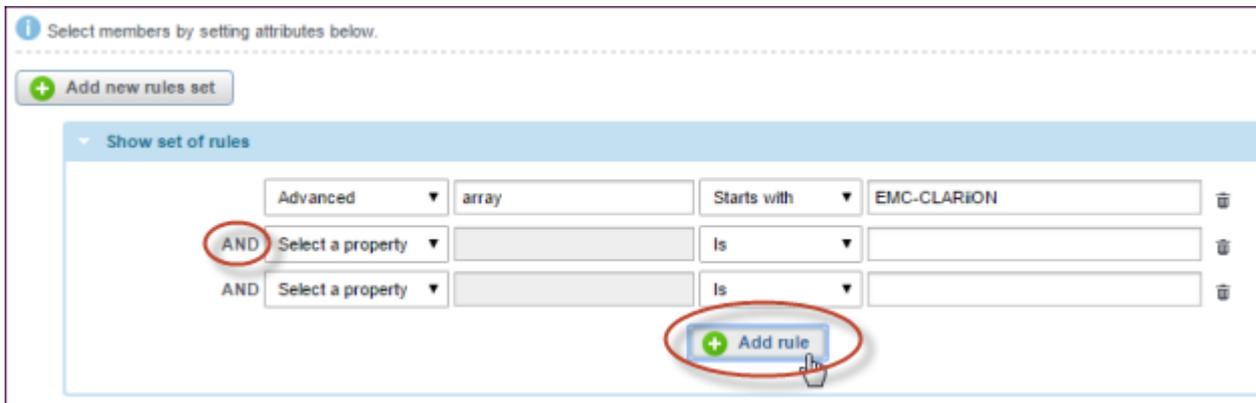


c. In the last text box, provide the property value.

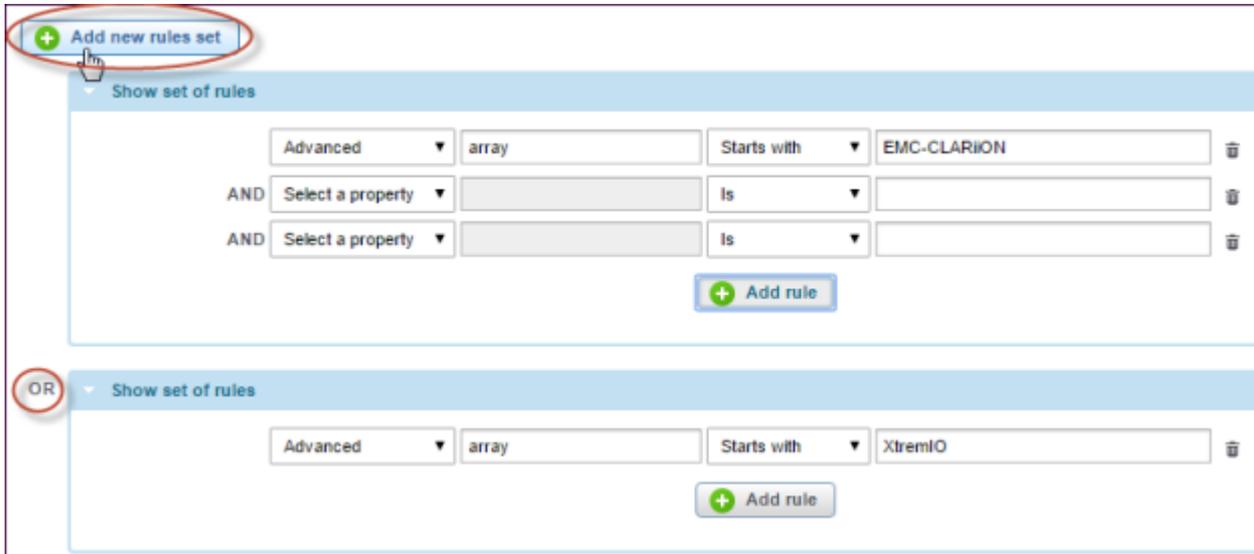
Click in the box to see suggested values based on the property named in the first text box. The suggestions are obtained from existing discovered values in your infrastructure.



d. To AND another rule, click **Add Rule** inside the ruleset box.



e. To OR another ruleset, click **Add new ruleset** outside of the existing ruleset box.



f. To verify defined membership, click **Show Members**.

7. For hierarchical groups, optionally add static members.

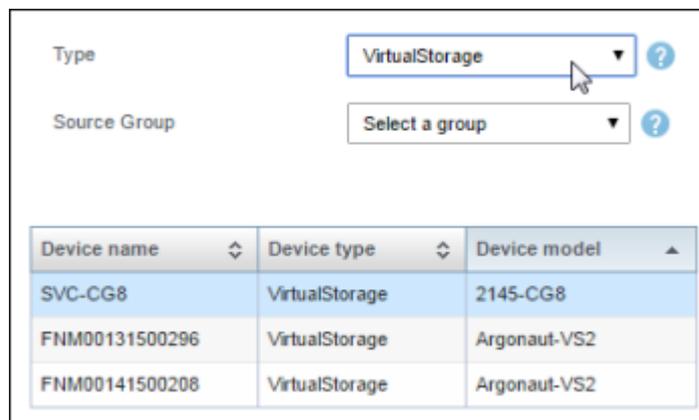
Static members are those that you select from a list and manually add to the membership.

Note

Static and dynamic members are combined to create the total membership for a group.

a. Click the **Static Members** tab.

b. Make selections from the **Type** or **Source Group** drop-down lists to populate the preview table. This creates a subset of items to choose from.



c. To add a static member to the group, select a row in the preview table and click **Add to group**.

The selection moves to the static member list.

8. Click **Save**.

Edit a group

Edit a group to change the membership rules for the group or change the new property values associated with the group.

Procedure

1. Go to **Administration > Centralized Management > Group Management > *group_type***.
2. In the right pane, select the group to edit by clicking it.
3. Click **Edit**.
4. Change the property values in the text boxes at the top of the page, if needed.
5. Change or add to the membership rules, as needed.
6. Click **Save**.

Results

If you change a group name, all new tagging after the change uses the new group name. Regarding existing records in the database:

- if you changed only the group name, the change is applied the next time the PTF process runs, and the old group name is overridden by the new group name.
- If you changed both the group name and the rule definition, you will see both the old and new names in reports for a while, until the old group name ages out of the reports.

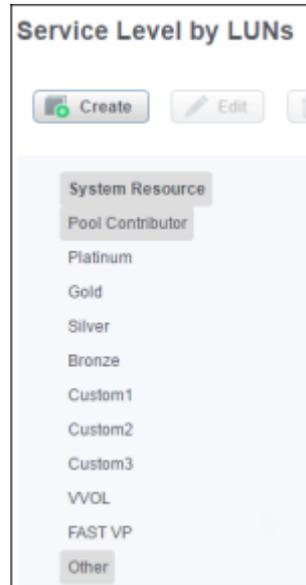
View and reorder rules

For flat group types, the order of the group names affects the order that the rules are processed and the tagging outcome. You can review all rules for all groups, and reorder the groups to change the processing order of the rules.

The **Reorder** and **View All Rules** buttons only appear on the dialogs for flat groups.

The new properties in flat groups can contain only one value. The first group name whose rules match the collected data is assigned. Make sure to order the groups so the most restrictive groups are first.

For example, in the list of service level groups, System Resource and Pool Contributor identify specialized storage uses. After that, Platinum rules are most restrictive. At the bottom of the list, Other is very general and acts as the catchall group when no other rules match.



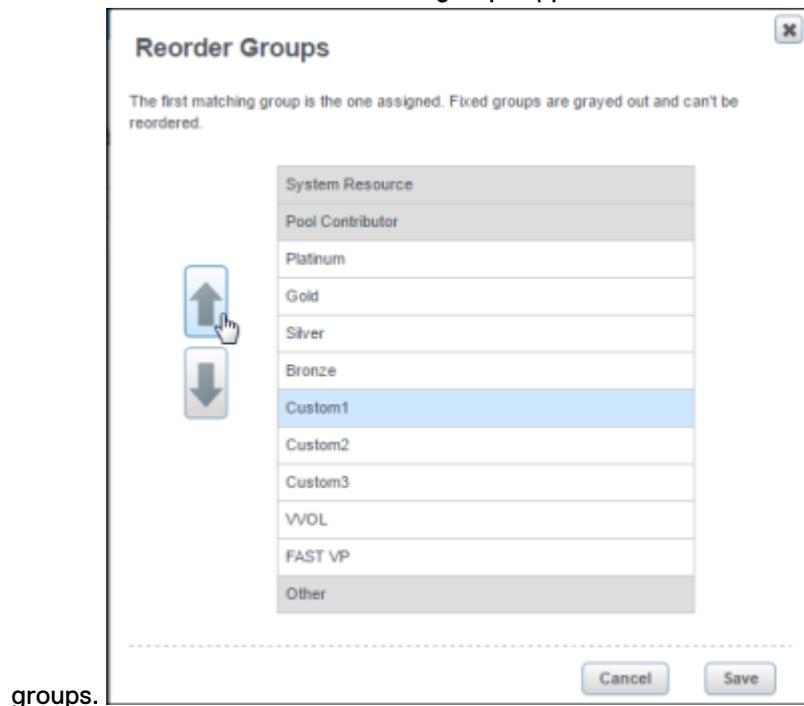
Procedure

1. Go to **Administration > Centralized Management > Groups Management > *group_type***.

This procedure is relevant to flat groups only.

2. To view all rules for all groups, click **View All Rules**.
3. To exit the list of rules, click **OK**.
4. To reorder the rules, click **Reorder**.
5. Click the group whose placement you want to change.
6. Use the arrows to change that group's place in the list.

Make sure that the more restrictive groups appear before the less restrictive



groups.

7. Click **Save**.

Create new group by copying an existing group

In hierarchical group types, you can quickly create a new group with almost the same rules as an existing group by copying and pasting the existing group.

The **Copy** and **Paste** buttons are only available on the dialogs for hierarchical group types.

Procedure

1. Go to **Administration > Centralized Management > Groups Management > *group_type***.
2. Click to select the existing group that you want to copy, and then click **Copy**.
3. Click to select the group in the hierarchy that will be the parent of the new group, and then click **Paste**.

The new group appears under the selected group. It has the same name as the copied group.

4. Click to select the new group, and then click **Edit**.
5. On the **Edit** dialog, change the group name and alter or add to the rules and static member selections, as needed.
6. Click **Save**.

Get group membership rules as regex expression

You can see the complete membership rules for a group as a single regex expression.

Procedure

1. Go to **Administration > Centralized Management > Groups Management > *group_type***.
2. Click to select the group name whose membership rules you want to see.
3. Click **Edit**.
4. Click **Click to show query** at the bottom of the **Edit** page.

Manage groups by exporting and importing files

As an alternative to creating and editing group names in the UI, developers or administrators can manage groups by exporting and importing files.

The Groups Management UI creates and maintains the underlying configuration files that support your group definitions. You can export and save those files to a location outside of ViPR SRM for backup purposes or to view and edit the settings. You can import an edited version into ViPR SRM.

Note

Editing files offline and importing them is an advanced activity. End users are recommended to edit group settings using the **Edit** and **Create** buttons on the Groups Management UI.

Procedure

1. Go to **Administration > Centralized Management > Groups Management > *group_type***.

For example, click the **Service Level by LUNs** group type.

2. On the group type page, click **Export** and follow your browser's prompts to save the zip file.
3. Extract the files with an unzip utility.

Each group type is supported by two files:

XML configuration file

The XML file defines the Property Tagging Filter (PTF) configuration for the group type. It includes:

- The `<files>` element defines the name of the associated data input file (the CSV file), and the characters used in the file for delimiters and wildcards.
- The `<key-properties>` element defines the structure of the rules that are used to match existing data in the database.
- The `<new-properties>` element defines the names of the new properties to add when a match occurs.

Typically, you do not change this file. Use it to interpret the format of the CSV file.

CSV data input file

The CSV file contains the group names and definitions in the acceptable format required by the PTF. Each line represents a group name, with all of the rules that define members of that group and the new property values for that group.

For example, for the **Service Level by LUNs** group type, each line contains the defining rules for membership in a group, and, at the end of each line, the chargeback cost basis, the minimum and maximum response rates, and the service level name.

4. To import files:
 - a. Make sure you have preserved an exported version of the existing settings.

Note

This step is important. You are about to overwrite the existing settings.

- b. Archive your edited CSV file and the XML file into a zip file. Both are required.
- c. Click **Import**, and choose the zip file to upload.
- d. Click **OK**.

Central Configuration Repository

The Central Configuration Repository provides a simple way to edit configuration elements that are shared between several SolutionPacks and SolutionPackBlocks.

The main page of the Central Configuration Repository contains all the editable configuration elements across the registered servers. Each element is described by its name, its parameters, and the SolutionPacks referencing it. Clicking on an item of the list opens the edition page of this particular element.

As the list size quickly increases with the installation of new SolutionPacks, it is recommended to use the search field to easily retrieve a specific configuration element.

The top of the edition page contains the element's parameters while the bottom contains the list of SolutionPackBlocks referencing this element.

Modifying a configuration element requires the following steps:

- Enter new parameters for the element.
- Select the SolutionPackBlocks that will be updated with the new parameters. The SolutionPackBlocks that are not selected will keep the previous configuration.
- Click the **Apply** button.

SolutionPacks

SolutionPacks provide an easy and quick way to add technology-specific dashboards and reports to your system. SolutionPacks are installed from the Centralized Management SolutionPack Center.

In SolutionPack Center, you can browse each of the SolutionPacks to learn more about what each SolutionPack monitors and the kinds of reports and dashboards that are provided with the SolutionPack. Before you can install a SolutionPack, you must have a license.

Installing a SolutionPack

You can install a SolutionPack from Centralized Management.

Before you begin

- Determine whether you need a SolutionPack license file by checking the feature names and expiration dates listed in **Centralized Management > License Management**.
- Core modules must be up-to-date in all servers because not all module dependencies are validated during the installation or update process.

Core modules are modules that are not installed by a SolutionPack. Examples of core modules include webapps, module-manager, license-manager, java, mysql, backend, tools, and so forth. Update these core modules: **Databases, Backends, Frontends, Miscellaneous** (except Blocks).

Procedure

1. Navigate to **Centralized Management**.
2. Select **SOLUTIONPACK CENTER**.
3. Select the SolutionPack in the **Browse and Install SolutionPacks** screen.
4. Read the summary information and click **Install**.
5. Select the components to install.
 - a. Type the instance name.
 - b. Select the server in one or more list boxes. For example, select the server in the **Data collection** and **Reports** list boxes.
 - c. Click **Next**.
6. For each list box you select, a screen appears.
 - a. Click **Next** after you complete each screen.

- b. Click **Install** after you complete the last screen.

The installation process begins.

7. Select the maximize arrow next to each component to view the installation process.

When the installation successfully completes, green checkmarks appear.

After you finish

After the installation is complete, select **Centralized Management > SolutionPack** to verify the installed SolutionPack.

Reconfiguring a SolutionPack

You can change the configuration of a SolutionPack.

Procedure

1. Go to **Administration > Centralized Management > SolutionPacks**.

A table that lists all of your installed SolutionPacks appears.

2. Click the SolutionPack name that you want to reconfigure.

A summary description about the SolutionPack appears, followed by a table of configurable components.

3. Click the **Edit** (pencil) icon for the component you want to reconfigure.
4. Make changes on the dialog that appears, and click **Reconfigure**.
5. Click **OK**.

Installing a SolutionPackBlock

A SolutionPackBlock adds features to the system that are not associated with any particular SolutionPack.

Similar to a SolutionPack, you configure a SolutionPackBlock when you install it. You can reconfigure it later.

For more information on configuration parameters in a SolutionPackBlock, see the *EMC M&R Advanced Administration Guide* and other guides in the EMC M&R portfolio.

Procedure

1. Navigate to **Administration > Centralized Management > SolutionPacks**.

The table shows the installed SolutionPackBlocks and their instance names.

2. Click **Add Component**.

This table shows all available components for installation.

Note

You can install multiple instances of the same component.

3. Click the SolutionPackBlock to install.
4. Work through all of the dialogs, completing each one and clicking **Next**.
The last dialog contains the **Install** button, rather than **Next**.
5. Click **Install**.
6. Click **OK**.

Reconfiguring a SolutionPackBlock

You can change the configuration of a SolutionPackBlock.

Procedure

1. Go to **Administration > Centralized Management > SolutionPacks**.
2. Scroll to the second table on the screen, called **Other Components**.
3. Click the **Edit** (pencil) icon for the component you want to reconfigure.
4. Make changes on the dialog that appears, and click **Reconfigure**.
5. Click **OK**.

Configuring server settings for online updates

Procedure

1. Click **Administration** .
2. Click **Centralized Management**.
3. Click **Configuration > Online Update**.
4. Ensure that you are on the **Settings** tab.
5. Type the username and password that you use for the EMC Online Support website.
6. Test server connectivity by clicking the Test Server Connectivity icon in the upper-right corner of the window.
If the test is successful, the icon becomes a green check mark.
7. Click **Save**.

Enabling the online update task

Enable the Online Update task to download the latest updates from the EMC Update server automatically.

Procedure

1. Click **Administration** .
2. Click **Centralized Management**.
3. On the **Physical Overview** page, click the **<host_name> - Front End** where the Online Update task needs to run.
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.

Disabling the online update task

You can disable the Online Update task if you prefer to manually download the updates from the EMC Update server.

Procedure

1. Click **Administration** .
2. Click **Centralized Management**.
3. On the **Physical Overview** page, click the *<host_name>* - **Front End** where the Online Update task needs to run.
4. Click **Tasks**.
5. Type `onlineUpdate` in the **Search** bar.
6. Click the **OnlineUpdate** scheduled task.
7. Click **Disable**.

Running the online update task manually

At any time, you can run the Online Update task manually to access the available updates.

Procedure

1. Click **Administration** .
2. Click **Centralized Management**.
3. On the **Physical Overview** page, click the *<host_name>* - **Front End** where the Online Update task needs to run.
4. Click **Tasks**.
5. Type `onlineUpdate` in the **Search** bar.
6. Click the **OnlineUpdate** scheduled task.
7. Click **Run Now**.

Troubleshooting

This section provides SolutionPack troubleshooting tips that can be used to resolve common SolutionPack issues.

- [What to do if data does not appear in any reports](#) on page 183
- [What to do if data does not appear in some reports](#) on page 184
- [Viewing collector errors in the Collector-Manager log files](#) on page 184

What to do if data does not appear in any reports

Procedure

1. After the completion of at least three collection cycles, verify if data is populating into the reports. If there is still no data in the reports, continue to the next step.
2. Run the scheduled task to import data into reports. If there is still no data in the reports, continue to the next step.

3. To view the log files for errors, go to **Centralized Management** and click **Logical Overview > Collecting > Collector-Manager::*<instance name>* > Log Files.**

What to do if data does not appear in some reports

Procedure

1. Run the scheduled task to import data into reports. If there is still no data in the reports, continue to step 2.
2. Search for the metric in the database.
3. To view the log files for errors, go to **Centralized Management** and click **Logical Overview > Collecting > Collector-Manager::*<instance name>* > Log Files.**
4. To view the log files for errors, go to **Centralized Management** and click **Events > Event-Processing_Manager::Instance name > Logs.**

To debug the errors, enable the Event-Spy for the Event Processing Manager.

Viewing collector errors in the Collector-Manager log files

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

Procedure

1. Click **Administration** .
2. Click **Centralized Management > Logical Overview.**
3. Expand **Collecting.**
4. Click the Collector-Manager for your collector instance.
Collector-Manager::*<Collector-Manager instance>* - *<host_ID>*
5. Expand **Log Files** and click the **View File** icon to review the error messages.

Data Enrichment

The Data Enrichment page enables you to quickly add tagging to data coming from a collector-manager or an event-processing manager. To access the data enrichment page, select the data enrichment root node from the tree or select the data enrichment node located under the physical node.

The collector manager and/or the event-processing manager must be running to tag data and there must be at least one property tagging filter (for the collector manager) and one event property tagger (for the event processing manager).

The tagging status provides an indication of how the module is running:

- Running - the module is running.
- Module is not running - the module is not running.
- Server not reachable - the target server might be down.
- WS not Supported - the module is running but cannot get a response from the webservice. The module might be too old and does not support the web service.
- Unknown - any other errors.

To add a new property tagging filter or a new event property tagger to the list of tagging, click **Register a new module** and select a server, category, and the modules to register.

The Data Enrichment page cannot be used to modify property tagging filters or event property tagger configurations that contain an accessor.

Values that are not correctly written in the property tagging filter or event property tagger data file appear in red italics. You cannot edit that value, but you can delete the row.

Using the Data Enrichment UI

Use the data enrichment GUI to create the new properties and define the rules for tagging data with the new property values.

Here is the sequence of tasks required to add a property, populate it, verify it, and make it visible in a report.

1. [Register a Collector Manager](#) on page 185
2. [Create new tagging structure](#) on page 187
3. [Populate the tag set rows](#) on page 193
4. [\(Optional\) Run the import property task and restart Collectors](#) on page 196
5. [Verify the new property](#) on page 197
6. [Add the new property to a report](#) on page 197

Register a Collector Manager

To enable a PTF to run, you must associate it, or register it, to a Collector Manager.

To associate a data enrichment PTF to a Collector Manager, you register the Collector Manager.

Typically, a collector server hosts many Collector Managers. Typically, you want all collectors to use the same PTF to consistently tag data and populate your new properties. A best practice is to register a single module, the `Collector-Manager :: Load-Balancer :: DataEnrichment` module, to minimize overhead and ensure that the same set of tagging rules is used on all data processed on a collector server.

If there is a reason to use different sets of rules for different Collector Managers, you can register each Collector Manager separately, and associate different tagging rulesets to each one.

Procedure

1. Log onto the Console and go to **Administration**  > **Centralized Management** > **Data Enrichment**.
2. In the right pane, click **Register a new module**.
3. In the **Server** column, select the server that hosts the data collecting module you want to register.

Typically, select the **Collector** server. However, other server types also host a few collecting modules.

4. In the **Categories** column, select **Collecting**.
5. In the **Modules** column, select one or more modules that you want to participate in the property tagging.

For typical applications, choose `Collector-Manager :: Load-Balancer :: DataEnrichment`.

Register a new module

Select the servers from the left column, then the category, then one or multiple modules from the 'module' column. The modules are represented in the form <module> :: <name> :: <id>. (To register a module, it must be running. Only the unregistered modules are listed.)

Servers

- lglak116.lss.emc.com
- Primary Backend
- lglak117.lss.emc.com
- Additional Backend
- lglak118.lss.emc.com
- Collector

Categories

- Collecting
- Event-Processing

Modules

ServiceLevel

- Collector-Manager :: Load-Balancer :: PTF-

Location

- Collector-Manager :: Load-Balancer :: DataEnrichment
- Collector-Manager :: Load-Balancer ::

User_Defined_Group_Management

- Collector-Manager :: Load-Balancer ::

ObjectServiceLevels-UI

Note

The list shows only running modules. If the desired module is not in the list, exit and check the status of the Collector Manager under **Administration > Centralized Management > Physical Overview > server_name > collector_manager_name**. Use the **Start** or **Restart** buttons on the page if needed.

6. Click **Register**.

The **Data Enrichment** page redisplay, and the newly registered module appears in the table. You can now associate property tagging rules to the module.

Data Enrichment

This is the registered modules and property tagging definitions used for data enrichment across all configured servers.

Showing 1 to 1 of 1 entries Search:

	Server ▲	Category ▲	Module ▲	Status ▾
<input type="checkbox"/>	lglak118.lss.emc.com - Collector	Collecting	Collector-Manager :: Load-Balancer :: DataEnrichment	Running

0 entries selected

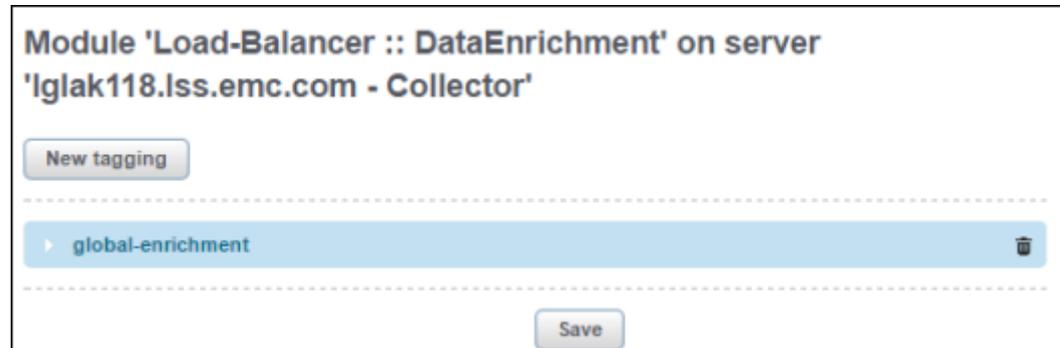
Create new tagging structure

Define the new tagging structure by defining the keys to match against and the new properties to add.

Procedure

1. Go to **Administration**  > **Centralized Management** > **Data Enrichment**.

The table shows the modules that are registered to participate in data enrichment.



2. Click a row for a registered module.

The tagging page appears.

This page contains a blue bar for each defined tag set.

Note

ViPR SRM is installed with a predefined tag set named **global-enrichment**.

3. (Optional) Click the **global-enrichment** entry to expand it.

Examine the structure of this predefined tagging structure as a preview to the next few steps. Notice the following:

- This tagging structure contains one key, `device`.
- It contains four possible properties to add to records when a collected `device` value matches a rule.
- To populate the tagging structure, you would create rows in the table that specify the collected `device` values to match and the resulting four property values to add.
- The bunit, customer, and location properties can also be managed using the Groups Management interface.

4. To proceed with defining a new ruleset structure, click **New tagging**.

A new blue bar appears, with input fields for defining the new tag set.

Module 'emc-watch4net-health :: DataEnrichment' on server 'Iglak118.Iss.emc.com - Collector'

New tagging

global-enrichment

new-tagging*

Name

No entries to show Search:

No records found

Previous Next Show 10 entries

Save and ...

5. For **Name**, create a unique name for your new tag set.
This name becomes the name of the CSV input file containing the rules you define in the tag set.
6. To define a key for the structure, click **Add new key** and complete the dialog.

Create a new key

* New column ?

after ?

Default value ?

Type ?

Delete after use ?

- a. For **New column**, start typing the APG database field name that you want to match on, and select the field from the list.

- b. For **after**, select the position of this new key in the key set. (For the first key, the position must be **first**.)
- c. For **Default value**, optionally enter a default value for this field.

The value must be in the format specified by the **Type** field in the next step. **Default value** is used as follows:

- A first record is created that uses all of the default values.
- If you add a new key into an existing structure that was already populated with records, the default value is used in all of the existing records.

- d. For **Type**, select the type of match value to use for this field when populating the ruleset.

The following table provides guidelines for choosing an appropriate type. See [Type attribute for PTF keys](#) on page 191 for descriptions and syntax of acceptable match values for each **Type**.

Type	Description	Notes								
Range	The match value must be a number range.	Used only for number range matches. Syntax examples for match values are: <ul style="list-style-type: none"> • [1;10] to match 1 through 10 inclusive • [0;100[to match 0 through 99.99999 								
String	The match value must be an unquoted, exact-match, case-sensitive string.	String matches have good performance, supporting long files (>100,000 lines) with minimal impact.								
Regex	The match value must be a Java regular expression.	Try to avoid using this type. Although a regex match is powerful, it is CPU-intensive. We recommend using one of the SQL wildcard matches instead.								
Wildcard	The match value must be a SQL pattern. Case-insensitive.	Wildcard and Wildcard-CI are the same. (Both perform case-insensitive SQL pattern matching.) Wildcards are:								
Wildcard-CS	The match value must be a SQL pattern. Case-sensitive.									
Wildcard-CI	The match value must be a SQL pattern. Case-insensitive.									
		<table border="1"> <thead> <tr> <th>Wildcard</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>% or *</td> <td>Any multiple characters</td> </tr> <tr> <td>_ or ?</td> <td>Any one character</td> </tr> <tr> <td>\</td> <td>Escape character if needed</td> </tr> </tbody> </table>	Wildcard	Description	% or *	Any multiple characters	_ or ?	Any one character	\	Escape character if needed
Wildcard	Description									
% or *	Any multiple characters									
_ or ?	Any one character									
\	Escape character if needed									

Type	Description	Notes
Smarts Wildcard	Used for EMC Service Assurance Suite/SMARTS integrations.	

- e. For **Delete after use**, indicate whether to delete the matched property from the collected data after tagging is complete.



Use this option with extreme caution. Deleting fields can break your data model. Usually leave unchecked.

The **Delete after use** option is useful if your intention is to replace a collected property and value with a new property and value.

- f. Click **Save**.
7. Optionally repeat the previous step to add additional keys to the structure.
8. To define a new property, click **Add new property** and complete the dialog.

Create a new property ✕

* New column ?

after first ▼ ?

Default value ?

- a. For **New column**, type a name for the new database property.

Note

APG database property names are limited to 8 characters. If you use more than 8, the property name is truncated to 8.

Note

If you reuse a name that already exists in the APG database, your rules will populate the existing property. It is generally recommended to create a new property.

- b. For **after**, select the position of this new property in this tag set. (For the first property, the position is always **first**.)
- c. For **Default value**, optionally enter a default value for this field.

It is used as follows:

- A first record is created that uses all of the default values.

- If you add a new property into an existing structure that was already populated with records, the default value is used in all of the existing records.

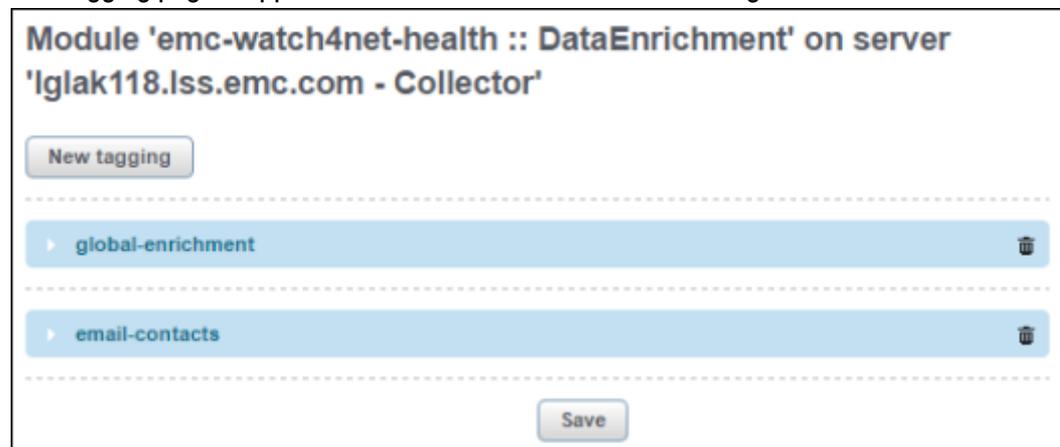
d. Click **Save**.

9. Optionally repeat the previous step to add additional new properties to the tag set structure.
10. After creating all keys and all properties, click **Save**.
11. On the **Save Data Enrichment** dialog, select the modules that should use this tagging structure.

If you have multiple registered modules, select all that apply.

12. Click **Update**.

The tagging page reappears with a new blue bar for the new tag set.

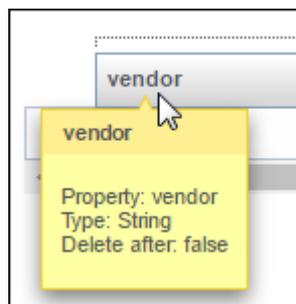


Type attribute for PTF keys

PTF keys are defined using a Type attribute. The Type defines the format of the match values.

View the defined Type for a key

You can view the Type that was defined for a key by hovering your cursor over the column header until a tooltip appears for that column. The following image shows that the vendor key was defined with a Type of string. This means that the values you provide in that column must be strings, and cannot include wildcards.



Type attribute values

The following list defines each Type value.

String

An unquoted string, where the string must match exactly.

Wildcards

Wildcard characters are accepted in a match value. There are multiple types that use wildcards:

Type	Description
Wildcard	Case insensitive. Same as Wildcard-CI .
Wildcard-CS	Case sensitive. For example, z does not match z.
Wildcard-CI	Case insensitive. For example, z matches z or z.

The following values are supported in the match values for keys with wildcard types:

Wildcard	Description
% or *	Any multiple characters
_ or ?	Any one character
\	Escape character if needed

Range

A range of numerical values is accepted as the match value. Decimal values are permitted. The following rules apply:

- Specify the beginning and ending integer or decimal values
- Separate the two values with a semi-colon
- Enclose the phrase with brackets, as follows:
 - [] enclosing brackets indicate inclusive
 -] [non-enclosing brackets indicate exclusive
 - combinations are permitted

Range value examples:

Value	Description
[5;20]	Matches any value 5 through 20, including 5 and 20.
]5;21[Matches any value greater than 5 and less than 21.
[5;20[Matches any value 5 up to but less than 20. 19.5 would match.
[.5;1]	Matches any value .5 to 1, inclusive.

Regex

Any Java regular expression is supported as the match value.

@ values

The following special values are available for keys:

Value in the GUI	Value in the XML/CSV files	Description
@DEFAULT	"**"	Matches all values.
@EMPTY	"" (no space)	Matches a property that exists but is empty.
@NULL	"@@"	Matches when the property does not exist.
@MATCHALL	"%%"	Always matches, including when the property does not exist.

Populate the tag set rows

Populate the tag set with the key values to match and new property values to add.

There are several ways to populate the rows:

- Add rows manually using the Console interface. This method creates a CSV file using the name of the tagging set.
- Import a CSV or XSL file

Note

Importing a file overwrites all entries currently in the table.

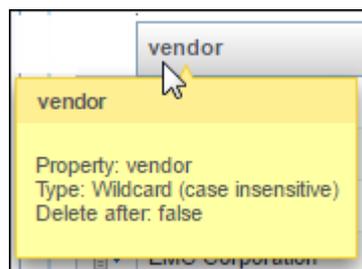
- Add a few rows using the Console, and then download the file, edit it to add more rows, and upload it.
- Add a few rows using the Console, and then open the file in the Console, edit it, and save it.

Key values must match the Type attribute that was configured for the key or use one of the special @ values.

Procedure

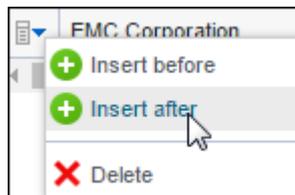
1. Go to **Administration** > **Centralized Management** > **Data Enrichment**.
2. Click a registered module associated with the tagging set that you want to populate.
A blue bar appears for each tagging set.
3. Click to expand the tagging set that you want to populate.
4. To enter the first row:
 - a. Click in the first key field and type or select a match value.
The value for a key field must be either:
 - A value that conforms to the key's Type attribute (Range, String, Regex, or Wildcard).
 - One of the @value selections.

To view the Type attribute for a key, hover the cursor over the column heading until the tooltip appears.



In the example above, a valid wildcard match value is `lgl1*` to match all device names that start with LGL or lgl.

- b. Click in each key field and provide a match value.
 - c. Click in each property (result) field and provide the desired value for the new properties.
 - d. Press **Enter** to commit a row.
5. To enter subsequent rows, right-click the list icon at the beginning of a row, and select either **Insert before** or **Insert after** to place the new row in the desired spot.



Rows are processed in the order in which they appear in this table (or the file). Row order can affect results, especially when you are using wildcards, regular expressions, or ranges. A subsequent row can overwrite a match by a previous row. Typically, you would start with the most general rule, such as a match on `@DEFAULT` or `@MATCHALL`, and end with more specific rules.

6. Click **Save**.
This command saves the table contents to a file. The filename is the tag set name with an extension of `.csv`.
7. On the **Save Data Enrichment** dialog, select all of the registered modules that the file applies to, and click **Update**.
The file is saved on each registered server, under `Collecting/Property-Tagging-Filter/module-name/conf/tagset_name`.
8. To export a tag set file, click **Export CSV** or **Export XSL**.
9. To import tagging set values from an external file:

Note

Importing a file overwrites all entries currently in the table.

- a. Click **Import CSV** or **Import XSL**.
- b. Follow your browser prompts to import the file.
- c. Click **Save** after importing the file.
- d. Select all of the registered modules that the file applies to, and then click **Update**.

View and edit the PTF input file

After creating a few rows using the Console interface, you can view the file to determine the required format, and then manually add additional rows or import a file in the proper format.

There are two ways to view and edit a saved PTF input file.

- Export the file, edit it in a text editor, and then import it.

Note

Importing a file overwrites all entries currently in the table.

- Open the file in the Console's text editor.

Either way, you will notice that the @*values* and wildcards used in your Console entry are translated into standardized wildcard values.

Procedure

1. To export and import a PTF input file:
 - a. Go to **Administration > Centralized Management > Data Enrichment** .
 - b. Click a registered modules that you associated to the tag set.
 - c. Expand the tag set name.
 - d. Click **Export** or **Import**.
2. To open the file on the Console:
 - a. Go to **Administration > Centralized Management > Physical Overview > *registered-server-name* > Modules > Collecting > Property-Tagging-Filter :: *registered-module-name***.

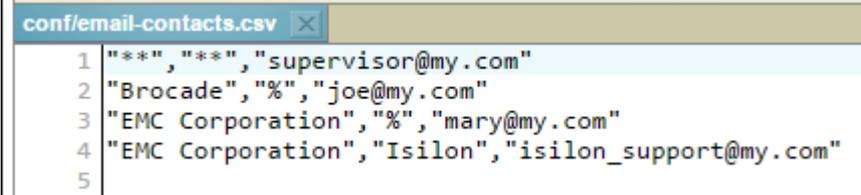
The *registered-server-name* is typically a **Collector** server. The *registered-module-name* is typically **Load-Balancer**.

- b. In the right pane, expand **Configuration Files**.
- c. Locate the `conf/tag-set-name.csv` file.
- d. Click **Edit** (pencil icon) next to the file.

The file opens in the Console file editor.

3. Add additional rows using the format of the existing file.

For example, here is the `conf/email-contacts.csv` file.



```

conf/email-contacts.csv
1 "***", "***", "supervisor@my.com"
2 "Brocade", "%", "joe@my.com"
3 "EMC Corporation", "%", "mary@my.com"
4 "EMC Corporation", "Isilon", "isilon_support@my.com"
5

```

Notice that @**DEFAULT** values entered on the Console were translated into "***". If you add rows, use the "***" notation, not the @**DEFAULT** notation.

4. After editing or importing an edited input file, apply the updates to all registered modules that use the tag set.

Note

If there is only one registered module that uses the tag set, this step is not needed.

- a. Go to **Administration > Centralized Management > Data Enrichment** .
- b. Select all modules that use the tag set associated with the edited file.
- c. Click **Update**.

Edit a ruleset

You can change the structure of the ruleset by adding or deleting keys and properties. You can edit the values in the rows,

Procedure

1. Logon to the Console and go to **Administration > Centralized Management > Data Enrichment**.
2. Click a registered module associated with the tagging ruleset that you want to edit.
A blue bar appears for each tagging ruleset.
3. Click to expand the tagging ruleset you want to populate.
4. To delete a key or a property from the structure, click **Delete** (trashcan icon) in the column header.
This action removes the column and all values.
5. To add a row, right-click the list icon in a row and select **Insert Before** or **Insert After**.
For entry instructions, see [Populate the tag set rows](#) on page 193.
6. To delete a row, right-click the list icon in the row and select **Delete**.
7. To edit a value in a row, click on the value, provide a new value, and press **Enter** to commit the change
8. Click **Save** to save the table, and then **Update** to apply the changes to selected registered modules.

(Optional) Run the import property task and restart Collectors

Data enrichment processing depends on the configured intervals for the database property import task and the Collectors. For quicker data enrichment results, you can run these manually.

It might take a day for all involved processes to run. The appearance of new properties depends on the following:

- The database import-properties task must run. This task updates the in-memory database cache on the Frontend.
- The collecting process must execute the Property-Tagging-Filters. For accurate tagging results, sometimes two collection cycles are required.

If needed, you can run these processes manually.

Procedure

1. To run the database property import task manually:
 - a. Under **Physical Overview**, go to *server_name* - Front End > **Tasks > Database > import-properties-Default** .
 - b. In the right pane, click **Run Now**.
2. To restart a Collector-Manager:
 - a. Click **Administration > Centralized Management > Physical Overview > collector_server_name** .
 - b. With the desired collector server selected, click the **Services** tab.
 - c. Select a Collector-Manager, and click **Restart**.

A Collector-Manager calls all of the individual collecting components in its configuration, such as the stream and XML collectors and the PTFs.

Verify the new property

After the new property has propagated through the system and the new PTF tag set is applied to collected data, you can see the property using Frontend tools.

Procedure

1. In the User interface, click **Administration > Advanced Search**.
2. In the **Expansion** field, click the icon for the **Property selection helper** .
3. Click the APG tab, and then use the search field at the bottom of the dialog to locate the new property.

When searching, remember that a long property name was truncated to the first 8 characters.

If the new property does not appear in the list, wait an hour and try again.

4. If the new property appears in the list, select it and exit the property selection helper.

The new property name should appear in the **Expansion** field. Delete any other property names.

5. Click **Apply**.

The search result lists the values for the property that have so far been applied to records in the database. It can take up to a day for all collectors to run, depending on settings at your site.

Add the new property to a report

To make the new property visible, add it to a report.

If you choose a report that includes the keys on which the tagging was based, you do not need to redefine the report filter. The following example adds a column to a table report, assuming that the existing filter captures records that include the new property.

Procedure

1. Navigate to the table report and click **Modifications**  > **Edit Reports**.
2. Click the **Report Details: Table** tab.
3. Click **+ Property**.
4. Scroll to the end of the property list (blue bars) and expand the **Property: *** bar.
5. Configure the new column:
 - a. For **Column Name**, type the column header.
 - b. (Optional) For **Description**, type the description that will appear in the tooltip if the user hovers the cursor over this column header.
 - c. For **Property**, type the new property name.
 - d. Click **Save**.
6. Click **BROWSE MODE** to return to the report.

Enabling Impact Analysis

Enable Impact Analysis to receive alerts that provide details about the effects of an event and to populate impact details on reports.

Before you begin

Impact Analysis uses the event change listener in the SolutionPack for Storage Compliance. The SolutionPack for Storage Compliance must be installed and configured before you can enable Impact Analysis.

Impact Analysis gives users additional information about effects of an event in the SAN Environment. For example, if a switch port goes down many paths from host to lun can be affected and become unavailable. Impact Analysis displays alerts for these paths. Similarly, if a disk fails or faults, the user receives alerts for all luns that are carved from the disk.

Procedure

1. Navigate to **Centralized Management**.
2. Click **SolutionPacks** in the tree.
3. Under **Other Components**, click the pencil icon for **Alert Consolidation**.
The **SolutionPack Reconfiguration Window** appears.
4. Click the **Enable Impact Analysis** checkbox.
To disable Impact Analysis, clear the checkbox.
5. Type the hostname or ip address of the host of the compliance change event listener. This is typically the Compliance Backend host.
6. Type the change event listener port number, typically 2060.
7. Click **Reconfigure**.
If the operation succeeds, an **OK** button is enabled.
8. Click **OK**.

Alerting

The Alerting module provides access to alerting administration and development features that are performed using the Web Portal.

This information is general to the EMC M&R core. For more information about alerting, see:

- Product documentation—Alerting implementations are product-specific. Your product documentation provides relevant alerting implementation details. See [Information you can find here](#) on page 10 for links to product documentation.
- Backend documentation—The Alerting features are flexible and offer many opportunities for customization. For information about the alerting infrastructure and supporting configuration files, see the *EMC M&R Advanced Administration Guide*, available on <https://support.emc.com>.

Alerting reports

Users access the alerting reports on the User Interface, in Browse mode.

Change the refresh rate on alerting reports

Alerting reports are installed with preconfigured automatic refresh interval set to 60 seconds. You can change that value. To disable automatic refresh, change the setting to 0.

At installations with a small number of alerts, you can reduce the refresh interval to as low as 15 seconds.

For larger events databases (with a total number of active alerts more than a few thousand), it may take a while for the alerting reports to load. In addition, if those reports have the automatic refresh interval set too low, the alerting reports become unusable. You can increase the refresh interval or disable automatic refresh.

If automatic refresh is disabled, you can always refresh manually with the browser's refresh feature to update the report as needed.

Note

New alerts that originate from SNMP can take up to four minutes to appear in the alerting reports. This delay is independent of the refresh interval setting.

Procedure

1. Go to the report whose refresh rate you want to change, and click **Modifications**  > **Edit Reports**.
2. Click the **Report Details: Table** tab.
3. Scroll to the bottom, and expand **Display Options**.
4. Change the value in **Refresh interval (secs)**.
To disable automatic refresh, enter 0.
5. Click **Save**.
6. Click **BROWSE MODE** to return to the report display.

Alerting infrastructure

The alerting module consists of various components installed on the Alerting Frontend and the Alerting Backend and a separate database just for alerts.

Alerting Frontend

In most installations, the Alerting Frontend is the Primary Frontend. The components on the Alerting Frontend are the user-facing components:

- Alerting reports show information from the database.
- Alert Definitions define all of the conditions that cause events to be written to the database. Alert definitions also define other actions that can occur, such as an email notification, an entry to a log file, or a trap to a third party application.
- Alert Templates are useful for consistency in creating alert definitions.
- Alert Adapters and Grouped Boxes customize the alerting module and are mostly for developer use.

Alerting Backend

In most installations, the Alerting Backend is the Primary Backend. The components on the Alerting Backend include all of the supporting processes for consolidating events from various sources into a single format and for writing the consolidated alerts to the Events database.

Alerting (Events) database

The Events database contains two tables.

genericevents_live table

This table stores all ACTIVE alerts. The **All Alerts** report on the Console, and all of the reports that are filtered versions of it, are based on the entries in this table.

An ACTIVE alert is one whose *Active* property has the value 1.

genericevents_archive table

This table stores INACTIVE alerts until they are purged from the system.

An INACTIVE alert is one whose *Active* property has the value 0.

Access the Alerting portal page

The alerting page lets you create, edit, enable, and disable alert definitions. It also accesses the templates, adapters, and grouped boxes nodes.

The URL to the Alerting page is `<Tomcat server address>/alerting-frontend`. For example:

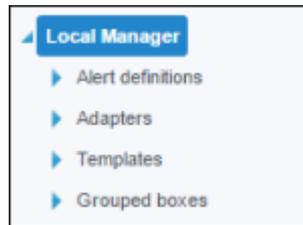
```
http://localhost:58080/alerting-frontend
```

From the user interface, use these instructions:

Procedure

1. If you are in the reporting interface, click **Administration** .
2. Click **Modules > Alerting**.

The first level nodes in the resulting tree are the Managers configured on the Frontend.



Alerting node tools

The tools above the alerting node tree manage the node hierarchy.

The tools can be unavailable depending on the node you select.

	Creates a new element under a list node and opens the element edit page. If you select another node before saving the new one, the new element is deleted.
	Copies the element.
	Pastes the copied node if it is a list that contains the same type of elements.
	Exports the element to an XML file.

	Imports an XML file to the Manager
	Deletes the element.

Alert definitions

An alert definition defines the actions that should occur when an event meets defined conditions.

For an event to result in an action, such as an email notification or getting added to the database, it must match a valid and ENABLED alert definition.

SolutionPacks install predefined alert definitions appropriate to their function and devices. You can create custom alert definitions.

An alert definition contains the following components:

- An entry point filter defines which events are handled by this alert definition. The filter operates on the fields (properties) in the event data. The filter identifies events by matching values in data fields, such as a MIB name and field, or a collector name and metric, or an event name and device type.
- Operations and conditions are optional components in an alert definition. Operations provide a way to manipulate the metric data in an event. Conditions test the data and provide alternate outcomes for different actions.
- Actions define what should occur as a result of the event, such as an email, logging file entry, SNMP trap to an Events database, or traps to other applications. Only events that are written to the Events database appear in the alerting reports.

Access alert definitions

You can configure and enable or disable alert definitions. You can also view a graphical representation of the alert definition, modify it (such as add or remove actions or change the filter), and add new alert definitions.

Procedure

1. Click **Administration > Modules > Alerting > Local Manager > Alert definitions**.
2. Expand the alert definition folders until the alert definition name appears.

In the right pane, the **Type** column indicates whether the entries are folders or alert definition names.

	Alert definition folder
	Contextualized alert definition
	Non-contextualized alert definition

3. In the table in the right pane, right-click an alert definition, and choose one of the following actions.

Edit	Provides a graphical view of the components in the alert definition, with access to each component. You can add,
-------------	--

	delete, and change the configuration of components. Use this option to add or remove components, such as actions or comparators.
Configure	Provides access to the configuration dialogs for each component in the alert definition. Use this option to change the configuration of a component, such as email recipients on Mail actions or threshold values in comparators.
Probe	Unavailable. Probes are implemented as a Test button when you configure an alert definition.
Enable/Disable	Activates or inactivates the alert definition.
Delete	Removes the alert definition from the system.

Note

If you click (as opposed to right-click) an alert definition, the UI invokes either the **Edit** or **Configure** action, depending on the last choice you made using the right-click menu.

4. To set the default action (**Configure** or **Edit**) for your current session when you click an alert definition, select the desired option one time on the right-click menu.

About configuring alert definitions

An alert definition defines the actions that should occur when an event meets defined conditions. The conditions are sometimes configurable.

The Alerting Frontend interface presents the configurable parts of an alert definition as fields that you can easily configure. For example, here is the form for configuring utilization thresholds for an alert related to switch ports.

Configurations of alert definition template Brocade FC Switch Alert Definitions/Threshold/Brocade-Switch-Port/Link Utilization

This page allows you to dynamically create versions of this alert definition by specifying some specific parameters. Those configurations will be preserved after an update of the alert definition.

+ Add

▼ Default ✔
Enabled ▼ ⚙ 🗑

Critical Threshold >= ▼

Minor Threshold >= ▼

Save

An alert definition must be ENABLED for the specified actions to occur.

Components in an alert definition

An alert definition contains at least one entry point filter and one action. It can contain multiple entry points and actions, and optional operations and conditions that manipulate and evaluate data elements in the event data.

Entry point

The entry point is the filter that defines which events are handled by this alert definition. The filter operates on the fields (properties) in the event data.

Note

At this point in the alerting process, the event data is normalized into a standard format, but the event is not yet in any database. Therefore, the filter wizard might not be able to prompt you with all valid values. To write a filter, you might need to research the source of the event and its normalization rules.

Here is an example entry point filter for the **Backend Processing Delay** alert definition in the **EMC M&R Health** folder. The source for this alert is the metric called **ProcessingDelay** collected by the **APG_Health** collector.

Figure 3 Filter in an alert definition



An alert filter typically includes a test for the source of the event. The following table shows some often used field names for filtering on event sources in ViPR SRM.

Table 15 Field names for filtering on event sources

Event type	Event source filter	Example
metric-based thresholds from collectors	<code>source=='collector_name'</code>	<code>source=='APG_HEALTH'</code>
events from the event listener	<code>source=='name_GenericEvent'</code>	<code>Source=='VMAX-GenericEvent'</code>
events from SNMP traps	<code>Source='vendor_in_MIB%' & (sourcedomainname==MIB_data_field</code>	<code>Source='Cisco%' & (sourcedomainname=='cieLinkDown'</code> <code> </code>

Table 15 Field names for filtering on event sources (continued)

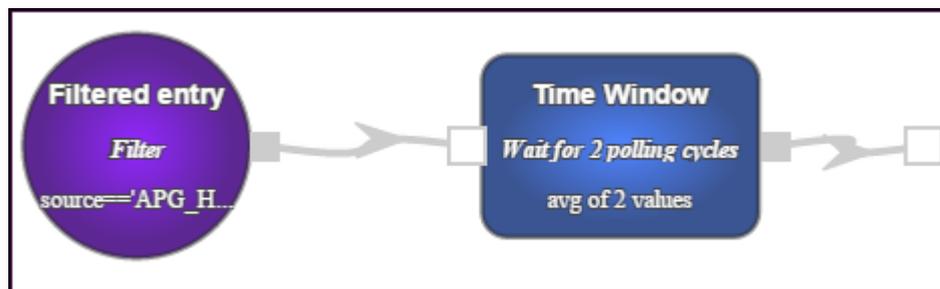
Event type	Event source filter	Example
		<pre>sourcedomainname== 'cieLinkUp')</pre>
events from scheduled reports	source=='report_name'	reportName=='apg report'

Operations

Operations evaluate incoming data to see if they should go to the next component or be dropped. A large set of operations are available, including case, arithmetic, aggregation, time window, date, and so on.

Here is a Time Window operation that calculates an average.

Figure 4 Operator in an alert definition

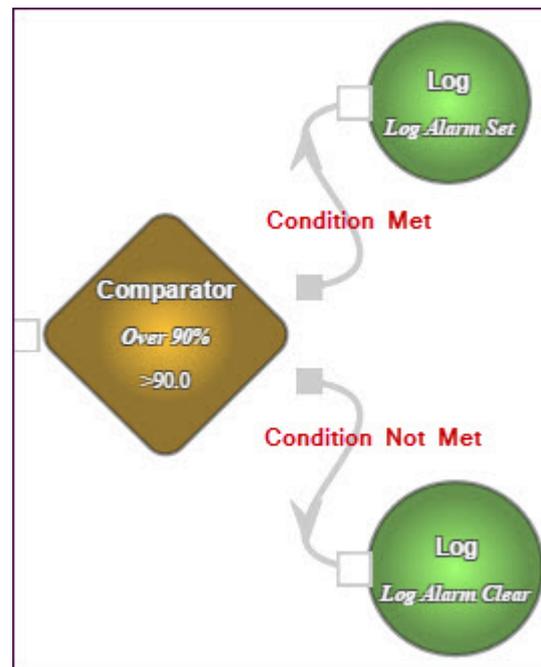


Conditions

Conditions let you set up multiple outcomes in the same alert definition. Each outcome is based on the results from a condition, and would typically have different actions associated with them. Two types of conditions are available.

- Comparators let you define value tests. A common technique in ViPR SRM is a chain of comparators, each one using a different comparator value. The outcomes are a series of SNMP traps creating Critical, Major, Minor, and finally a Clear Trap, all on the same property, and in the same alert definition.
- Counters let you count occurrences within a specified time range

Here is an example comparator that uses Log actions as outcomes. Events processed by this alert definition do not make it into the Events database because the Alert Trap action is missing.

Figure 5 Comparator with two outcome paths**Actions**

The following actions are available:

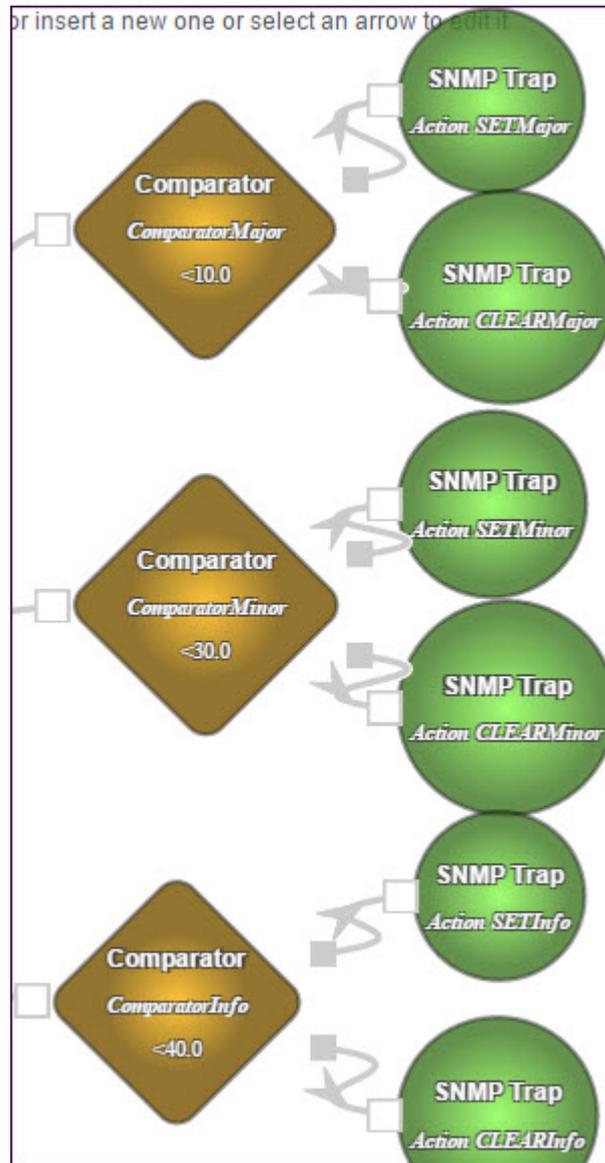
- An SNMP Alert Trap or Clear Trap, which gets the event into the alerting database or subsequently clears it.
- A Log action, which writes the event to a log file.
- A Mail notification, which sends an email to preconfigured email addresses.
- An SNMP trap to a third party, such as another management system.

Note

- Do not add the ESRS Event action to any alert definition. The ESRS Event action is for EMC use only.
- The Send to SAM action is for use by customers who have EMC Service Assurance Suite installed.

Here is a set of comparators, each one associated with two outcome paths. The Alert Trap contents in each action are very similar. The difference would be the severity that is set in the trap contents.

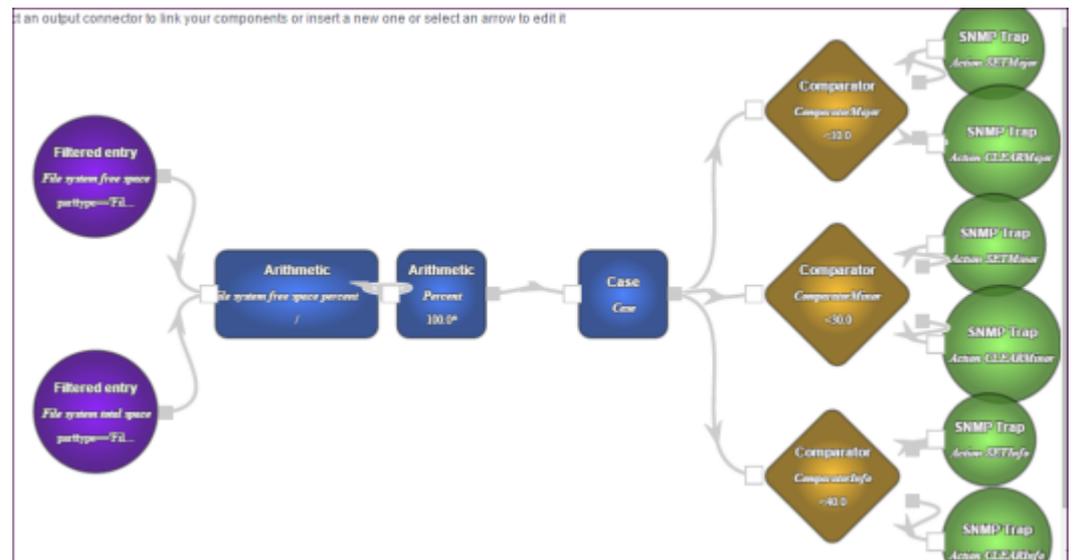
Figure 6 Actions in an alert definition



Example: Putting it all together

Here is an alert definition that contains two entry points to define two data fields, and then performs several operations to derive a new value. There are three conditional tests on that value, and the outcome actions are a series of Alert Traps with different severities and corresponding Clear Traps.

Figure 7 Example alert definition



This is the **Percentage of File System Free Space Alert** in the **EMC M&R Health** folder, in case you want to explore how all of the components are configured.

Contextualized and non-contextualized alert definitions

Contextualized alert definitions allow for more than one set of configuration values. Different named contexts can be configured differently.

Contextualized alert definitions provide the following advantages:

- If future SolutionPack upgrades change the alert definition, your customized parameter settings are not affected because the contexts are saved separately from the alert definition.
- You can define several contexts, or sets of values, for the same parameter set, and enable/disable the contexts separately. For example, you might have different settings for production and testing, or you might preserve the installed values in the default context and create a new context for your customized values.

When you configure a contextualized alert definition, you select the context that you are providing values for. You can create and name new contexts.

Non-contextualized alert definitions can have only one set of configuration values.

Contextualizing an alert definition

You can add contextualization to components in an alert definition and set the values for the default context.

This procedure describes how to add contextualization to a component in an alert definition. This is a one-time process that sets up a component and parameter for contextualization. Thereafter, users can add additional contexts by right-clicking the alert definition and selecting **Configure**.

Many alert definition components, including all threshold Comparators, are contextualized out of the box. Administrators can use this procedure to contextualize those that are not.

Note

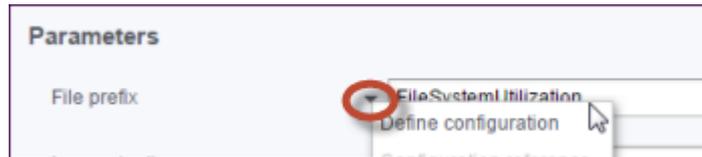
The Action components, such as traps, logging, and email notifications, do not need to be contextualized. Those configurations are not overwritten.

Procedure

1. Right-click the alert definition and choose **Edit**.
2. Double-click the component in the alert definition that you want to contextualize.

A configuration dialog appears for the component.

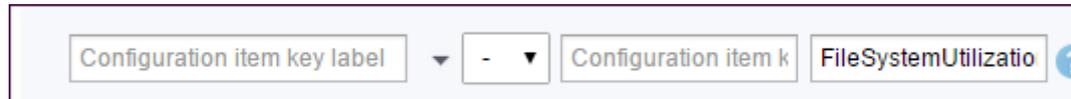
3. Click the small down arrow next to the configurable element, such as a parameter or filter, and choose **Define configuration**.



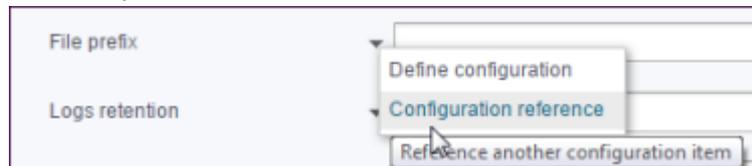
Here is the arrow for a filter:



A set of context configuration fields appears.



4. Complete the fields:
 - For **Configuration item key label**, type a name for this parameter that the user sees on the configuration window.
 - For ordering, optionally select the order in which the parameters appear on the configuration window.
 - For **Configuration item key**, type a key used to uniquely identify the context and reference the context in other contexts within the same alerting definition. For example, an alert definition might have three log actions for different conditions. By using a key, you can configure the log action component one time and then reference it in the other components by choosing **Configuration Reference** when you configure the second and third components.



- For the last field, type the default configuration that appears in the field when the user adds a context.
5. Click **OK**.
 6. Repeat steps 3 and 4 for additional parameters in the component.
 7. Click **Save** below the alert definition whiteboard.

8. As a best practice, add the first context.

Otherwise, the component remains unconfigured.

- a. Return to the alert definition list, right-click the alert definition, and choose **Configure**.
- b. Click **Add**.
- c. Provide a name for the context. A best practice is to name the first context **Default**.
- d. Click **Save**.

Results

The alert now appears in the alert definition list with an icon indicating that it is contextualized. Users can configure it to add additional contexts, and enable or disable the contexts.

Configuring an alert definition

The Configure action lets you change the filter and configurable settings in an alert definition, add new contexts and configure them if appropriate, and test or probe the alert definition results.

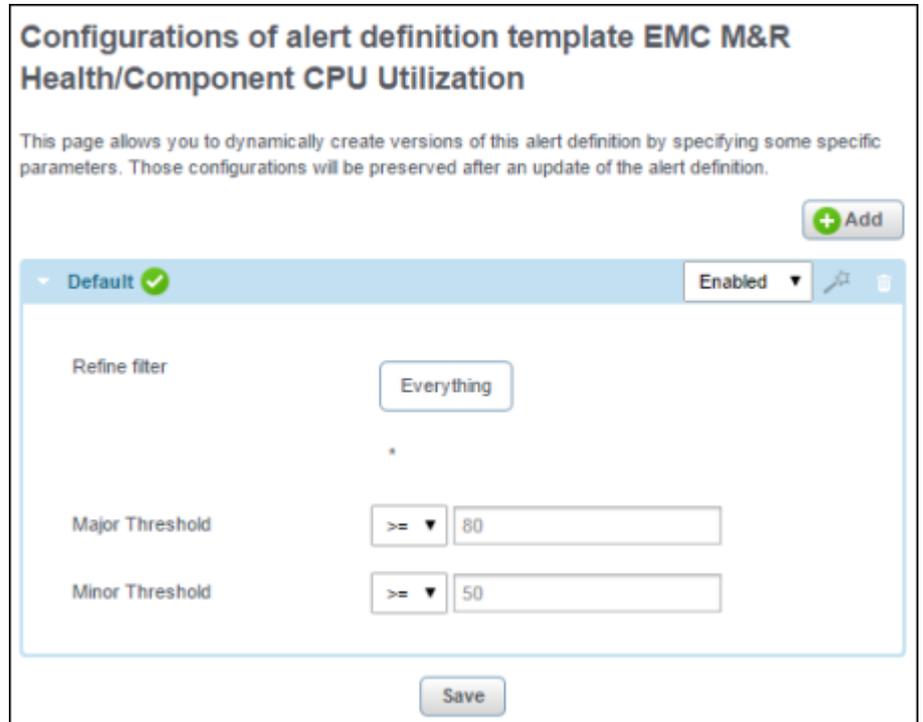
For example, alert definitions often compare values in an event to a configurable threshold setting. A filter defines which events to consider. If an alert definition is contextualized, you can you define different filters or threshold values and name them.

Note

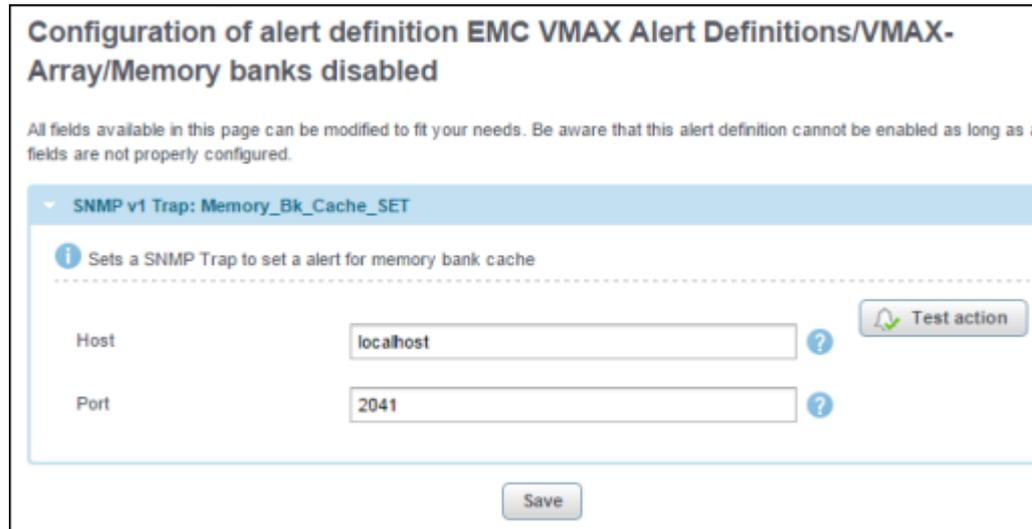
Not all alert definitions are contextualized.

Procedure

1. Click **Administration > Modules > Alerting > Alert Definitions**.
2. In the table of alert definitions, expand the folders to find the alert definition to configure.
3. Right-click the alert definition and choose **Configure**.
 - If the alert is contextualized, you see one or more named contexts, with the configurable parameter names and the values for each context. Out of the box, one context is named **Default**.



- If the alert is not contextualized, you see the configurable parameters and the current settings.



4. For a contextualized alert, you can:
 - Change the values in any of the existing contexts, including the **Default** context or any other named context.
 - Create a new context, as follows:
 - a. Click **Add**.
 - b. Name the new context, and click **OK**.
 - c. Type values in the new context.
 - Choose **Enable** or **Disable** for each context. When multiple contexts are enabled, the incoming events are processed multiple times, once for each enabled context.

5. For a non-contextualized alert, change the values in one or more fields.
6. Click **Save**.

If the alert definition is enabled, the change takes effect immediately.

Editing an alert definition

You can change or add additional components to an alert definition.

By editing an alert definition, you can change the filter or add or delete the actions or operations in the alert definition.

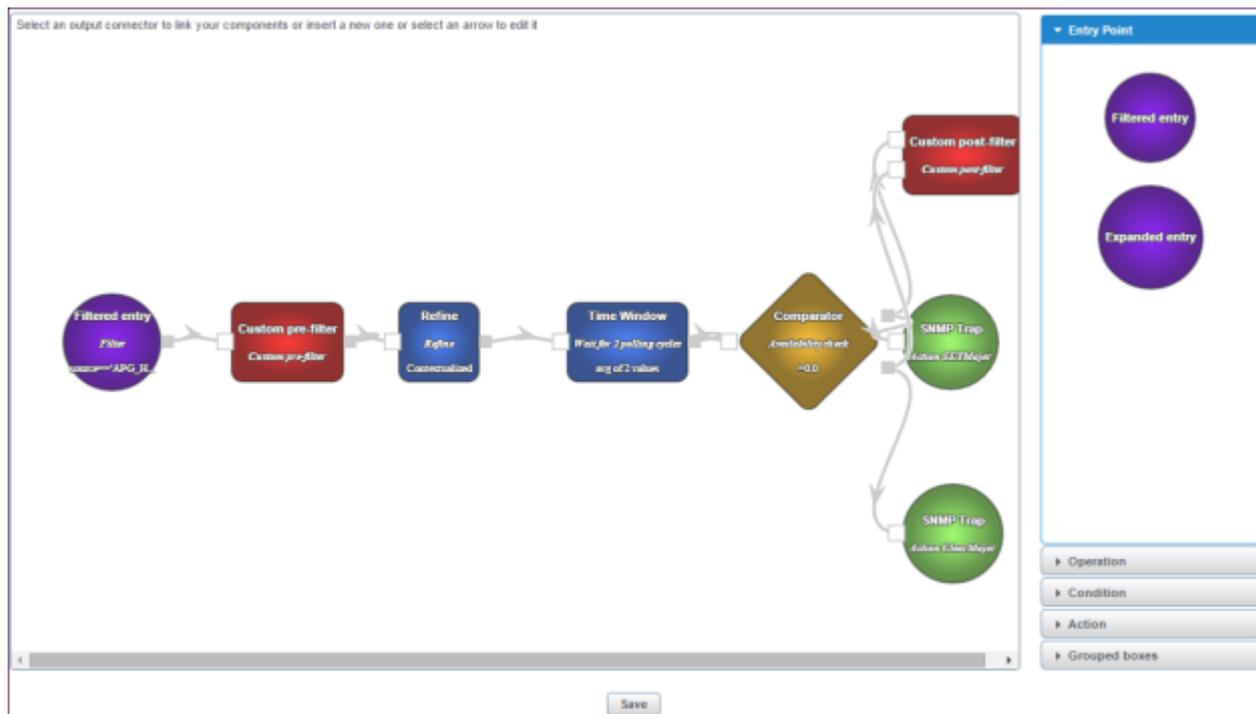
Note

Do not add the ESRS Event action to any alert definition. The ESRS Event action is for EMC use only.

Procedure

1. Click **Administration > Modules > Alerting > Alert Definitions**.
2. Expand nodes until the alert definition you want to edit appears in the table in the right pane.
3. In the right pane, right-click the alert definition and choose **Edit**.

A graphical representation of the alert definition appears on a whiteboard, with an expanding list of available components on the right.



4. To reconfigure an existing component:
 - a. Hover the cursor over the component until a menu of small icons appears above the component, and click **Edit** (pencil icon).
For example, to refine the filter in the Filtered Entry, hover the cursor over the Filtered Entry component in the alert definition.



- b. Make changes to the component's configuration in the dialog that appears.
- c. Click **OK** to preserve the changes and exit the dialog.

5. To add a new component:

- a. In the list of components on the right, click the type of component you want to add.

The type expands to show all available components in that category. For example, click **Action** to show all of the action components.



- b. Drag the symbol for the desired component from the list onto the whiteboard.

For example, to add an Email notification action, drag the Mail object.

- c. Complete the configuration dialog that appears when you release the drag, and click **OK**.

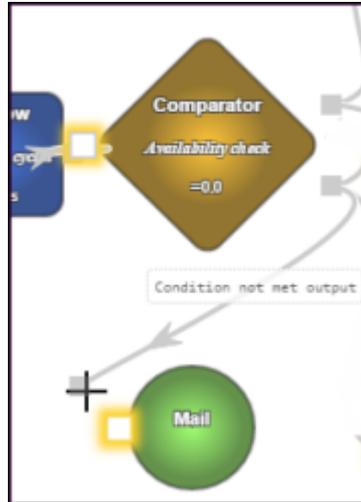
To reconfigure the component later, hover the cursor over the symbol and click the **Edit** icon that appears.

- d. Add the new component into the alert definition flow by dragging a line between connection points on the symbols.

- Start the drag from an exit connection point.
- Drag to an entry connection point.

Exit connection points are on the right side of a symbol. Entry connection points are on the left.

The following example connects the "Condition not met" exit point of a Comparator to the Mail action. (The alert definition sends an email when the condition is not met.)



6. Click **Save** at the bottom of the whiteboard to save all changes to the alert definition.

Enable/disable alert definitions

Many alert definitions are installed in the disabled state. You need to enable them if you want the events defined by their filters to generate alerts.

You can enable/disable alert definitions at the folder level, sub-folder level, or individually.

Procedure

1. Go to **Administration > Modules > Alerting > Alert definitions**.
2. To enable/disable alert definitions:
 - a. Click folder rows in the right pane until the folders or individual alert definitions that you want to enable or disable appear as a row in the table.
 - b. Select one or more rows in the right pane using the checkboxes in the first column.
 - c. Click **Enable** or **Disable** at the bottom of the page.

If you selected a folder, all alert definitions under that folder are affected, including the contents of any sub-folders.

3. In the right pane, use the **State** column to determine the status of individual alert definitions.

State column	Description
empty	This is a folder. Click it to show its contents.
	The alert definition is enabled.
	The alert definition is disabled.
	The alert definition requires some configuration before you can enable it. Right-click the row and choose Configure .

Creating alert definitions

You can use alert definition examples to help you get started with constructing alert definitions.

Procedure

1. From the Alerting page, click **Alert definitions**.

2. Click the **New element**  icon above the tree.

A whitespace opens with components listed on the right.

3. Drag an **Entry Point** component onto the whitespace and configure it.

When you release the drag, a dialog for configuring the component opens.

The Entry Point filter works like report filters do in the main user interface.

The difference between using a filter in the Alerting Frontend and the Backend is the suggestions shown in the wizard. The Alerting Backend is not bound to a database, so the list of known properties and values is created at runtime.

Because of this, it is normal to have no suggestions when you use the wizard immediately after starting the Alerting Backend.

4. In the component list on the right, expand the type of component you want to work with next.

5. Drag a component onto the whitespace and configure it.

You can rearrange components later by dragging them.

6. Continue to drag components and configure them, ending with an Action component.

7. Connect components to define the processing flow for an entry that is captured by the entry point filter.

Every component has connectors, which are yellow squares. To link components, click an output connector on the right side of a component, and drag to an input connector on the left side of another component. When the input connector turns green, release the drag.

Note

To delete a connection, drag the line off of the input connector point and release. The connecting line disappears.

There is no limit to the number of objects you can link to the same output.

8. When you are finished adding components, click **Save**.
9. In **Name**, enter a name for the alert definition.
10. In **Description**, enter a description for the alert definition.
11. Click **Save and enable** or **Save and disable**.

About the ESRS alert definitions

Do not change or reconfigure the EMC Secure Remote Support (ESRS) alert definitions.

The SolutionPack for EMC M&R Health includes an optional feature that forwards alerts about certain health conditions to EMC. The feature is installed and configured during SolutionPack for EMC M&R Health installation by the alerting block question.

Note

Unless you configure the ESRS settings in the alerting block question, no information is sent to EMC.

A set of ESRS alert definitions is associated with this feature. Note the following:

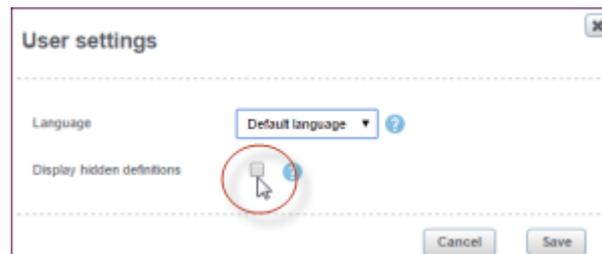
- The ESRS alert definitions are hidden by default. See the next section to unhide them.
- Do not change or reconfigure the ESRS alert definitions.
- Do not add the ESRS Event action (used in the ESRS alert definitions) to other alert definitions.
- The alert definition **Enable** and **Disable** commands do not apply to the ESRS definitions. The ESRS definitions are always enabled when the SolutionPack for EMC M&R Health alerting block is installed, and always disabled otherwise.

Unhide the ESRS alert definitions

You can optionally unhide the ESRS alert definitions to view them.

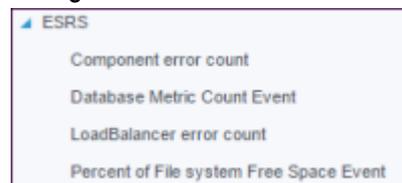
Procedure

1. Go to **Administration > Modules > Alerting > Alert Definitions**.
2. Click **Profile > View Profile** in the banner area of the alerting page.
3. Check the **Display Hidden Definitions** box.



4. Click **Save**.

The ESRS alert definitions appear in the tree. They are for viewing only. Do not change them.

**Alert definition component reference**

An alert definition is a sequential set of components that filters events and takes appropriate action. There are four types of alert definition components.

Entry point filter

Defines the events and metrics to examine. Every alert definition requires an entry point filter.

Operations

Define operational actions on the filtered event data. Operations are optional.

Conditions

Define test conditions on the event data. Different actions can depend on conditional outcomes. Conditions are optional.

Actions

Define the actions to take, such as creating an SNMP trap to forward to the alerting backend or write a log entry. At least one action component is required for the alert definition to have any results.

Entry Point filter

The Entry Point filter works like report filters in the main user interface.

The difference between using a filter in the Alerting Frontend and the Backend is the suggestions shown in the wizard. The Alerting Backend is not bound to a database, so the list of known properties and values is created at runtime. Because of this, it's normal to have no suggestions when you use the wizard immediately after starting the Alerting Backend.

Operations

Operations evaluate incoming data to see if they should go to the next component or be dropped.

No Operation

This operation does not do anything. Use it if an entry needs separate senders in its process when all the entries come from the same operation, such as a Previous Polling Operation. It's a good way to make a used connector (red) of a grouped box available as an output if needed.

Case

This operation creates logical output groups that send data to each next operation until one has its condition met, such as a group of alert levels that send an alert only once for the highest level of alert. The order of criticality of the next components is the order of link creation, so the next component link is the highest level and if its condition is met when the next data is received, the other next components won't receive data.

Constant Arithmetic

This operation applies an operator on each entry data with a constant value.

Value

Value used by the operator

Operator

The arithmetic operator (+, -, / or *)

Constant is left operand

Select this if the constant value must be left operand

Data Arithmetic

This operation applies an operator on two data lines defined by the filter.

Operator

The arithmetic operator (+, -, / or *)

Time limit

Maximum delta between the two values in minutes

Left operand

Optional and needed only if you use an operator that produces a different result depending on the order of the operand, such as - and /. This parameter must contain the name of the sender who will send the data that must be the left side of the operation.

Group by

Properties, each separated by a white space, that create a logical group of values for the operation

Value filter

Optional and is necessary only if the entries of this operation come from the same input, as it defines two different elements of the same group

Constant Math

This operation applies a complex mathematical operation to the data.

abs

The abs function which returns the absolute number

ceil

The ceil function which returns the smallest value but greater than the argument

exp

The exp function which returns the exponential value raised to the power of a double value

floor

The floor function which returns the largest value but less than the argument

Configurable Math

This operation applies a complex mathematical operation using a parameter to the data.

Operation

The mathematical operation

Value

Value applied to the second parameter of the selected operation

History Operation

This sliding window operation collects a range of values bounded by the number of values or the time range of the values. The output of the collected values can be their minimum, maximum, sum or average.

Wait for period

Optional and for a time only bounded buffer, this flag is set if the buffer must wait a full period before computing values

Time range

Optional and is the time window of the collected values in minutes

Window length

Optional and is the values number to collect in the sliding window

Output

Min, max, avg or sum

Previous Polling Operation

This sliding window operation keeps a series of previous polling for each data ID and distributes it to the outputs.

Time range

Optional and is the time window of the collected values in minutes

Window length

Number of previous polling to keep

Data Repeater

This operation blocks or allows output for every data that have their timestamps outside the time frame.

Update timestamps

Select this to have the timestamps of each repetition be the current timestamp; don't select this to have the actual timestamp used

Repetitions

Repetition number for each ID. 0 = unlimited

Repeat period

Time in minutes between repetitions

Date

This operation blocks or allows the output of data that doesn't have a timestamp within the time frame.

Block if in range

Select this to block data if the timestamp is in the range you set here

Start time

Inclusive start time of a day, following a hh:mm 24 hour format

End time

Exclusive end time of a day, following a hh:mm 24 hour format.

End time must be later than start time

Start day**End day**

Optional. select this if the time range can be repeated for successive days, such as 8:00 to 17:00 Monday to Friday.

Properties To Values Creator

This operation takes a list of properties and creates a new value for each property that contains a numerical value. The original data is not sent to the next components. Enter a list with each property separated by a white space.

Timestamp To Value Creator

This operation creates data with the data timestamp as its value. The original data is not sent to the next components.

Alive Timer Operation

This operation sends the last received data to the next components only when no data for the same ID has been received before the timer expires.

Inform only once

Select this to send non-updated data to the next components only once rather than at the end of each period.

Timer

Period of the timer. This value cannot be lower than 5, as every 5 minutes a check is made to see if any ID isn't updated.

Conditions

A condition verifies a condition of the data before sending it to its actions or operations. Conditions can be stateful or stateless. In a stateless condition, new data is sent to the next component. A stateful condition sends the data only when the state is changed. For example, if state of ID A = true, data is sent only when the state is false.

Constant Comparator Operation

Compares data with a constant value and adds the ADDED_THRESHOLD property containing the comparator in the output data. Data that meets the conditions is sent to the true actions/operations. Data that does not is sent to the false actions/operations.

Stateless

Select this to send the result for each piece of data received, or do not select this to send data if a state changes for an ID.

Constant value

Constant value used to compare data

Operator

Operator of the comparison (<, <=, >, >=, = or !=)

Property Comparator Operation

Compares data with a numeric property contained in the data.

Stateless

Select this to send the result for each piece of data received, or do not select this to send data if a state changes for an ID.

Property

Must be a valid number to compare data with its value

Operator

Operator of the comparison (<, <=, >, >=, = or !=)

Data Comparator Operation

Compares two pieces of data coming from two entries and defined by a grouping filter. Each comparison ends with the piece of data on the left sent to the true actions/operations when the condition is met, or to the false actions/operations if the condition is not met.

Stateless

Select this to send the result for each piece of data received, or do not select this to send data if a state changes for an ID.

Operator

Operator of the comparison (<, <=, >, >=, = or !=)

Left operand

Must contain the name of the sender who will send the data that must be the left side of the operation.

Time limit

Maximum delta between the two values in minutes

Group by

Properties that create a logical group of values for this operation

String Comparator Operation

Compares one of the fields in the data with a custom string. Each piece of data with a condition met is sent to the true actions/operations, or to the false actions/operations if the condition is not met.

Stateless

Select this to send the result for each piece of data received, or do not select this to send data if a state changes for an ID.

Field name

Name of the field in the data to use in the comparison

Comparator

Type of comparison (equals, contains, starts with or ends with)

Compare to

Custom string used in the comparison

Data Counter Operation

This operation counts the number of data with the same ID and sends it to its true actions/operations when the count reaches the counter in the time limit.

Stateless

Select this to send the result for each piece of data received, or do not select this to send data if a state changes for an ID.

Time range

Time window boundary of the collected values in minutes

Count

Counter

Grouped Data Counter Operation

This operation is useful to do an AND of previous conditions for data related by properties with different IDs. If the count of received related data matches the counter before the time limit, the last data is sent to its true actions/operations.

Stateless

Select this to send the result for each piece of data received, or do not select this to send data if a state changes for an ID.

Time range

Time window boundary of the collected values in minutes

Count

Counter

Group by

A list of properties separated by white spaces that create a logical group of values for the operation

Sender Counter Operation

This operation counts the number of data from the sender and sends them to their true actions/operations when the count reaches the counter in the time limit.

Stateless

Select this to send the result for each piece of data received, or do not select this to send data if a state changes for an ID.

Time range

Time window boundary of the collected values in minutes

Count

Counter

String Counter Operation

This conditional operation counts the times that data with the same ID has a property equal to a comparator, and sends them to their true actions/operations when the count reaches the counter in the time limit. To make a counter with more string operations, use a stateless String Comparator with a Sender Counter.

Stateless

Select this to send the result for each piece of data received, or do not select this to send data if a state changes for an ID.

Time range

Time window boundary of the collected values in minutes

Count

Counter

Field name

Name of the field in the data to use in the comparison

Compare to

Custom string that must match the field value

Actions

Actions inform users of new or resolved problems in their system or run an auto resolving process.

About action data

Each alert data comes with a series of properties. Since most of the time they are coming from the Watch4net collectors, they are the same properties that you find when editing a report. When the Alerting Backend is running, you can have a view of those properties either in an entry point filter or when probing an active definition.

For example, suppose you want to send an email reporting the current CPU utilization for a specific CPU in a device based on the alert data below:

```
AlertData: {
  VALUE: 25
  TMST: 123456789 (unix timestamps)
  ID: uniqueId for this metric
  properties: {
    devtype: Array
    device: VPLEX1
    parttype: CPU
    part: CPU0
    name: CurrentUtilization
  }
}
```

You would add the following as the content of the email message:

```
Device PROP.'device' currently has a PROP.'name' on its PROP.'part'
of VALUE.
```

This would translate dynamically as:

```
Device VPLEX1 currently has a CurrentUtilization on its CPU0 of 25.
```

Log action

This action writes an entry in a log file. If two actions write to the same file, use the same parameters except for the content, because the Alerting Backend can't guarantee which configuration will be used. The default directory for the log files is Backends/Alerting-Backend/Default/logs.

Parameter	Description
File prefix	Complete directory path and prefix of each log file. The prefix is followed by <code>_index.log</code> , where <code>index</code> is incremented until it reaches the retention parameter. Newest log files have the lowest index, so the current log is always <code>_0.log</code> .
Log retention	Maximum number of log files kept.
Rotation time	Optional if rotation-entry is used. Maximum time frame in minutes of the log content prior to its creation.
Rotation entries	Optional if rotation-time is used. Maximum entries in a log file.

Parameter	Description
Entry content	Content of the log entry. To add the value, date or ID of the current data, use the keywords VALUE, TMST or ID. To add a property of the current data, use the keyword PROP.'propertyName'.

Mail action

This action sends an email when triggered. For successful emails, an SMTP server must be configured as described [here](#).

Parameter	Description
To	Comma-separated list of recipients
Subject	Subject of the email. To add the value, date or ID of the current data, use the keywords VALUE, TMST or ID. To add a property of the current data, use the keyword PROP.'propertyName'.
Message	Content of the email. To add the value, date or ID of the current data, use the keywords VALUE, TMST or ID. To add a property of the current data, use the keyword PROP.'propertyName'.

SNMP Trap action

This action sends an SNMP trap when activated. SNMPv1, SNMPv2, or SNMPv3 traps are supported.

Note

You can not change the SNMP version on an existing trap action. Instead, drag a new SNMP action to the whiteboard, define it as needed, and delete the older action.

The following header fields appear on a new SNMP component.

Field name	Description
Name	Unique name for this component in the alert definition.
Type	Select SNMP v1, SNMP v2, or SNMP v3. The fields in the Parameters section change to match the version specification.
Description	Customized description for this trap action.

The following fields appear in the **Parameters** section of an SNMP component.

Parameter	Description
Host	Receiver of the traps
Port	Receiving port
Community	Community of the traps

The following parameters change to match the SNMP version selected in the header.

SNMPv1	SNMPv2	SNMPv3
Enterprise OID	Trap OID	Trap OID
Trap OID		Engine ID
Generic ID		Username

Parameter	Description
SNMPv1	SNMPv2 SNMPv3
Enterprise specific ID	Authentication Protocol (MD5 or SHA)
Variables OID format (legacy or RFC)	Authentication Password
	Private Protocol (DES or AES)
	Private Password
Trap content	Comma-separated list of the content for each variable of the trap. A comma is escaped if it's preceded by a period. To add the value, date or ID of the current data, use the keywords VALUE, TMST or ID. To add a property of the current data, use the keyword PROP.'propertyName'. There's no limit to the number of variables in the trap content.

External Process action

This action runs an external process on the Alerting Backend when triggered. To add the value, date, or ID of the current data, use the keywords VALUE, TMST or ID. To add a property of the current data, use the keyword PROP.'propertyName'.

Parameter	Description
Command	Command to run
Command parameters	Optional. Comma-separated parameters of the command, for example -n 200. A comma is escaped if it's preceded by a period.
Environment parameters	Optional. Comma-separated list of environment parameters of the terminal in key=value format (ex: JAVA_HOME=/opt/...). A comma is escaped if preceded by a period.

Send to SAM action

This action sends an alerting notification to a Smarts SAM Console.

Parameter	Description
Host	SAM hostname or IP address
Port	SAM port
Domain name	Name of domain where EMC Smarts is installed
Domain username	Username to connect to the EMC Smarts domain
Domain password	Password to connect to the EMC Smarts domain
Broker username	Username to connect to the SAM broker
Broker password	Password to connect to the SAM broker
Operation type	Type of operation that is involved (clear or notify)

Parameter	Description
Source	Name(s) of the domain or domain group(s) that have originally diagnosed and notified, directly or indirectly, current occurrences of this event. If there is more than one original domain, the attribute lists each separated by a comma. When the notification is cleared, the last clearing domain stays in the value. Type: string
User	Notification user
EventName	Name of the event. This attribute, along with ClassName and InstanceName uniquely identified this event. Type: string
EventType	Indicates the nature of the event (MOMENTARY or DURABLE). A MOMENTARY event has no duration. An authentication failure is a good example. A DURABLE event has a period during which the event is active and after which the event is no longer active. An example of a durable event is a link failure. Type: string
Severity	Enumerated value that describes the severity of the event from the notifier's point of view: <ul style="list-style-type: none"> • Critical Action is needed NOW and the scope is broad (for example, an outage to a critical resource) • Major Action is needed now • Minor Action is needed but the situation is not serious at this time • Unknown Element is unreachable, disconnected, or in an otherwise unknown state • Normal Event is purely informational Type: unsigned int
Notification properties	ICS_Notification property from the Smarts Java API.

ESRS event action

Note

Do not add the ESRS event action to any alert definitions.

The SolutionPack for EMC M&R Health includes an optional feature that forwards alerts about certain health conditions to EMC. The ESRS event action is for the exclusive use by EMC for the ESRS alert feature.

Creating adapters

An adapter takes data and translates it to alerting data.

Creating a Values Socket Listener Adapter

This adapter waits for values from the socket connector in the collecting.xml file. The data has the properties adapterName containing the name of this adapter and adapterType containing RVSocketListener.

Procedure

1. From the Alerting page, click **Adapters**.
2. Click the **New element** icon above the tree.
3. From the **Type** list, select **Values Socket Listener Adapter**.
4. In **Name**, type a descriptive name for the adapter.
5. In **Port**, type the port that the adapter will take its incoming raw data from.

Creating an Events Adapter

This adapter waits for events from a generic-event-writer in the processing.xml file. The data has the properties adapterName containing the name of this adapter and four parameters.

Procedure

1. From the Alerting page, click **Adapters**.
2. Click the **New element** icon above the tree.
3. From the **Type** list, select **Events Adapter**.
4. In **Name**, type a descriptive name for the adapter.
5. In **Value**, optionally enter the name of the field containing the decimal value of the event to be used in the operation, such as bytes for flow events. You can leave this box blank for a strings only event.
6. In **Properties**, enter a comma-separated list of the properties used in the entry point or special components.

If the default data type of a property does not give the best display format, enter the data type of the field with the field name separated by a comma, for example SRC_PORT:INT. The data types can be Boolean, byte, bytes, double, float, int, long, numeric, short, object, or string. The data type is implementation-specific for the type of data the adapter is parsing.

7. In **Timestamps**, optionally enter the field that contains the timestamps of the events. Otherwise, the current timestamps are used.
8. In **Port**, type the port that the adapter will take its incoming raw data from.

Creating a Report Data Adapter

This adapter parses exported reports. The properties of the data are the report displayed properties and the report name (property = reportName). Table reports contain the non value columns and the name of the value column (property = name). Graph reports have the legend of each series (property = name). Each file parsed is deleted after parsing. The resulting data has the properties adapterName containing the name of this adapter and adapterType containing APGXmlReportAdapter.

Procedure

1. From the Alerting page, click **Adapters**.
2. Click the **New element** icon above the tree.
3. From the **Type** list, select **Report Data Adapter**.
4. In **Name**, type a descriptive name for the adapter.
5. In **Time check**, enter the time in minutes between each check.
6. In **Directory**, optionally enter the subdirectory under the Alerting custom directory that contains the files to parse.
7. In **File name regex**, enter the file name regex. You can use the * wildcard character.

Creating a File Data Adapter

This adapter parses each file corresponding to a file name regex in a directory. Each file parsed is deleted after parsing.

Procedure

1. From the Alerting page, click **Adapters**.
2. Click the **New element** icon above the tree.
3. From the **Type** list, select **File Data Adapter**.
4. In **Name**, type a descriptive name for the adapter.
5. In **Directory**, optionally enter the subdirectory under the Alerting custom directory that contains the files to parse.
6. In **Time check**, enter the time in minutes between each check.
7. In **File name regex**, enter the file name regex. You can use the * wildcard character.
8. Upload any parser or formatter files you want to use.

Creating a Kafka adapter

This adapter retrieves raw values from a Kafka cluster and pushes the data to the alerting engine. The data is processed by the alerting engine to produce alerts.

Before you begin

This Kafka adapter is useful only if the following supporting elements are configured in the EMC M&R infrastructure:

- A Kafka collector is collecting and inserting values into a Kafka topic called rawvalues.
- Alerting definitions exist that filter for data from this adapter and from the rawvalues topic.

Note

Not all M&R products support the Kafka collector and data bus.

Procedure

1. From the Alerting page, click **Adapters**.
2. Click the **New element** icon above the tree.
3. In the **Type** list, select **Watch4net Values Kafka Adapter**.

4. In **Name**, type a descriptive name for the adapter.
5. In **Configuration type**, select one of the following:
 - **Configuration file** to specify a custom configuration file.
 - **Explicit settings** to use a system-provided configuration file and provide only a few installation-specific parameters.
6. If you selected **Configuration file** above, use **Watch4net Values Kafka Adapter configuration file** to select a custom configuration file. The file must conform to the `kafka-consumer.xsd` specification. For parameter information, see <http://kafka.apache.org/documentation#newconsumerconfigs>
7. If you selected **Explicit settings**, the following parameters are required:

Kafka cluster addresses

Provide a comma-separated list of Kafka cluster addresses in the following format:

```
hostname:port, hostname:port, ...
```

At least one entry is required for explicit configuration. Two or more entries are recommended for redundancy and required for high availability. With multiple entries, the Alerting Backend continues to receive data even if one of the Kafka processes stops.

Topic

Use the literal value `rawvalues`.

This is one of several consumer topic-names defined in the Kafka collector's configuration file. The Alerting Backend expects values as collected in the `rawvalues` topic.

8. Click **Create**.

Creating alert templates

Templates enable you to pre-configure values and add them to an alert definition.

Creating alert templates

Alert Manager comes with preconfigured templates. You can edit the existing templates or create new ones.

Procedure

1. From the Alerting page, click **Templates**.
2. Click the **New element** icon above the tree.
3. In **Family**, select the type of template to create.

Option	Description
Entry Point	The Entry Point filter works like report filters on the APG Frontend. It filters the flow of incoming data for alert definition.

Option	Description
Operation	Operations evaluate incoming data to see if they should go to the next element or be dropped.
Condition	A condition verifies a condition of the data before sending it to its actions or operations. Conditions can be stateful or stateless. In a stateless condition, new data is sent to the next component. A stateful condition sends the data only when the state is changed. For example, if state of ID A = true, it sends data only when the state is false.
Action	Actions inform users of new or resolved problems in their system or run an auto resolving process.

4. In **Name**, enter a name for the template.
5. In **Description**, enter a description for the template.
6. Use the Filter Wizard to select the properties you want to filter on.
7. Click **Create**.

Using preconfigured alert templates and examples

Alert Manager comes with preconfigured templates and example alert definitions. You can edit the existing ones or create new ones based on the templates and examples.

There are several example alert definitions. Alert definition components can reference a template. Log and Traps are the most frequently used actions in alert definitions. We provide a log template and two trap templates.

Procedure

1. To view an example alert definition:
 - a. From the Alerting page, click **Alert definitions > Examples**.
 - b. In the right pane, right-click an alert definition name and select **Edit**.
2. To view a preconfigured template:
 - a. From the Alerting page, click **Templates > *template_name***.
 - b. Review and edit the template contents before associating it to alert definitions.
3. To edit the default trap templates to work in your environment:
 - a. From the Alerting page, click **Templates > *default trap template***.
 - b. In the **Template Parameters** section, examine the **Trap Content** field.

The preconfigured trap template lists all possible alert properties, including some properties that are user-defined. It is probable that your MIB does not contain fields for the entire list of properties. In that case, any SNMP trap action component that references the default alert trap template will not work.

- c. Replace any property in the list that does not exist in your MIB with the value `null`.

The order of properties is important in a template.

Note

To preserve the out-of-the-box template for future use, create a new template with an edited list of properties.

Here is an example edited property list.

```
PROP. 'id', PROP. 'Name', null, PROP. 'count', PROP. 'eventstate', PROP
. 'Source', PROP. 'parttype', PROP. 'part', PROP. 'eventname', PROP. 'p
arttypedisplayname', PROP. 'partdisplayname', PROP. 'eventdisplayn
ame', PROP. 'fullmsg', PROP. 'devtype', PROP. 'device', PROP. 'sourcei
p', PROP. 'sourcedomainname', PROP. 'sourceeventtype', PROP. 'value'
, PROP. 'active', PROP. 'timestamp', null, null, null, null, null, PROP.
'acknowledged', PROP. 'eventtype', PROP. 'category', PROP. 'eventtex
t', PROP. 'severity', PROP. 'impact', PROP. 'certainty', PROP. 'inmain
tenance', PROP. 'troubleticketid', PROP. 'owner', PROP. 'systemdefin
ed1', null, null, null, null, PROP. 'userdefined1', null, null, null, nu
ll, null, null
, null, null, null
```

d. Click **Save**.

4. To associate a template to an SNMP trap action:
 - a. Under **Alert Definitions**, click an alert definition folder.
 - b. In the table in the right pane, right-click the alert definition name and select **Edit**.
The diagram of components for this alert opens.
 - c. Click the **Edit** (pencil) icon above the SNMP trap action component (a green circle).
 - d. In the **Edit Component** dialog that opens, in the **Template** field, select the appropriate template name from the drop down list.

Setting up a log template

This example shows creating a new template based on the Log template. Logs are frequently used actions in alert definitions.

Procedure

1. From the Alerting page, expand **Templates**.
2. Click **Log Template**.
3. In **Description**, enter a description for the new template.
4. In **File prefix**, enter the directory path and prefix for each log file.

The log files are created in the Alerting-Backend/instance/Alerts folder and named log_INDEX.log, where index is incremented until it reaches the log retention limit. The newest log files have the lowest index, and the current log is the 0 log.

5. In **Logs retention**, enter the number for how many logs you want to keep.
6. In **Rotation time**, enter a value for the maximum time in minutes that log data is gathered before the log is created.

This value is optional if you are setting a value for rotation entries. In this example, the logs files are rotated once a day (1440 minutes).

7. In **Rotation entries**, enter a value for the maximum number of entries per log file.

This value is optional if you are setting a rotation time. In this example, the log is rotated when the 0 log file has 5000 entries.

8. In **Entry content**, enter the properties for the log entry.

To add the value, date or ID of the current data, use the keywords VALUE, TMST or ID. To add a property of the current data, use the keyword PROP.'propertyName'. In this example, each line of the log contains: DEVICE with PART -> NAME = VALUE.

9. Click **Save**.

Setting up a Trap Notify and Trap Clear template

This example shows how to create a new template based on the Trap Notify template. Traps are frequently used actions in alert definitions. A Trap Notify alert sends an alert to your fault management system, such as EMC Ionix. The Trap Clear deactivates the alert.

Procedure

1. From the Alerting page, expand **Templates**.
2. Click **Trap Notify Template**.
3. In **Description**, enter a description for the new template.
4. In **Host**, enter the host that will receive the traps.

The host is the address of the SNMP trap manager and should be changed if it's not on the same host as the alerting backend.

5. In **Port**, enter the trap manager port.162 is the default for this protocol.
6. In **Community**, enter the trap community. Public is the default.
7. In **Generic ID**, enter the generic ID of the trap.
8. In **Enterprise specific ID**, enter the specific ID of the trap format for the generic type.

9. In **Trap content**, enter the content for each variable of the trap, separated by commas.

Commas will be escaped if they are preceded by a \. To add the value, date or ID of the current data, use the keywords VALUE, TMST or ID. To add a property of the current data, use the keyword PROP.'propertyName'. It is recommend that you do not change the content of the traps as the order of the properties match our trap configuration. You can change the last three properties to ones you define. The only difference between Trap Notify and Trap Clear is the number value in the Trap content: 2 is notify and 4 is clear.

10. Click **Save**.

Grouped boxes

Grouped boxes are containers of standard alerting components, linked in the same way as alert definitions, that you can use in an alert definition or other grouped boxes.

With grouped boxes you can:

- Create a subset of boxes such as baseline computations that you can reuse in multiple definitions. Grouped Boxes can contain any alert definition component and can easily be reused.

- Create a single box that has its instances updated whenever it is modified, unlike templates that do not have a link. When you make a change to a grouped box, the change is immediately propagated to all the alert definitions that use it.

Unlike alert definitions, grouped boxes have free inputs and outputs on any number of components, whether they are linked or not. The free connectors appear as inputs or outputs on the group box that you can use to link to other components. You can mouse over a free connector to see its source component, which is its name, and the input/output type, for example Condition met output.

Grouped boxes are created like alert definitions with these exceptions:

- Entry points are not required to save a grouped box.
- Connectors have three possible states: red, yellow, and green.
- When you save a grouped box, the resulting box in the definitions is a red box and its connectors are the unused input/output connectors of its internal components. No configuration is needed when using a grouped box in a definition. The description of the configuration is the description you enter when you save the grouped box.

Creating a Grouped Box

You can create and edit Grouped Boxes and use them in alert definitions. Grouped Boxes can contain any alert definition component and can be easily reused.

Procedure

1. From the Alerting page, click **Grouped boxes**.
2. Click the **New element** icon above the tree.
3. Edit the Grouped Box definition by dragging entry points, operations, conditions and actions on the page, and editing their parameters. Unlike alert definitions, entry points are not mandatory for grouped boxes.
4. Link the components in a logical flow for how you want the data to be evaluated.

Connectors have three states:

Green represents free input or output connectors. You can link within the scope of the Grouped Box you are working on or leave the connector free, in which case it appears as a free connector on the Grouped Box when it is used in an Alert Definition.

Red represents a connector that is linked with another element within the Grouped Box. The data flow between these connectors takes place in the Grouped Box, and the connectors are not visible and cannot be linked to in the Grouped Box.

Yellow represents connectors that have been used within an Alert Definition. You can mouse over a yellow connector and see a list of the alert definitions that use it, including the component that links to it. A yellow connector cannot be linked to in the source Grouped Box as it is already occupied in one or more Alert Definitions, and it cannot be deleted until it is freed in the Alert Definitions it is linked to.

You can use a Grouped Box within another Grouped Box but not within itself. Grouped Boxes used within a Grouped Box display their available connectors.

5. Click **Save**.

Probing alerts

Probing helps you confirm the behavior of an alert definition. Do a probe to test the elements in an alert definition in real-time to see how data is evaluated at each component, and how many times actions have been triggered.

Probes help you spot problems with how you construct your alert definitions, so that you can adjust them and retest until your alerts are satisfactory. This helps you to construct ever more complex alerts to match your network's increasing complexity.

Probing is available only to users who have write permission.

Preparing for probing

Before you start probing, you can make these changes to improve the success of probing your alert definitions.

- Make sure the alerting sources, such as timeseries, events, text, or report data, are properly configured to be sent to the Alerting Backend.
- Reduce the polling intervals of one or more collectors, or how often source data is sent to an Alerting Backend, so that your probing yields more results that correspond to the time frame. If, for example, the default polling for the Smarts Collector uses four minutes, you would have to wait four minutes for every probe result. You can change the Smarts Collector to 10 seconds for the duration of the probing, and then switch it back for production.
- Reduce the sampling property that determines how many properties per data sample are preserved. This is controlled by the `alerting.suggestion.sample` property in the `module.properties` file located at `<APG>/Backends/Alerting-Backend/<instance>/conf`. By default this property does not appear, and 1 in 100 properties are kept. For example if polling takes place every ten seconds, it can take a long time to accumulate enough properties for effective probing. You can change this to 1:1 for the duration of your probing session.
- Filter hints and sampled properties are dependent on the data coming into the Alerting Backend. You can increase the number of hints and sampled properties by adding and adjusting the `alerting.suggestion.sampling` property in the `module.properties` file located in the conf directory of the Alerting Backend.

Probing an alert definition

Use probing to test an alert definition in real-time.

Procedure

1. From the Alerting page, click **Adapters**.
2. Enable the adapter you want to use.
3. In the tree, click **Alert definitions**.
4. Enable the alert definitions that you want to use.
5. In the **Alert definitions** list, click the alert definition you want to probe and click **Probe**.

The alert definition whiteboard appears.

6. Edit the probe settings.

Option	Description
Filter	Applies a filter to the probe results. This affects only the probe results and not the alert definition. Use a standard filter. Only a subset of properties may appear in the wizard depending on how many samples were cached. You can adjust this by changing the value of the <code>alerting.suggestion.sampling</code> property, in the <code>module.properties</code> file, at <code><APG>/Backends/AlertingBackend/<instance>/conf</code> .
Show in probes	<p>Selects the property to show in all probe results. The properties appear at the beginning of each result line followed by the selected sampled properties. They appear only when you probe a link.</p> <ul style="list-style-type: none"> • Sender name: component the line of alerting data originated from. • Destination name: component the line of alerting data is flowing to. • Variable ID: report node the line of alerting data originated from. The variable ID format depends on the adapter type. • Variable value: value from the line of alerting data to be evaluated by the alert definitions, unless a string in the data is selected for evaluation, such as a string comparator. • Variable timestamps: timestamp associated with the line of alerting data. It is converted from a Unix timestamp to a human-readable form for clarity in the probes. For category graphs this is when the report was generated in the main user interface.
Properties	Selects properties whose names and values will appear in probe results. You can select All to display all the available properties.

7. Test or clear a component or link by clicking it. The component is highlighted and its information appears in an area of the same color.
8. To disable the Probe UI, for example to edit the definition to refine the probe results, click anywhere in the tree, and then click on the alert definition you were working with. From here you can edit it.

Probing a report

You can test how a report is alerted on. This procedure has three major steps: set up a scheduled report to be alerted on and generate it, create an alerting definition for the report, and probe the alert definition.

Notes about report probe results:

- Some components, such as entry points and operations, may be less interesting to probe, as they return only their descriptions and settings.
- Actions, besides listing all their settings, show the number of times they have been executed. You can see in real-time the effect the report data had on the number of times the action executed by watching this number during the probing session.
- Enabling links displays the information in each line of alert data according to the selections in the Probe window. Since there is no filter all data is parsed and every property with its value appears. Depending on where the link is located in the alert

definition, a subset of alerting data lines appear according to how they are routed from previous components in the chain.

Procedure

1. Set up a scheduled report to be alerted on and generate it:
 - a. In the main user interface, navigate to a report that contains alerting data.
 - b. Using the **Display** menu, set the time range to 1 month.
 - c. Using the **Tools** menu, select **Schedule this Report**.
 - d. Set the schedule to **Every month**.
 - e. Click the **Alert** tab and select the Alerting Backend that contains the alert definitions to test with the data from this scheduled report.
 - f. In the tree, click **Scheduled Reports**, right-click the report, and then click **Launch now**.
 - g. The Report Data Adapter, which parses reports from the Frontend, is automatically created the first time it receives a report from the Frontend to be alerted on. The Report Data Adapter is named `AUTO_CREATED_APG_REPORT_ADAPTER`. On the Alerting interface, change its Time check parameter to 1 (minute), so that reports received from the Frontend are parsed and can be probed. You can change this back to a larger interval after the alert definitions are ready for production, to reduce the load on the server caused by needless polling.
2. Create an alerting definition for the report:
 - a. On the Alerting Frontend, click **Alert Definitions > Examples > Multiple Severities Comparator** and above the tree click the **Copy** and then the **Paste** icon. Rename it Utilization Over Month Alert.
 - b. Click the newly created alert definition in the tree to display its components.
 - c. Mouse over its entry point and click the **pen** icon to edit it.
 - d. In the **Edit Components** dialog that appears, delete the contents of the filter. You can change its name and description.
 - e. You can add a filter later, but first test the unfiltered content using a probe, and include a filter from the Probe UI for iterative testing purposes. Alert definitions pass their data to the next component only if the state data has changed from a previous data line in a series. This is useful, for example, if instead of using the Report Data Adapter, you use the Values Socket Listener Adapter to directly input UtilizationPct metric data about the file systems received by the Smarts-Collector through the Collector-Manager, and do not want to continually trigger an alert when a file system is over a certain value, but only the first time it crosses the threshold, and subsequently when it is under the threshold. In this example the report will be parsed once a month and all the data will be parsed to see which thresholds are crossed for each line of alerting data. Therefore, edit each Comparator condition to enable their Stateless check boxes.
 - f. Save and enable the alert definition.
 - g. To make sure the Report Data Adapter named `AUTO_CREATED_APG_REPORT_ADAPTER` is enabled, click it in the adapter list and click **Enable**. You can disable any other adapters so that you can concentrate on the report data being received from the APG Frontend, if this does not impact other required running alerts.

3. Probe the alert definition:

- a. From the Alerting page, click **Alert Definitions**.
- b. Click **Utilization Over Month Alert > Probe**.

The alert definition appears with the Probe window on the left, and space at the bottom of the screen for the probe results.

- c. Enable a few of the components and links of the alert definition to see the data that will be flowing through them by clicking each one.

For each component/link clicked, a probe result window appears at the bottom of the screen.

- d. Open up another browser window, return to the main user interface and on the **Scheduled Reports Management** page right-click the scheduled report you configured to be alerted upon and click **Launch now**.

Create alerts from scheduled reports

The metrics that appear on reports can be the source of an alert.

Process flow to alert on a scheduled report

An alert might enter the system as a metric from a scheduled report.

You can configure a report for alerting using the Schedule a report tool. Data from a report scheduled for alerting is converted into XML form and then parsed into alerting data by the APG Report Data adapter.

The alerting data is sent to the Alerting Backend, where it is evaluated by alert definitions. This alerting method depends on an alert definition that filters on the report name and one or more report metric names. The alert definition defines the conditions that will generate an alert, such as a value that is exceeded.

You can evaluate any metric that appears on the report. You can also use operations in the alert definition to create a new metric value from the data on the report, and generate alerts on the calculated value.

Prerequisites for creating alerts from scheduled reports

To alert on reports, all of the following elements must be in place.

- The Alerting Frontend and the Alerting Backend must be configured as described in [Configure the alerting engine](#).
- Communication between the APG Frontend and the Alerting Backend must be configured if they are on different servers. See [Configure alerting communication](#).
- An APG Report Data Adapter must be defined in the Alerting Frontend. This adapter is created automatically the first time a report is parsed by alerting.
- Alert definitions must exist in the Alerting Frontend to gather data from the reports. The procedure [Create a scheduled report alert](#) on page 238 includes how to create the appropriate alert definition.

Configure the alerting engine

Configure the service of the Alerting Backend to make alerts accessible to the Alerting Frontend.

This is an Administrator task. Configure the Alerting Engine with the following set of JVM arguments.

File	Required	Description
alerting.config.file	Yes	Configuration file containing the list of definitions, adapters, components and templates.
alerting.snmp.oid	Yes	SNMP OID used by the SNMP trap action.
alerting.max.crunching.thread		Number of threads used for the crunching of incoming data. Default is 5.
alerting.conf.dir		Directory containing the configuration schema files. Default is AlertingBackendDir/conf.
alerting.custom.conf.dir		Directory of the custom configuration files for some components. Default is AlertingBackendDir/custom.
alerting.suggestion.sampling		Sampling of properties catching for the suggestions. Default is 100.
alerting.suggestion.limit		The number of values kept for each property. Default is 1000.
alerting.mail.from		The From address of a mail action. Default is noreply@mybusiness.com.
alerting.probe.size		Probe buffer size. Default is 5MB.
alerting.probe.ttl		Probe time to live in seconds. Default is 60 seconds.

Although not mandatory, alerting is usually used with the APG values. To enable collecting these values, edit the Alerting connector in the `Collecting/Collector-Manager/Default/conf/collecting.xml` file from the APG installation directory. The connector should look like this:

```
<!-- This collecting component connects the collecting process to
the Alerting itself, using a
plain text socket. -->
<connector enabled="false" name="Alerting" conf="Collector-Manager/
Default/Cache-Connector/conf
/alertingconnector.xml" />
```

The file `alertingconnector.xml` creates a socket on port 2010, which is the listening port of the default APG Values Socket Listener adapter of the Alerting Backend. By enabling both the connector and the adapter, Alerting is able to trigger on APG values.

Configure alerting communication

If the APG Frontend and the Alerting Frontend are not on the same server, add lines to the `APG.xml` file so that the Frontend knows which Alerting Backend instances it can send reports to.

This procedure assumes that the APG Frontend and the Alerting Frontend and Backend are configured.

Procedure

1. On the server where the Alerting Frontend resides, open `alerting-frontend.xml`, located at `<APG>/Web-Servers/Tomcat/<Instance>/conf/Catalina/localhost`.
2. Copy the alerting properties.

For example, if you have two Alerting Backends configured, you copy these lines:

```
<Resource name="manager/Local Manager" auth="Container"
type="com.watch4net.alerting.jmx.AlertManagerFactory"
factory="org.apache.naming.factory.BeanFactory" user="admin"
pass="changeme" url="service:jmx:rmi://
//jndi/rmi://localhost:52569/jmxrmi" />
<Resource name="manager/Second Manager" auth="Container"
type="com.watch4net.alerting.jmx.AlertManagerFactory"
factory="org.apache.naming.factory.BeanFactory" user="admin"
pass="changeme" url="service:jmx:rmi://
//jndi/rmi://localhost:52570/jmxrmi" />
```

3. On the server where the Alerting Frontend resides, open `APG.xml`, located at `<APG>/Web-Servers/Tomcat/<Instance>/conf/Catalina/localhost` and paste the lines you copied from the `alerting-frontend.xml` file anywhere in the `<context>` tags.
4. If the Alerting Frontend references the local Alerting Backend, replace `localhost` in each copied entry with the IP address of the Alerting Backend.
5. Restart Tomcat.

Create a scheduled report alert

Follow these steps to configure alerts from scheduled report data.

Procedure

1. Schedule the report for alerting:
 - a. On the Console, navigate to the report that has the metric you want to alert on and click **Tools > Schedule this report**.
 - b. On the **Scheduling** tab, define how often and when to run this report.
 - c. On the **Alert** tab, select **Local Manager**.
 - d. Click **Save**.
2. Enable the APG Report Data Adapter.
 - a. Go to **Administration > Modules > Alerting > Local Manager > Adapters**.
 - b. Look for **APG Report Data** adapter. If it is in the list of adapters, use the table in the right pane to make sure it is enabled.
If the adapter is not listed, perform the following step.
3. (One-time step) Add the APG Report Data adapter:
 - a. On the User Console, go to **My Reports > Scheduled Reports**.
 - b. In the right pane, right-click the report that you scheduled for alerting, and choose **Launch now**.

This action adds the APG Report Data adapter to the list of adapters if it is not already there, and enables it.

4. Create an alert definition that includes:
 - a. A filter that identifies the report name and one or more metrics.
 - `'reportname==report name'`
 - `'name==metric name'`
 - b. Conditions that evaluate the metric and define outcome paths. You can optionally include operations that create new metric values from the entry metric data.
 - c. To make the alert appear in alerting reports, include SNMP Alert Trap and Clear Trap actions on the outcome paths. Other actions are also valid, such as log, email, or SNMP notifications.
5. Enable the alert definition.
6. Launch the report again, or wait for the next scheduled run.

Each time the report runs, the Alerting Backend receives the data and generates or clears alerts when conditions in the alert definition are met.

You can use probing to simulate conditions and test the results.

Report data parsing

The APG Report Data adapter parses the data in the scheduled reports that it receives.

The APG Report Data adapter generates alerting data from scheduled reports. The data and related properties depend on the report type.

Each data point in a report, such as a point in a graph or a table row with a value column, results in a line of data that can be evaluated by an alert definition.

Common properties

Each line of alerting data output by the APG Report Data adapter has the following properties.

```
timestamp
value
name
reportName
```

The displayed properties with their values from the header of the report are included with the lines of alerting data from the report. If it is a mixed report, the lines of alerting data from each child report have these common properties, although usually properties are used only in individual reports for a single device.

Tables

Each row of a table report results in one line of alerting data per value column that can be evaluated by an alert definition.

- The value column of the table is used for alert definition evaluation, unless something else is specified for evaluation, such as a field in a string operation.
- The column name of the value column is used as the property value of the name property. The name is what the report editor named the column.
- Data from other table columns are input as additional properties, with their column names used as the property names. The name is what the report editor named the column.

Graphs

In a standard time-series graph, each timestamp (point on the graph), generates a line of alert data that can be evaluated. In a graph with four lines (metrics), at each timestamp there are four lines of data.

The value of the property name is derived from the legend, even if the legend is hidden as it is in Mixed Reports.

For horizontal bar and pie charts, each category in the graph results in a line of alerting data that contains the following information:

- Relative value in percentage of the metric compared to the other categories.
- Value of the name property, which is the name of the report.
- Timestamp of each line of alerting data, which, for these types of reports, corresponds to when the report was generated rather than the time period the report covers. We recommend you do not use these types of reports for time-based alerting.
- Section property for each line of alerting data, which corresponds to the legend for the category reported on in the data.
- The absoluteValue, which is the absolute value of the metric for the category for the report period.
- The totalValue, which is the sum of the absolute values for every category in the graph.

Baseline graphs

Baseline graphs are like other timeseries graphs with these additional characteristics:

- One alerting data line is created for each timestamp with the value of the metric and its name, according to the contents of the legend, using the value of the name property.
- The baselineValue property is added, which is an average for the metric at the specified time over the last month.
- The baselineMin is added, which is the minimum value recorded for the metric at the specified time over the last month.
- The baselineMax is added, which is the maximum value recorded for the metric at the specified time over the last month.

You can use a formula to vary the length of the baseline instead of using the default, which is the past four weeks.