Hewlett Packard Enterprise

HPE Storage Operations Manager

Software Version: 10.20
Windows® and Linux® operating systems

User Guide

Document Release Date: August 2016
Software Release Date: August 2016

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

AMD is a trademark of Advanced Micro Devices, Inc.

© 2012 Google Inc. All rights reserved. Google™ is a trademark of Google Inc.

Intel®, Intel® Itanium®, and Intel® Xeon® are trademarks of Intel Corporation in the U.S. and other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

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

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

Red Hat® is a registered trademark of Red Hat, Inc. in the United States and other countries.

SAP®, SAP® BusinessObjects™, and SAP® BusinessObjects™ Web Intelligence® are the trademarks or registered trademarks of SAP SE in Germany and in several other countries.

UNIX® is a registered trademark of The Open Group.

Oracle Technology – Notice of Restricted Rights

Programs delivered subject to the DOD FAR Supplement are 'commercial computer software' and use, duplication, and disclosure of the programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, programs delivered subject to the Federal Acquisition Regulations are 'restricted computer software' and use, duplication, and disclosure of the programs, including documentation, shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

For the full Oracle license text, see the `open_source_third_party_license_agreements.pdf` file in the `license-agreements` directory in the SOM product download file.

Acknowledgements

This product includes software developed by the Apache Software Foundation.
(http://www.apache.org)

This product includes software developed by the Indiana University Extreme! Lab.
(http://www.extreme.indiana.edu)

This product uses the j-Interop library to interoperate with COM servers.
(http://www.j-interop.org)

Documentation Updates

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

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

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hpe.com/group/softwaresupport/search-result?keyword=.

This site requires an HP Passport account. If you do not have one, click the **Create an account** button on the HP Passport Sign in page.

Support

Visit the HPE Software Support web site at: **https://softwaresupport.hpe.com**

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

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

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

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to **https://softwaresupport.hpe.com** and click **Register**.

To find more information about access levels, go to:
https://softwaresupport.hpe.com/web/softwaresupport/access-levels

Contents

Chapter 1: Getting Started with SOM

The following topics introduce you to the main features of the Storage Operations Manager (SOM) console, tasks you can perform from the console and how to configure your browser for SOM.

- "Configuring Web Browsers for SOM" below
- "Configuring the Console" on page 22
- "Features of the Console " on page 23
- "Using the Console" on page 38

Configuring Web Browsers for SOM

Configure your web browser according to the information included here.

- "Configure Mozilla Firefox for SOM" below
- "Configure Mozilla Firefox Timeout Interval" on the next page
- "Configure Microsoft Internet Explorer for SOM" on the next page
- "Configure the Microsoft Internet Explorer Title Bar" on page 22

Configure Mozilla Firefox for SOM

By default, the SOM help opens in a new browser window. With Mozilla Firefox, you can sign in to only a single SOM session on each client system.

In the main SOM console window,  (the **Show View in New Window** / **Show Form in New Window** icon) opens a duplicate of the current view or form in a new browser window.

The number of windows generated can be controlled by configuring Mozilla Firefox so that SOM responds to requests in a new tab within the current Firefox window.

> Note: The browser context menu might be displayed on right-click from the SOM Console. However, these options do not work and you can ignore them.

To configure how Mozilla Firefox responds to SOM links:

1. In the Mozilla Firefox address bar, type `about:config` and then press **Enter**.
2. At the top of the displayed form, in the **Filter** field , type `newwindow`. A list of relevant attributes appears.
3. Double-click **browser.link.open_newwindow**.

4. In the **Enter integer value** dialog box, type one of the following choices:
   - **1** = Replace the current Firefox window/tab.
   - **2** = Open a new Firefox window.
   - **3** = Open a new tab within the current Firefox window.
5. Click **OK** to save your changes and close the dialog box.
6. Double-click **browser.link.open_newwindow.restriction**.
7. In the **Enter integer value** dialog box, type one of the following choices:
   - **0** = Use settings in **browser.link.open_newwindow**.
   - **1** = Ignore settings in **browser.link.open_newwindow**.
   - **2** = Use settings in **browser.link.open_newwindow** unless the URL contains other window instructions.
8. Click **OK** to save your changes and close the dialog box.

Configure Mozilla Firefox Timeout Interval

If you use the Mozilla Firefox browser and have timeout issues (for example, being prompted to continue before a map appears), try resetting the Mozilla Firefox timeout value:

1. In the Mozilla Firefox address bar, type `about:config`.
2. Select the **dom.max_script_run_time** entry from the list.
3. Increase the value displayed. For example, enter 0 (zero) to set the timeout value to infinity.

Configure Microsoft Internet Explorer for SOM

By default, the SOM help opens in a new browser window. You can sign in to multiple SOM sessions with Microsoft Internet Explorer. Use a different user name for each browser session.

In the main SOM console window, the **Show View in New Window** / **Show Form in New Window** icon opens a duplicate of the current view or form in a new browser window.

To control the number of windows generated, you can configure Microsoft Internet Explorer so that SOM responds to requests in a new tab within the current Explorer window.

> Note: The browser context menu might be displayed on right-click from the SOM Console. However, these options do not work and you can ignore them.

To configure how Microsoft Internet Explorer responds to SOM requests, follow these steps:

1. From the Microsoft Internet Explorer browser, select **Tools** → **Internet Options**.
2. Select the **General** tab.
3. Under the **Tabs** section, click **Settings**.
4. In the **Tabbed Browsing Settings** dialog, locate the radio box group labeled **When a pop-up is encountered**.
5. Make your selection:
   - **Let Internet Explorer decide...**
   - **Always open pop-ups in a new window**
   - **Always open pop-ups in a new tab**
6. Click **OK** to save your configuration and close the dialog box.
7. Click **OK** to close the **Internet Options** dialog and return to the browser window.

Configure the Microsoft Internet Explorer Title Bar

When using Internet Explorer, the browser settings determine whether the name of an SOM view or form displays in the title bar.

To configure Microsoft Internet Explorer to display form and view titles, follow these steps:

1. Open the Internet Explorer browser and click the **Tools** menu.
2. Select **Internet Options**.
3. Navigate to the **Security** tab, **Trusted Sites**, **Custom Level**, **Miscellaneous** section.
4. Disable the **Allow websites to open windows without address or status bars** attribute.

Configuring the Console

You can configure the following user interface features:

- The console timeout interval.
- The initial view to display in the SOM console.

To configure user interface features, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **User Interface** > **User Interface Configuration**. The User Interface Configuration form is displayed.
2. Make your Global Control configuration choices. (See "Attribute" on the next page column in the table.)
3. Click **Save and Close** to apply your changes.

To apply your Console Timeout or Initial View configuration changes, sign out of the SOM console. Your changes should take effect after restarting the console.

Attribute	Description
Console Timeout	Use this attribute to change the timeout interval in days, hours, and minutes. The default session inactivity timeout value is 18 hours. The minimum timeout value is 1 minute. After this period, if no mouse movement occurs, the console locks and the user is prompted to sign in again.
Initial View	Use this attribute to specify the initial view to be automatically displayed in the console by default. If you do not want the default view, select **None (blank view)**. Select a view from the drop-down menu list.

Features of the Console

The SOM console is the graphical user interface of the SOM application. The main components of the SOM console are labeled in the following screen capture and explained with a brief description.

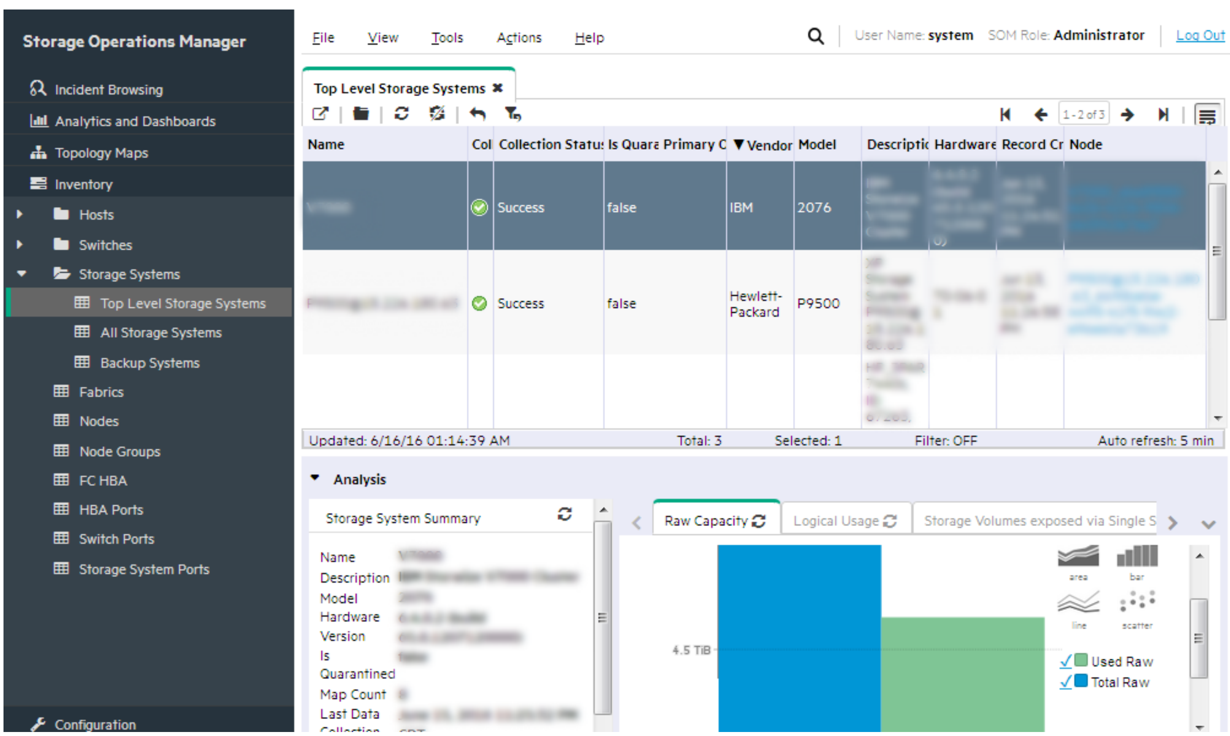

Title Bar

Used to identify the application you are running. The top-right corner contains the standard browser buttons for closing and resizing the SOM console window.

Menu bar

Menus available in the SOM console:

- File
- "View Menu" on page 37
- "Tools Menu " on page 36
- "Actions Menu" on page 36
- "Access Help" on page 41

> **Tip:** To expand SOM menus, click the menu or press Ctrl+Shift and the underlined character in the menu name (if any). SOM uses Ctrl+Shift (instead of Alt) to avoid the browser's main menu behavior. For example, in SOM, press Ctrl+Shift+H, then u for **Help → Using the SOM Console**. If the SOM menu does not expand as expected, your browser configuration already overrides the SOM configuration for that keyboard combination of Ctrl+Shift+<*ASCII character*>.

Workspace navigation panel

Helps you navigate between workspaces and views. For more information, see "Display Views" on page 43 and "Workspaces" on the next page.

Workspace

A context that represents your current scope of interest and work. Workspaces provide a means of grouping views for a related purpose or task flow. Multiple views are available in each workspace. See "Workspaces" on the next page.

Console message bar

Alerts about any problems with the application.

User, Role, and Sign Out button

Your current user name, and role assignment. Your role assignment determines what you can see and do within the console.

Breadcrumb trail

Title of the view you selected from the workspace navigation panel and the breadcrumb trail. Each view provides access to a group of objects. More details about each object are available when you double-click the object to display that object's form. The breadcrumb trail appears in the title of the view, so you can easily navigate to previously accessed views and forms.

View Toolbar

Tools available within the current view or form. These tools enable you to remove any data filters that you previously applied, restore any columns that you previously hid, and manipulate objects within the view. The drop-down selectors enable you to modify the default filter values applied to the visible data. See "Toolbars" on page 32 for more information.

Content Pane

Displays the currently selected view or form.

Status Bar

In table views, the status bar shows the following information:

- Updated: The date and time when the view was last refreshed.
- Total: The current number of objects in the database that match the criteria for this table (each row displays data about one object).

  > **Tip:** To reduce the number of objects displayed, so that you see only the objects of interest, use filters.

- Selected: Indicates the number of rows selected in the table.
- Filter: Indicates if the currently displayed data is a filtered subset of available objects.
- Auto Refresh: Indicates the current refresh time interval.

In map views, the status bar shows the following information:

- The number of nodes displayed on the map.
- Auto status refresh: Automatic refresh rate for the Refresh Status option.

In both table and map views, the status bar displays the Last Updated time to indicate the time at which the view was last refreshed.

Analysis Pane

Displays information dynamically about the object selected in the content pane. Additional information can include information such as capacity utilization, performance metrics, member nodes and child node groups.

> **Note:** This pane remains blank until an object is selected.

Workspaces

A workspace is a collection of views that represent a scope of interest and work. Workspace is a collection of views with a related purpose or task flow.

When you click a workspace, the views associated with that workspace display are displayed. After you select a view, the view display panel shows the requested data.

The views within workspaces provide convenient access to information associated with each object type represented. A view displays all objects of a given type that meet the filter criteria specified for that view.

SOM includes the following workspaces:

- Incident Browsing

  The **Incident Browsing** workspace provides views that contain information that SOM considers important to bring to your attention regarding your storage environment.

- Dashboards

  Use the **Dashboards** workspace to view at-a-glance information about your storage network. Dashboard views enable you to easily compare and quickly isolate the information you need to manage your storage environment.

- Topology Maps

  The **Topology Maps** workspace includes the system topology map view by default.

  > **Tip:** The following changes are not automatically visible in the **Topology Maps** workspace folders:
  >
  > - Add one or more node groups
  > - Delete one or more node groups
  > - Modify a node group hierarchy
  >
  > To view any of these changes, click **Refresh** in the upper right-hand corner of the workspace. **Refresh** collapses the node group maps folders. Expand each folder of interest to view the updated node group map list.

- Inventory

  Each view in the Inventory workspace contains information related to the object listed. For example, the Nodes view contains information related to the node objects.

  > **Note:** If your role includes Administrator privileges, you can access the Configuration workspace.

- Configuration

  The Configuration workspace enables you to configure SOM for your storage environment.

Views

Views contain information about the objects in your storage environment. A view is a collection of related objects that are depicted as a table or map with graphical representation of connectivity information.

- Table views – Presents summary information for a list of objects in a sorted order. Examples include viewing a list of storage systems sorted by collection status and filtered by vendor.

  Table views display data in tabular format. Each row displays data about one object. If there are more rows than fits on a single screen, you can scroll through the table view using the scroll bar.

- Map Views – Displays a graphical representation of connectivity information that shows relationships among objects. A map view is a powerful way to display information about your storage environment.

Views Available from the Console

The following views are available from the console:

- "Analytics and Dashboards View" on the next page
- "Topology Maps View" on page 29
- "Inventory Views" on page 29

Incident Browsing Views

The **Incident Browsing** workspace includes views for incidents filtered on the incident's Lifecycle State. For example, the Open Incidents view displays all Incidents with Lifecycle State values other than **Closed**. The Closed Incidents view displays all incidents that have a Lifecycle State of Closed. This workspace also includes a view for SNMP trap incidents and another view for displaying all incidents.

The following views are available in the Incident Browsing workspace.

Views for Incident Browsing

View Title	Description
Open Incidents	Lists all incidents with lifecycle state values other than Closed.
Closed Incidents	Lists all incidents that have a lifecycle state of Closed.
All Incidents	Shows all incidents. This view is useful for determining all of the incidents that have been generated within a certain time period.
SNMP Traps	Lists all of the traps that were received from devices in your network environment. Your

Views for Incident Browsing , continued

View Title	Description
	SOM administrator must configure specific traps before they are displayed within SOM incident views.

Analytics and Dashboards View

The **Analytics and Dashboards** workspace comprises dashboards that display data pertaining to the entire storage environment, a specific object (for example, storage system), or a group of objects.

Dashboard panels contain a variety of tables and charts that provide at-a-glance information, enabling you to easily compare and quickly isolate information.

There are two types of dashboard views:

- Dashboards available in the Analytics and Dashboards workspace
  These dashboards contain information for an entire set of objects managed by the SOM management server.

  To view a dashboard, click **Analytics and Dashboards** and select a dashboard.
- Dashboard of an object
  An individual dashboard is specific to a selected object (hosts, storage systems, switches, and so on). For example, the Host Capacity panel in the Host Dashboard, displays capacity information for the selected host.

  To view an individual object dashboard, select **Open Dashboard** from the context menu.

  Object-based dashboard views include the same data available in the Analysis pane for a selected object.

  > **Tip:** Use the breadcrumbs in the title bar to return to the previous screen from a dashboard view.

The following dashboards are available:

Dashboard	Description
Environment Capacity	Overall capacity utilization in the environment.
Asset Dashboard	Number of discovered devices based on Device Family, Device Vendor, or the OS Type of a device.
Collection Status Dashboard	Data collection statuses and quarantined devices in the environment (storage systems, hosts, and switches).
Storage Systems DTT Analytics	Number of days to threshold values for the capacity utilization (Raw Used, Actual Allocated, and Actual Used) of storage systems or pools.

Dashboard	Description
Storage Pools DTT Analytics	Number of days to threshold values for the capacity utilization (Actual Allocated, and Actual Used) of storage pools.
ThP Analytics for Virtual Servers	Number of datastores that are susceptible to an outage of physical disk space in a VMware® virtual environment.
Storage Systems Unused Volumes Analytics	Storage that can be reclaimed at storage system and volume levels.

Topology Maps View

In a map view, nodes, ports, and connectors are represented as symbols on the map. The lines between the nodes represent the connection or relationship between these objects.

You access map views from the Topology Maps workspace. You can also open maps from table views using the **Actions** menu. For more information, see "Actions Menu" on page 36.

The following map view is available from the console.

View	Description
System Topology	Displays the physical connectivity of all the storage elements in your network. The topology shows the fabric and network connections among the discovered devices. The map view changes dynamically as new devices are discovered in the environment.

Inventory Views

The following views are available from the Inventory workspace.

View	Description
Hosts	Lists the hosts discovered by SOM. Hosts are further categorized as discovered hosts, virtual servers, virtual machines, inferred hosts, created hosts, and host clusters. Sorting this view by status lets you see all of the nodes that are down or somehow disabled.
Switches	Lists the information about all the switches discovered by SOM including the physical and virtual switches.
Storage Systems	Storage systems discovered by SOM. The view is categorized broadly as top level storage systems (includes file storage, block storage, cluster storage), all storage systems that includes physical storage systems and their underlying nodes, and backup systems that provide data protection and disaster recovery.

View	Description
Fabrics	List of fabrics associated with the switches discovered by SOM.
Nodes	Lists the nodes associated with the elements discovered by SOM.
Node Groups	Lists the node groups provided by SOM and created by the administrator.
FC HBA	Displays the total list of host bus adapter cards that are discovered and managed by SOM in the environment.
HBA Ports	Displays the entire list of host bus adapter ports that are discovered and managed by SOM in the environment.
Switch Ports	List of switch ports in the environment that are discovered and managed by SOM.
Storage System Ports	Lists the storage system FC ports in the environment that are discovered and managed by SOM.

Configuration Views

View	Description
Discovery Addresses View	Lists the status for all IP addresses hosted on Nodes that are discovered by SOM. (The Node form: IP Addresses tab also displays relevant addresses from this view.)
Discovery Ranges View	Displays list of IP address ranges that are configured for discovery. You can start scanning a selected address range from this view. After scanning the status is displayed in the view.
Discovery Credentials View	Displays list of discovery credentials.
Tenants View	Displays list of tenants configured by the administrator. You can delete a selected tenant from this view.
Data Collection Policies View	Lists the policies configured for data collection.

View	Description
BlackOut Periods View	Lists the blackout periods that are configured.
Data Collection Control View	Lists the default data collection level based on device profiles.
Monitoring Policies View	Displays list of policies configured for collecting performance metrics.
Collectors View	Displays the default collectors provided by SOM.
Node Groups View	Lists the node groups provided by SOM and created by the administrator.
Monitoring Groups View	Displays monitoring groups that are created for collecting performance metrics.
Host Inference Rules View	Displays rules created for inferring hosts without discovering them. You can use this view to manually run a rule and delete hosts inferred by a rule.
User Accounts View	List of user accounts that are created by the administrator.
User Groups View	List of default user groups provided by SOM and user groups that are created by the administrator.
User Account Mappings View	Displays the mapping between user accounts and user groups.
Security Groups View	List of default security groups provided by SOM and security groups that are created by the administrator.
Security Group Mappings View	Displays the mapping between various user groups and security groups.

Toolbars

The following toolbars are available from the SOM interface:

- Table Views Toolbar
- System Topology Map Views Toolbar
- "The Form Toolbar" on page 35

Table View Toolbar

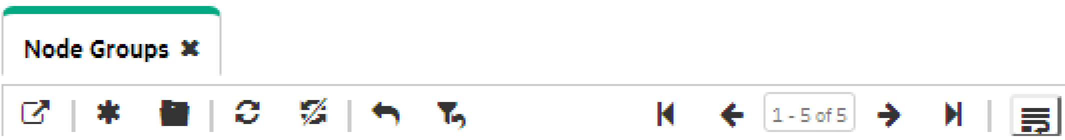

Use the table view toolbar to perform the following tasks within the displayed view.

View Toolbar Icons

Icon	Description
	Show View in New Window. Displays the current view in a new window.
	New. *SOM Administrators only*. Opens the form to create a new object instance.
	Open. Displays the form for the selected object. See "Use Forms and Analysis Panes to Access More Information About an Object" on page 68.
	Refresh. Refreshes the current view. See "Refresh a View" on page 44 for more information. Restarts periodic refresh if it has been disabled.
	Stop Periodic Refresh. Temporarily disables the periodic refresh of a view. See "Stop Periodic Refresh of a View" on page 48 for more information.
	Restore Default Settings. Resets default settings, including the resizing of table columns, sort selections, and filters. Any hidden columns are restored to the view. See "Hide a Column" on page 46.
	Restore Default Filters. Clears any currently applied filters. See "Filter a Table View" on page 48.

View Toolbar Icons, continued

<table>
<tr><th>Icon</th><th>Description</th></tr>
<tr><td></td><td>Delete. If your role permits, deletes the selected object instance and any objects contained in that object. For example, deleting a host also deletes the card and port instances associated with that host, and the history of those objects.</td></tr>
<tr><td></td><td>Close. Close the current view.</td></tr>
<tr><td colspan="2">The page controls only appears when viewing tables. They let you page through table information by rows.

Use Previous or [Page Up] to move up one page.

Use Next or [Page Down] to move down one page.

Use First or [Home] to move to the top of the table.

Use Last or [End] to move to the end of the table.

Use the [↑] up arrow key to scroll up one row.

Use the [↓] down arrow key to scroll down one row.

The page control displays the total number of rows in the current table, as well as which group of rows within that total is currently visible.

If the page control displays <code><maximum_table_size value></code>, then the table row count exceeds the maximum table size specified by SOM.

To view the actual table size, look for the Total value in the table status bar. SOM displays the total number of rows for the table, followed by the display limit set for the table.

When the table size exceeds the maximum table size value, also note the following:
SOM recomputes the actual number of rows in the table each time you refresh the table view or update the table filter.
When you scroll to the last row of the table, SOM displays a dialog box explaining that the table is larger than the specified limit and recommending that you filter the table view. See "Filter a Table View" on page 48 for more information about how to filter a table view.</td></tr>
<tr><td>or</td><td>Toggles text-wrap on or off.</td></tr>
</table>

System Topology Map Toolbar

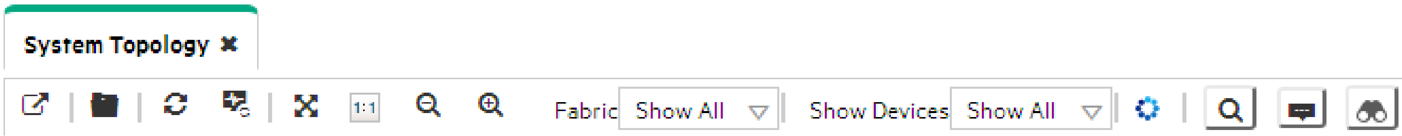

The System Topology Map view toolbar lets you perform the following tasks within the displayed map.

System Topology Map Toolbar Icons

Icon	Description
	**Show View in New Window**. (Only available from the main Console .) Displays the current view in a new window.
	**Open**. Displays the form for the selected object. See "Use Forms and Analysis Panes to Access More Information About an Object" on page 68.
	**Refresh**. Refreshes the current view. See "Refresh a View" on page 44 for more information.
	**Refresh Status**. Refreshes only the status of each node in the map. See "Refresh Node Status on a Map" on page 65 for more information.
	**Fit Content**. Adjusts the size of the node symbols so that all members of the node group fit within the current window. See "Adjust the Zoom Factor" on page 63 for more information.
	**Actual Size**. Cancels any current zoom setting. See "Adjust the Zoom Factor" on page 63 for more information.
	**Zoom Out**. Zooms out 25% of current size. See "Adjust the Zoom Factor" on page 63 for more information.
	**Zoom In**. Zooms in 25% of current size. See "Adjust the Zoom Factor" on page 63 for more information.
	**Close**. Close the current view.
 or 	**Find**. Toggles on or off highlighting the identified Node in the current map and ensures that node is in the map's display area. See "Find a Node in a Map" on page 61.
 or 	**Tool Tips**. Toggles on or off Tool Tips information that pops up when the mouse cursor is placed over an object on a map. See "Control Tool Tips Information on a Map" on page 64 for more information.

The Form Toolbar

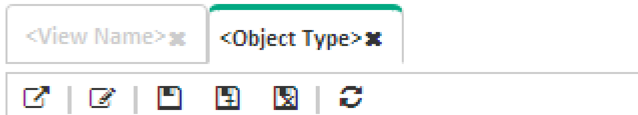

If your role permits, the toolbar lets you perform the following tasks within the form. The group of available actions can change from form to form:

> **Tip:** You can also right-click any object in a table or map view to access the items available within the **Actions** menu.

Form Toolbar Icon Actions

Icon	Action
	**Show Form in New Window**. Displays the current form in a new window. **Note:** SOM closes the current form before displaying the form in a new window.
	**Show Analysis**. Displays the analysis pane information for the current form. See "Analysis Pane" on page 36 for more information.
	**Save**. Saves the current form.
	**Save and New**. Saves the current form, and opens a new empty form where you can create a new object instance.
	**Save and Close**. Saves and closes the current form.
	**Refresh**. Refreshes the data in the current form.
	**Delete**. Deletes the selected object instance and any objects contained in that object. For example, deleting a host also deletes the card and port instances associated with that host, and the history of those objects. **Note:** When you delete an object instance that is created using a filter, such as a node group, SOM deletes only the node group filter. SOM does not remove the nodes that belong to the selected group.
	**Close**. Closes the current form.

Analysis Pane

The analysis pane displays related details about a selected object. SOM performs the appropriate analysis on a selected object and displays the most significant information in the Analysis pane. Hyperlinks within the tabs of the Analysis pane allow you navigate to additional information.

The Analysis pane is available below most workspace views, such as, inventory, topology, configuration views, and so on. From a workspace view, select an object or an element, to see its analysis information at the bottom of the window.

> **Note:** If the Analysis pane is not visible, point to the bar at the bottom of the window to view ↕ resize icon, drag the bar to view the Analysis pane.

Note the following:

- The analysis pane remains empty until an object is selected.
- To update the displayed information, click **Refresh** in the tabs of the analysis pane.

Tools Menu

SOM Tools Menu Options

Tool	Description
Find Element	Searches the SOM database for elements using the string of characters (*not case-sensitive*) that you provide. For more information about finding an element, see "Find Elements" on page 433.
Signed In Users	View a list of the SOM users who are currently signed in to SOM.

Actions Menu

The actions available to you depends on the object selected. If no actions are available for a particular object, the Actions menu is empty.

To perform an action, select an object, and then select an action from the **Actions** menu. The **Actions** menu is accessible from the SOM console main menu toolbar and from the menu toolbar in any view or form that is opened in a new window.

- **To invoke an action from a table or map view**:
  a. If you do not have a view displayed, from the workspace navigation panel, select a view.

b. Do one of the following:
   - In a table view, click a row.
   - In a map view, click the object of interest.

   Tip: For multiple selections, use Ctrl+click.

c. On the menu toolbar, select the **Actions** menu.

   Tip: You can also right-click any object in a table or map view to access the items available within the **Actions** menu.

d. Select the action you want to perform from the list of available actions.

- **To invoke an action from a form**:

  a. If you do not have a form open, from the Workspaces navigation panel, select a table view.

  b. From the table view, double-click the row representing an object instance (for example, **node groups**).

  c. From the **Actions** menu, select an action. For example, select **Actions** → **Node Group Details** → **Show Members (Include Child Groups)** to view the members of a node group.

When invoking actions, note the following:

- If you are running an action that modifies attributes on a form, the action takes effect immediately. You need not click Save.
- An action might cause a new window to open.
- If you selected the wrong number of objects for an action, you can cancel the selection of all objects by clicking twice in the row. (The first click selects the object and the second click cancels the selection of the object.)

View Menu

The View menu from the console provides the following options:

- "Refresh a View" on page 44
- "Select Multiple Objects in a View" on page 44

Using the Console

The main window of the SOM console is the starting point for navigation. From the main window, you can perform the following tasks:

- "Display Information About SOM" on page 39
- "Tools Menu " on page 36
- "Actions Menu" on page 36
- "Display Views" on page 43
- "Customize a Dashboard View" on page 45
- "Customize a Table View" on page 45
- "Use Map Views" on page 61
- "Use Forms and Analysis Panes to Access More Information About an Object" on page 68
- "Work with Objects" on page 67
- "Customize Charts" on page 70

Display Information About SOM

Two menu items that provide current information about your installed SOM:

- **Help → System Information**

  The **System Information** window provides current information about SOM.

  > **Note:** The information available depends on your assigned SOM role.

  Within the **System Information** window, click the icon for access to the help information.

  The **System Information** window has three tabs as follows:

  a. "System Information: Product Tab" below.

  b. "System Information: Server Tab" on the next page.

  c. "System Information: Extensions Tab" on page 41.

- **Help → About HPE Storage Operations Manager Software**

System Information: Product Tab

To display the SOM system information, click **Help → System Information**.

The **Product** tab displays the following information about SOM:

- Product name and version number
- Locale Information (language) for the current SOM session:
  - Client locale
  - Server locale
  - SNMP string encodings
  - Web browser
- The current status of SOM System Health:
  - Status
  - Last updated

  The following table describes the possible SOM health status values:

Status	Description
Normal	Indicates that SOM is not experiencing any problems.
Warning	Indicates performance issues that are not significantly affecting SOM.
Minor	Indicates problems that might result in out of date data.
Major	Indicates problems that are significantly affecting the SOM management server's operations, but are not yet critical. Major status usually indicates that some action is required.
Critical	Indicates the SOM is not functioning. For example, SOM is out of memory, all database connections are lost, or a major SOM component has failed.

- User Information about the current SOM user:
  - User Name that you used when logging into SOM.
  - SOM role to which you are currently assigned.
  - User groups to which you currently belong.
- For license information, click **View Licensing Information**.

  The **Type** of license. It can be one of the following:
  - Instant-On
  - Premium
  - Ultimate

  For information about license types and to purchase additional licenses, contact your HPE Sales Representative.

System Information: Server Tab

To display the SOM system information, click **Help** → **System Information**.

The **Server** tab displays the following information about the SOM server:

- Hostname
- IP Address
- Official Fully Qualified Domain Name (FQDN)
- User Account and User Group information obtained from (either the SOM database or a directory service using LDAP)
- Operating System

- Install Directory
- Data Directory
- Available Processors
- SOMs Free / Allocated Memory (% Free)
- SOMs Maximum Attemptable Memory

System Information: Extensions Tab

To display the SOM system information, click **Help** → **System Information**.

The **Extensions** tab lists the SOM extensions deployed on your SOM management server.

Display SOM License Information

Select **Help** → **About HPE Storage Operations Manager Software** to display the following information:

- The current version number of SOM.
- The **Type** of license. It can be one of the following:
  - Instant-On
  - Premium
  - Ultimate

  For information about license types and to purchase additional licenses, contact your HPE Sales Representative.

Access Help

To access the help

1. In the SOM Console main menu, click **Help**.

   Note: Within the SOM user interface, pressing [F1] does not access context-sensitive help.

2. From the **Help** menu, you can access all information described in the following table.

Help → Links

Help Link	Description
SOM Documentation Library	Provides access to the documentation landing page that contains resources for the SOMdocumentation.

Help → Links , continued

Help Link	Description
HPE Software Support Website	Provides access to the HPE Software support web site.
System Information	Displays product, server, and health information about SOM.
About HPE SOM Software	Provides access to the current licensing information about SOM.

Search the Help Topics

To search for specific information across all help topics, follow these steps:

1. In the navigation pane of the Help window, click the **Search** tab.
2. Type in a search string (see table).
3. Click the **Search** button. The order of the resulting list of topics is based on a ranking order, with highest ranking topics at the top of the list.

Search Variables

Description	Variable	Example
Search for one or more words. When you enter a group of words into the search field, "or" is inferred.		host switch
Search for a phrase.	" " (wrap a text string in quotes)	"navigation pane"
Search for "either of" or "any of" specific strings.	OR (case insensitive) | (pipe symbol)	host OR switch OR asset "host capacity"|"switch capacity"
Search for two or more specific strings.	AND (case insensitive) + (plus symbol) & (ampersand)	presented AND storage AND host "presented storage"+host "presented storage"&"host"

Search Variables , continued

Description	Variable	Example
Search for all topics that do not contain something.	NOT (case insensitive) ! (exclamation mark)	NOT switch ! switch
Search for all topics that contain one string and do not contain another.	^ (carat symbol)	host ^ switch
Combinations of the above.	( ) parenthesis	capacity and (host or switch) host or node (!group)

Note: Results returned are case insensitive. However, results ranking takes case into account and assigns higher scores to case matches. Therefore, a search for "templates" followed by a search for "Templates" would return the same number of help topics, but the order in which the topics are listed would be different.

Mark Your Favorite Help Topics

Use the Favorites tab in the help system to set favorites for your commonly used help topics.

When using this feature, note the following:

- This feature is not related to the Favorites option in your web browser.
- The help topic favorites list is deleted when you delete the cookies in your web browser.

Work with Views

The following topics provide information on how you can work with views:

- "Display Views" below
- "Refresh a View" on the next page
- "Select Multiple Objects in a View" on the next page

Display Views

Views contain information about the objects in your network. A view can be a table view or a map view.

Note the following about accessing views:

- When you select another view from the workspaces navigation panel, the selected view replaces the current view.
- If you open a view using the **Show View in New Window** icon, the view opens in a new window.
- If the view has more than one page of information, use the scroll bar or the page controls to navigate through each page of the view.

To display a view, follow these steps:

1. Click a workspace name in the Workspaces navigation panel. The workspaces provided by SOM are:
   - Analytics and Dashboards
   - Topology Maps
   - Inventory
   - Configuration
2. Select a view.

Refresh a View

You can manually refresh a view at any time so that you are viewing the latest set of information.You cannot change the automatic refresh rate that is set by SOMfor each view.

To refresh a view, do one of the following:

- Display any view, then select the **View** →.**Refresh**.
- To refresh a table view, click the Refresh icon in the table view.

  The table view status bar displays the refresh rate and whether the refresh rate is enabled or disabled. (If disabled, clicking the Refresh icon enables periodic refresh.)
- To refresh a map view, click the Refresh icon in the map view toolbar to update changes in node placement, nodes added, and nodes deleted.

Select Multiple Objects in a View

You can select or cancel the selection of multiple objects when using a table or map view. This feature is useful when you want to access details or invoke an action on multiple objects, such as nodes, hosts, or storage systems.

Multiple Objects in Table Views

> **Tip:** Look in the status bar of each table view to see the number of objects currently selected as well as

the total number of objects in the view.

To select multiple objects in a table view:

Press Ctrl-Click to select the row for each object you want to select.

To cancel the selection of an object in a table view:

Select the row for each object again to cancel the selection.

Multiple Objects in Map Views

To select multiple objects in a map view:

Do the following:

- Use Ctrl-Click to select each object of interest on the map.

Each object you select changes to indicate it has been selected.

To de-select an object in a map view:

Re-select an object of interest on the map.

Each object you de-select returns to normal on the map.

Customize a Dashboard View

To minimize or maximize a dashboard panel, click the ▾ (minimize) button or ▸ (maximize) button available in the upper-left corner of a given panel.

Customize a Table View

Table views display data in tabular format. Each row displays data about one object. If there are more rows than fits on a single screen, you can scroll through the table view using the scroll bar. If a table contains more rows than the maximum limit set for a table, filter your table view to reduce the number of rows.

Note: SOM displays a message indicating an error in fetching data if you leave an inventory view idle for more than five minutes. Refresh the browser to resolve this issue.

From a table view, in addition to the tasks accessed with the view display panel toolbar, you can perform the following tasks:

- "Resize a Column" on the next page
- "Hide a Column" on the next page
- "Display a Hidden Column" on the next page
- "Select all Rows in a Table" on page 47

- "Sort Column Data " on the next page
- "Stop Periodic Refresh of a View" on page 48
- "Filter a Table View" on page 48
- "Restore Table View Defaults" on page 59
- "Export Table Information" on page 59
- "Limits to View Settings" on page 60

The following customizations are saved across browser sessions:

- Column width
- Hidden columns
- Sort column and order
- Column filters
- Quick filter value

Resize a Column

You can resize columns using your mouse.

To resize a column in the table:

1. Mouse over the edge of the column until you see a ↔ resize icon.
2. Drag the column edge to the width you want.

Hide a Column

If you find you no longer want to include a column of information in your view, you can hide a specified column.

To hide a table column:

1. Right-click the column of interest.
2. Select **Visibility**.

   The list of column names appears.

3. Click to clear the check box ☐ that precedes the name of the column you want to hide.

Display a Hidden Column

If you want to display a hidden column.

To display a hidden table column:

1. Right-click the column of interest.
2. Select **Visibility**.

The list of column names appears.

3. Click to check the check box that precedes the name of the column you want to display.

Select all Rows in a Table

If you want to select all rows in a table.

To select all rows in a table:

1. Select any row in the table.
2. Do one of the following:
   - Press **Ctrl-a**.
   - Right-click any row in the table view, and select **Select All**.
3. The table view data appears highlighted.

Sort Column Data

By sorting columns, you can get the most important information at the top of your table. For example, at times you might want to view all storage systems for which data collection is not currently enabled.

To sort by columns:

1. Right-click the column heading or data cell on which you want to sort.
2. To sort the column in ascending order, select **Sort:** → **Ascending**
3. To sort the column in descending order, select **Sort:** → **Descending**

When sorting column data, note the following:

- You can click the column header to initiate a sort on the column values. Clicking the column heading again, reverses the sort direction.
- SOM might provide table views in which sorting is disabled for one or more columns.
- Sorting tables that contain large amounts of data (for example, viewing all interfaces or incidents), can sometimes result in slow response times. In this case, it is better to first filter the table information so that it contains only the values of interest before sorting the remaining data.

More About Sorting

When sorting table columns, note the following:

- You can sort on only one column heading at one time.
- Uppercase letters are sorted separately from lowercase letters.
- SOM sorts some table columns using lexicographical ordering. This might produce unfamiliar orders for strings such as object IDs that contain numbers. For example you might expect the following order when sorting the hardware version data type:

- 1.3.6.1.4.1.1
- 1.3.6.1.4.1.3
- 1.3.6.1.4.1.20

Using lexicographical ordering, these system object ID values are ordered as follows:

- 1.3.6.1.4.1.1
- 1.3.6.1.4.1.20
- 1.3.6.1.4.1.3

- Your sort choices are saved across user sessions.

Stop Periodic Refresh of a View

You can manually stop the periodic refresh of the group of items displayed in a table view at any time.

> **Note:** The status of this group of objects is always updated on a regular basis, you are only stopping updates to the SOM console based on network objects being added or deleted from the SOM database.

SOM's status bar displays the refresh rate and whether the refresh rate has been disabled.

To stop the periodic refresh of a table view:

1. Click the Stop Periodic Refresh icon. In the bottom right corner of the SOM console, the following message displays:

   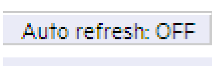

2. If you want to restart the refresh rate, click the Refresh icon in the view display panel toolbar. In the bottom right corner of the SOM console, the following message displays:

   

   > **Note:** You cannot change the refresh rate. SOM sets that default rate for each view.

Filter a Table View

When using table views, you can reduce the amount of information displayed by filtering a view using one of the object's attribute values.

Filtering a table view is also useful to reduce the number of rows when a table contains more rows than the maximum limit set for a table. See "Toolbars" on page 32 for more information about how SOM indicates the table has exceeded the maximum limit specified.

When a view is first displayed, it displays a set of filtered columns based on the view definition provided by SOM.

> **Note:** The view status bar indicates if one or more filters have been set for the view. FILTER:ON indicates that one or more filters have been set. These are filters that you can modify. FILTER:OFF indicates no modifiable filters have been set for the view. These views might have default filters.

When specifying filters, you can perform the following tasks:

Column Selection Filters

- "Filter by Attribute Value" on the next page
- "Modify a Table View Filter" on page 54
- "Remove a Filter" on page 57
- "Restore Default Filters" on page 58
- "Display Current Filter Settings" on page 58

When using filters, note the following:

- You can filter on multiple table columns. The resulting filter is a logical AND of the filters for all of the columns.
- SOM might provide table views in which filtering is disabled for one or more columns.
- SOM restricts certain filter operations or options per data type. Only the filter options that apply to the attribute data type appear. The data types and valid filter options are described in the table below.
- A different subset of these filter options appears depending on whether you are clicking a data cell, a column header, or a blank row. A data cell filter menu includes filters that use the value of the selected data cell.

Filter Options Available

Data Type	Valid Filter Options
All data types	Equals this value Not equal to this value See "Filter by Attribute Value" on the next page for more information.
Text (String) Numeric	Create Filter...

Filter Options Available, continued

Data Type	Valid Filter Options
Boolean data types	Is true Is false
Numeric (Integer, IP address, and date)	Greater than or equal to this value Less than or equal to this value See "Filter by Attribute Value" below for more information.

Filter by Attribute Value

When specifying filters based on attribute (column) values, some filter options require that you specify the value by selecting the value in an object instance and some require that you open a Create Filter... dialog to specify the value. You can also specify a filter based on whether an attribute contains a value.

- "Select Filter Values" below
- "Create a Filter to Specify Values" on the next page
- "Use Null Value Filters" on page 54

Select Filter Values

The following filter options require an attribute value as the basis for the filter.

> **Note:** When using the filter options listed below, first right-click the value in the table on which you want to filter.

Filter Settings

Filter Option	Description
Equals this value	SOM displays only instances that contain the attribute value you specify.
Not equal to this value	SOM displays only instances that have an attribute value that does *NOT* contain the value specified. You can specify multiple values for this filter; but can only supply one value each time you select this option. **Note:** For certain attributes, this option can also filter values that are "empty" or null.

Filter Settings , continued

Filter Option	Description
Greater than or equal to this value	SOM displays only instances that contain the attribute values that are greater than or equal to the value you specify.
Less than or equal to this value	SOM displays only instances that contain the attribute values that are less than or equal to the value you specify.
Is true	SOM displays only instances that have an attribute value that contains the value **true**.
Is false	SOM displays only instances that have an attribute value that contains the value **false**.

You can change a filter at any time. SOM saves filters per user so that the filters you specify are maintained during subsequent user sessions.

To filter your view by selecting an attribute value in the table:

1. Right-click the attribute value on which you want to base your filter.
2. Select one of the following filter options:
   - **Equals this value**
   - **Not equal to this value**
   - **Greater than or equal to this value**
   - **Less than or equal to this value**
   - **Is true**
   - **Is false**

   SOM displays a table view of all instances that have been selected based on the filter option and attribute values you specified or selected.

   Each filtered column is indicated using the filter icon.

Create a Filter to Specify Values

SOM allows you to provide attribute values on which you want to filter. You can provide attribute values by creating a filter for any of the following types of values:

- Text (String)
- Numeric

To filter information by specifying one or more values:

1. Right-click the column or attribute on which you want to filter.
2. Select **Create filter...**

   SOM displays the filter dialog that is appropriate for the selected column's data type.
3. In the Create Filter ... dialog:
   a. Select a Filter Option (see "Text (String) Filter Options " below) .
   b. Specify one or more valid values (see Valid Filter Values)
4. Click **Apply**.

   SOM displays a table view of all instances that have been selected based on the filter option and attribute values you specified or selected.

   Each filtered column is indicated using the filter icon.

Text (String) Filter Options

Filter Option	Description
starts with	SOM displays only instances that have an attribute value that starts with the text string value specified. Use this option when you are looking for an entry that begins with a specific string value.

contains	SOM displays only those instances that have an attribute value that contains the text string value you enter.
	You can use the wildcard character ('*') to match one or more characters within the contains value. For example, **c*m** matches the following values:
	`3com9000`
	`callmgr1`

matches	SOM displays only those instances that have an attribute value that matches the text string value you enter. Use the question mark (?) to match one character. Use the asterisk (*) as a wildcard character to match zero or more characters.
less than or equal to	SOM performs an alphabetical (lexicographical) comparison and displays all text string values that are before the text string value you enter.
greater than or equal to	SOM performs an alphabetical (lexicographical) comparison and displays all text string values that are after the text string value you enter.

Numeric Filter Options

Filter Option	Description
equals	SOM displays only instances that contain the numeric value or values you specify.
not equals	SOM displays only instances that have an attribute value that does *NOT* contain the numeric value or values specified. **Note:** For certain attributes, this option can also filter values that are "empty" or null.
greater than or equal to	SOM performs an alphabetical (lexicographical) comparison and displays all text string values that are after the text string value you enter.
less than or equal to	SOM displays all values that are less than or equal to the numeric value you enter

Valid Filter Values

Data Type	Description
Text	Enter the value for which you want SOM to search. Text (string) filters are case sensitive.
Numeric	Enter the numeric value or values for which you want SOM to search. To enter more than one numeric value, enter a comma-separated list.
Enumerated List	Select one or more values from the enumerated list.
IP Address	Enter an IP Address or range of addresses using two IP addresses separated by a '-' or using Classless Inter-Domain Routing (CIDR) notation: IPv4 examples: 10.168.0.1 - 10.168.13.1 10.2.120.0/21
Date and Time	You must enter either a date or time or both. When entering a date only the day is required. When entering the time only the minutes are required. **Note:** SOM uses a 24-hour clock, beginning at midnight (which is 0000 hours). For example, 1:00 AM is 0100 hours, 2:00 AM is 0200 hours, and 11:00 PM is 2300 hours.

Use Null Value Filters

SOM provides the following filter options to filter your view based on whether the attribute contains a value. These filter options appear for data types that do not require a value:

- Is not empty
- Is empty

To filter your view based on null values:

1. Right-click the column or attribute value on which you want to filter.
2. Select from the filter options described in the table below.
3. SOM displays a table view of all instances that have been selected based on the filter option and any attribute values you specified or selected.

   Each filtered column is indicated using the filter icon.

   You can change a filter at any time. SOM saves filters per user so that the filters you specify are maintained during subsequent user sessions.

Filter Choices

Filter Option	Description
Is not empty	SOM displays only the instances that contain a value for this attribute.
Is empty	SOM displays only the instances that do not have a value for this attribute.

Modify a Table View Filter

You can change a filter for a table view at any time. SOM saves filters per user so that the filters you specify are maintained during subsequent user sessions.

To modify a filter :

1. Right-click the column or attribute on which you want to filter.
2. Select **Modify filter...**

   Note the following:

   - If a filter is not created for the selected table column, the **Modify filter...** option does not appear.
   - If you have used an existing attribute value with the **Not equal to this value** filter, you can select an additional attribute value and select **Not equal to this value also**.
   - SOM displays the filter dialog that is appropriate for the selected column's data type. If the filter was created using the “Not equal to this value” option with the Text (String) data type, SOM does not include the values for the current filter.

3. In the Modify Filter ... dialog:

a. Select a Filter Option (see "Text (String) Filter Options " below) .

b. Specify one or more valid values (see Valid Filter Values)

4. Click **Apply**.

   SOM replaces the previous filter with the new filter values. SOM displays a table view of all instances that have been selected based on the filter option and attribute values you specified or selected.

   Each filtered column is indicated using the filter icon.

Text (String) Filter Options

Filter Option	Description
start with	SOM displays only instances that have an attribute value that starts with the text string value specified. Use this option when you are looking for an entry that begins with a specific string value. For example if all of your Cisco devices started with the text string "Cisco", and you wanted to find all Cisco devices, you could use the value string "Cisco".
contains	SOM displays only those instances that have an attribute value that matches the text string value you enter. You can also use a wildcard character (*) within your string value. If a wildcard is not specified, this filter option finds those values that exactly match the value string you enter. For example, if you wanted to find only the Cisco1 device from the following list of values, you would use "Cisco1" as your value string: • Cisco1 • Cisco12 • Cisco123 In this example, SOM would not include Cisco12 and Cisco123.
matches	SOM displays only those instances that have an attribute value that matches the text string value you enter. **Note:** Do not use the asterisk (*) within your string value unless you want SOM to match * (asterisk).
less than or equal	SOM performs an alphabetical (lexicographical) comparison and displays all text string values that are before the text string value you enter.
greater	SOM performs an alphabetical (lexicographical) comparison and displays all text string values

Text (String) Filter Options , continued

Filter Option	Description
than or equal	that are after the text string value you enter.

Numeric Filter Options

Filter Option	Description
equal	SOM displays only instances that contain the numeric value or values you specify.
not equal	SOM displays only instances that have an attribute value that does *NOT* contain the numeric value or values specified.
greater than or equal	SOM performs an alphabetical (lexicographical) comparison and displays all text string values that are after the text string value you enter.
less than or equal	SOM displays all values that are less than or equal to the numeric value you enter.

Enumerated List Filter Options

Filter Option	Description
equal	SOM displays only instances that contain the value or values you select.
not equal	SOM displays only instances that have an attribute value that does *NOT* contain the value or values you selected.

IP Address Filter Options

Filter Option	Description
equal	SOM displays only instances that contain the IP Address value you specify.
range	SOM displays only instances that are within the IP Address range you specify.

Date and Time Filter Options

Filter Option	Description
on or after	SOM displays only instances that have date and time values that occur on or after the date and time you specify.

Date and Time Filter Options, continued

Filter Option	Description
on or before	SOM displays only instances that have date and time values that occur on or before the date and time you specify.
between	SOM displays only instances that have date and time values that occur after the first date and time you specify and before the second date and time you specify. Use the between operator when you want to filter on instances that occur within a specified hour or day.

Valid Filter Values

Data Type	Description
Text	Enter the value for which you want SOM to search. Text (string) filters are case sensitive.
Numeric	Enter the numeric value or values for which you want SOM to search. To enter more than one numeric value, enter a comma-separated list.
Enumerated List	Select one or more values from the enumerated list.
IP Address	Enter an IP Address or range of addresses using two IP addresses separated by a '-' or using Classless Inter-Domain Routing (CIDR) notation: IPv4 examples: 10.168.0.1 - 10.168.13.1 10.2.120.0/21
Date and Time	You must enter either a date or time or both. When entering a date only the day is required. When entering the time only the minutes are required. **Note:** SOM uses a 24-hour clock, beginning at midnight (which is 0000 hours). For example, 1:00 AM is 0100 hours, 2:00 AM is 0200 hours, and 11:00 PM is 2300 hours.

Remove a Filter

You can remove a filter for a selected column at any time.

To remove a filter

1. Right-click the column that has a filter that you want to remove.
2. Select **Remove filter** from the drop-down menu.

SOM removes all existing filters that have been set for that column.

Restore Default Filters

You can restore the default filters for a view. This option removes any filters that you have defined for the current view.

To restore the default filter settings:

In the table toolbar, click the **Restore Default Filters** icon.

> **Note:** If you are viewing a table that appears in the form, open the table in a new window using the **Show View In New Window** icon, and then click the **Restore Default Filters** icon.

All filters are reset to the default values and any filters that you have created are removed for the current view.

> **Note:** You will lose any selections you have in the view.

Display Current Filter Settings

To keep track of the filters you have created, you can view the filter that has been set on a column-by-column basis.

To view the filter for a column

Mouse over the column of interest.

The following table explains the symbols used for each filter option.

Filter Definitions

Symbol	Filter Option
<=	Less than or equal to this value
>=	Greater than or equal to this value
=	Equals this value
NOT IN	Not equal to this value
IS NULL	Is empty
IS NOT NULL	Is not empty
LIKE	Contains string... Starts with string... Matches string...

Note: The percent sign (%) represents the wildcard character.

Restore Table View Defaults

You can remove the types of table view customizations described in the following table. Remove customizations when you no longer find them useful or when you reach the limit for the maximum number of customizations that can be stored.

If you are concerned about reaching your table view settings limit, remove view settings for tables that are not important.

Note: The **Restore Default Filters** icon does not appear on table views in forms. If you want to clear settings from a table view that appears in a form, use the **Show View in New Window** icon to open the table view in a new window, and then perform these operations.

Restore Default Table View Settings

What You can Remove	How	Description
Only table view filters	From the view display panel toolbar, select **Restore Default Filters**.	Clears all customizations to filters for the table view and refreshes the view with the view's default filter settings.
All customizations for all table views	Select **View** → **Restore All Default View Settings**	Clears all customizations (and cookies) for all table view customizations. See "Limits to View Settings" on the next page for more information about how SOM uses cookies to store table view customizations.

Export Table Information

You can export the contents of a table view for use in other applications. You can choose to export only the rows that are selected or all of the rows in your table.

Note: You must have a minimum of Operator Level1 Role to export table information.

When printing table information, note the following:

- The first column of table information does not appear in the exported version. The contents of this column is for selection purposes only.

- You can copy and paste the table data to other applications, such as Microsoft Excel, for additional editing and manipulating, such as getting a list of hostnames.

To export selected table rows

1. Right-click any cell or column header within the table.
2. Select **Export to CSV**.

   SOM displays the Export to CSV dialog.
3. Select **Selected Rows**.
4. To include the table column headings in the exported data, select the **Include Column Headings** check box.
5. SOM stores date and enumerated values in both localized (Jul 12, 2010 10:07 AM) and raw (1278950859739) format. By default, SOM exports only the localized (human readable) format for date and enumerated values.

   Select the **Raw Data** option to include only the raw (computer readable) format for both date and enumerated values.

   Select the **Localized Data** option to include only the localized (human readable) format for both dates and enumerated values.
6. In the **File Download** dialog, select one of the following options.
   a. **Open** to view the file contents.
   b. **Save** to save the file to a specified file name.

To export all table rows

1. Right-click any cell or column header with the table.
2. Select **Export to CSV**

   SOM displays the Export to CSV dialog.
3. Select **All Rows**.
4. To include the table column headings in the exported data, select the **Include Column Headings** check box.
5. In the **File Download** dialog, select one of the following options.
   a. **Open** to view the file contents.
   b. **Save** to save the file to a specified file name.

Limits to View Settings

SOM automatically saves the following kinds of table view settings:

- Column width
- Hidden columns

- Sort column and order
- Filters

When customizing table views, note the following:

- All settings for each table are stored in a corresponding cookie for that table.
- The number of tables that have settings that can be saved, and the behavior when the limit is reached depends on the browser.

  Microsoft Internet Explorer discards your oldest cookie and lets you continue creating customizations. If you are using Microsoft Internet Explorer, the cookie limit is 48.

  If you are using Mozilla Firefox, the cookie limit is 48. If you reach the 48 limit, Mozilla Firefox removes the next to last cookie before saving the latest cookie created.
- Table view settings are specific to each operating system user. Therefore, if you log on as a different operating system user, you will have *different* view settings. If you sign into SOM as a different SOM user, but as the same operating system user, you will have *the same* view settings.

> **Note:** SOM displays a warning message when your last cookie and subsequent table settings are being saved.

Use Map Views

You can perform the following operations in a map view:

- "Use Forms and Analysis Panes to Access More Information About an Object" on page 68
- "Change the Map Layout" on page 64
- "Adjust the Zoom Factor" on page 63
- "Pan Around the Map" on the next page
- "Set the Location of the Overview Pane" on the next page
- "Refresh Node Status on a Map" on page 65
- "Control Tool Tips Information on a Map" on page 64

If you use the Mozilla Firefox browser and have any timeout issues; for example, being prompted to click **Continue** before a map appears, see "Configure Mozilla Firefox Timeout Interval" on page 21.

Find a Node in a Map

You can easily find a particular Node in a map.

To find a Node in the current map:

1. In the Map toolbar, toggle the Find button on:

   

2. In the text box , do one of the following:
   - Click the ▽ drop-down icon to display the complete list of choices within this map. Select any item from the list.
   - Type the Name attribute value (*not case-sensitive*) of the node you want to find in the map. This is the Name attribute value from the Node form, that becomes the map icon label.

     As you type, the auto-complete feature displays a list of possible matches. Select any item from the list.
3. SOM selects the Node that has a Name attribute value that matches your choice. If necessary, SOM pans to the appropriate area in the Map.
4. To navigate elsewhere in the map, toggle the Find button off:

   

Set the Location of the Overview Pane

You can choose which corner of the map contains the Overview Pane. Alternatively, you can hide the Overview Pane .

To set the location of the Overview Pane:

1. In the Map toolbar, toggle the Overview Location button on:

   

2. In the menu, select one of locations

   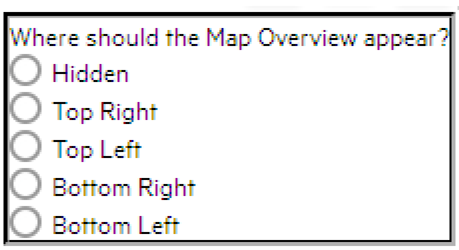

3. To close the menu, toggle the Overview Location button off:

   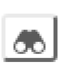

Pan Around the Map

If the node of interest is not easily visible on the map, you can move to other sections of the map.

To pan to other sections of the map:

Do one of the following:

- Use the arrow keys (UP ARROW, DOWN ARROW, RIGHT ARROW, LEFT ARROW).
- Use the Focus Area of the Overview Pane to pan to other sections in the map. The Overview Pane provides a view of the entire map you selected. SOM indicates the Focus Area using a gray rectangle as shown in the following example:

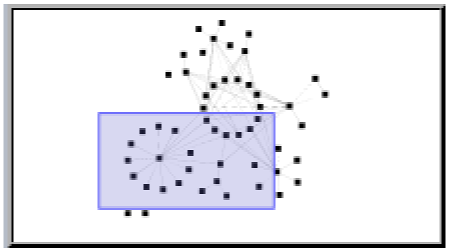

To pan around the map, drag the Focus Area.

Adjust the Zoom Factor

You can adjust the zoom factor in a map in several ways.

To zoom in a map:

Do one of the following:

- In the map, rotate the mouse wheel button forward.
- Click .
- Press **+** (PLUS SIGN).
- In the Overview Pane, double-click the Focus Area that is indicated with the blue rectangle as shown in the following example:

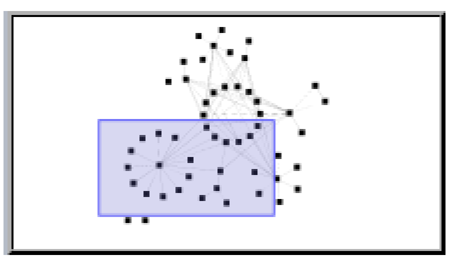

> **Tip:** If SOM does not indicate a Focus Area, double-click anywhere in the Overview Pane.

The Overview Pane provides a view of the entire map you selected. The Focus Area indicates the portion of the map that SOM displays in the larger map view.

To zoom out in a map:

Do one of the following:

- In the map, rotate the mouse wheel button backward.
- Click .

- Press **-** (MINUS SIGN).
- In the Overview Pane, double-click outside the Focus Area that is indicated with the blue rectangle.

To fit the map to the screen size:

Do one of the following:

- Click .
- Press **=** (EQUALS SIGN).

To display the map at 100 percent (actual size):

Do one of the following:

- Click
- Press 1 (ONE).

Change the Map Layout

If you prefer a different layout for the symbols on your map, you can change the placement.

To move a single object:

Drag any map symbol to the location you want on the map.

To move multiple objects:

Use Ctrl-Click to select each object of interest on the map.

Each object you select changes to indicate it has been selected. Drag any one of the selected nodes to move the group of objects to the location you want on the map.

Tip: If you drag in a space between the nodes, all objects are de-selected.

This placement persists until you refresh or otherwise reload the map.

Control Tool Tips Information on a Map

When you place the mouse cursor over an object on a map, SOM displays Tool Tips information for the current object. Tool Tips information is a subset of the information contained in the object's form.

Each time you open a map, the **Tool Tips** button is enabled.

To disable Tool Tips popups, click the **Tool Tips** button. SOM closes any open Tool Tip dialog boxes.

Refresh Node Status on a Map

In map views, you can refresh the node status on a map while maintaining the layout of the nodes. This feature enables you to quickly refresh a view when you are only concerned with status updates for one or more nodes on your map.

SOM automatically refreshes node status on a map every 60 seconds. The refresh counter begins after the completion of the last status refresh.

> **Note:** If a map is taking longer than expected to refresh, the maximum number of nodes that are displayed on a map might be set too high. The SOM administrator can set this value.

You can refresh the node status manually.

To refresh only the node status:

Click the **Refresh Status** icon in the toolbar to manually refresh node status on your map view.

The last update time is changed and the status of each node is refreshed on the map. Nodes are not added, deleted, or rearranged. Connectivity is not recalculated.

Some SOM users can delete nodes and other objects from the SOMdatabase (depending on the assigned SOM role). Any node that has been deleted appears as a transparent icon to all SOM users until their map is refreshed using the **Refresh** icon. After refresh, the deleted node is removed from the map. SOM does not automatically refresh the connectivity or set of nodes in a map view.

Access Maps

You can access maps in the following ways:

- From table views using the **Actions** menu
- From the **Topology Maps** workspace.

Table Views

To display a map view from a table view:

1. Select the table view you want from the workspaces navigation panel. (For example, select the **Inventory** workspace, Click to expand the **Storage Systems** folder and then click **All Storage Systems**.)
2. In the table view, click the row that contains the object of interest.
3. Click the **Actions** menu in the main toolbar and then select **Launch Topology**.

   > **Note:** You can also access map views from the **Actions** menu in a form.

> **Tip:** You can also right-click any object in a table or map view to access the items available within the **Actions** menu.

Topology Maps Workspace

To display the topology map of an entire storage infrastructure:

1. From the workspace navigation panel, click the **Topology Maps** workspace.
2. Click **System Topology**.

The System Topology pane displays the physical connectivity of all the storage elements in your network. You can access storage element nodes, and filter the view by fabrics and element types.

Work with Objects

Objects are database records of information about your environment. Each type of object represents a particular kind of information.

An object is defined by its attributes. Different object types have different numbers and types of attributes. Some attribute values are simple, such as numbers and text strings. Other attribute values are more complex, such as a reference to a related object.

If more than one of a certain type of object can be related to the selected object, the form contains a tab that displays a table with the entire list of related objects.

A *view* is a collection of related objects that are depicted graphically as a table or map. A *form* provides all stored attributes about a selected object. The attributes on the form can be attributes of the selected object or related objects.

Operations that can be performed on objects are called actions. Actions are shortcuts to simple or complex tasks. A particular action can be associated with a specific object type. For example, when displaying the hosts table view, you might want to open a map showing the storage elements connected to a host.

> **Tip:** You can also right-click any object in a table or map view to access the items available within the **Actions** menu.

To access an object's form from a table view:

Double-click the row representing an object.

SOM displays the form for the selected object.

To access an object's form from a map view:

Select the node, and then click **Open**.

OR

Double-click the node.

> **Note:** If the map object is a child node group, double-clicking the child node group object replaces the current map with a map of the nodes in the child node group. To access a child node group form, select the child node group object and click the Open icon.

SOM displays the form for the selected object.

> **Tip:** A red asterisk (*) that precedes an attribute on a form indicates the attribute requires a value.

From an object form, you can:

- "Use Forms and Analysis Panes to Access More Information About an Object" below
- "Access a Subset of the Available Information" on the next page
- "Access All Information" on page 70

Use Forms and Analysis Panes to Access More Information About an Object

You can use forms and analysis panes to access more information about an object. For example, if you want to obtain more information about a physical or virtual switch, you can access the information from the physical or virtual switch form.

Access all object attributes and related objects by viewing the form:

> **Tip:** A red asterisk (*) that precedes an attribute on a form indicates the attribute requires a value.

- **To open a form using Tools → Find Element**:

  For more information, see the "Tools Menu " on page 36.

- **To open a form from a table view**:

  Double-click the row representing an object.

  The form appears, containing the details about the object.

- **To open a form from a map view**:

  Select the node and then click **Open** on the toolbar.

  OR

  Double-click the node.

  > **Note:** If the map object is a child node group, double-clicking the child node group object replaces the current map with a map containing each of the nodes in the child node group. To access a child node group form, use the Open icon in the toolbar.

  The form appears, containing the details about the object.

Access more details about an object using the analysis pane:

The Analysis pane is available below most workspace views, such as, inventory, topology, configuration views, and so on. From a workspace view, select an object or an element, to see its analysis information at the bottom of the window.

Note: If the Analysis pane is not visible, point to the bar at the bottom of the window to view ↕ resize icon, drag the bar to view the Analysis pane.

Note the following:

- The analysis pane remains empty until an object is selected.
- To update the displayed information, click **Refresh** in the tabs of the analysis pane.

Access a Subset of the Available Information

You can access information about the attributes of an object and the objects related to it, from its form. The related objects are indicated by a **Lookup** icon. For example, when viewing information for a node object, you can access information about the device profile associated with that node.

Tip: A red asterisk (*) that precedes an attribute on a form indicates the attribute requires a value.

This is an example Lookup Field:

To display a subset of information about a related object from within a form:

1. Locate the field for the related object that you want to learn more about.
2. Click the **Lookup** icon, and then select **Show Analysis**.

   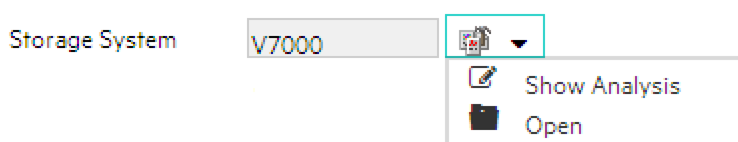

3. The analysis pane appears showing information about the related object. For more information, see "Analysis Pane" on page 36 .

   Note: SOM displays only the information that the SOM security configuration permits you to access.

4. To see the date and time when the details were last refreshed, mouse over **Refresh**.

   Click **Refresh** to gather the most recent data.

Access All Information

While investigating the details for one object using a form, you can access information about a related object. For example, when viewing all information stored for a node, you can access all available information for the associated device profile.

Tip: A red asterisk (*) that precedes an attribute on a form indicates the attribute requires a value.

From a form, you can open another form associated with an object. Such objects are indicated using a **Lookup** icon in the form.

This is an example Lookup Field:

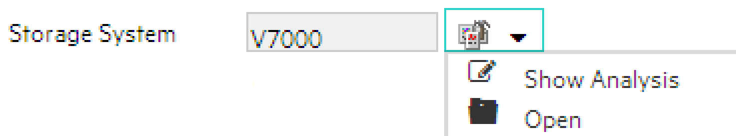

To open another form from a form:

1. Locate the field of an object about which you want to see more information.
2. Click the Lookup icon, and then select **Open**.

A new form appears showing all of the attributes for that object. Any default values specified for the object are pre-populated in the form.

Customize Charts

Certain dashboard panels and tabs in the analysis pane have charts that can be customized.

For example, you can customize these charts:

- area
- bar
- line
- scatter

To customize a chart, do any or all of the following:

- Change the chart type by clicking one of the controls provided in the upper-right portion of the panel.

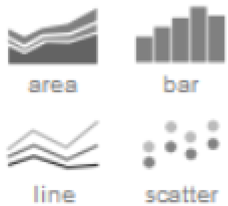

- Select the desired components from the component list on the right side of the panel.
- Hover the cursor over the component list to instantly view data for one particular component.
- Hover the cursor over the chart to view a value for a given component at a particular time.
- Isolate the view to a particular time frame by clicking and dragging the buttons under the time captions beneath the chart.

Chapter 2: Configuring SOM for your Storage Environment

You must perform the following configurations before you can manage your storage environment with SOM.

Task	Description	User Role	License
"Summary of Security Tasks" on page 77	Configure user accounts, user groups, and security groups to control access to the managed storage infrastructure.	Administrators Only	All
"Create a Node Group" on page 107	Define node groups based on device category, vendor, family and profile and assign nodes to node groups.	Administrators Only	All
" Discovery Tasks" on page 155	Configure IP address, IP address range, credentials for discovery, and tenant associations.	Administrators Only	All
"Inferring Hosts Based on Rules" on page 167	Create rules based on host security groups, zones or zone aliases to infer hosts from storage systems and fabrics in your environment.	Administrators Only	All
"Configuring Data Collection Settings" on page 176	Configure policies and blackout periods for collecting data from managed elements.	Administrators only	All
"Monitoring Performance" on page 190	Collect performance metrics from managed elements.	Administrators only	Ultimate Perf-Pack Only
"Managing Storage Tiers" on page 196	Configure automated rules-based assignments for categorizing storage systems, volumes and pools into storage tiers.	Administrators only	All

Configuring Emails

You can configure SOM to send emails to specified email addresses at the occurrence of pre-defined error conditions and/or at scheduled intervals.

The pre-defined conditions for which SOM can trigger an email are:

- **At data collection failure**: The email will contain the details of the element for which the data collection has failed.
- **At CSV export failure**: The email will contain the details of the element for which the csv export has failed.

If you choose to schedule an email at specified intervals, the email will contain SOM status reports with the following information:

- **Data collection failures**: A zip file with details about the elements for which the data collection fails and the time stamps of their data collection failure.
- **Stale element**: An attachment with details about the elements which are stale.
- **Invalid or empty observer schedules**: An attachment with details about the invalid observer schedules.
- **CSV failures**: An attachment with details about the CSV failures.

To configure SOM emails, follow the appropriate procedure:

- "Configure Emails for SMTP Servers That Do Not Require Authentication (Manual Configuration)" below
- "Configure Emails for SMTP Servers That Require Authentication (Configuration Using the Command Line)" on the next page

Configure Emails for SMTP Servers That Do Not Require Authentication (Manual Configuration)

To configure SOM email for Simple Mail Transfer Protocol (SMTP) servers that do not require authentication, follow these steps:

1. Create an `email.properties` file in the following directory:
   - *Windows*: `%OvDataDir%\Conf\som`
   - *Linux*: `/var/opt/OV/conf/som`
2. In the `email.properties` file, add the following properties:

Properties	Description
`email.host=<FQDN of the Email Server>`	Type the FQDN of the SMTP server. For example, `email.host=myserver.example.com`
`email.port=25`	Type the port number for SMTP. The default port number is 25.
`email.auth=false`	Set this property to false. If set to true, the SMTP server requires authentication.
`email.host.user=admin`	Type the user name of the SMTP server.

Properties	Description
`email.schedule=true`	Set this property to true if you want emails to be sent at scheduled intervals.
`email.schedule.period=5`	Type the time interval in hours between emails. By default it is set to 24 hours.
`email.dcfailure=true`	Set this property to true if email has to be sent for every data collection failure. By default it is set to false.
`email.csvfailure=false`	Set this property to true if email has to be sent on every csv export failure. By default it is set to true.
`email.to=<Recipient's Email Address>`	Type the recipient's email address. For example, `email.to=someone@example.com`
`email.from=<Sender's Email Address>`	Type the sender's email address. For example, `email.from=someone@example.com`

Note: SOM supports only one set of SMTP configurations.

3. Save the `email.properties` file.
4. Restart the `somjboss` service for the changes to take effect.

Configure Emails for SMTP Servers That Require Authentication (Configuration Using the Command Line)

While configuring email for Simple Mail Transfer Protocol (SMTP) servers that require authentication, you must supply the password only through the command line to ensure that the password is encrypted.

To configure SOM email for SMTP servers that require authentication, follow these steps:

1. Create a text file (for example, `EmailConfiguration.txt`) in your local host server with the following properties:

Properties	Description
`email.host=<FQDN of the Email Server>`	Type the FQDN of the SMTP server. For example, `email.host=myserver.example.com`
`email.port=25`	Type the port number for SMTP. The default port number is 25.

Properties	Description
`email.auth=true`	Set this property to true. If set to false, the SMTP server does not require authentication.
`email.host.user=admin`	Type the user name of the SMTP server.
`email.schedule=true`	Set this property to true if you want emails to be sent at scheduled intervals.
`email.schedule.period=5`	Type the time interval in hours between emails. By default it is set to 24 hours.
`email.dcfailure=true`	Set this property to true if email has to be sent for every data collection failure. By default it is set to false.
`email.csvfailure=false`	Set this property to true if email has to be sent on every csv export failure. By default it is set to true.
`email.to=<Recipient's Email Address>`	Type the recipient's email address. For example, `email.to=someone@example.com`
`email.from=<Sender's Email Address>`	Type the sender's email address. For example, `email.from=someone@example.com`

Note: SOM supports only one set of SMTP configurations.

Note: You must manually set the properties `email.csvfailure` and `email.dcfailure` in the `email.properties` file later. The command line does not set these properties.

2. Save the `EmailConfiguration.txt` file.
3. Change to the following directory:
   - *Windows*: `%OvDataDir%\bin`
   - *Linux*: `/var/opt/OV/bin`
4. At the command prompt, run the following command:

```
somemailconfiguration.ovpl -configure <full path of the EmailConfiguration.txt file>
[-hostpassword <password>]
```

Note: This command writes the contents of the `EmailConfiguration.txt` file to the `email.properties` file and encrypts the password.

5. To restart the `somjboss` service, type 'yes'.
6. Verify the contents of the `email.properties` file located in the following directory:
   - *Windows*: `%OvDataDir%\Conf\som`
   - *Linux*: `/var/opt/OV/conf/som`
7. Manually set the properties `email.csvfailure` and `email.dcfailure` in the `email.properties` file.

   If any properties are missing, manually add the properties or use the command line to supply passwords.
8. When you finish making the changes, save the `email.properties` file.
9. Restart the `somjboss` service for the changes to take effect.

Configuring Security

The SOM security model provides user access control to the objects in the SOM database. The model can be configured to meet the needs of your environment. To configure security, you need an understanding of user accounts, user groups and security groups and how they can be mapped to meet the security needs of your environment.

You can configure the following components of security to meet your environment's security needs:

- Users – Identifies users of the system.
- User Groups – Groups of users based on their roles and control the access to the SOM console.
- User Group Mapping – Determines the access level to the SOM console for each user group.
- Security Group – Identifies set of nodes that the user can access.
- Security Group Mapping – Controls what the users can do with the nodes.

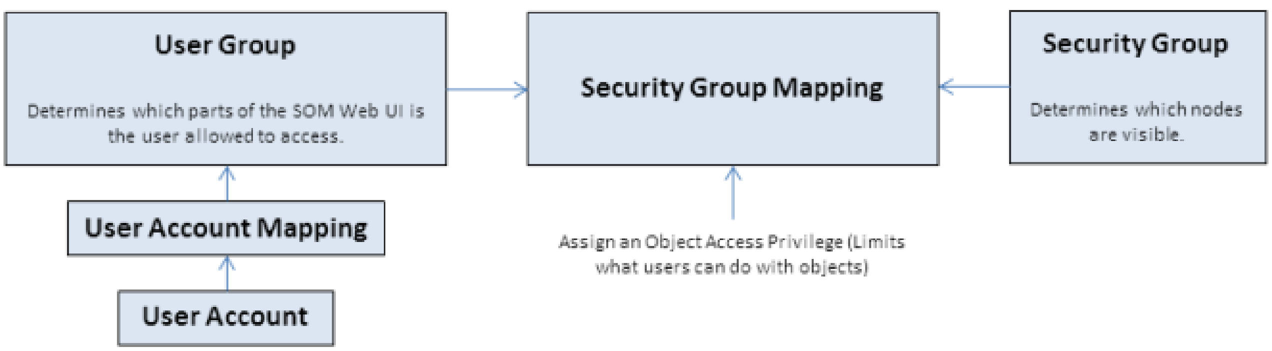

Summary of Security Tasks

The following table lists all possible choices for configuring security. The tasks vary based on the type of user authentication mode you choose.

Task	Description
"Choose a Mode for User Authentication " on page 79	Choose the type of user authentication – SOM Console Access or LDAP.

Task	Description
"Create a User Account" on page 82	You must create a user account for each SOM user.
"Create a User Group" on page 87	The administrator can create any number of user groups to meet the needs of your network environment. Examples of the need for additional user groups include the following: • When you need a subset of users to access only a subset of nodes. • When you need to divide node access between two or more user groups (such as multiple shifts or multiple sites that share responsibilities).
Map user accounts to the default user groups	A particular user cannot access the SOM console until their user account is mapped to at least one of the default user groups. For more information, see" Predefined User Groups " on page 86.
"Create a User Account Mapping" on page 89	If you created additional user groups, map the appropriate user accounts to each user group you created.
"Create a Security Group" on page 91	By default, all operators can access all nodes discovered by SOM. However, you can limit the visibility to a subset of nodes for some or all operators by using user groups and security groups. **Note:** Each node can be mapped to one and only one security group. Examples of the need to create additional Security Groups to limit node access include the following: • When you need a subset of users to access only a subset of nodes. • When you need to divide node access between two or more user groups
"Configure Security Group Mappings " on page 94	After creating any additional user groups, you map each user group to a security group and assign the *Object Access Privilege* for this security group mapping. The *Object Access Privilege* determines the level of access that each user group has to the nodes that are visible. Users can view a node only if one of the user groups to which they belong is associated with the security group of that node.
"Methods for	By default, all SOM user groups have access to nodes assigned to the default security

Task	Description
Assigning Nodes to Security Groups" on page 93	group. If you create security groups to limit node access, you must assign nodes to the appropriate security group. Each node is associated with one and only one security group.

Choose a Mode for User Authentication

SOM can integrate with a directory service using LDAP for consolidating storage of user names, passwords, and optionally, user groups. You can choose to use any of the following authentication methods best suited for your environment.

Option 1: SOM Configuration settings

User names, passwords and user group memberships are defined within the SOM database.

Option 2: Lightweight Directory Access Protocol (LDAP)

SOM communicates with the directory service using LDAP. You can use the LDAP external mode with LDAP password and user account mapping with SOM User Group membership assignments.

Note: You must choose *one* user authentication method and configure all SOM users with the same approach.

If you choose option 2, you must have already configured SOM to integrate with the directory service using LDAP.

Option 1: SOM Configuration Settings

Configure the user names, passwords, and user group membership assignments in the SOM database. The following table describes the SOM security settings:

Mode	User Authentication Method	User Account Definitions in SOM	User Group Membership in SOM	User Groups Mapping
Internal	SOM Password	Yes	SOM	Yes

Security Configuration Tasks	
Task 1	"Create a User Account" on page 82

Security Configuration Tasks	
Task 2	"Create a User Group" on page 87
Task 3	"Create a User Account Mapping" on page 89
Task 4	"Create a Security Group" on page 91
Task 5	"Map User Groups to Security Groups" on page 94

Option 2: Lightweight Directory Access Protocol (LDAP)

SOM administrators can configure SOM to rely on your environment's directory service to provide any of the following:

- **Mixed**: SOM password
- **External**: SOM password plus SOM User Group membership assignments

Mode	User Authentication Method	User Account Definitions in SOM	User Group Definitions in SOM	User Group Membership Method
Mixed	LDAP Password	Yes	Yes	SOM User Account Mappings
External	LDAP Password	No	Yes	LDAP

For information about integrating with a directory service for the user name and password, see *"LDAP-Based Authentication"* in the *SOM Deployment Guide*.

If you are using the LDAP External mode , note the following:

- Do not create user accounts in the SOM console.

  Note: If you are a new user, you might not be able view the following:
  - System Topology
  - HBA ports and FC ports for inferred hosts.
  - Any data for hosts (Presented Storage tab)

- Do not create user account mappings in the SOM console.
- To modify user account information such as user name, password or user group assignment, you must use the LDAP directory service software. You cannot modify the user account from the SOM console.
- You can choose to configure the user display name value to be one or more LDAP properties rather than the name used to sign in to SOM.

Different Ways to Configure Security

You can configure security using the following methods:

SOM Console Forms

The forms for individual security and multi-tenancy objects in the console are useful for configuring one aspect of the security at a time. The following views are available under Security folder in the Configuration:

- User Accounts

  Each User Account form enables you to configure one user and shows the user accounts to which the user belongs. If you are storing user group membership in a directory service, user accounts are not visible in the console.

- User Groups
  The User Group form enables you to configure user groups.
- User Account Mapping
  With a User Account Mapping form you can configure user account-to-user group association. If you are storing user group membership in a directory service, user account mappings are not visible in the console.
- Security Groups
  The Security Group form enables you to create security groups and shows the nodes currently assigned to the security group. The node assignment information is read-only.
- Security Group Mapping
  The Security Group Mapping form enables you to configure a user group-to-security group association.

The Security Wizard

The Security Wizard is useful for visualizing the security configuration. It is the easiest way to assign nodes to security groups within the Storage Operations Manager console. The View Summary of Changes page in the wizard presents a list of unsaved changes from the current wizard session. It also identifies potential problems with the security configuration.

Note: The Security Wizard does not include any information about tenants.

Configure User Accounts

Each user account represents a user. You can perform the following tasks for a user account:

- "Create a User Account" on the next page
- "Modify a User Account" on page 83
- "Delete a User Account" on page 84

Create a User Account

User Account configurations include creating user name and password settings. It also involves specifying if SOM should use an external resource for password information.

> **Note:** If you are a new user, you might not be able to view the following:
>
> - System Topology
> - HBA ports and FC ports for inferred hosts
> - Any data for hosts (Presented Storage tab

To configure a user account, follow these steps:

1. From the workspaces navigation panel, select **Configuration**> **Security** > **User Accounts**. The User Accounts view is displayed.
2. Click **New** on the view toolbar. The User Account form is displayed.
3. Specify the user account details. (See the User Account attributes below.)

   > **Tip:** You can filter the User Accounts view by User Group or Security Group.

4. Click one of the following icons to save the user account:
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

Attribute	Description
Name	Enter a string that identifies a user uniquely. The name can be up to 40 alpha-numeric characters. Do not use punctuation, spaces, or underline characters.
Directory Service Account	☐ indicates that user name and password are stored in the SOM database. For more information, see "Option 1: SOM Configuration Settings " on page 79. ☑ indicates that SOM uses Lightweight Directory Access Protocol (LDAP). Additional steps are required. For more information, see "Option 2: Lightweight Directory Access Protocol (LDAP)" on page 80.
Password	Enter the **Password** value. Type any combination of alpha-numeric characters, punctuation, spaces, and underline characters. Reenter the **Password** value.

Attribute	Description
	Note: If you have enabled **Directory Service Account** ☑, do not provide a password. **Tip:** When SOM is configured with **Directory Service Account** ☐, SOM users who are assigned to the Object Access Privilege Mapping can change their SOM password at any time using **File** → **Change Password.** *Object Access Privilege* = one of the following: • Object Administrator • Object Operator Level 2 • Object Operator Level 1 (with more limited access privileges than Level 2)

Modify a User Account

Use the instructions in this topic only if you have configured SOM to store user names and passwords in the SOM database.

If you have configured SOM to use an external User Authentication Method (passwords stored outside of the SOM database) such as LDAP, see "Option 2: Lightweight Directory Access Protocol (LDAP)" on page 80 .

To change the user name:

You must delete a user account (see "Delete a User Account" on the next page, and then recreate the account mapping (see "Option 1: SOM Configuration Settings " on page 79).

To change the password:

1. From the Workspaces navigation panel, select the **Configuration**>**Security**> **User Accounts**. The User Accounts view is displayed.
2. Double-click the user account row that you want to edit.
3. Locate the **Password** attribute and change the **Password** value. Type up to 40 alpha-numeric characters, punctuation, spaces, and underline characters.
4. Reenter the new password.
5. Click **Save and Close**. SOM immediately implements your changes.

To change the user group to user account assignment:

> **Note:** To change a user group to user account assignment, you first delete the user account mapping. If you change the user account or user group configuration for a user who is currently signed into the SOM console, the change does not take effect until the next time the user signs in. By default, the SOM timeout limit is 18 hours. If a user has not signed out within 18 hours, SOM forces the user to sign out.

1. From the Workspaces navigation panel, select the **Configuration** > **Security** > **User Accounts**. The User Accounts view is displayed.
2. Select the user account mapping that you want to change.
3. Delete the user account mapping by clicking the Delete icon.
4. Select the **New** icon to configure the new user account mapping.
5. Make your configuration choices. (See the User Account Mapping Attributes table.)
6. Click **Save and Close**.

User Account Mapping Attributes

Attribute	Description
User Group	In the **User Group** attribute, click **Lookup**. • To create new user group, click **New** and provide the required information. (See "Create a User Group" on page 87 for more information.) • To select an SOM user group configuration, click the **Quick Find** icon and make a selection.
User Account	In the **User Account** attribute, click **Lookup**. • To create a new user account, click **New** and provide the required information. For more information, see "Create a User Account" on page 82 • To select an SOM user group configuration, click **Quick Find** and make a selection. **Note:** If you map a user account to two or more SOM User Groups, SOM gives the user account the privileges associated with each mapped SOM user group.

Delete a User Account

Ignore this topic if SOM is configured to access LDAP information for user group assignments. When SOM is configured to access LDAP information, to disable a user's access to SOM, you must use the appropriate process required by your environment's directory service software (see "Option 2: Lightweight Directory Access Protocol (LDAP)" on page 80).

Caution: If you delete the last SOM user assigned to the SOM Administrators User Group, no one can access the Configuration workspace. For more information about how to recover from this mistake, see "Restore the Administrator Role" on page 105.

To delete a user account, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** > **User Accounts**. The User Accounts view is displayed.
2. Select the user account that you want to delete from the table view.
3. Do one of the following:
   - Click **Delete**. The message "Are you sure you want to perform this action on the selected items?" is displayed.
   - Click **Open**. The user account is displayed in the User Account form view. Click . The message "Are you sure you want to delete this item? This will also delete all contained objects and references." is displayed.
4. Click **OK** to delete the user account.

   The user account configuration is removed from the User Accounts view.

Note: If you remove the User Account for a user who is currently signed into the SOM console, the change does not take effect until the next time the user signs in. By default, the SOM timeout limit is 18 hours. If a user has not signed out within 18 hours, SOM forces the user to sign out.

Configure User Groups

User groups enable you to group users and control the access to the SOM console.

SOM provides predefined user groups. Users cannot access the SOM console until their user account is mapped to at least one of the predefined user groups. You can create additional user groups to fine tune access to SOM based on your environment. For more information, see " Predefined User Groups " on the next page

You can perform the following tasks for a user group:

- "Create a User Group" on page 87
- "Modify a User Group" on page 88
- "Delete a User Group" on page 88

Predefined User Groups

The following predefined SOM user groups determine a user's access to the SOM console workspaces and forms. Each user account must be mapped to one of these predefined SOM user groups to enable access to the SOM console:

- Administrators
- Level 2 Operators
- Level 1 Operators (with more limited access privileges than Level 2 Operators)
- Guest Users

SOM provides two additional predefined user groups that do not provide access to the SOM console:

- Global Operators

  Provides access to all topology objects (the nodes in all security groups) but does not change a user's access to the SOM console.

  > **Tip:** User accounts mapped to the Administrators user group can access all topology objects and do not need to be mapped to the Global Operators user group or any security groups.

- Web Service Clients

  Provides access for software that is integrated with SOM. Do not use any other user group for software integrations.

You cannot delete the predefined user groups.

If you map a user account to two or more user groups, SOM gives the user account the privileges associated with each user group to which the user account is assigned.

In addition to the default user groups, administrators can create additional user groups. Creating user groups enables you to fine tune user group access when using security groups. For example, you might want one user group to have Level 2 Operator access to the nodes in one security group and Level 1 Operator access to the nodes in another security group.

Determine the User Group

Before configuring SOM sign-in access for your team, determine which default user group is appropriate for each team member. The user groups are hierarchical, meaning the higher level user groups include all privileges of the lower level user groups in the hierarchy (Administrator is the highest level and Guest is the lowest level).

As SOM administrator, you can change the "Control Menu Access" on page 99 (restrict access to certain SOM Actions menu items and Tools menu items) to provide tighter security than those enforced by the default settings.

The following table lists the User Group required to access SOM worskspaces. You cannot modify User Group settings for workspaces. For more information about workspaces, see "Workspaces" on page 25.

Access to Workspaces

Workspaces	Guest Users	Level 1 Operators	Level 2 Operators	Administrators
All views in the Topology workspace	Yes	Yes	Yes	Yes
All views in the Monitoring workspace	Yes	Yes	Yes	Yes
All views in the Troubleshooting workspace	Yes	Yes	Yes	Yes
All views in the Inventory workspace	Yes	Yes	Yes	Yes
All views in the Configuration workspace				Yes

The following table provides some examples of how User Groups control permission for modifications to certain forms:

Access to Forms (some examples)

Forms	Guest Users	Level 1 Operators	Level 2 Operators	Administrators
Node forms	Read-Only	Read-Write	Read-Write	Read-Write
IP Address forms	Read-Only	Read-Write	Read-Write	Read-Write
Node Group forms	Read-Only	Read-Only	Read-Only	Read-Write
Configuration Forms				Read-Write

You cannot modify User Group settings for forms.

Create a User Group

User groups enable you to control the access to the SOM console. In addition to the predefined user groups, you can create additional user groups to fine tune access to the SOM console. Each user account must be mapped to one or more user group.

To configure a user group, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** > **User Groups**. The User Groups view is displayed.
2. Click ✱ **New** on the view toolbar. The User Group form is displayed.
3. Specify the user group details. (See the User Group Attributes table.)
4. Make your additional configuration choices.

5. Click one of the save options to save the user group:
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The user group is displayed in the User Groups view.

User Group Attributes

Attribute	Description
Name	Enter the name that uniquely identifies the user group. Enter a maximum of 40 alpha-numeric characters. Spaces are not permitted.
Display Name	Enter the name that should be displayed in the SOM console to identify this User Group. Enter a maximum of 50 characters. Alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _+ -) are permitted.
Directory Service Name	*Optional*. When Lightweight Directory Access Protocol (LDAP) defines this User Group, enter the group's Distinguished Name. See "Option 2: Lightweight Directory Access Protocol (LDAP)" on page 80.
Description	Type a maximum of 2048 characters to describe this user group. Alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _+ -) are permitted.

Modify a User Group

To modify a user group, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** > **User Groups**. The User Groups view is displayed.
2. Select the user group that you want to modify from the table view.
3. Click **Open**. The user group is displayed in User Group view.
4. Make the required changes to the user group.
5. Click to save changes to the user group. The User Group View is refreshed to display the changes to the user group.

Delete a User Group

To delete a user group, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** > **User Groups**. The User Groups view is displayed.
2. Select the user group that you want to delete from the table view.
3. Do one of the following.
   - Click **Delete**. The delete confirmation message is displayed. Click **OK** to delete the user group.
   - Click **Open**. The user group is displayed in the User Groups view. Click . The delete confirmation message is displayed. Click **OK** to delete the user group.

The user group configuration is automatically removed from the User Groups view.

Configure User Account Mapping

User Account Mappings enable you to assign a User Account to one or more User Groups to control SOM console access.

Each user account must be mapped to at least one predefined user group to access the console. A user account can be mapped to two or more user groups.

A User Account Mapping is a separate object in the SOM database. Therefore, when you create or delete a User Account Mapping, you create or delete only the User Account Mapping, not the User Account or User Group.

The following tasks are associated with a user account mapping:

- "Create a User Account Mapping" below
- "Delete a User Account Mapping" on the next page

Create a User Account Mapping

To assign a user account to a user group, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** > **User Account Mappings**. The User Account Mappings view is displayed.
2. Click **New** on the view toolbar. The User Account Mapping form is displayed.
3. Make your configuration choices. (See the User Account Mapping Attributes table.)
4. Click one of the save options to save the mapping:
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

The user group account mapping is displayed in the User Account Mapping view.

> **Note:** If you create a user account to user group mapping for an SOM user who is currently signed into the SOM console, the change does not take effect until the next time the user signs in. By default, the SOM timeout limit is 18 hours. If a user has not signed out within 18 hours, SOM forces the user to sign out.

User Account Mapping Attributes

Attribute	Description
User Group	In the **User Group** attribute, click the **Lookup** icon. • To create new user group, click the **New** icon and provide the required information. For more information, see "Create a User Group" on page 87. • To select an SOM user group configuration, click the **Quick Find** icon and make a selection.
User Account	In the **User Account** attribute, click the **Lookup** icon. • To create new user account, click the **New** icon and provide the required information. For more information, see "Create a User Account" on page 82. • To select an SOM user group configuration, click the **Quick Find** icon and make a selection. **Note:** If you map a user account to two or more SOM user groups, SOM gives the privileges associated with each mapped SOM user group.

Delete a User Account Mapping

When you remove a user account from a user group, only the mapping between the two gets deleted. The user account or the user group does not get deleted from the SOM database.

To remove a user account mapping from a user group, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** > **User Account Mappings**. The User Account Mappings view is displayed.
2. Select the row that contains the User Account and User Group mapping that you want to delete.
3. Do one of the following.

   - Click **Delete**. The delete confirmation message is displayed. Click **OK** to delete the mapping.

- Click **Open**. The mapping is displayed in the User Group Mapping form view. Click . The delete confirmation message is displayed. Click **OK** to delete the mapping.

Configure Security Groups

Security Groups define sets of nodes within your network environment. Each node is assigned to only one Security Group. Your security strategy determines the number of security groups required for your network environment. By default, all nodes are assigned to the **Default Security Group** and all the users see all the nodes. You can create additional security groups to group nodes that require the same access level.

The following tasks are associated with a security group:

- "Create a Security Group" below
- "Modify a Security Group" on the next page
- "Delete a Security Group" on page 93

Recommendations for Planning Security Groups

- Map each user account to only one default user group.
- Do not map the default user groups to security groups.
- Because any user account mapped to the administrators user group receives administrator-level access to all objects in the SOM database, do not map this user account to any other user groups.
- In general, related elements should be configured as part of the same security group. Some examples of related elements include the following:
  - If a virtual machine is part of a security group, then its virtual server also needs to be part of the same group.
  - Arrays where the storage volumes are part of remote replication pairs need to be part of the same group.
  - The array which provides backend storage needs to be part of the security group as the storage virtualizer
  - Cluster members and the cluster should be part of the same group.
  - When host is presented storage from an array, the host , array, and fabric elements in path need to be part of the same group.
  - Virtual switches that are part of the physical switch should also be mapped to the same security group.

Create a Security Group

Required only for Operator or Guest users:

To create a security group, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** > **Security Groups**. The Security Groups view is displayed.
2. Click **New** on the view toolbar. The Security Group form is displayed.
3. Make your configuration choices. (See the Security Group Attributes table.)
4. Click one of the save options to save the security group:
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The security group is displayed in the Security Groups view.
5. See "Methods for Assigning Nodes to Security Groups" on the next page.

Security Group Attributes

Attribute	Description
Name	Enter the name that uniquely identifies this Security Group. Type a maximum of 255 characters. Alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _+ -) are permitted.
UUID	SOM assigns a Universally Unique Object Identifier to the security group. This UUID is unique across all databases.
Description	Type a maximum of 2048 characters. Alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _+ -) are permitted.

Modify a Security Group

To modify a security group, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** > **Security Groups**. The Security Groups view is displayed.
2. Select the security group that you want to modify from the table view.
3. Click **Open**. The security group is displayed in the Security Group view.
4. Make the required changes to the security group.
5. Click to save changes to the security group. The Security Group View is refreshed to display the changes to the security group.

Delete a Security Group

To delete a security group, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** >**Security Groups**. The Security Groups view is displayed.
2. Select the security group that you want to delete from the table view.
3. Do one of the following.
   - Click **Delete**. The delete confirmation message is displayed. Click **OK** to delete the security group.
   - Click **Open**. The security group is displayed in the Security Group view. Click . The delete confirmation message is displayed. Click **OK** to delete the security group.

   The security group is removed from the Security Groups view.

Methods for Assigning Nodes to Security Groups

You can assign nodes to security groups using any of the following:

- "Configure Security Using the Security Wizard" on page 97
- Node form
  However, until you define at least one security group in addition to the default security groups, the security group attribute does not appear in the Node form and the Security Group column does not appear in the Nodes view.

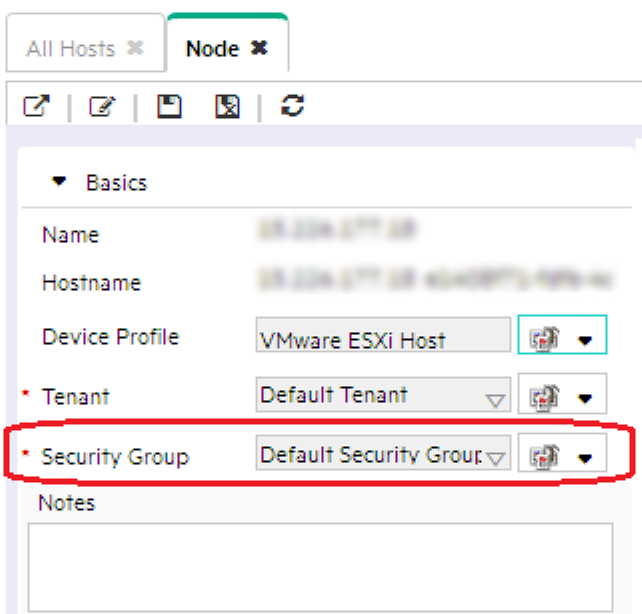

Tip: Administrators can use security groups in node group definitions that become filters in SOM views. If an SOM user cannot access any nodes in a particular node group, that filter dynamically disappears from the filter selection list in the SOM views.

Configure Security Group Mappings

Required only for Operator or Guest users:

Security Group Mappings control the nodes that visible to operators and guests, and the actions that the operators and guests can perform with those visible nodes. (Security Group Mappings are irrelevant to users assigned to the Administrators User Group. Administrators automatically see all nodes and have full access rights.)

SOM provides the *default* Security Group Mappings that allow all SOM operators and guests to see all nodes. Administrators can delete these *default* mappings and create new mappings that provide more limited control. (Deleting a security group mapping does not delete the associated user group or security group. Therefore, administrators can map the associated user groups and security groups in other ways with more limited control.)

SOM provides predefined *Object Access Privileges*. The Object Access Privilege determines the level of access that each User Group has to the visible nodes. Level of node access includes the actions that can be performed on the nodes.

For example, if an SOM operator is mapped to a User Group with **SOM Level 2 Operators**, but their Security Group Mapping's *Object Access Privilege* is **Object Operator Level 1** (with more limited access privileges than Level 2), that SOM operator *sees* all of the actions available to SOM Level 2 Operators, but can run only those *actions allowed* for SOM Level 1 Operators.

If an SOM operator or guest is assigned to Multiple predefined SOM User Groups, the SOM console displays all the parts of SOM that are available to the highest User Group. However, If an SOM operator or guest is assigned to Multiple *Object Access Privileges*, actions available for each node are determined by the node's Security Group Mapping. If mapped to the same security group multiple times, the highest access level is available.

Administrators can map user groups to security groups using the following methods:

- "Configure Security Using the Security Wizard" on page 97
- "Map User Groups to Security Groups" below

Map User Groups to Security Groups

(*Required only for Operator or Guest users*)

To assign a user group to a security group, follow these steps:

1. From the workspaces navigation panel, select the **Configuration**> **Security** > **Security Group Mappings**. The Security Group Mappings view is displayed.
2. Click **New** on the view toolbar. The Security Group Mapping form is displayed.
3. Make your configuration choices. (See the Security Group Mapping Attributes table.)
4. Click one of the save options to save the security group mapping:
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The security group mapping is displayed in the Security Group Mappings view.

Security Group Mapping Attributes

Attribute	Description
User Group	Specify the user group to be assigned to the security group. In the **User Group** attribute, click the Lookup icon. • To create new User Group, click the **New** icon and provide the required information. (See "Create a User Group" on page 87 for more information.) • To select a User Group configuration, click the **Quick Find** icon and make a selection.
Security Group	Specify the security group to be assigned to the user group. In the **Security Group** attribute, click the Lookup icon. • To create new security group, click the **New** icon and provide the required information. For more information, see "Create a Security Group" on page 91. • To select a security group configuration, click the **Quick Find** icon and make a selection.
Object Access Privilege	Determines the level of access each user account in the user group has to the nodes assigned to its security group. In the **Object Access Privilege** attribute, select a privilege level from the drop-down list. SOM provides the following privileges: • Object Administrator • Object Operator Level 2

Security Group Mapping Attributes, continued

Attribute	Description
User Group	Specify the user group to be assigned to the security group. In the **User Group** attribute, click the Lookup icon. • To create new User Group, click the **New** icon and provide the required information. (See "Create a User Group" on page 87 for more information.) • To select a User Group configuration, click the **Quick Find** icon and make a selection.
Security Group	Specify the security group to be assigned to the user group. In the **Security Group** attribute, click the Lookup icon. • To create new security group, click the **New** icon and provide the required information. For more information, see "Create a Security Group" on page 91. • To select a security group configuration, click the **Quick Find** icon and make a selection.
	• Object Operator Level 1 (with more limited access privileges than Level 2) • Object Guest

Default Object Access Privileges

When you map user groups to security groups, you also determine the Object Access Privilege.

The Object Access Privilege determines the level of access each User Account in the User Group has to the nodes associated with the assigned Security Group. For more information, see "Control Menu Access" on page 99.

SOM provides the following Object Access Privileges. Each can be used in any number of security group mappings:

- Object Administrator
- Object Operator Level 2
- Object Operator Level 1 (with more limited access privileges than Level 2)
- Object Guest

You cannot change the Object Access Privileges definitions that SOM provides.

Configure Security Using the Security Wizard

The Security Wizard enables you to configure User Accounts, User Groups, and Security Groups. You can access the pages of the wizard in any order.

Notes before you begin using the wizard:

- You can choose to perform all the security configuration tasks using the wizard or you can access individual pages of the wizard specific to any task.
- Your configuration changes are not saved until you click **Save and Close** in the wizard.

To configure security using the Security Wizard, follow these steps:

1. From the workspaces navigation panel, select the **Configuration** > **Security** > **Security Wizard**. The Welcome page of the Security Wizard is displayed.
2. Click **Next**. The **Map User Accounts and User Groups** page is displayed.
3. Create User Accounts with the following steps:
   a. Click **New**. The **Create User Account** dialog box is displayed.
   b. Enter the following information.

Username	Enter the user name. You can use up to 40 alpha-numeric characters. Do not use punctuation, spaces, or underline characters.
Password	Type any combination of alpha-numeric characters, punctuation, spaces, and underline characters.

Note: The Security Wizard is unable to create accounts for use with LDAP . These accounts may be created using the User Accounts Form. For more information, see "Create a User Account" on page 82.

   c. Click **Close** to add the user account and close the dialog box.
   d. Click **Add** to add more user accounts.
   e. Repeat Step 3 (a) and 3 (b) for each User Account that you want to create.
4. Create User Groups with the following steps:
   a. Click **New**. The **Create User Group** dialog box is displayed.
   b. Enter the following information:

Name	Enter the name that uniquely identifies the User Group. You can use up to 40 alpha-numeric characters. Do not use spaces.

Display Name	Enter the name that you want to be displayed in the SOM console to identify this User Group. You can use up to 50 characters with any combination alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _+ - are permitted).
Directory Service Name*(Optional)*	When a directory service defines this User Group, enter the group's Distinguished Name. SOM communicates with the directory service using Lightweight Directory Access Protocol (LDAP).
Description	Type a maximum of 2048 characters to describe this User Group. Alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _+ -) are permitted.

c. Click **Close** to add the user group and close the dialog box.

d. Click **Add** to add more user groups.

e. Repeat Step 4(a) and 4 (b) for each User Group that you want to create.

5. Map User Accounts to User Groups with the following steps:

   a. Select a row in the **User Accounts** table.

   b. In the **User Groups** table, click the left arrow in the row of the User Group you want to assign to the selected User Account.

      The User Account and User Group names appear in the **User Account Mappings** table.

   c. Repeat steps 5(a) and 5(b) for each User Account Mapping.

6. Map User Groups to Security Groups with the following steps:

   a. Select a row in the **User Groups** table.

   b. In the **Security Groups** table, select the left arrow in the row of the Security Group you want to assign to the selected User Group.

      The User Group and Security Group names appear in the **Security Group Mapping** table.

   c. Repeat steps 6(a) and 6(b) to assign each User Group to a Security Group.

7. Assign Nodes to Security Groups.

   a. Select a row in the Security Groups table. The nodes that are already assigned to the selected Security Group is displayed.

   b. Select a row in the **Available Nodes** table. The selected node to be assigned is displayed. Use Ctrl + click for multiple selections.

   c. Click to assign the selected node to the selected node group.

   d. Repeat step 7(a) to step 7(c) to assign more nodes to the security groups.

8. Click **Next**. The View Summary of Changes page is displayed.
9. Click **Save & Close** to save the Security Configuration.

Control Menu Access

Access to the Tools and Actions menu items is controlled by Security Group Mapping configuration settings User — Group, Security Group, and *Object Access Privilege*.

> **Tip:** You can also right-click any object in a table or map view to access the items available within the **Actions** menu.

Note the following:

- User groups determine access to SOM console workspaces, views and forms. User groups also determine the Tools and Actions that the users in the User Group can access.
- You must assign each user account to one of the predefined user groups before that user can access SOM. For more information, see " Predefined User Groups " on page 86.
- If you map a User Account to two or more SOM User Groups, SOM gives the User Account the privileges associated with each User Group to which the User Account is assigned.
- Security Groups are optional and control (through User Groups) which Users can access a node and its hosted objects, such as an interface. Each node is associated with only one Security Group.

  > **Note:** Users see only those members of an object group (for example, Node Group or Router Redundancy Group) for which they have access. If a user cannot access any nodes in the group, the group is not visible to that user.

- Object Access Privileges are associated only with security groups and their associated User Groups. Object Access Privileges determine the Tools and Actions that the User Group can access for the nodes they are permitted to view.
  - If a user account is assigned to an SOM user group with *more privileges* than the Object Access Privilege, the user sees all of the actions available for the User Group (not restricted because of the Object Access Privilege setting). For example, if a user account is assigned to the user group **SOM Level 2 Operators** and has an Object Access Privilege of **Object Operator Level 1** (with more limited access privileges than Level 2 Operators) for a set of nodes, the operator sees all actions available to Level 2 Operators.

  - If a user account is assigned to an SOM user group with *fewer privileges* than the Object Access Privilege, the user cannot see all of the actions available for the Object Access Privilege. For example, if a User Account is assigned to the User Group **SOM Level 1 Operators** (with more limited access privileges than Level 2 Operators) and has an Object Access Privilege of **Object Operator Level 2** for a set of nodes, the operator can see only those actions that are available to Level 1 Operators. As an administrator, you must do either of the following:

- Configure the **Menu Item Context Basic Details** to change the **Required SOM Role** for the menu item
- Assign the operator User Account to the **SOM Level 2 Operators** User Group.

- All menu items are visible to users, but an *Access Denied* message displays when any user with insufficient privileges tries to use a menu item. For example, both Level 1 or Level 2 Operators are denied access to the Communication Settings action.
- If the menu item does not require node access, (for example, **Status Details** for a Node Group) SOM uses the privileges assigned to the SOM User Group that is mapped to the User Account.

User Group and Object Access Privilege Required for the Tools Menu:

Access to the SOM Tools menu items is determined by User Group and the Security Group Object Access Privilege that is set for the node. For information about Tools Menu Access Limitations Click here .

SOM Tools Menu Access Limitation

Tools Menu Item	SOM User Group	Object Access Privilege
Find Element	SOM Guest Users	Object Guest
Signed In Users	SOM Administrators	Object Administrator

User Group and Object Access Privilege Required for the Actions Menu:

Access to the SOM Actions menu is determined by User Group and the Security Group Object Access Privilege that is set for the node.

URL Action Access Limitations

Action Menu Item	Submenu Item	SOM User Group	Object Access Privilege
Configuration Details	Communication Settings	SOM Administrators	Object Administrator
Configuration Details	Monitoring Settings	SOM Level 1 Operators	Object Operator Level1
Custom Attributes		SOM Administrators	Object Administrator
Graphs		SOM Level 1 Operators	Object Operator Level 1
Management Mode		SOM Level 2 Operators	Object Operator Level 2

URL Action Access Limitations, continued

Action Menu Item	Submenu Item	SOM User Group	Object Access Privilege
Node Group Details	Show Members (Include Child Groups)	SOM Level 1 Operators	Object Operator Level 1
Node Group Details	Preview Members (Current Group Only)	SOM Level 1 Operators	Object Operator Level 1
Node Group Details	Status Details	SOM Level 1 Operators	Object Operator Level 1
Node Group Membership		SOM Administrators	Object Administrator

Note: Each Tools and Action menu item provided by SOM is also associated with a *predefined SOM Role*. If you change the setting for a menu item provided by SOM to a Role that is a *lower level Role* than the *predefinedSOM Role* assigned to the menu item, SOM ignores that change. Any user group with the lower level role than the *predefined SOM Role* cannot access the menu item.

Check Security Configuration

Each SOM user can be assigned to multiple Security Group Mappings. The *Object Access Privilege* determines what SOM users can do with a node object. For example, if their User Group is **SOM Level 2 Operators**, but the Object Access Privilege is **Object Operator Level 1** (with more limited access privileges than Level 2), each user assigned to the Security Group Mapping *sees* all of the actions available to a Level 2 Operator, but can run only those *actions allowed* for Level 1 Operators. If an SOM user is assigned to multiple Security Group Mappings, that user sees all the parts of SOM that are provided to the highest User Group setting and access for each node is determined by the node's Security Group Mapping.

Post-configuration Tasks

After configuring user passwords and roles, communicate the following information to your team:

- "Opening the SOM Console" below
- "Configuring Sign-In to the SOM Console" on page 103
- "Signing Out from the SOM Console" on page 103

Opening the SOM Console

Provide each user with the following information:

`http://<serverName>:<portNumber>/som/main`

<serverName> = the fully-qualified domain name of the SOM management server

<portNumber> = the SOM HTTP port number, which can be configured during the installation.

When a SOM management server has more than one fully-qualified domain name, SOM chooses one during the installation process. There are two ways to find out which domain name that is in use by SOM in your network environment:

- Click **Help** → **System Information** and navigate to the **Server** tab. Locate the **Official Fully Qualified Domain Name (FQDN)** attribute value.
- Use the `somofficialfqdn.ovpl` command. For more information, see the CLI Reference Pages.

To determine the current port number configuration, look at the line `#HTTP Ports` in the `nms-local.properties` file (see table for the location of this file).

Determine the SOM console Port Number

Operating System	Identify Current Port Number
Windows	`<Install_Dir>\HP\HP BTO Software\conf\nnm\props\nms-local.properties`
Linux	`<Install_Dir>/var/opt/OV/conf/nnm/props/nms-local.properties`

The browser requirements to use the SOM console are as follows:

- Pop-ups, cookies, and JavaScript must be enabled.
- Each user's screen resolution must be 1024x768 pixels or higher.
- When using Microsoft Internet Explorer as your browser, you can access multiple browser sessions of SOM.
- When using Mozilla Firefox as your browser, multiple browser sessions point to the same window.

> **Note:** Users can bookmark the URL for the SOM console. Use the URL for the SOM console rather than the SOM Welcome page.

To open the console:

1. Type the following URL (Uniform Resource Locator) into your browser navigation bar:
   `http://<serverName>:<portNumber/som/main/`
2. Sign in with the following name and password:
   <name you configured>
   <password you configured>
3. Click the **Sign In** button.
4. The console opens in a new window.

5. *Optional*. Close the SOM Welcome page.

> **Note:** If you do not close the SOM Welcome page or sign out, you can relaunch the console from the SOM Welcome Page without signing in again.

To refresh the console window:

Click the Refresh icon in the tool bar of any SOM window.

Configuring Sign-In to the SOM Console

After entering the URL to access the SOM console, SOM prompts you to sign into the console:

1. At the **User Name** prompt, enter the user name that was provided by your administrator.
2. At the **Password** prompt, enter the password that was provided by your administrator.
3. Click the **Sign In** button.

After you access the SOM console, the user account name and the highest associated object access privilege appear in the upper right corner of the console.

Signing Out from the SOM Console

To sign out from the console, follow these steps:

1. Select **File → Sign Out**.
2. Click **OK**.

Note the following:

- Sign in is not preserved across user sessions. After signing out, each user must sign in again.
- You must sign out of each browser session that is running SOM. For example, if you have signed in twice with two different browsers, signing out in one browser does not cause you to lose access in the other browser.
- By default, SOM automatically signs out any user after 18 hours of inactivity. An administrator can configure the timeout period.

Troubleshoot Access

> **Tip:** Select **Help → System Information** to view the User Name, SOM Role, and User Group for the current SOM session.

SOM provides several tools to help you troubleshoot and monitor SOM access:

- "Check Security Configuration" below
- "View the Users who are Signed in to SOM" below
- "Restore the Administrator Role" on the next page
- "Restore SOM Access to the System User" on the next page

Check Security Configuration

Each SOM user can be assigned to multiple Security Group Mappings. The *Object Access Privilege* determines what SOM users can do with a node object. For example, if their User Group is **SOM Level 2 Operators**, but the Object Access Privilege is **Object Operator Level 1** (with more limited access privileges than Level 2), each user assigned to the Security Group Mapping *sees* all of the actions available to a Level 2 Operator, but can run only those *actions allowed* for Level 1 Operators. If an SOM user is assigned to multiple Security Group Mappings, that user sees all the parts of SOM that are provided to the highest User Group setting and access for each node is determined by the node's Security Group Mapping.

View Summary of Changes in the Security Wizard

Use the Security Wizard **View Summary of Changes** option to view your recent configuration changes, including the following:

- The user accounts created
- The user groups created
- The security groups created
- The user accounts and user groups mappings
- The user groups and security groups mappings
- The security groups that have new nodes assigned to them

To view the summary of security configuration changes:

From the **Security Wizard** main page, select the **View Summary of Changes** option.

SOM displays a summary of the configuration changes made since you last saved your changes.

View the Users who are Signed in to SOM

You can view the users who are currently signed in to SOM. This option is useful when you want to determine which users and systems are available. For example, you might want to view the users who are signed in before shutting down a system.

To see the list of users who are currently signed in to SOM:

- Select **Tools → Signed In Users**.

SOM displays the number of users currently signed in to SOM as well as each user name, IP address of the client that is running the SOM console, and the sign in time of the user.

Restore the Administrator Role

If you have accidentally configured SOM so that zero SOM users are mapped to the SOM user group (preventing anyone from being able to access the Configuration workspaces), then an administrator can access the SOM console as the `system` user to correct the problem.

Sign in to the console using the password that was configured for the `system` user when SOM was first installed.

If you do not remember the password assigned to the `system` user, use the `somchangesyspw.ovpl` command to reset the `system` user's password.

Restore SOM Access to the System User

SOM provides an `nms-roles.properties` file that stores part of the `system` user configuration. Do not modify this file.This file is located in the following directory:

- **Windows**:
  `Install_Dir\HP\HP BTO Software\nmas\NNM\conf\props\nms-roles.properties`
- **Linux**:
  `Install_Dir/var/opt/OV/nmas/NNM/conf/props/nms-roles.properties`

To verify the contents of this file:

1. With a text editor, open the `nms-roles.properties` file.
2. Verify that the following required line is present:

   ```
   system = system,admin
   ```

3. Save and close the file.

Configuring Node Groups

A node group is a collection of nodes (elements) or child node groups that have the same device filter criteria. You can use node groups to categorize elements for easier administration and monitoring. Elements can be categorized based on filters such as devices vendor, model, profile, category, and so on. Node groups act as filters and provide you with filtered views or help you limit access to a set of nodes through security mappings.

Elements are automatically assigned to node groups based on predefined attributes. SOM provides default node groups. For more information, see "Node Groups Provided by SOM" on the next page.

You can create additional node groups based on your environment and requirements. You can define attributes to determine node group membership. Each node group is defined using one or more of the following options:

- **Device Filters**: Provides filters such as **Device Category**, **Device Vendor**, **Device Family**, and **Device Profile**. Nodes must match at least one specification to belong to the node group.

- **Additional Filters**: Provides option to specify additional filters using Boolean expressions based on a list of object attributes.
- **Additional Nodes**: Enables you to add additional nodes to the node group based on the *hostname* attribute of the node.
- **Child Node Groups**: Enables you to add node groups to the node group to establish hierarchical containers.

It is recommended that you do not modify the default node groups as the data collection policies are preconfigured to automatically collect data from discovered elements that are categorized into these node groups.

SOM combines the results of all node group configuration settings in the following manner:

- SOM first evaluates **Device Filters**. If any exist, nodes must match at least one specification to belong to this node group.
- SOM then evaluates any **Additional Filters**. Nodes must pass all additional filters specifications to belong to this node group.
- Any nodes specified as **Additional Nodes** are always included in the node group, regardless of any filters.
- Any child node group results are treated the same as **Additional Nodes**.

Node Groups Provided by SOM

SOM provides default node groups to automatically group newly discovered elements based on device filter criteria. In addition to filtered views, node groups provide default security mappings and data collection policies for the elements in each group. You can view the node groups to which an element belongs in the summary section of the analysis pane.

It is recommended that you do not modify the default node groups as the data collection policies are preconfigured to automatically collect data from discovered elements that are categorized into these node groups.

Name	Description
All Elements	This node group includes all elements discovered by SOM. It includes Hosts, Switches, Storage Systems, and Fabrics as child groups.
FC Fabrics	Any fabric discovered within your management domain are automatically included in this node group.
FC Switches	This node group is populated with a list of categories for storage switches. Any switch, physical or virtual within your management domain is included in this node group.
Hosts	Any host, physical or virtual within your management domain is included in this node group.

Name	Description
Storage Systems	Any storage devices discovered within your management domain are automatically included in this node group.

Recommendations for Planning Node Groups

Some key points to consider while planning node groups for your environment:

- Keep in mind that node groups add overheads to the system. Therefore, ensure that you have valid use cases based on your needs when creating node groups.
- Create node groups that cater to a definite purpose. Identify your topmost use cases before you begin planning your node groups. For example, you could create node groups for managing Windows hosts, Linux hosts or storage devices based on vendor, model or the device profile. You could then attach data collection or monitoring policies to these node groups.
- Use different node groups for different purposes. Not all node groups created for data collection makes sense for filtering views or restricting node access. So you will need to configure them independently based on the purpose.
- Find a balance by creating a rich set of groups for monitoring purpose and viewing purpose without overloading the system with a large number of superfluous node groups that will never be used.
- Do not use the Additional Nodes tab extensively to add nodes to a node group as it consumes excessive resources on the management server. As a rule of thumb, node group definitions should be filter-driven and this feature should be used as an exception.

Create a Node Group

You can create any number of node groups in addition to the default node groups provided by SOM.

To create a node group, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Object Groups** > **Node Groups**. The Node Groups view is displayed.
2. Click ✱ **New** on the view toolbar. The Node Group form is displayed.
3. Enter node group details. (See the Node Group attributes below.)
4. Configure a device filter to the node group with the following steps:
   a. Under the Device Filters tab in the right pane, click ✱ **New**. The Node Device Filter form is displayed.
   b. Select the device filter options. (See the Device Filter options below.)
   c. Click one of the **Save** options.

- **Save** – To save the form.
- **Save and New** – To save and open a new form.
- **Save and Close** – To save and close the form.

Note: Repeat Step 4 to configure more device filters.

5. Associate additional filters to the node group using the Filter Editor. See ""Specify Additional Filters for Node Groups " on page 110" for more information.
6. Associate additional nodes based on the *hostname* attribute of the node with the following steps:
   a. Under the Additional Nodes tab in the right pane, click **New**. The Additional Node form is displayed.
   b. Enter the fully-qualified, host name of the node as it appears in the Nodes view.

   Note: This entry is case-sensitive. The name you provide must match the host name attribute as it appears in the Nodes view (**Inventory** > **Nodes** view).

   c. **Note**: Additional nodes must be manually deleted as these associations are not automatically generated by SOM.

   d. Click one of the **Save** options.
      - **Save** – To save the form.
      - **Save and New** – To save and open a new form.
      - **Save and Close** – To save and close the form.

Note: Repeat Step 6 to associate additional nodes to the node group.

7. Add child node groups to the node group with the following steps:
   a. Under the Child Node Groups tab in the right pane, click **New**. The Node Group Hierarchy form is displayed.
   b. Enter the child node groups details. (See "Child Node Group Attributes" on page 110 below.)

   Note: To create a new child group, follow Steps 1 through 8 in this procedure.

   c. Click one of the **Save** options.
      - **Save** – To save the form.

- **Save and New** – To save and open a new form.
- **Save and Close** – To save and close the form.

8. Click one of the **Save** options to create the node group.

- **Save** – To save the form.
- **Save and New** – To save and open a new form.
- **Save and Close** – To save and close the form.

Node Group Attributes	Description
Name	The name of the node group. The text string can be alpha-numeric with a maximum of 255 characters and can include spaces and special characters (~ ! @ # $ % ^ & * ( ) _+ -).
Notes	Can be used to document information about the node. Information might include why the node is important, if applicable, or to what customer, department, or service the node is related. Additional information might include where the node is located, who is responsible for it, and its serial number. You might also track maintenance history using this attribute. A maximum of 1024 characters, alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _+ -) are permitted. **Note**: You can sort your node group table views based on this value. Therefore, you might want to include keywords for this attribute value.

Device Filters	Description
Device Category	*Optional*: Select from the drop-down list that has the available device categories. SOM provides four predefined categories – Hosts, Storage Systems, FC Switches, and FC Fabrics.
Device Vendor	*Optional*: Select from the drop-down list that displays the available device vendors.
Device Family	*Optional*: Select from the drop-down list that displays the available device families.
Device Profile	*Optional*: Select from the drop-down list to choose from the predefined device profiles or click Lookup for additional options.

Child Node Group Attributes	Description
Child Node Group	Select the node group from the drop-down list or click lookup for additional options. • Show Analysis – Displays Analysis Pane information for the selected object. • Quick Find – Displays a list of valid choices for populating the current attribute field. • Open – Opens the form for the related object instance that is currently selected in the lookup field. You can use this option to make changes to the selected object. • **New** – Opens a new form to create a new instance of the object.
Expand Child in Parent Node Group Map	Used to indicate whether all the nodes contained in a Child Node Group are displayed in the Parent Node Group Map as a part of the parent node group map. Select this option for each child node group if you want to have an expanded view of all the child nodes displayed in the parent node group view. If enabled, each node in the Child Node Group appears on the Parent Node Group Map. If disabled, a hexagon represents a Child Node Group on the Parent Node Group Map. Multiple child node groups, if any, are also displayed in the same manner. If a child node group is also a parent, its member nodes and child groups are displayed in the parent node group map if the Expand Child in Parent Node Group Map option is selected for each child node group. **Note**: This attribute appears in the Child Node Groups tab of the Node Group Form.

Specify Additional Filters for Node Groups

You can specify additional filters for node groups using Boolean expressions based on a list of object attributes. Use the **Additional Filters Editor** to create expressions that refine the requirements for membership for a node group.

Read the following topics to create additional filters for a node group:

- "Guidelines for Creating Additional Filters for Node Groups" on the next page
- "Add Boolean Operators in the Additional Filters Editor" on page 113
- " Create an Additional Filters Expression" on page 115

Guidelines for Creating Additional Filters for Node Groups

The **Additional Filters Editor** enables you to create expressions to further define the nodes to be included in a node group. Make sure to design any complex additional filters offline as a Boolean expression first. This method can help to minimize errors when entering expressions using the **Additional Filters Editor**.

Notes on additional filters for a node group:

- SOM treats each set of expressions associated with a Boolean Operator as if it were enclosed in parentheses and evaluated together rather than in order of grouping as the nesting implies. Therefore, when using the AND operator to combine expressions that include Custom Attributes, include only one `customAttrName/customAttrValue` pair in the expression. Otherwise, if you use multiple `customAttrName` and `customAttrValue` pairs with the AND operator, the results might not be as expected.
  Click here for an example.

  In the following example, because the AND Boolean operator indicates that the system should evaluate all of the `customAttrname` and `customAttrvalue` pairs together, it is not possible for any nodes to match this Additional Filters expression:

  Additional Filter Expression Example 1

  ```
  ((customAttrName = capability) AND (customAttrValue = com.hp.som.capability.card.fru))
  AND ((customAttrName = location) AND (customAttrValue = datacenter1))
  ```

  This is because `customAttrName` would need to match both capability and location at the same time. However, if you use the OR operator to combine the `customAttrName` and `customAttrValue` pairs as shown in the following example, the filter should work as expected.

  Additional Filter Expression Example 2

  ```
  ((customAttrName = capability) AND (customAttrValue = com.hp.som.capability.card.fru))
  OR ((customAttrName = location) AND (customAttrValue = datacenter1))
  ```

  Using the Node values listed in the following table, all three nodes (nodeA, nodeB, and nodeC) pass the filter in Example 2 because each of these nodes has either the value `com.hp.som.capability.card.fru` for capability or the value `datacenter1` for location.

  Example Data

Node Name	capabilty	customAttrName	customAttrValuee
nodeA	com.hp.som.capability.card.fru	location	datacenter1
nodeB	com.hp.som.capability.card.fru	<undefined>	<undefined>
nodeC	<undefined>	location	datacenter1

- Use the EXISTS and NOT EXISTS operators when you want the system to consider nodes that either do or do not have any Capabilities or Custom Attributes when evaluating the Filter String.
- View the expression displayed under Filter String to see the logic of the expression as it is created.

- The AND and OR Boolean Operators must contain at least two expressions as shown in the example below.

```
AND
sysName like cisco*
sysName != cisco2811
OR
sysLocation = Boston
sysContact In (Johnson,Hickman)
```

SOM evaluates the expression above as follows:

```
sysName like cisco* AND sysName != cisco2811 AND (sysLocation = Boston OR sysContact in
(Johnson, Hickman))
```

- SOM finds all nodes with a (system name) sysName beginning with cisco, except not cisco2811.
- Of these nodes, SOM then finds all nodes with a (system location) sysLocation of Boston or (system contact name) sysContact of Johnson or Hickman.
- SOM evaluates only those nodes that contain values for all of the attributes included in the Additional Filter expression.
  Click here for an example.

  If your Node Group filter expression includes the capability and customAttrName attributes, then SOM evaluates only nodes that have a value defined for both capability and customAttrName. For example, if you create a Node Group using the following Additional Filters expression, then SOM evaluates only those nodes that have a value defined for capability and a value defined for customAttrName:

  ```
  (capability = com.hp.som.capability.card.fru) OR (customAttrName = location)
  ```

  Using the Node values listed in the following table, SOM only evaluates nodeA. This is because nodeA contains a value for capability and a value for customAttrName. SOM does not evaluate nodeB because it does not have a value for customAttrName. SOM does not evaluate nodeC because it does not have a value for capability. NodeA also passes Node Group Additional Filter because its capability value of com.hp.som.capability.card.fru matches the value specified in the Additional Filter expression. Therefore, only nodeA is included in this example Node Group.

Node Name	capabilty	customAttrName	customAttrValuee
nodeA	com.hp.som.capability.card.fru	location	datacenter1
nodeB	com.hp.som.capability.card.fru	<undefined>	<undefined>
nodeC	<undefined>	location	datacenter1

  Tip: You can populate a placeholder value, such as "none" or "undefined" for any attribute that you want to use in an Additional Filter.

- The placement of your cursor and the subsequent text that is selected is important when performing operations using the Additional Filters Editor. For example, you append to or replace, the expression that is selected.
- The placement of your cursor and the subsequent text that is selected is especially important when adding your Boolean operators.
- You can drag any of the following items to a new location in the Filter String:
  - Filter Editor Options: AND, OR, NOT, EXISTS, NOT EXISTS
  - Filter Expression (Attribute, Operator and Value)
- When moving items in the Filter String, note the following:
  - Click the item you want to move before dragging it to a new location.
  - As you drag a selected item, an underline indicates the target location.
  - If you are moving the selection up, SOM places the item above the target location.
  - If you are moving the selection down, SOM places the item below the target location.
  - If you attempt to move the selection to an invalid target location, SOM displays an error message.

Add Boolean Operators in the Additional Filters Editor

Note the following when adding or deleting Boolean Operators using the **Additional Filters Editor**:

- Add your highest level Boolean operator first. For example, **AND** is the highest level Boolean operator in the following expression:

  (sysName like cisco* OR sysName like hp*) **AND** (sysLocation = Boston OR sysContact in Johnson,Hickman)
- Add each additional Boolean Operator before the expressions to which it applies.
- Select the appropriate Boolean Operator in the expression before you add the expressions to which the Boolean Operator applies.
- When a Boolean Operator is selected and you click **Delete**, any expressions that are associated with the Boolean Operator are also deleted.

  In the example expression below, If you select **AND** and then click **Delete,** the **Additional Filters Editor** deletes the entire expression.

Click here for an example for creating additional filters for node groups.

Node Group Additional Filters Expression Example

```
((sysName like cisco* OR sysName like hp*) AND (sysLocation = Boston OR sysContact in
(Johnson, Hickman)))
```

To add the expression above, after you are in the Additional Filters Editor, follow these steps:

1. Click **AND**.
2. Click **OR**.
3. Select the **OR** you just added to the expression.
4. In the **Attribute** field select **sysName** from the drop-down list.
5. In the **Operator** field, select **like** from the drop-down list.
6. In the **Value** field, enter **cisco***.
7. Click **Append**.
8. In the **Attribute** field, select **sysName** from the drop-down list.
9. In the **Operator** field, select **like** from the drop-down list.
10. In the **Value** field, enter **hp***.
11. Click **Append**.
12. Select the **AND** that you previously added to the expression.
13. Click **OR**.
14. Select the **OR** you just added to the expression.
15. In the **Attribute** field, select **sysLocation** from the drop-down list.
16. In the **Operator** field, select **=** from the drop-down list.
17. In the **Value** field, enter **Boston**.
18. Click **Append**.
19. In the **Attribute** field, select **sysContact** from the drop-down list.
20. In the **Operator** field, select **in** from the drop-down list.
21. In the **Value** field:
    a. enter **Johnson** and press **<Enter>.**
    b. On the new line, enter **Hickman.**
22. Click **Append**.
23. Click **Save** to save your additional filters.
24. Select **Actions** > **Preview Members (Current Group Only)** to view the members of the Node Group that is a result of this filter.

    Tip: To test the effects of your node group definition on child node groups, in the Node Group form,

> select **Save**, then **Actions** > **Node Group Details** > **Show Members (Include Child Groups)**. SOM displays the members of the current node group members as well as the members of each associated child node group. Depending on the complexity of your node group hierarchy, SOM might take some time to complete updating the results. Click **Refresh** to check for the most recent changes to the contents of the node group.

25. Click Refresh to check for the most recent changes to the contents of the node group.

Create an Additional Filters Expression

Use the **Additional Filters Editor** to create expressions that refine the requirements for membership in a node group. Make sure to design any complex additional filters offline as a Boolean expression first. This method can help to minimize errors when entering your expressions using the **Additional Filters Editor**.

If any additional filters are created, SOM combines any **Device Filters** and **Additional Filters** using the AND Boolean operator as follows:

- SOM first evaluates any Device Filters. Nodes must match *at least one* Device Filter specification to belong to this node group.
- SOM then evaluates the Additional Filters expression. Nodes *must also match all* Additional Filters expression specifications to belong to this node group.

To create an Additional Filters expression:

1. Establish the appropriate settings for the Additional Filters you need (see the Additional Filters Editor Choices and Additional Filters Editor Buttons table).
   a. Plan out the logic needed for your Filter String.
   b. Use the buttons on the bottom half of the Additional Filters Editor to establish the logic structure. See "Add Boolean Operators in the Additional Filters Editor" on page 113 for more information.

      For example, to establish the following structure, select **Insert**, then click **AND**, then **NOT**, and then **AND** a second time:

      `(( ) AND NOT ( ))`
   c. Now place your cursor in a location within the displayed Filter String, and use the top half of the filter editor to define the parameters of the selected filter requirement.

      For example, select a set of parentheses and use the Insert button to specify the filter requirement within those parentheses:

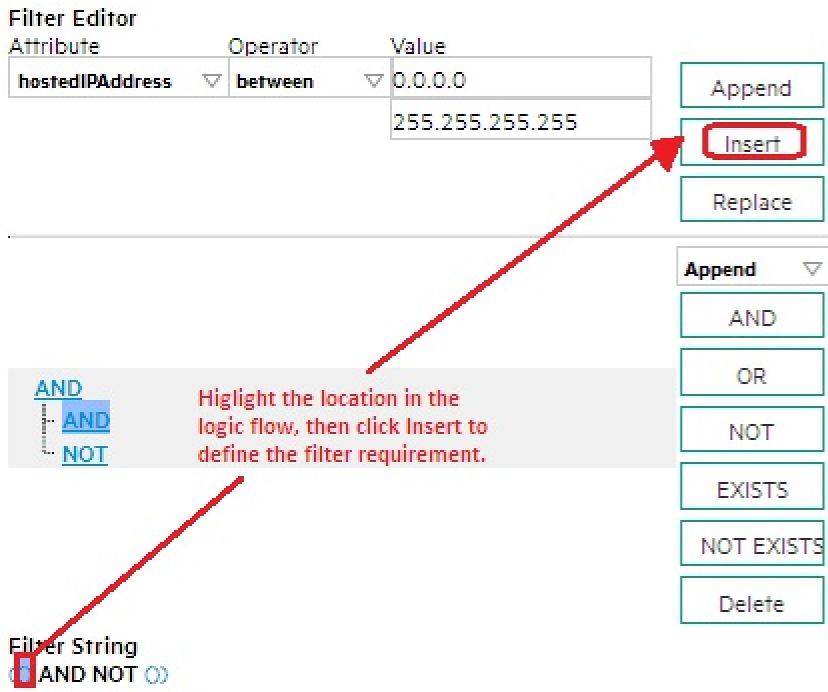

2. Click **Save and Close**.

Additional Filters Editor Choices for Node Groups

Attribute	Description
Attribute	SOM provides Additional Filters codes for a subset of the following object attributes: • tenantName (Name) • securityGroupName (Security Group) • sysName (System Name) • sysLocation (System Location) • sysContact (System Contact) • hostname (Hostname, case-sensitive) • hostedIPAddress (Address) • mgmtIPAddress (Management Address) • nodeName (Name)
Operator	The standard query language (SQL) operations to be used for the search.

Additional Filters Editor Choices for Node Groups, continued

Attribute	Description
	Note: Only the `is null` Operator returns null values in its search. Valid operators are described below. • **=** Finds all values equal to the value specified. Click here for an example. Example: `sysName = cisco2811` finds all devices with system name equal to **cisco2811**. • **!=** Finds all values not equal to the value specified. Click here for an example. Example: `sysName != cisco2811` finds all system names other than **cisco2811**. • **<** Finds all values less than the value specified. Click here for an example. Example: `mgmtIPAddress < 15.239.255.255` finds all IP address values less than **15.239.255.255** • **<=** Finds all values less than or equal to the value specified. Click here for an example. Example: `mgmtIPAddress <= 15.239.255.255` finds all IP address values less than or equal to **15.239.255.255**. • **>** Finds all values greater than the value specified. Click here for an example. Example: `mgmtIPAddress > 15.238.0.0` finds all IP address values greater than **15.238.0.0** • **>=** Finds all values greater than or equal to the value specified. Click here for an example. Example: `mgmtIPAddress >= 15.238.0.0` finds all IP address values greater than or equal to **15.238.0.0**. • **between** Finds all values equal to and between the two values specified. Click here for an example. Example: `mgmtIPAddress between 15.238.0.10 15.238.0.120` finds all IPv4 address values equal to or greater than **15.238.0.10** and equal to or less than **15.238.0.120**. • **in** Finds any match to at least one value in a list of values. Click here for an example. Example: `sysName in` 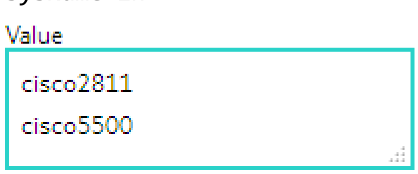

Additional Filters Editor Choices for Node Groups, continued

Attribute	Description
	finds all systems with names that are **cisco2811** or **cisco5500**. **Note:** Each value must be entered on a separate line as shown in the example. SOM displays the list of attributes using comma-separated values enclosed in parentheses, for example (**cisco2811, cisco550**). However, the comma-separated list is used only for display purposes. The actual delimiter is the new line. • **is not null** Finds all non-blank values. Click here for an example. Example: `sysName is not null` finds all systems that have a name value. • **is null** Finds all blank values. Click here for an example. Example: `sysName is null` finds all systems that do not have an assigned name value. • **like** Finds matches using wildcard characters. Click here for more information about using wildcard characters. The following attributes cannot be used with the `like` operator: • mgmtIPaddress The asterisk (*) character means *any number of characters of any type at this location*. **Note:** For optimum performance, avoid beginning your search string with an asterisk (*). The question mark (?) character means *any single character of any type at this location*. Examples: • `sysName like cisco*` finds all system names that begin with **cisco**. • `sysName like cisco??*` finds all system names that *start with* cisco **followed by two characters**. • `sysName like rtr??bld5*` finds all system names that have *specific characters at an exact location*, positions 1-3 (rtr) and 6-9 (bld5). • **not between** finds all values except those between the two values specified. Click here for an example. Example: `mgmtIPAddress not between 15.238.0.10 15.238.0.120` finds all IP address values less than **15.238.0.10** and greater than **15.238.0.120**. • **not in** Finds all values except those included in the list of values. Click here for an example.

Additional Filters Editor Choices for Node Groups, continued

Attribute	Description
	Example: `sysName not in` 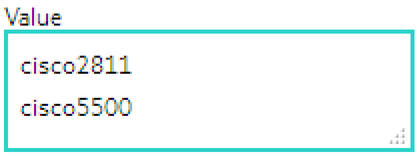 finds all system name values other than **cisco2811** and **cisco5500**. **Note:** Each value must be entered on a separate line as shown in the example. SOM displays the list of attributes using comma-separated values enclosed in parentheses, for example, (**cisco2811, cisco550**). However, the comma-separated list is used only for display purposes. The actual delimiter is the new line. • **not like** Finds all that do not have the values specified (using wildcard strings). Click here for an example. The following attributes cannot be used with the `not like` operator: • mgmtIPaddress The asterisk (*) character means *any number of characters of any type at this location*. The question mark (?) character means *any single character of any type at this location*. Examples: • `sysName not like cisco*` finds all system names that do not begin with **cisco**. • `sysName not like cisco??*` finds all system names that do not *begin with* cisco **followed by two characters**. • `sysName not like rtr??bld5*` finds all system names that do not have *specific characters at an exact location*, positions 1-3 (rtr) and 6-9 (bld5).
Value	The value for which you want SOM to search. Note the following: • The values you enter are case sensitive. • SOM displays a variable number of value fields depending on the Operator selected. For example, the `between` Operator causes two value fields to be displayed. • The `in` and `not in` operators require that each value be entered on a separate line.

Additional Filters Editor Choices for Node Groups, continued

Attribute	Description
	• When entering a value for the Capability attribute, copy and paste the Unique Key value from the Node form: Capability tab.

Additional Filters Editor Buttons

Button	Description
Append	Appends the current expression (Attribute, Operator,and Value) to the selected expression already included in the Filter String.
Insert	Inserts the current expression (Attribute, Operator,and Value) in front of the cursor location within the Filter String.
Replace	Replaces the selected expression with the expression displayed in the Attribute, Operator, and Value fields.
AND	Inserts the AND Boolean Operator in the selected cursor location. **Note:** View the expression displayed under **Filter String** to see the logic of the expression as it is created.
OR	Inserts the OR Boolean Operator in the current cursor location. **Note:** View the expression displayed under **Filter String** to see the logic of the expression as it is created.
NOT	Can be used in any part of the Filter String to specify that SOM should exclude nodes with values that pass the expression that immediately follows the NOT. For example, when evaluating the following Filter String, SOM includes nodes with a hostname that contains **router**, followed by any number of characters, followed by **hp.com** and excludes any nodes with a Device Profile that includes **Cisco** as the Vendor value: `(hostname like router*.hp.com OR NOT (devVendorNode = Cisco))`
EXISTS	Used for filters that include Capabilities or Custom Attribute names and values in the Filer String. Indicates that you want SOM to consider nodes that have Capabilities or Custom Attributes when evaluating the Filter String. **Tip:** When creating complex Filter Strings that include `customAttrName` and `customAttrValue` pairs as one component of an "***or***" statement, to prevent SOM from

Additional Filters Editor Buttons, continued

Button	Description
	excluding nodes that have zero Custom Attributes, use **EXISTS** or **NOT EXISTS** criteria for the `customAttrName` and `customAttrValue` pair definitions. Otherwise nodes that do not have any Custom Attributes are automatically excluded even if they have values that pass other aspects of your filter. For example, when evaluating the following Filter String, SOM includes nodes with a hostname that includes **router**, followed by any number of characters, followed by **hp.com** as well as any nodes that have the Custom Attribute **edge** and that edge value is **true**: `(hostname like router*.hp.com OR EXISTS((customAttrName=edge AND customAttrValue=true)))`
NOT EXISTS	Used for filters that include Capabilities or Custom Attribute names and values in the Filer String. Indicates that you want SOM to consider nodes that do not have any Capabilities or Custom Attributes when evaluating the Filter String, but exclude the nodes that match the expression that follows the NOT EXISTS. **Tip:** When creating complex Filter Strings that include `customAttrName` and `customAttrValue` pairs as one component of an "***or***" statement, to prevent SOM from excluding nodes that have zero Custom Attributes, use **EXISTS** or **NOT EXISTS** criteria for the `customAttrName` and `customAttrValue` pair definitions. Otherwise nodes that do not have any Custom Attributes are automatically excluded even if they have values that pass other aspects of your filter. For example, when evaluating the following Filter String, SOM includes nodes with a hostname that includes **router**, followed by any number of characters, followed by **hp.com** and excludes any nodes with Custom Attribute **edge** and that edge value is **true**. `(hostname like router*.hp.com OR NOT EXISTS((customAttrName=edge AND customAttrValue=true)))`
Delete	Deletes the selected expression. **Note:** If the Boolean Operator is selected, the Additional Filters Editor deletes all expressions associated with the Boolean Operator.

Modify a Node Group

To modify a node group, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Object Groups** > **Node Groups**. The Node Groups is displayed.
2. Select the node groups that you want to modify from the table view and click **Open**. The Node Group form is displayed.
3. Make the necessary changes to the node group. You can modify any of the attributes of the node group.

   Note: You can modify the default node groups provided by SOM.

4. Click to save changes to the node group.

   The Node Groups view is refreshed to display the changes made to the node group.

Delete a Node Group

You cannot delete the following:

- The default node groups — All Elements, FC Fabrics, FC Switches, Hosts, and Storage Systems.
- Node groups that have nodes associated with them.

To delete a node group, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Object Groups** > **Node Groups**. The Node Groups view is displayed.
2. Select the node group that you want to delete, right-click, and select **Delete Node Group.** The selected node group is deleted.

Discovering Devices

The devices that comprise your Storage Area Network (SAN) must be discovered so that they can be monitored and managed by SOM. The discovery process involves specifying the IP address of the management proxies to access devices associated with the management proxy and collect data. Management proxies are CIM agents and device managers such as EMC Solutions Enabler or Hitachi HiCommand Device Manager. In some cases, the management proxy may be the device itself.

SOM can discover hosts, storage systems, and switches. An administrator must configure a device for discovery by providing the IP address. Some devices might also require the authentication credentials to access the device.

An element and its node are created automatically after a storage device is discovered. The management server collects data from the devices and stores it in a CIM, WBEM or SMI standards-based database. This enables you to view and manage your multi-vendor storage infrastructure uniformly.

The default data collection policy, is automatically triggered when a new node is created. Inventory details of the discovered device and its components are collected and displayed in the relevant element category of the Inventory workspace. For example, if the discovered device is a host, then it appears in the Hosts view in the Inventory workspace.

Recommendations for Planning Discovery

Key points to consider when you plan discovery for your storage environment:

The maximum number of addresses for which you can start discovery from the user interface at a time is 1000. To configure addresses beyond this number, use the `somdiscoveryconfigexportimport.ovpl` command.

- To configure bulk discovery, set the following two properties in the `ovjboss.jvmargs` file.
  - da.bulkDiscoveryQueueSize default: 100
  - da.bulkDiscoveryIntervalInSeconds default: 20

  The file is located at `<Install_Dir>\HP\HP BTO Software\shared\nnm\conf\props\ovjboss.jvmargs`
- Plan sequence of discovery such that you discover switches first, storage systems followed by hosts. This helps reduce time to value in realizing connectivity information.
- Use the Queue Discovery option to automate the discovery process rather than manually discover each address.
- SOM relies on a healthy database and sufficient disk space to function properly. If you include the management server address for discovery and discover the management server, SOM will monitor its own health. You can review the product health using the Product tab on the System Information page (**Help** > **System Information**).
- Each discovered node (physical or virtual) counts toward the license limit. The capacity of your license might influence your approach to discovery.

Prerequisites for Discovering a Device

You need the following to discover a storage device:

- The IP address or the FQDN of the device
- A tenant — use the default tenant if you have not created one
- Device-specific prerequisite, if any

Note: If required, provide the discovery credential for the IP address or the FQDN.

Prerequisites for Agentless Discovery of Linux, Solaris, HP-UX, and AIX Hosts

SOM uses Secure Shell (SSH) to discover the following hosts:

- Linux
- Solaris
- HP-UX
- IBM AIX

SSH uses the default port (port 22) to establish a connection between SOM and the remote host.

To ensure agentless discovery of these hosts, do the following:

- Provide the IP address or DNS name of the host.

  For a non-default port, append the port number to the IP address or DNS name. For example, `hostname.domain.com:36` or `ipaddress:36` for port 36.
- Access the host using an appropriate user account.
  - **Root user account** – Accessing a host using a root user account provides SOM with access to all the information about the host.
  - **Non-root user account** – Accessing a host using a non-root user account provides SOM with access to limited information about the host.
    Example
    - For a Linux host - information related to the serial number, manufacturer, disk drives, disk partitions, and Veritas DMP devices, is not available.
    - For a Solaris host - information related to the disk drives, disk partitions, Mpxio Multipath, HBA adapter, ports and target mappings is not available.
    - For an AIX host - the `MaxMediaSize` of disk drives and information related to the HBA adapter, ports and target mappings, is not available.
  - **Sudo user account**
    Sudo enables users (permitted by the administrator) the ability to run commands as a root user. Using a non-root user account, information about disk drives, disk partitions, volume management, multipath path information, and the serial number of a host is not available to SOM. Discovery and data collection of a host is based upon the privileges of the sudo user account as configured in the `/etc/sudoers` configuration file. For more information about creating and configuring a sudo user, see "Creating and Configuring a Sudo User" on page 523.

    > **Note:** If a host is discovered using a sudo user account and if that sudo user account is later deleted from the `sudoers` file, the existing information related to the disk drives, disk partitions,

> volume management, multipath path, and the serial number of the host is lost.

> **Note:** If an HP-UX host is discovered using a sudo user account, you cannot delete the sudo user account from the host console because the SSH session remains open. This issue is specific to the openssl version installed on the HP-UX host.

- Configure SSH on the host.

For a Linux Host

- Ensure that at least one of the following is true:
  - The `lsb` package is available on the Linux host.
  - The `/etc/issue` file on the host is not modified manually.

  SOM runs the `lsb_release -d` command to identify if the discovered host is a Linux host. The output of the command also identifies the distribution of the Linux system, that is whether the host runs on a Redhat or a SUSE distribution of Linux. If the `lsb_release -d` command is not available on the discovered host, the management server fails to identify the type of the host. In this case, SOM uses the `/etc/issue` file to identify the discovered host. However, it can use this file only if it is not modified manually.

  > **Note:** If at least one condition mentioned above is not satisfied, SOM fails to discover the Linux host.

- Install the rpm, `sysfsutils` to ensure agentless data collection.

  SOM collects HBA details using the commands in the `sysfsutils` package. Therefore, in the absence of HBA Port and target mapping information, the LUNs presented to the host are shown as local disks.

> **Note**: During agentless discovery of a Linux host, SOM uses hostname to uniquely identify a host. If SOM discovers hosts which have a default hostname, that is `localhost.localdomain` or `localhost`, SOM displays the IP address as the name of the host.

Commands for Data Collection

SOM runs a set of commands to collect data from a host based on the access rights of the user account. You can also log on to a host, with the appropriate user account, and run the commands at the command line interface to get the required information.

See the following for the set of commands:

- "Commands for a Linux Host as a Sudo User" on the next page
- As a Root User

  - "Commands for a Linux Host as a Root User" on the next page
  - "Commands for a Solaris Host as a Root User" on page 128
  - "Commands for an AIX Host as a Root User" on page 131
  - "Commands for an HP-UX Host as a Root User" on page 130
- As a Non-Root User
  - "Commands for a Linux Host as a Non-Root User" on page 131
  - "Commands for a Solaris Host as a Non-Root User" on page 133
  - "Commands for an AIX Host as a Non-Root User" on page 134

Note: If data collection fails with an internal error or provider time-out error, check the privileges of the user account or any commands running for a long time. Commands may sometimes stop responding and not collect any data.

Commands for a Linux Host as a Sudo User

Use the following commands at the command line interface to collect data from a Linux host based upon the privileges of the sudo user account. Sudo may be configured to require the root password or no password at all.

Note: While running the following commands, provide the absolute path of a command. For example: /usr/sbin/fdisk –l

Command	Description
`dmidecode`	Determines the serial number and name of the hardware manufacturer.
`fdisk`	Collects information about the disks, disks partitions, and capacity details of the Device Mapper partitions.
`vgdisplay -v`	Provides the details of all the volume groups.
`vgdisplay --version`	Provides the version of LVM on the host.
`lvdisplay -vm`	Provides the LVM extent details of the host.

Command	Description
`/usr/sbin/vxprint`	Provides information on Veritas Volume Manager's disk groups and their associations.
`/usr/sbin/vxdg free`	Provides information on Veritas Volume Manager disk information and also determines available space in the disk group.
`/usr/sbin/vxdisk -q list <diskname>`	Provides the details of the disk controlled by the Veritas Volume Manager.
`/usr/sbin/vxdisk -q list | cut -f1 -d`	Collects information on Veritas Volume Manager's disks and sub-path information.
`/sbin/dmsetup info`	Provides the Device Mapper partition details.
`/sbin/dmsetup ls`	Provides the Device Mapper device and partition details.
`/sbin/dmsetup --version`	Determines the Device Mapper version and multipath device details.
`/sbin/multipath -ll`	Provides multipath disk details.
`udevinfo`	Collects SCSI information on all other supported distributions of Linux except SUSE 11.
`udevadm info -a`	Collects SCSI information about SUSE Linux.

Commands for a Linux Host as a Root User

Use the following commands at the command line interface to collect data from a Linux host with the root user account:

Command	Description
`dmidecode -t system`	Determines the serial number and name of the hardware manufacturer.
`fdisk -l <diskname>`	Collects information about the disks, disks partitions, and capacity details of the Device Mapper partitions. **Note:** If a disk is not enabled on the host, then this disk information will not listed in the `fdisk`

Command	Description
	`-l <diskname>` output.
`udevadm info -a`	Collects SCSI information about SUSE Linux.
`/usr/sbin/vxprint`	Provides information on Veritas Volume Manager's disk groups and their associations.
`/usr/sbin/vxdg free`	Provides information on Veritas Volume Manager disk information and also determines available space in the disk group.
`/usr/sbin/vxdisk -q list | cut -f1 -d`	Collects information on Veritas Volume Manager's disks and sub-path information.
`vgdisplay --version`	Provides the version of LVM on the host.
`vgdisplay -v`	Provides the details of all the volume groups.
`lvdisplay -vm`	Provides the LVM extent details of the host.
`vgcfgbackup -f`	Provides the mirror volume extent details of the host.
`/sbin/dmsetup --version`	Determines the Device Mapper version and multipath device details.
`/sbin/dmsetup ls`	Provides the Device Mapper device and partition details.
`/sbin/multipath -ll`	Provides multipath disk details.
`/sbin/dmsetup info`	Provides the Device Mapper partition details.
`/usr/sbin/vxdisk -q list <diskname>`	Provides the details of the disk controlled by the Veritas Volume Manager.

Commands for a Solaris Host as a Root User

Use the following commands at the command line interface to collect data from a Solaris host with the root user account:

Command	Description
`uname -t system`	Verifies if it is the Solaris operating system.

Command	Description
`uname –n`	Provides the node name or system name.
`uname -X`	Provides the node name, machine type, number of processors, and the OS version.
`uname -i`	Provides the machine type.
`prtconf`	Collects the RAM or physical memory size.
`ifconfig –a`	Provides the machine IP address that is used only if the IP resolution fails.
`kstat –p cpu_info`	Collects the number of processors and the processor type.
`df -k`	Provides file system capacity details.
`df -an`	Provides the file system type.
`/usr/sbin/zfs list -H -t filesystem -o name, used, avail, mountpoint, recordsize, compression`	Provides details about the zfs filesystem.
`/usr/sbin/fcinfo hba-port`	Collects HBA and HBA port information.
`/usr/sbin/fcinfo remote-ports –sp <portwwn>`	Collects HBA target mapping information.
`echo | format`	Collects information about the disks, disks partitions, and capacity details. The Echo command is used to ensure that command output is used without modifications.
`/usr/sbin/metastat`	Provides information on Solaris native volume manager disks and their associations.
`/usr/sbin/metaset`	Provides information on Solaris native volume manager disk sets.
`pkginfo -l SUNWmdu`	Determines the Solaris native volume manager version.

`cat /etc/driver/drv/fp.conf | grep "mpxio-disable"`
or
`cat /kernel/drv/scsi_vhci.conf | grep`
	Determines if Solaris native Mpxio is enabled.

Command	Description
"mpxio-disable"	
/usr/sbin/luxadm probe	Provides the Native Mpxio Multipath device names.
/usr/sbin/luxadm display <rdisk>	Provides multipath disk details.
cat /kernel/drv/scsi_vhci.conf | grep "load-balance"	Determines Mpxio multipath type/algoritm.
pkginfo –l SUNWcsu or pkginfo –l SUNWmdi	Provides native Mpxio multipath version.

Commands for an HP-UX Host as a Root User

Use the following commands at the command line interface to collect data from a HP-UX host with the root user account:

Command	Description
uname -s	Verifies if it is the HP-UX operating system.
Hostname	Provides the node name or system name.
uname -rv	Provides the node name, machine type, number of processors, and the OS version.
model	Provides the machine type.
getconf MACHINE_SERIAL	Provides the machine serial number.
machinfo	Collects the RAM or physical memory size.
grep "Physical:" /var/adm/syslog/syslog.log	Collects the RAM or physical memory size. **Note:** This command is run only if the `machinfo` command is not available.
ifconfig –a	Provides the machine IP address that is used only if the IP resolution fails.
ioscan –fnC processor	Collects the number of processors and the processor type.
"ioscan –kfnc" (Versions older than 11.31) and "ioscan –kfNnc" (11.31)	Collects information about the disks, and disks partitions. (read-only command, do not modify).
ioscan –m hwpath	Collects the hardware path for agile disks (11.131

Command	Description
	format).
"diskinfo –v" and "idisk –p"	Collects block size, bytes per sector, number of blocks for disks and disk partitions.
df -p	Provides file system capacity details.
df -n	Provides the file system type.
fstyp –v </dev/…partition>	Provides file system, block size, and number of blocks.
/tmp/hbatest -v	Collects HBA adapter, ports and target mapping information.
scsimgr -p lun_map -D </dev/rdiskxx>	Collects NativeMultipath Disk Lunmap.
"vgdisplay", "pvdisplay", and "lvdisplay"	Provides information about HP-UX native Volume Manager, disks, volumes and their associations.

Commands for an AIX Host as a Root User

Use the following commands at the command line interface to collect data from an AIX host with the root user account:

> **Note**: `hbatest` is an SOM provided binary, copied over the SSH channel. It is deleted after the operation is complete.

Command	Description
bootinfo –s <disk-name>	Collects the MaxMediaSize of the disk drive.
hbatest	Collects the HBA and HBA port information.

Commands for a Linux Host as a Non-Root User

> **Note:** For Redhat Linux version 7 and onwards, use `cat /etc/redhat-release` command to identify the Linux distribution on the discovered host when the `lsb_release -d` command fails.

Use the following commands at the command line interface to collect data from a Linux host with the non-root user account:

Commands	Description
`uname -nsrm`	Identifies if the discovered host is a Linux host. Also, provides information related to the node name, kernel release, and model details of the discovered hosts.
`lsb_release -d`	Identifies the Linux distribution on the host.
`cat /etc/issue`	Identifies the Linux distribution on the discovered host from the /etc/issue file, in case the `lsb_ release -d` command fails.
`cat /etc/redhat-release`	Identifies the Linux distribution on the discovered host for RHEL version 7 and onwards, in case the `lsb_release -d` command fails.
`ps -aef | grep "com.appiq.cxws.main.LinuxMain" | grep -v "grep"`	Identifies if the CIM Extension is running on the host.
`rpm -q APPQcime`	Identifies if the CIM Extension is installed on the host.
`cat /proc/meminfo`	Collects memory information about the host.
`cat /proc/cpuinfo`	Collects information about host processor count.
`cat /proc/partitions`	Determines information about the disks and disk partitions of the host. The output of this command is used by the `fdisk -l` command.
`udevinfo -a -p`	Collects SCSI information about Redhat Linux.
`ls -l`	Determines permission and ownership details. Also, includes the permission details for the LXM volumes.
`rpm -qa VRTSvxvm-common`	Identifies if Veritas Volume Manager is installed on the host.
`/usr/sbin/vxprint -lr`	Provides information on the Veritas Volume Manager's sub-disk details.
`/usr/bin/systool -c fc_host -v`	Collects information related to HBA.
`/usr/bin/systool -c scsi_host -v`	Collects information related to HBA ports.

Commands	Description
/usr/bin/systool -c fc_remote_ports -v	Provides the target port information.
/usr/bin/systool -c scsi_disk -v	Provides detailed information of the LUNs presented to the host.
df -PT	Provides file system details of the host.
/bin/df	Collects information related to Device Mapper disks mounted on the File Systems.
cat /proc/scsi/scsi	Used for collecting SCSI information.

Commands for a Solaris Host as a Non-Root User

Use the following commands at the command line interface to collect data from a Solaris host with the non-root user account:

Command	Description
uname -t system	Verifies if it is the Solaris operating system.
uname -n	Provides the node name or system name.
uname -X	Provides the node name, machine type, number of processors, and the OS version.
uname -i	Provides the machine type.
prtconf	Collects the RAM or physical memory size.
ifconfig -a	Provides the machine IP address that is used only if the IP resolution fails.
kstat -p cpu_info	Collects the number of processors and the processor type.
df -k	Provides file system capacity details.
df -an	Provides the file system type.
/usr/sbin/zfs list -H -t filesystem -o name, used, avail, mountpoint, recordsize, compression	Provides details about the zfs filesystem.
/usr/sbin/metastat	Provides information on Solaris native volume manager disks and their associations.

Command	Description
`/usr/sbin/metaset`	Provides information on Solaris native volume manager disk sets.
`pkginfo -l SUNWmdu`	Determines the Solaris native volume manager version.

Commands for an AIX Host as a Non-Root User

Use the following commands at the command line interface to collect data from an AIX host with the non-root user account:

Command	Description
`uname -s`	Verifies the AIX OS.
`hbatest`	Provides the node name or system name.
`uname -rv`	Provides the OS version.
`lsconf | grep 'Machine Serial'`	Provides the serial number.
`uname -M`	Provides the machine type/model.
`lsattr -El sys0 -a realmem`	Collects the RAM or physical memory size.
`ifconfig -a`	Provides the machine IP address that is used if the IP resolution fails.
`odmget -q"PdDvLn LIKE processor/*" CuDv`	Collects the number of processors and the processor type.
`df -tMk`	Provides file system capacity details.
`lsdev -Cc disk`	Lists all AIX disk drives.
`lscfg -l <disk-name>`	Provides disk drive information.
odmget –q "attribute=unique_id" CuAt odmget –q "attribute=lun_id" CuAt odmget –q "attribute=scsi_id" CuAt odmget –q "attribute=wwn" CuAt	Collects the UniqueID, OS LunID, ScsiID, and WWN attributes of a disk from the CuAt ODM object.
`lsvg -o`	Lists the active AIX native Logical Volume Manager (LVM) volume groups.

Command	Description
`lspv`	Lists all the AIX physical volumes.
`lsvg -l <volume-group>`	Lists LVM physical volumes in a specified volume group.
`lspv -l <physical-volume>`	Lists the physical and Logical Extents (LV) in a specified physical volume. This is modeled as a disk partition (Volume manager partition).
`lqueryvg -sp <physical-volume>`	Provides the size of an extent (in powers of 2) for a specified physical volume.
`odmget -q "PdDvLn=logical_ volume/lvsubclass/lvtype" CuDv`	Lists the LVM logical volumes of the CuDv ODM object.
`lslv <logical-volume>`	Provides details of the specified LVM volume.
`lslpp -l <package-name>`	Provides the version of the package used against supported MPIO ODM packages.
`lspath -F name:parent:connection:path_status_ status`	Provides the native MPIO Multipath device names, parent scsi/fscsi device, connection (WWN), path and MPIO status.
`lsattr -F "attribute=value" -El <mpio-device> -a reserve_policy, algorithm`	Provides the MPIO reserve policy and load balancing algorithm for the specified MPIO device.

Agentless Discovery of Windows Hosts

SOM uses the Windows Management Instrumentation (WMI) service to discover remote Windows hosts without the CIM extension. WMI uses the default port 135 to establish a connection between the Windows host and SOM. The operating system on the SOM management server can be either Linux or Microsoft Windows.

Prerequisites

On the SOM management server

- Provide the IP address or DNS name of the host.
- Provide a user account with administrator privileges.
- Use the Discovery Hint "Windows - WMI (Agentless)" (Configuration > Discovery > Discovery Addresses).
  If the Discovery Hint is not specified, the `somwindowsAgentlessDiscovery.ovpl` command still attempts to copy the files configured in the script.

> **Note:** This is a one-time configuration. Subsequent discoveries happen automatically.

- On Windows SOM
  Download and install the binary `psexec` from Microsoft's Windows SystInternals website:
  http://live.sysinternals.com
  `psexec` needs the following setup:
  - TCP ports 135, and 445 must be open.
  - The `Admin$` and `IPC$` shares must be enabled.
  - The environment variable `PATH` must include the `psexec` installation path.
    For example, if `psexec.exe` is located at `C:\Temp`, type the following command:
    `set PATH=%PATH%;C:\Temp`
- On Linux SOM
  - TCP port 139 must be open.
  - Ensure that you can mount the remote admin share, `admin$/Temp` using a CIFS mount. For example,
    `mount -t cifs -o username=<UserName>,password='<password>' //<remote-host-name>/admin$/Temp /mnt/<mount-point>`
- Run the `somwindowsAgentlessDiscovery.ovpl [-i]` command.
  On Windows SOM
  [*drive*:]`\Program Files (x86)\HP\HP BTO Software\bin\somwindowsAgentlessDiscovery.ovpl –i`
  On Linux SOM
  `/opt/OV/bin/somwindowsAgentlessDiscovery.ovpl –i`
  The option "`-i`" ensures that subsequent discoveries run automatically after the script is completed.
  If there are more than 150 hosts to be discovered, it is recommended that you do not use the "`-i`" option. Instead, select hosts (less than 50 at a time) manually, and trigger discovery in time intervals of 10 minutes.
  For more information about the `somwindowsAgentlessDiscovery.ovpl` command, see the *HPE Storage Operations Manager CLI Reference Page*.

On the remote Windows host

- Port 135 must be open.
- Enable the WMI service.
- Enable port 135 through the firewall (**Start** > **Administrative Tools** > **Windows Firewall with Advanced Security**)
- If the host has the HPMPIO multipathing software, specify the HP MPIO binaries in the system path to collect multipathing information. The path variable must be set to point to the HPMPIO DSM install location,
  [*drive*]:`\Program Files (x86)\Hewlett-Packard\HP MPIO DSM\P6000 DSM\AMD64`

You can also log on to a Windows host using the administrator account, and run the commands at the command line interface to get the required information. For more information about the set of commands, see "Commands as an Administrator" below.

Commands as an Administrator

Use the following commands at the command line interface to collect data from a Windows host as an administrator:

Commands	Description
`vxdisk list`	Provides information about the disks used by Veritas DMP on a managed server or on a specified disk group.
`vxdisk diskinfo`	Provides disk information for a Veritas DMP device.
`vxvol -v volinfo`	Provides volume information of a storage volume for Veritas DMP device.
`vxdmpadm pathinfo`	Provides information on path details, path status, load balance policy, port, target, and LUN numbers for a multipathing device.
`hpdsm devices`	Provides multipath device details related to HP MPIO device.
`hpdsm paths device=<APPIQ_Win32HPMPIOAll.NUMBER>`	Provides detailed information about HP MPIO paths to the specified device.
`mpclaim -s -d`	Provides multipath device details related to Microsoft MPIO device.
`mpclaim -s -d <MicrosoftMPIOMultipathing.NUMBER>`	Provides detailed information about Microsoft MPIO paths to the specified device.
`reg query <path_to_the_registry_key> /v DisplayVersion`	Provides version information for HP MPIO and Veritas DMP.

Prerequisites to Discover Hosts with CIME Agent

SOM can use a CIM extension installed on a remote host to collect detailed information about the host. The remote host can be any of the following:

- Windows Host
- Linux host

- HP-UX
- Solaris
- IBM AIX

To discover a host with a CIM agent, follow these steps:

1. Install the CIM extension on the host.
2. Provide the IP address or DNS name of the host.
3. Provide the authentication credentials (user name and password) of the host.

   If you change the password of a host after you discover it, stop and restart the CIM extension running on the host, and change the host password in the discovery list. You must rediscover the host.

Note: A host discovered using a CIM extension, cannot be subsequently discovered using agentless discovery. To discover such a host using agentless discovery, delete the host from SOM and re-discover.

The `AppStorWin32Agent` service is automatically enabled when the CIME agent is installed. If you specify a discovery hint (Configuration > Discovery > IP Addresses) for a host, the discovery results are as mentioned in the following table:

Service	Discovery Hint	Discovery Result
Enabled	CIM Extension	Pass, discovered with CIM agent
Enabled	Windows Agentless **Note**: If the registry key is modified to enable the agentless mode, discovery will pass with the Windows Agentless hint.	Fail
Enabled	No Hint	Discovered with CIM agent
Disabled	No Hint	Discovered as agentless
Disabled	CIM Extension	Fail
Disabled	Agentless	Pass, Discovered as agentless
Enabled	Agentless	Pass, Discovered as agentless

Prerequisites to Discover Host Clusters

A host cluster is automatically discovered if its member nodes are discovered. However, a CIM extension must be installed on the member node used for discovery.

The following cluster services support automatic discovery:

- HP Serviceguard Cluster on HP-UX
- Microsoft Cluster Services (MSCS) on Windows
- VMware Clusters (through vCenter)

Prerequisites to Discover IBM Virtual IO

To discover an IBM Virtual IO host and obtain the underlying connectivity, you must discover the following components:

- IBM VIO Server
- IBM VIO Client
- IBM HMC

Note: There is no particular order to discover the above three components. Once all the three components are discovered, the SOM correlates the connectivity among them.

To discover IBM VIO Server and IBM VIO Client with CIME agent, follow the procedure below:

- Intsall the CIM extension on the host.
- Provide the IP address or DNS name of the host.
- Provide the authentication credentials (username and password).
- Select **IBM AIX CIME** discovery hint.

Note: IBM VIO Server and Client are discovered using the same CIM extension as is necessary to discover any IBM AIX host.

SOM uses Secure Shell (SSH) to discover an IBM HMC. To discover IBM HMC without CIME agent, follow the procedure below:

- Provide the IP address or DNS name of the host.
- Provide the authentication credentials (username and password) of the host.
- Select **IBM HMC** from the **Discovery Hint** list.

Note: IBM HMC is not represented in the system topology by itself, it is discovered to get relevant information for connectivity between IBM VIO Server and client.

Prerequisites to Discover VMware ESX Servers and Virtual Machines

VMware® ESX® Servers and Virtual machines can be discovered through a VMware® vCenter™ or through individual ESX Servers.

A Virtual Machine can also be discovered as an individual host (either agentless or with a CIM agent).

To discover via a vCentre:

1. Install and run VMTools on each virtual machine.

   If VMTools is not running, the virtual machine is unmanaged and only limited data is available. For example, you cannot view the element topology of the associated discovered host for an unmanaged virtual machine.

   > **Note**: SOM checks the status of the `VMTools` property in the **Properties** pane of a virtual machine. If `VMTools=GuestToolsRunning`, then VMTools is running on the virtual machine.

2. Provide the user name and password of the vCenter account that can view or access the ESX Servers or virtual machines to be discovered.
   The vCenter account must have "Datastore Browse" privileges.
3. Discover the vCenter.

   > **Notes:**
   > - All ESX Servers and virtual machines that the vCenter account can access are discovered automatically.
   >   Use the custom property `discovery.exclude.vmware.vm=true`, in the `%OvDataDir%\Conf\som\custom.properties` file to disable the automatic discovery.
   > - You can discover a DRS cluster and its details only via a vCenter.
   > - For the reconciliation of a VM either with a CIM agent or as an agentless host, VMTools must be running on the VM while it is being discovered through a vCenter or ESX server.

To discover via an ESX Server:

1. Install and run VMTools on each virtual machine.

   If VMTools is not running, the virtual machine is unmanaged and only limited data is available. For example, you cannot view the element topology of the associated discovered host for an unmanaged virtual machine.

   > **Note**: SOM checks the status of the `VMTools` property in the **Properties** pane of a virtual machine. If `VMTools=GuestToolsRunning`, then VMTools is running on the virtual machine.

2. Provide the user name and password of an ESX server.
3. Discover the ESX Server.

> **Notes:**
> - All VMs that are hosted on an ESX server are discovered automatically.
>   Use the custom property `discovery.exclude.vmware.vm=true`, in the `%OvDataDir%\Conf\som\custom.properties` file to disable the automatic discovery.
> - For the reconciliation of a VM either with a CIM agent or as an agentless host, VMTools must be running on the VM while it is being discovered through a VC or ESX server.

Prerequisites to Discover Brocade Switches

SOM discovers Brocade switches through a Brocade Network Advisor (BNA) server.

Specify the following to discover Brocade Switches:

- The IP address or DNS name of the BNA server (Configuration > Discovery > Discovery Addresses)
- Authentication credentials of the BNA server

Brocade switches that operate in the Access Gateway mode are also discovered using the Brocade Network Advisor (BNA) server. This default behavior is enabled in the `custom.properties.sample` file. If required, disable automatic discovery using the `brocade.discoverAccessGateway` property in the `custom.properties` file.

Note the following, before you discover a Brocade Access Gateway:

- If a Brocade Access Gateway is connected to multiple edge switches, the Access Gateway and all other connected edge switches must be managed by the same BNA server.
- If a BNA server that is managing both Brocade Access Gateways and HP Virtual connect switches is discovered, only the Brocade Access Gateways are discovered.

Discovery of a Brocade switch results in the discovery of the Brocade fabric that contains the switch. For details about a Brocade Fabric and its related components, see the "Fabrics View" on page 294.

> **Note:**
> - Access Gateways are shown in SOM as N-Port ID virtualizers and are visible in the Physical Switches inventory view.
> - If the switch dependency details are not displayed correctly, re-run data collection.

Prerequisites to Discover Cisco Switches

SOM discovers top level Cisco physical switches using SNMPv2 or SNMPv3.

SNMPv2

SOM uses SNMPv2 as the default method to discover Cisco switches that have the community string enabled on the switch.

> **Note:** On Cisco switches, the default configuration does not enable SNMP community strings.

1. Type the following commands to set the community string on a Cisco Switch:
   a. `cisco_switch# show snmp`

      To display the Cisco SNMP configuration settings.
   b. `cisco_switch# config t`

      To enter the configuration mode.
   c. `cisco_switch(config)# snmp-server community public ro`

      To enable the read only community string.
   d. `cisco_switch(config)# exit`

      To exit the configuration mode.
   e. `cisco_switch# copy run start`

      To save.
2. Specify the IP address (Configuration > Discovery > Discovery Addresses) to discover the switch. You do not need to provide a password.

SNMPv3

SOM uses SNMPv3 to discover Cisco switches that support the following:

- Authentication: **MD5** or **SHA**
- Encryption: **DES**, **AES**, or **None**

To enable the discovery of Cisco switches using SNMPv3:

1. **Modify the Custom Properties File**
   a. Create and edit the `custom.properties` file using the `custom.properties.sample` file in the following location:
      - *Windows*: `%OvDataDir%\conf\som`
      - *Linux*: `/var/opt/OV/conf/som`
   b. Set the following properties in the `custom.properties` file:

- `cisco.useSNMPv3=true`
- `cisco.snmp.authenticationProtocol=MD5` (Message Digest 5) If the switches use the Secure Hash Algorithm -1 (SHA), replace `MD5` with `SHA`.
- `cisco.snmp.privacyProtocol=DES`
  If the switches use a privacy protocol other than DES, replace `DES` with `AES` or `None`.

c. Restart the somjboss service.

2. **Create an Account on a Switch**

   To discover a switch, you must create an account on a Cisco switch using one of the following:

   - **Command Line Interface (CLI)**

     Use the following commands at the command prompt:

     i. To enter the configuration mode
     `cisco_switch# config t`

     ii. To create a user account
     `cisco_switch(config)# username <user> password <password>`

     In this instance *<user>* is the user name of the new account and *<password>* is the password for the corresponding account.

   - **Cisco Device Manager**

     To create an account on one switch.

     > **Note**: All Cisco switches with the same credentials are discovered in a fabric. For switches with different credentials, repeat the discovery process for each switch.

   - **Cisco Fabric Manager**

     To create an account on all the switches in a fabric with the same credentials and security settings.

     > **Note:** If the switch dependency details are not displayed correctly, re-run data collection.

Customize Switch Properties

For enhanced functionality, you can customize Cisco switch properties in the `custom.properties` file. These properties may also be copied from the `custom.properties.sample` file.

The `custom.properties` file exists in the following location:

- *Windows*: `%OvDataDir%\Conf\som`
- *Linux*: `/var/opt/OV/conf/som`

Description	Property
To set the time-out in milliseconds for discovery.	`cisco.snmp.timeout=15000`
To set the number of retries for discovery	`cisco.snmp.retries=2`
To optimize and ensure discovery of all switches in the environment.	`discovery.snmp.timeout=10000`
To set the number of retries for discovery	`discovery.snmp.retries=3`
To enable the discovery of ports that do not have a connector.	`cisco.showPortsWithNoConnector=true`
To enable the discovery of non-SAN ports.	`cisco.showNonSanPorts=true`
To enable the discovery of port-channel ports.	`cisco.showPortChannelPorts=true`
To exclude remote switches from discovery.	`cisco.snmp.switch.exclude=switchWWN`
To discover disabled VSANs. By default, disabled VSANs are not discovered.	`cisco.showDisabledVsans=true`

Cisco VSANs and Fabrics

A Cisco VSAN is a collection of FC switch ports from a set of connected FC switches that comprise a fabric. The ports of a switch can be members of multiple VSANs. Likewise, ports of multiple switches can be grouped to form a single Cisco VSAN. The fabrics and VSANs that a switch belongs to are automatically discovered when a switch is discovered. For details about a Cisco Fabric and its related components, see the "Fabrics View" on page 294.

Prerequisites to Discover HPE XP/P9500 Arrays

SOM discovers HPE XP arrays through the service processor (SVP) on the array using the RMI-API.

To discover HPE XP arrays:

1. **Create a user account to access the SVP**
   The user account must have the "View Only" privilege. The authentication credentials can be defined by the user.
2. **Provide access to the SVP**
   Specify the IP address (Configuration > Discovery > Discovery Addresses) of the SVP.

By default, the 1099, 51100, and 51099 ports are used.

The discovery hint (Configuration > Discovery > Discovery Addresses), **HDS/XP –Native API**, applies to both HPE XP/P9500 and HDS arrays.

Discovery or data collection might fail for XP 24000 or P9500 arrays due to the following reasons:

- The limit of 32 simultaneous open connections has been crossed for a XP 24000 array.
- Multiple users are simultaneously accessing a P9500 RMI-server/Web-console through the SVP on the array at a given point in time.

  In such instances retry discovery or data collection after a while.

Prerequisites to Discover HPE 3PAR Arrays

SOM uses the 3PAR SMI-S server to discover a 3PAR array.

To discover a 3PAR array:

1. **Start the 3PAR SMI-S server**
   By default, the SMI-S server is not started on the array. To start the SMI-S server, follow this step:
   - Open the InForm CLI interface on the array and run the following command:

     ```
     startcim
     ```

     This command starts the SMI-S server.

2. **Provide access to the HPE 3PAR array**
   Provide the discovery address or DNS of the HPE 3PAR array along with the user name and password. To discover and collect data from an HPE 3PAR, ensure that the user account has either the Super or Browse role. The audit role does not authorize data collection.
3. **(Optional) Enable Secure Shell (SSH)**
   For the collection of deduplication and Adaptive Optimization (AO) data, enable SSH on the HPE 3PAR array. SOM collects the following deduplication properties:
   - Dedup
   - Compaction
   - Provisioning Type

   Note: To collect deduplication and AO data, the user account must have either the Super or Browse role (not the audit role).

   SOM runs a set of commands to collect deduplication and AO data from HPE 3PAR based on the access rights of the user account. For information about commands to collect deduplication and AO data, see "Data Collection Commands Used for HPE 3PAR Arrays" below.

Data Collection Commands Used for HPE 3PAR Arrays

SOM uses the following commands to collect deduplication and AO data from HPE 3PAR storage systems:

Note: To collect deduplication and AO data, ensure that the user account has either the Super or Browse role. The audit role does not authorize data collection.

Commands	Description
`showsys -space`	Collects information including dedup and compaction on the storage system.
`showvv -space`	Collects information including dedup, compaction, and provisioning type on the virtual volumes.
`showcpg -space`	Collects information including dedup and compaction on the storage pool.
`showaocfg`	Collects information about adaptive optimization configuration in the storage system.

Prerequisites to Discovering IBM SAN Volume Controller or IBM Storwize V7000 Arrays

SOM uses the SMI-S server to discover the IBM SAN Volume Controller or IBM Storwize V7000 arrays.

To discover an IBM SVC or IBM V7000 array:

- Provide the IP address and authentication details of the device.
- Use the Discovery Hint "IBM SVC- SMIS" (**Configuration** > **Discovery** > **Discovery Addresses**).

Prerequisites to Discover HP StorageWorks EVA Arrays

SOM discovers HP StorageWorks Enterprise Virtual Arrays (EVA) arrays using the default TCP port number 5989 (CIM XML transaction over HTTPS) of the Command View (CV) proxy server and its SMI-S provider over a SSL fiber channel connection.

Prerequisites

- HP StorageWorks CV EVA must be installed on a server that is not running SOM before you discover an HP EVA storage system.
- The IP address, user name, and password of the active Command View EVA server that manages the EVA system.
- If the active and standby CV EVA proxy machines exist, both the proxies must be discovered.
  SOM does not discover the EVA if only the CV EVA server that is passively managing the array is discovered. If the passive CV EVA server does not have active management of any EVAs at the time discovery is run, no EVA is listed for the discovered passive CV EVA server. If at some point in time an

EVA becomes managed by the passive CV EVA server, you must start discovery and data collection of both the active and standby CV EVA proxies.

Note: If storage arrays discovered through a proxy server are moved to another proxy server, you must discover the new proxy server and either re-run data collection manually on the storage arrays or wait for the next data collection cycle, configured as per the data collection policy.

Prerequisites to Discover Hitachi HUS and AMS Arrays

SOM uses the Hitachi HiCommand Device Manager (HDvM) and the built-in HDS provider to discover and collect data from an HUS or AMS array. The HiCommand Device Manager must be installed on a server (proxy host). The proxy host can be used to discover multiple arrays. This proxy host can run Windows, Linux, or the HP-UX operating system.

To discover anHitachi HUS or AMS array:

1. **Provide access to the HiCommand Device Manager**:
   Specify the IP address (**Configuration** > **Discovery** > **Discovery Addresses**), user name and password of the server running the HiCommand Device Manager. Do not point to the disk array. The default authentication credentials of the HDvM are system/password.
2. **Open port 2001**:
   SOM accesses the port that the HiCommand Device Manager listens to. By default, the HiCommand Device Manager listens on port 2001, and the management server assumes this configuration during discovery. If the HiCommand Device Manager uses a different port, specify the other port number separated by a colon in the IP Address/DNS Name box (**Configuration** > **Discovery** > **Discovery Addresses** > **New**).

   The HiCommand Device Manager can also listen to other default ports, based on its version. Hence ensure that the following ports are also open:

   - 1099
   - 51099
   - 51100

   While scanning an IP address range, the management server discovers only those instances of the HiCommand Device Manager that are configured for default ports.

The SOM management server communicates with the HiCommand Device Manager through both non-secure and secure connection. To enable a secure connection using a secure port (default secure port is 2443), perform the following steps:

1. Stop the SOM JBosss application server.
2. Set the following property to true in the `%OvDataDir%\Conf\som\custom.properties` file:

```
provider.hds.useSecureConnection=true
```

3. Start the SOM JBosss application server.

   The built-in HDS provider uses secure connection to discover the HDS arrays.

For more information about discovering HDS arrays, see " Discovery Tasks" on page 155.

Prerequisites to Discover an EMC Isilon Clusters

To discover and collect data from EMC Isilon devices, SOM uses SSH to connect to any node within the cluster.

Prerequisites to discover EMC Isilon devices

- Ensure that the SSH service is enabled on the node that is used to discover the cluster.
- Do one of the following to provide access details of the Isilon cluster:
  - Either specify the IP address or DNS name of any node in the cluster.
    SOM communicates with an EMC Isilon cluster using the default SSH port number 22 configured on the node. If the node is configured for a port other than the default port, enter the port number separated by a colon along with the IP address or DNS name of the node in the IP Address/DNS Name box (**Configuration** > **Discovery** > **Discovery Addresses**).

    > **Note**: Irrespective of the number of nodes in an Isilon cluster, you can enter the IP Address or DNS Name of any one node to discover all the nodes in the cluster.

  - Specify the EMC Isilon SmartConnect zone name instead of the IP of a node within the cluster.
- Provide User Credentials
  SOM can discover an Isilon cluster by using either the root user account or a sudo user account. The information obtained using either of the accounts is the same.

  > **Note:** Ensure you delete all the existing access points before discovering the EMC Isilon cluster using a non-root user.

Prerequisites to Discover EMC VNX Filer

SOM uses the EMC® VNX™ Series XML API interface to remotely manage and monitor an EMC VNX Filer using HTTPS.

To discover a VNX Filer storage system:

1. **Specify the IP address or DNS name**
   SOM communicates with the storage system using the default SSL port number 443 configured on the device. If a port other than the default port is configured on the device, enter the port number separated

by a colon in the IP Address/DNS Name box (Configuration > Discovery > Discovery Addresses) along with the IP address or DNS name of the Control Station of the device.

2. **Provide User Credentials**
   Specify a device user that belongs to the nasadmin group, with the "XML API v2 allowed" client access role. The storage system has a default user id **nasadmin** with password **nasadmin** that can be used to discover the device.

Prerequisites to Discover EMC Symmetrix Arrays

To discover and collect data from EMC Symmetrix (DMAX/VMAX) arrays, SOM uses the SMI-S provider on a proxy server.

To discover the array:

- Ensure that the SMI-S provider is running on the proxy server.
- Provide the IP address or FQDN (Configuration > Discovery > Discovery Addresses) and the authentication credentials of the proxy server. The default credentials are: admin/#1Password
- Ensure that the default ports for HTTP (port 5988) and HTTPS (port 5989) are open.

Before you discover EMC VMAX array using Solution Enabler 8.x, ensure that you delete the existing access point discovered using SMI-S provider 4.x.

Ensure you discover the fabrics before you discover EMC VMAX and the backend storage arrays.

> **Note:** If storage arrays discovered through a proxy server are moved to another proxy server, you must discover the new proxy server and either re-run data collection manually on the storage arrays or wait for the next data collection cycle, configured as per the data collection policy.

Prerequisites to Discover HPE StoreEasy Storage

SOM uses the Windows Management Instrumentation (WMI) service to discover supported HPE StoreEasy Storage systems.

Prerequisites to discover HPE StoreEasy Storage:

- Provide the IP address or the DNS name and authentication details of the cluster.
- Enable WMI service on SOM and StoreEasy.
- Enable port 135 on SOM and StoreEasy. In StoreEasy, enable port 135 through the firewall (**Start** > **Administrative Tools** > **Windows Firewall with Advanced Security**).
- You must register the SMI-S provider on the StoreEasy system and update the WMI classes before you start data collection. If you do not register the SMI-S provider, the backend storage system dependency is not shown.

Command	Description
`Register-SmisProvider –ConnectionUri http://<3PAR StoreServ systemip address>:<port>`	Registers the SMI-S provider

`Get-StorageProvider`
The following result is displayed after running the command:
`Type Name Manufacturer`
`---- ---- ------------`
`SMP Storage Spaces Management Provider Microsoft`
`Corporation`
`SMI-S 192.85.142.152 HP 3PAR`
	Verifies if the SMI-S provider is listed

`Update-StorageProviderCache –DiscoveryLevel full`	Updates the WMI classes
`Unregister-SmisProvider -ConnectionUri http://<3PAR StoreServ systemip address>:<port>`	Removes the registered SMI-S provider

- Execute the StoreEasy discovery script. For information about executing StoreEasy discovery script, see "Executing StoreEasy Discovery Script" below.
- Use the Discovery Hint "HP StoreEasy - WMI" (Configuration > Discovery > Discovery Addresses).

Executing StoreEasy Discovery Script

- **On Windows SOM**
  Download and install the binary `psexec` from Microsoft's Windows SystInternals website:
  http://live.sysinternals.com
  `psexec` needs the following setup:
  - TCP ports 135, and 445 must be open.
  - The `Admin$` and `IPC$` shares must be enabled.
  - The environment variable `PATH` must include the `psexec` installation path.
    For example, if `psexec.exe` is located at `C:\Temp`, type the following command:
    `set PATH=%PATH%;C:\Temp`

    [*drive*:]`\Program Files (x86)\HP\HP BTO Software\bin\somwindowsAgentlessDiscovery.ovpl -i`
- **On Linux SOM**
  - TCP port 139 must be open.
  - Ensure that you can mount the remote admin share, `admin$/Temp` using a CIFS mount. For example,
    `mount -t cifs -o username=<UserName>,password='<password>' //<remote-host-name>/admin$/Temp /mnt/<mount-point>`

- Run the `somwindowsAgentlessDiscovery.ovpl [-i]` command.

  > **Note:** Same command is applicable for agentless discovery of Windows host.

- `/opt/OV/bin/somwindowsAgentlessDiscovery.ovpl –i`

  The option "`-i`" ensures that subsequent discoveries run automatically after the script is completed.

> **Note:** This is a one-time configuration. Subsequent discoveries happen automatically.

For more information about the `somwindowsAgentlessDiscovery.ovpl` command, see the *HPE Storage Operations Manager CLI Reference Page*.

> **Note:** If SOM does not display the IP address of a file server because of DNS resolution issues, add the IP address of the file server in the `hosts.ini` file.

Prerequisites to Discover IBM XIV Storage System

SOM uses SMI-S to discover the IBM XIV Storage System.

To discover an IBM XIV Storage System:

- Provide the IP address and authentication details of the device.
- (Optional) Use the discovery hint "IBM XIV- SMIS" (**Configuration** > **Discovery** > **Discovery Addresses**).

Prerequisites to discover EMC CLARiiON and VNX Block Storage Systems

SOM uses the EMC Solutions Enabler with the SMI-S provider on a proxy server to discover CLARiiON and VNX storage systems.

To discover a CLARiiON and VNX storage system:

1. Install EMC Solutions Enabler with the SMI-S package on a proxy server.
2. Provide details to access the proxy server.
   Specify the IP address of the Solutions Enabler server and the user name and password of the EMC SMI-S provider (ECOM).

> **Note**: If an EMC ClARiiON/VNX block array is managed by multiple SMI-S proxies and these are discovered by SOM, the proxies are reconciled and a single Access Point (Analysis pane > Summary tab) is retained.

3. Ensure that port 5989 is open.

> **Note:** If storage arrays discovered through a proxy server are moved to another proxy server, you must discover the new proxy server and either re-run data collection manually on the storage arrays or wait for the next data collection cycle, configured as per the data collection policy.

Prerequisites to Discover EMC VPLEX Clusters

SOM discovers EMC VPLEX clusters using the VPLEX Element Manager (REST) API via HTTPS. It enables you to discover VPLEX Local, VPLEX Metro and VPLEX Geo clusters.

To discover a VPLEX cluster:

1. Provide the IP address and authentication details of the VPLEX management console.
2. Keep port 443 open.

> **Note**: In VPLEX Metro and VPLEX Geo configurations, you can discover either or both clusters. However, it is recommended that you discover both VPLEX systems. This is because if both clusters are discovered, information from each is collected through the local management station and is therefore not susceptible to inter-cluster communication failures.

Prerequisites to Discover NetApp Devices

SOM uses the ONTAP API to discover NetApp devices.

To discover a NetApp 7-Mode or C-Mode device:

1. Provide the IP address and authentication details of the NetApp device.

   > **Note:** For a C-mode device, specify the IP address of the cluster and for a 7-mode device, specify the IP address of the device.

2. Keep port 443 open.

Prerequisites to Discover EMC Data Domain Storage System

The EMC Data Domain system is a backup storage device. EMC Data Domain system provides data protection and disaster recovery in an enterprise environment. SOM uses SNMPv2C or SNMPv3 protocols and Data Domain CLI (over SSH) to discover and collect data from an EMC Data Domain system.

Discovery using SNMPv2C

1. In the EMC Data Domain system, configure SNMPv2C community string (with read-only access) using the EMC Data Domain System Manager web interface. For information about configuring SNMP settings in EMC Data Domain storage system, see the device documentation.
2. Specify the IP address of the EMC Data Domain system in SOM (**Configuration** > **Discovery** > **Discovery Addresses**).
3. Create a discovery credential using the read community string as username. Leave the password field blank.
4. (Optional) Use the **Discovery Hint** "EMC Data Domain - SNMP" (**Configuration** > **Discovery** > **Discovery Addresses**).

Discovery using SNMPv3

1. Configure SNMPv3 settings in the EMC Data Domain system using the EMC Data Domain System Manager web interface. For information about configuring SNMP settings in EMC Data Domain storage system, see the device documentation.

   Ensure that the SNMPv3 user, used to communicate with SOM, has the same key for privacy and authentication protocols.
2. Stop the SOM JBoss service.
3. Edit the `%OvDataDir%\Conf\som\custom.properties` file to add the following properties from `%OvDataDir%\Conf\som\custom.properties.sample` file:

   ```
   dd.useSNMPv3=true
   dd.snmp.timeout=<time in seconds>
   dd.snmp.retries=<number of retries>
   #specify the privacy protocol that was configured on the device side in above step
   dd.snmp.privacyProtocol= <privacy protocol configured on device side>
   #allowed values for privacy protocol are: NOPRIV, AES, and DES.
   dd.snmp.authenticationProtocol= <authentication protocol configured on device side>
   #allowed values for authentication protocol are: NOAUTH, MD5, and SHA1.
   ```

   > **Note:** The default privacy protocol is AES and the default authentication protocol is MD5.
   >
   > These properties apply to all EMC Data Domain systems that SOM discovers and collects data from. If SOM needs to discover and collect data from multiple EMC Data Domain systems, ensure that all devices have SNMPv3 user configured with same authentication and privacy protocols. For

> information about configuring SNMPv3 settings in EMC Data Domain storage system, see the device documentation.
>
> You can copy these properties from `%OvDataDir%\Conf\som\custom.properties.sample` file.

4. Start the SOM JBoss service.
5. Specify the IP address of the EMC Data Domain system in SOM (**Configuration** > **Discovery** > **Discovery Addresses**).
6. Create a discovery credential with SNMPV3 user name as username and key as password. (**Configuration** > **Discovery** > **Discovery Addresses**).

   Ensure that keys for privacy and authentication protocols are same while configuring the keys in EMC Data Domain storage system. For information about configuring SNMP settings in EMC Data Domain storage system, see the device documentation.
7. (Optional) Use the **Discovery Hint** "EMC Data Domain - SNMP" (**Configuration** > **Discovery** > **Discovery Addresses**).

Data Collection

To collect endpoint information and snapshots, SOM requires Data Domain local user (with the role user) credentials to execute CLI commands through SSH. For information about the CLI commands, see "Data Collection Commands Used for EMC Data Domain Storage System" below.

If you discover EMC Data Domain system using SNMPv2C protocol, ensure that local user name in EMC Data Domain system is same as the SNMPv2C community string. If you discover EMC Data Domain system using SNMPv3 protocol, ensure that local user name and password in EMC Data Domain system is same as the SNMPv3 user name and key.

Data Collection Commands Used for EMC Data Domain Storage System

SOM uses the following CLI commands to collect endpoint information and snapshots from EMC Data Domain system:

Commands	**Description**
`scsitarget endpoint show detailed`	Collects the endpoint information.
`scsitarget group show detailed service ddboost`	Collects endpoint information for DD boost service.
`snapshot list mtree <mtree_name>`	Collects snapshots for an MTree.

Discovery Tasks

To discover a new device, follow these steps:

1. Configure an address for discovery.

   or

   "Steps to Configure a Range for Discovery" on page 157
2. (*optional*) "Configure Credentials for Discovery" on page 160
3. "Create a Tenant" on page 163
4. "Start Discovery" on page 165

Note: If you rename a host or device element after its discovery, you need to delete that element and discover the renamed element, though there is no change in its DNS name or IP Address. If you delete an element, SOM loses the historical data about the element.

Configure Addresses for Discovery

Use the Discovery Address form to configure a new IP address of a storage element that you want to discover.

To configure an address for discovery, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Addresses**. The Discovery Addresses view is displayed.
2. Click **New** on the view toolbar. The Discovery Address form is displayed.
3. Specify the discovery address details. (See the "Attributes" below table.)
4. Click one of the options to save the address.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The address is displayed in the Discovery Addresses view.

Attributes	Description
IP	Type the IP address or the FQDN of the device to be discovered.

Attributes	Description
Address/DNS Name	If a device is not configured for the default port, specify the port number, separated by a colon. For example, if you enter `proxy2:1234` • `proxy2` is the name of the proxy server or the IP address of the device. • `1234` is port number.
Credentials	The discovery credentials, if required, of the device. Select an existing discovery credential or click lookup for additional options. • **Show Analysis** to display details of the current selection. • **Quick Find** to access the list of existing items. • **Open** to view the details of the current selection. • **New** to create a new item, for example a new tenant or a new discovery credential.
Tenant	The tenant associated with the IP address. By default, the IP address is associated to the default tenant. Select an existing tenant or click lookup for additional options. • **Show Analysis** to display details of the current selection. • **Quick Find** to access the list of existing items. • **Open** to view the details of the current selection. • **New** to create a new item, for example a new tenant or a new discovery credential. Note: If a tenant is not created, select the default tenant.
Discovery Hint	Discovery hint is a combination of device bundle name, vendor name and the discovery mechanism. When you select a value, it serves as a hint to invoke only the selected provider for discovering the device instead of invoking all the providers. Use this option to reduce discovery time. Select a value from the drop-down list based on the service provider.
Comments	Type any additional notes for the IP address.
Queue Discovery	Enabled by default. If selected, this option enables you to perform automatic discovery by queuing the elements for discovery.

Delete an Address

To delete an address configured for discovery, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Addresses**. The Discovery Addresses view is displayed.
2. Select the address that you want to delete from the table view.
3. Do one of the following.
   - Click **Delete**. The delete confirmation message is displayed. Click **OK** to delete the address.
   - Click **Open**. The address is displayed in Address Form view. Click . The delete confirmation message is displayed. Click **OK** to delete the address.

Configure Address Ranges for Discovery

Use the Discovery IP Range form to configure a new IP address range or modify an existing range.

Considerations for Defining an Address Range

Before you define an address range, consider the following points:

- Enter a range within the same subnet. The management server cannot scan IP ranges across subnets.
- The discovery process behaves as if an IP range is in the same subnet even if the IP range includes more than one subnet. For example, if you specify the range 172.16.190.10–172.16.191.20, it will discover 172.16.190.10–172.16.190.20.
- In the IP range, include a proxy server that has a direct connection or a SAN connection to the SOM management server, such as the EMC Solutions Enabler. Make sure that the proxy service has started. For Microsoft Windows systems, check the status of the proxy service in the Services window.
- The management server does not scan port numbers in an IP range. For example, you cannot discover an instance of the HiCommand Device Manager that listens on a port other than 2001.
- The management server displays duplicate discovery addresses for an element in the following scenario:
  - You add an IPv4 address for an element to be discovered, and then run an IP range scan that includes the IPv4 address of the previously added element.

Steps to Configure a Range for Discovery

To configure an address range for discovery, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Ranges**. The Discovery Ranges view is displayed.
2. Click **New** on the view toolbar. The Discovery IP Range form is displayed.

3. Specify the address range details. (See the "Attribute" below details.)
4. Click one of the options to save the range.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The address range is displayed in the Discovery Ranges view.

Attribute	Description
From IP Address	The first IP address in the address range.
To IP Address	The last IP address in the address range. **Note**: Make sure that the first and last address belong to the same subnet. That means only the last part of the IP address must be different.
Credentials	The discovery credentials, if required, to discover the IP address range. Select an existing discovery credential from the list or • **Show Analysis** to display details of the current selection. • **Quick Find** to access the list of existing items. • **Open** to view the details of the current selection. • **New** to create a new item, for example a new tenant or a new discovery credential.
Tenant	The tenant associated with the IP address range. SOM provides a predefined Tenant, the *Default Tenant* that is mapped to the SOM *Default Security Group*. An administrator can create additional tenants as needed. If a tenant is not created, select the SOM *Default Tenant*. An administrator can change the tenant assignment at any time. Select an existing tenant from the list or • **Show Analysis** to display details of the current selection. • **Quick Find** to access the list of existing items. • **Open** to view the details of the current selection. • **New** to create a new item, for example a new tenant or a new discovery credential.

Attribute	Description
Comments	Any additional notes the administrator adds related to the particular IP address range.

Scan an Address Range

After you add an address range, check for valid addresses in the range by scanning the range. The valid addresses are added to the Discovery Addresses view from where you can start discovery.

To scan an address range for valid address, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Ranges**. The Discovery Ranges view is displayed on the right pane.
2. Select an address range from the table view.
3. Right-click on the selected address range and click **Start Scanning** to scan the addresses in the range.

Modify an Address Range

To change an address range configured for discovery, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Ranges**. The Discovery Ranges view is displayed.
2. Select the address that you want to modify from the table view.
3. Click **Open**. The address range is displayed in Discovery IP Range Form view.
4. Make the necessary modifications to the address range.
5. Click to save changes to the address range. The Discovery Ranges view is refreshed to display the changes in the address range.

Delete an Address Range

To delete an address range configured for discovery, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Ranges**. The Discovery Ranges view is displayed.
2. Select the address range that you want to delete from the table view.
3. Do one of the following.
   - Click **Delete**. The delete confirmation message is displayed. displayed. Click **OK** to delete the address range.

- Click **Open**. The address range is displayed in the Discovery IP Range Form view. Click . The delete confirmation message is displayed. Click **OK** to delete the address range.

Configure Credentials for Discovery

Use the Discovery Credentials form to add a new discovery credential.

To configure a credential for a device, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Credentials**. The Discovery Credentials view is displayed.
2. Click **New** on the view toolbar. The Discovery Credentials form is displayed.
3. Specify the credentials for the device. (See the "Attributes" below table.)
4. Click one of the options to save the credentials.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The credentials are displayed in the Discovery Credentials view.

Attributes	Description
Name	Type a unique string to distinguish the credential from others in the list. For example, two Windows hosts may have the same user name as "Administrator", but the credential names could be "Payroll server user" and "HR Server user".
User name	The identifier used to log in to the proxy processes running on the specified IP address or range during discovery.
Password	The password for the user name, if required. Ensure the password length does not exceed 32 characters.

Note: When you select a discovery credential, you can see a list of all associated IP Addresses/DNS names and discovery ranges in the analysis pane.

Modify a Discovery Credential

To modify a discovery credential, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Credentials**. The Discovery Credentials view is displayed.
2. Select the credential that you want to modify from the table view.
3. Click **Open**. The credential is displayed in Discovery Credentials Form view.
4. Make the necessary modifications to the credential.
5. Click to save changes to the discovery credential. The Discovery Credentials view is refreshed to display the changes to the credential.

Note: To restart the discovery of all IP addresses related to the discovery credential that has been modified, right-click the credential and select **Start Discovery of Related Addresses**. After successful discovery, you can either re-run the data collection manually or let SOM collect data as per the scheduled data collection policy.

Delete a Discovery Credential

To delete a discovery credential, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Credentials**. The Discovery Credentials view is displayed.
2. Click to select the discovery credential that you want to delete from the table view.
3. Do one of the following.
   - Click **Delete**. The delete confirmation message is displayed. Click **OK** to delete the discovery credential.
   - Click **Open**. The address range is displayed in the Discovery IP Range Form view. Click . The delete confirmation message is displayed. Click **OK** to delete the discovery credential.

Configure Tenants

Tenant settings help you to accomplish the following:

- Identify overlapping address domains in your network so SOM can avoid duplicate address problems.
- Assign the *Initial Discovery Security Group* to elements after discovery.

  Note: Devices within the Default Security Group are visible from all views. To control access to a device, assign that device to a Security Group other than Default Security Group.

- Identify logical groups of nodes for any purpose, for example to identify the resources assigned to a specific customer or to identify specific areas of your network or to identify company sites.

- Create Node Groups based on Tenant attribute values. See " Create an Additional Filters Expression" on page 115 for more information about Node Group filters.

Use the Tenant form to create an association between a tenant and a security group. When you configure the IP address or the FQDN of an element to be discovered, the element inherits the security settings of the security group that is associated with the selected tenant.

An administrator create additional tenants as needed and can change a node's tenant or security group assignment at any time. See "Change Tenant Assignment for a Node" on page 164 for a Node for more information.

Related Topics

"Configure Security Groups " on page 91

Tenant and Initial Discovery Security Group Assignments

When SOM discovers elements in your storage network environment, Tenant and Security Group settings are established in the following manner:

Discovery Addresses: You can specify a tenant for each discovery address. A node is automatically created for an IP address that is discovered successfully . When administrators define a tenant, they specify an **Initial Discovery Security Group**. A newly created node associated with a defined tenant is mapped to the security group (the Initial Discovery Security Group) that is associated with the selected tenant. An administrator can change either the node's tenant or security group assignment or both at any time.

Nodes assigned to the *Default Security Group* are visible from all views. To control access to a device, assign that device to a security group other than Default Security Group.

Nodes within one tenant can each be assigned to different security groups, and nodes within one security group each be assigned to different Tenants.

Consider setting up your security configuration so that all newly-discovered nodes belong to a security group that is mapped to User Group = SOM Administrators . Those nodes will be visible only to administrators until an administrator intentionally moves the node into a security group that is also visible to the appropriate SOM operator or guest.

Tenant assignments are useful for identifying groups of nodes—node groups—within your network environment. Security Group assignments enable administrators to restrict the visibility of nodes within the SOM console to specific User Groups. For more information, see "Configuring Security" on page 77.

Recommendations for Planning Tenants

Consider the following recommendations while planning tenant configuration:

- Configuring tenants during discovery reduces administration overheads of assigning discovered elements to respective tenants manually.
- For a small organization, a single security group per tenant is probably sufficient.

- You might want to subdivide a large organization into multiple security groups.
- To prevent users from accessing nodes across organizations, ensure that each security group includes nodes for only one tenant.

Create a Tenant

To create a tenant, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Tenants**. The Tenants view is displayed.
2. Click ✱ **New** on the view toolbar. The Tenant form is displayed.
3. Make your configuration choices. (See the Tenant Attributes table.)
4. Click one of the options to save the tenant.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The tenant is displayed in the Tenants view.

Tenant Attributes

Attribute	Description
Name	Type a name that uniquely identifies this tenant. **Note**: You must enter a name value.
UUID	SOM assigns a Universally Unique Object Identifier (UUID) to the Tenant. This UUID is unique across all databases.
Description	Description of the tenant. Type a maximum of 2048 characters to describe this Tenant. Alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _+ -) are permitted.
Initial Discovery Security Group	The *Initial Discovery Security Group* specifies the security group assigned to an *IP Address* or *IP Address Range* associated with the tenant before discovery. For more information, see "Tenant and Initial Discovery Security Group Assignments" on the previous page. **Caution:** Devices within the *Default Security Group* are visible from all views. To control

Tenant Attributes, continued

Attribute	Description
	access to a device, assign that device to a Security Group other than Default Security Group. Administrators can assign each node within one tenant to a different security group. Select an existing security group from the list or • **Show Analysis** to display details of the current selection. • **Quick Find** to access the list of existing items. • **Open** to view the details of the current selection. • **New** to create a new item, for example a new tenant or a new discovery credential.

Change Tenant Assignment for a Node

After discovery you can change the tenant of a node. However, you must have defined at least one tenant in addition to the default tenants.

If you have not created any tenant, then

- The Tenant attribute does not appear on any Node form.
- The Tenant column does not appear in the Nodes view.

To change the tenant of a node, follow these steps:

1. Navigate to the Node form.

   You can access the Node form from the table view of the element in the Inventory workspace. For example, if you want to access the node form of a host, navigate to **Inventory** > **Hosts** and click on the node column in the table view. The Node form is displayed.

2. To change the tenant, do one of the following:
   - Select the drop-down list and choose a different tenant.
   - Click **Lookup** and select **New** to create a new tenant.

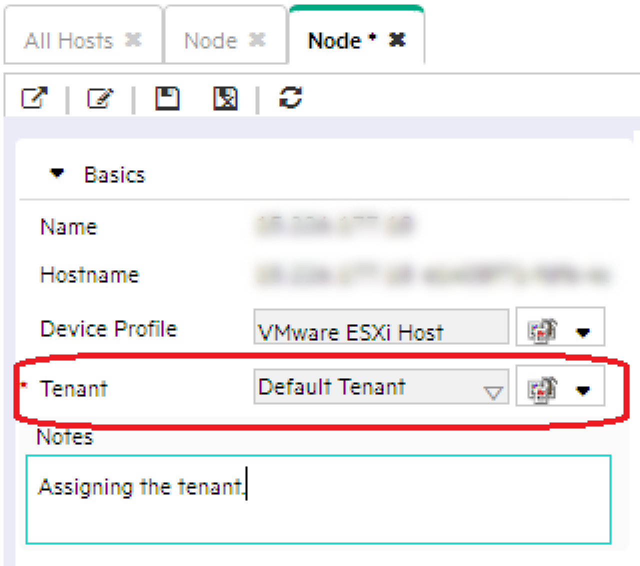

3. Click Save to change the tenant.

Start Discovery

To start discovery of a device, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Discovery** > **Discovery Addresses**. The Discovery Addresses view is displayed.
2. Select the address of the device for which you want to start discovery.
3. Right-click and select **Start Discovery**. The discovery starts and the result of the discovery is displayed in the Status column. See "Status of Discovery" below for the complete list of discovery statuses.

Note: The Start Discovery option does not restart the process for devices that are being discovered. However, if auto-unquarantine is enabled and the selected address (access point) has related elements that are quarantined, this option unquarantines the elements and initiates discovery. For more information about auto-unquarantine, see "Quarantine/Unquarantine Elements" on page 434.

Status of Discovery

The Status column in the Discovery Address view displays the discovery status of elements such as hosts, storage systems, switches, and fabrics. The following is the list of discovery statuses and their description.

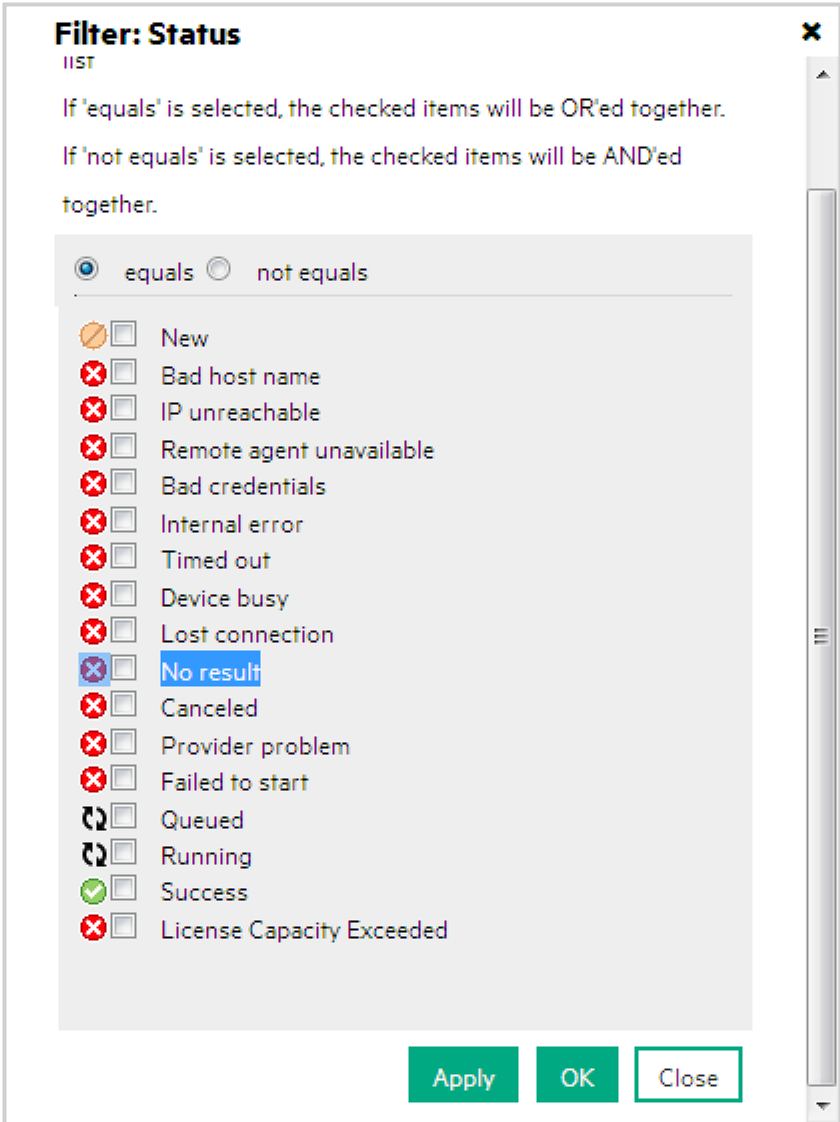

Discovery Views

To access discovery views, from the workspace navigation panel, click **Configuration** > **Discovery**. Select the view that you want to display. For example, select Discovery Addresses to display the Discovery Address view.

The Discovery folder in the Configuration workspace provides the following views:

Discovery Addresses

The Discovery Addresses view displays the list of addresses that are configured for discovery. Double-click on the address to open the address in the form view. The analysis pane provides a link to the Inventory view that show the top level elements that are discovered using the address.

You can perform the following tasks through this view:

- "Start Discovery" on page 165
- "Status of Discovery" on page 165

Discovery Range

The Discovery Ranges view displays the list of address ranges that are created so that they can be scanned valid discovery addresses. You can start scanning the addresses from this view.

Discovery Credentials

The Discovery Credentials view displays the list of discovery credentials that can be used to authenticate the discovery of new IP addresses from the Discovery Addresses list or new IP address ranges from the Discovery Ranges list.

Tenants

The Tenants view displays the list of tenants created and available in SOM. The default tenant is also listed in the view.

Inferring Hosts Based on Rules

SOM can gather and display information from hosts without discovering them. You can infer hosts by creating rules based on host security groups, zones, zone aliases, or device aliases created for Cisco switch ports. Rules probe your switch and storage configurations based on specific search parameters using regular expressions to infer connected hosts. When the scope is a zone, zone alias or a Cisco device alias, the hosts are inferred only after the fabric connectivity information is available.

> **Note:** If a zone or zone alias is used to infer a host, SOM considers the port WWNs configured in a zone or zone alias. SOM does not consider the Node WWNs to infer hosts.

Storage paths are obtained from the host security groups (HSG) configured for a storage system. Therefore if an HSG is used to infer a host, data collection must be rerun for the connected storage system to see information about the storage dependencies of a host.

The following functionality is not available for hosts inferred through rule-based host inference:

- Automatic cluster membership detection
- Associated virtual machines

When hosts are inferred after running the rules, the hosts are listed in the Inferred Hosts view (**Inventory** > **Hosts** > **Inferred Hosts**). Inferred hosts are associated with the host node group by default.

After inferring hosts, you can discover the inferred hosts by providing the credentials. If the discovery is successful, the hosts are reconciled and the inferred hosts become managed hosts.

When multiple rules are executed concurrently, it might be possible that multiple hosts with the same name are inferred. This can be ignored.

Regular Expressions in Rules

Consider the following best practices while creating regular expressions for inferring hosts:

- Consider the naming convention of the Cisco device aliases, zones, zone aliases, and host security groups in the environment so that the hosts can be detected. You might need multiple rules for different naming conventions.
- Use a capturing group that is used to display the host name. A capturing group is the characters within a set of parentheses.

Example

Assume that the hosts that you want to infer are prefixed with boston_, but you want to display only the host names without the boston_ prefix. In this case, you can use the following expression: boston_(.*)

Any host with a prefix of boston_ is inferred, but only the text after boston_ is displayed as the host name.

If you wanted boston_ to be displayed in the host name and you still want only hosts with the prefix boston_ inferred, you could change the expression so that boston_ is included in the capturing group, as shown in the following expression: (boston_.*)

If you are not sure where to begin, consult the following examples to see if any match your environment. Try entering some of the basic expressions, such as .*_.*_.*, and see what is inferred. You can always add additional rules to narrow the range to detect a particular naming convention.

Examples of Regular Expressions

Environment	Regular Expression	Result
`Boston_HostName_hba1`	`.*?_(.*?)_.*`	Strings that match the pattern of text_text_text will be scanned. The text between the first and second underscores will be displayed as the host name.
`Boston-HostName-disk`	`.*?-(.*?)-.*`	Strings that match the pattern of text-text-text will be scanned. The text between the first and second dashes will be displayed as the host name.
`Boston-HostName_com`	`.*?-(.*?)_.*`	Strings that match the pattern of text-text_text will be scanned. The text

Examples of Regular Expressions, continued

Environment	Regular Expression	Result
		between the first dash and second underscore will be displayed as the host name.
`Boston_storage_HostName`	`Boston_storage_(.*)`	Strings that match the pattern of Boston_storage_ text will be scanned. The text after the second underscore will be displayed as the host name.
`Boston___HostName_disk`	`.*?___(.*?)_.*`	Strings that match the pattern of text___text_text will be scanned. The text between the third and fourth underscores will be displayed as the host name.
`uhcHostName` HostName is always the fourth character.	`...(.*)`	Strings that have four or more characters will be scanned and any characters after the third character spot will be displayed as the host name.
`HostName:hba`	`(.*?):.*`	Strings that match the pattern of text:text will be scanned. Any text before the first colon will be displayed as the host name.

`boston_HostName_hba1`
`boise_HostName_hba1`
`marlborough_HostName_hba1`
but you do not want to infer `zebra_HostName_hba1`
	`[a-q]_(.*?)_.*`	Strings that begin with any lowercase letter from a to q and matches the pattern of text_text_text will be scanned. Any text between the first and second underscore will be displayed as the host name.

Examples of Regular Expressions, continued

Environment	Regular Expression	Result
		For uppercase letters use [A-Q]. You can change the range to match your environment; for example, a-s or N-Z.
boston1_HostName_hba1 boston3_HostName_hba1 but you do not want to infer boston9_HostName_hba1	.*[1-3]_(.*?)_.*	Strings that have number 1, 2 or 3 before the first underscore and that match the pattern. Any text between the first and second underscores will be displayed as the host name. You can change the range to match your environment; for example, 23 to 54.
HostName1_HostName2_HostName3		Strings that have two underscores will be scanned. Text before, after, and between the underscores will be displayed as host names.
Boston_HostName_hba1Boston-HostName-hba1	.*_(.*)_.*|(.*-(.*)_.*)	
MRO_HostName_disk My naming convention requires all zone names to begin with MRO, but I know a few have been created incorrectly and I want to capture those. For example, if I want to find any rogue zone names that do not start with "M" because my naming convention requires that all zones begin with "MRO," I would attempt to infer hosts with an expression like ([a-ln-zA-LN-Z]*).	([a-ln-zA-LN-Z]*)	This expression displays strings that begin with any letter except for the lowercase or uppercase letter M. The entire string would be displayed as the host name, so you could find the rogue zone names.

Examples of Regular Expressions, continued

Environment	Regular Expression	Result
`HostNameNN`	`(HostName.*)`	Strings that begin with HostName will be scanned.The text having same prefix will be displayed as host name.

The notation used in the expressions are defined as follows.

Definition of Common Notation Used in Expressions

Expression	Definition
`()`	Capturing group. Any expression within a set of parenthesis is displayed for the host name. If you do not provide a capturing group, no host name will be displayed from the hosts that were detected from the expression.
`?`	The reluctant quantifier. It starts search from the beginning of the input string, then reluctantly consumes one character at a time looking for a match. Finally, it tries the entire input string. Reluctant quantifiers are specifically used to extract host names from specific patterns like, all characters between the first underscore and the second underscore, as illustrated in the examples in the preceding table.

`.*`	Any character zero or more times. Use this expression carefully. For example, the following expression matches any element that has the boston_ prefix:
	`boston_.*`
	If you want HPE Storage Operations Manager to display any character after the boston_ prefix, add a capturing group as follows:
	`boston_(.*)`
	Assume though that you do not want to display all the characters after the boston_ prefix. If there is a character after .*, the wild card attribute will stop. For example, the following expression displays the characters that appear after boston_ and before _companyname:
	`boston_(.*)_companyname`
	Assume that all of your hosts do not end in _companyname. You can replace _ companyname with _.* as follows:
	`boston_(.*)_.*`
	The expression matches all hosts with the prefix of boston_, and displays any character that is after boston_ but before the second underscore.

Definition of Common Notation Used in Expressions, continued

Expression	Definition
	Note: Regular expressions are Java regular expressions and you must take care about using the greedy and reluctant quantifiers, as appropriate.

`.`	Any character. For example, assume the hosts in your environment all have different naming conventions, but contain three characters before the host name. You could provide an expression as follows:
	`...(.*)`
	Hosts with the name BosHost1 or LasHostA would appear as follows in the topology:
	`Host1 and HostA`

`[a-q]`	Lowercase letter between a and q
`[A-Q]`	Uppercase letter between A and Q
`[0-7]`	Digits between 0 and 7

`|`	The OR operator. Use the OR operator when you have different naming conventions in your environment. For example, assume you want to match hosts prefixed with boston_ or boise_. You could use the following expression to match those hosts:
	`boston_(.*) | boise_(.*)`
	You could also use the OR operator to find hosts when the naming convention differs between host names. For example, assume you have some hosts that have underscores in their name and others that have dashes. You could use the following expression to match those hosts:
	`.*_(.*) | .*- (.*)`

For more information about regular expressions, go to:

http://docs.oracle.com/javase/1.5.0/docs/api/java/util/regex/Pattern.html

Create a Rule

To create a rule for inferring hosts, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Rule Based Host Inference** > **Host Inference Rules**. The Host Inference Rules view appears.
2. Click **New** on the view toolbar. The Host Inference Rule form appears.
3. Specify the host inference rule details. (See the Host Inference Rule attributes below.)

4. Click one of the save options.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The host inference rule is displayed in the Host Inference Rules view.

   The table displays the attributes of a host inference rule.

Host Inference Rule Attributes	Description
Rule Name	Type the name of the rule.
Description	Type the description of the rule.
Priority	Type the priority of the policy. This can be any positive integer.
Run After Data Collection	The check box is selected by default and the rule is run after every successful data collection to infer new hosts and update information. For example, if the scope of the rule is host security group and you have new storage systems discovered, the rule is run after successful data collection of the storage systems. If you do not select this option, the system runs the rule based on its priority or you can choose to run the rule manually.
Add Hosts Inferred to Discovery Addresses	The check box is selected by default and the details of the hosts that are inferred from this rule are added to the Discovery Addresses view (**Configuration** > **Discovery** > **Discovery Addresses**) from where you can start discovery of the inferred host by adding its credential. If the discovery of the inferred host is successful, the host becomes a managed host and is no longer an inferred host. **Note**: If you clear this option, you cannot add inferred hosts to Discovery Addresses view later.
Scope	Select a scope from the following options: • **Zone** The rule searches the zone name for hosts on the fabrics. The discovery of the fabrics must be complete. • **Zone Alias**

Host Inference Rule Attributes	Description
	The rule searches the zone alias name for hosts on the fabrics. The discovery of the fabrics must be complete. Keep in mind the following when selecting Zone or Zone Alias as a scope: • You can run the rule from a management server where you have discovered only fabrics. You will be able to infer host names, but you will not obtain any storage details if no storage has been discovered. • Orphan zones and orphan zone aliases could return false inferences. • **Host Security Group** The rule searches the host security group names on the storage systems for hosts. The discovery and data collection for the storage systems must be complete for the rule to run. • **Cisco Device Alias** The rule uses a Cisco device alias along with the specified regular expression to infer a host name. SOM then uses the host name and the WWN of the device alias to collect details of the inferred host. Connectivity information of the Cisco fabric must be available for this rule to run. **Note:** If the symbolic name of a host is listed in the `%OvDataDir%\conf\som\custom.properties` file as a value for the `hostSymbolicNames` property, SOM considers the port is connected to a host and infers the host.
Regular Expression	Select the regular expression from the list. You can modify the regular expression as required. The regular expression determines what element will be inferred. See "Regular Expressions in Rules" on page 168 for more information.

Modify a Rule

To edit a rule, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Rule Based Host Inference** > **Host Inference Rules**. The Host Inference Rules view appears.
2. Select the rule that you want to modify and click Open . The rule is displayed in the Host Inference Rule form.

3. Modify the rule as necessary.
4. Click to save the changes to the rule.

The changes that you made to the rule become effective when the rule is run the next time, that is either after successful data collection or when you run the rule manually.

Delete a Rule

To delete a rule, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Rule Based Host Inference** > **Host Inference Rules**. The Host Inference Rules view appears.
2. Select the rule that you want to delete.
3. Do one of the following.
   - Click **Delete**. The delete confirmation message is displayed. Click **OK** to delete the rule.
   - Click **Open**. The rule is displayed in Host Inference Form view. Click . The delete confirmation message is displayed. Click **OK** to delete the rule.

When you delete a rule, the hosts inferred with the rule are not deleted and will continue to appear in the Inferred Hosts inventory view. However, the Host Inference Rule column is blank in the view as the rule is deleted and no longer exists in the system.

Run a Rule Manually

Before running a rule manually, the discovery and data collection of fabrics and storage systems must be complete based on the scope of the inference rule.

To run a rule manually to infer hosts, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Rule Based Host Inference** > **Host Inference Rules**. The Host Inference Rules view appears.
2. Click a rule to select it, right-click and click **Run a Rule**.

   The inferred hosts are displayed in the Inferred Hosts inventory view.

Note: SOM does not generate management events if you execute host inference rule manually. SOM generates management events only when host inference rule is executed during data collection of storage arrays.

View Inferred Hosts

You can view inferred hosts using the Inferred Hosts (**Inventory** > **Hosts** > **Inferred Hosts**) view. You can delete an inferred host from this view.

A rule must have run at least once for the hosts associated with the rule to be displayed. The view is refreshed whenever a rule is run and changes in the host topology are recalculated .

Delete an Inferred Host

To delete an inferred host, follow these steps:

1. From the workspace navigation panel, click **Inventory** > **Hosts** > **Inferred Hosts**. The Inferred Host view appears.
2. Click the inferred host that you want to delete. Right-click and select . The delete confirmation message appears. Click **OK** to delete the inferred host.

The hosts reappear in the list when the rule that was used to infer the deleted host runs again.

Delete Hosts Inferred by a Rule

To delete all hosts inferred by a rule, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Rule Based Host Inference** > **Host Inference Rules**. The Host Inference Rules view appears.
2. Select a rule from the table view, right-click and select . The delete confirmation message appears. Click **OK** to delete the hosts inferred by the selected rule.

Reconciliation of Hosts

SOM performs reconciliation of hosts when you discover the inferred hosts after providing their credentials.

Reconciliation of hosts results in the following:

- The ports and cards are associated with the managed host after reconciliation.
- Inferred hosts that are reconciled with managed hosts are deleted.

Configuring Data Collection Settings

A data collection policy is a set of rules that determine the elements from which data is collected and the schedule for data collection. After an element is discovered, SOM automatically creates a node for the element and associates it with one of the default node groups. A data collection policy is created with the following parameters:

- **Node Group** – Determines the target set of devices from which the data is to be collected.
- **Freshness Interval** – Specifies the number of hours after which data collection is to be triggered. After the specified interval, the data collected from the element is considered stale and data collection is triggered again.
- **Blackout Period** – Specifies the time interval during which data collection should not run. This is optional and can be useful in situations when you do not want to disrupt scheduled system activities, such as maintenance.
- **Priority** – Determines the collection policy that applies to a node group. Lower priority value means higher priority. For example, a policy with priority 1 is run before a policy with a higher priority value such as 2.

Data collection policies can be associated with multiple node groups. Therefore, if an element belongs to multiple node groups, it can have multiple effective polices. In such cases, priority of the policy determines when data is collected. Policy with the lowest priority value takes precedence. For example, if an element is simultaneously associated with policy P1 (Priority value 1) and policy P2 (Priority value 2), policy P1 takes effect first. When policy P2 is implemented, data is not collected again from the elements already part of policy P1.

SOM comes with a default data collection policy that is triggered automatically when a new element is discovered. The policy is defined with the following default values:

- Freshness schedule: 24 hours
- Blackout Period: None
- Priority: Zero
- Node Groups: Default Node groups (Hosts, Storage Systems, FC Switches, and FC Fabrics)

Recommendations for Configuring Data Collection

Key points to consider for data collection configuration:

- For effective data collection with minimal overload on the system, set the blackout period to less than or equal to half of the freshness interval. For example, if the freshness interval is 24 hours, the blackout period should not be more than 12 hours.
- It is good to ensure that data collection are not failing because of some very basic reasons such as provider problem, bad credentials, network issues, and such others. These failures add unnecessary overload to the system since there is at least one more data collection retry before the element is quarantined. After such elements are quarantined, visit the "Failure" pie in the collection dashboard to look for elements that report these errors. Take appropriate action to ensure that future data collections are successful and then manually un-quarantine the elements.

- When you assign priorities to policies, do not use numbers in a continuous sequence such as 0, 1, 2, 3, 4, 5, and so on. Ideally use multiples of a positive integer to set the priorities. For example, if you use multiples of 5 as the priority such as 5, 10, 15, 20, and so on. And suppose you want to modify the policy which has a priority of 10. You can change the priority to any number such as 12. This practice is helpful as you don't have to change priorities of all policies that have priorities in immediate succession.

Create a Data Collection Policy

Use the Data Collection Policy form to create a new data collection policy.

To configure a data collection policy, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Data Collection Settings**> **Data Collection Policies**.
2. Click **New** on the view toolbar. The Data Collection Policy form is displayed.
3. Make your configuration choices. (See the Data Collection Policy attributes table below.)
4. Associate a node group to the policy with the following steps:
   a. Under the Node Group Settings tab in the right pane, click **New**. The Node Group Settings form is displayed.
   b. Select the node group from the drop-down list or click Lookup for additional options.
      - Show Analysis – Displays Analysis Pane information for the selected object.
      - Quick Find – Displays a list of valid choices for populating the current attribute field.
      - Open – Opens the form for the related object instance that is currently selected in the lookup field. You can use this option to make changes to the selected object.
      - **New** – Opens a new form to create a new instance of the object.

      See "Create a Node Group" on page 107 for information on how to create a node group.
   c. Click one of the save options.
      - **Save** – To save the form.
      - **Save and New** – To save and open a new form.
      - **Save and Close** – To save and close the form.

      The associated node group appears in the right pane under the Node Group Settings tab.

> **Note**: Repeat Step 4 to associate more node groups to the policy.

5. (*optional*) Associate a blackout period with the policy with the following steps:
   a. Under the Blackout Settings tab in the right pane, click **New**. The Blackout Settings form is displayed.
   b. Select a blackout period from the drop-down list or click Lookup for additional options. See "Create a Blackout Period" on page 181 for more information.
   c. Click one of the save options.
      - **Save** – To save the form.
      - **Save and New** – To save and open a new form.
      - **Save and Close** – To save and close the form.

   The associated blackout period appears in the right pane under the Blackout Settings tab.
6. Click one of the save options.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The policy appears in the Data Collection Policies view.

Data Collection Attributes	Description
Policy Name	The name of the data collection policy.
Freshness Interval (in Hrs)	The maximum number of hours within which at least one data collection will be triggered for that element. After this time, the element will be declared as stale .
Priority (Integer >=0)	A number that is greater than or equal to zero. When multiple policies are applicable, the policy with the lowest priority value takes effect. **Note**: Priorities for data collection policies are set globally. Hence you cannot have multiple policies with the same priority.
Active	Indicates the policy is currently active. De-select this to disable a policy.
TimeOut (In Minutes)	The time in minutes that the SOM management server waits for a response from an element that is queried for data collection. The default value is 180 minutes if no value is specified.

Data Collection Attributes	Description
Description	A general description about the data collection policy.

Modify a Data Collection Policy

To change an address configured for discovery, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Data Collection Settings**> **Data Collection Policies**.
2. Select the policy that you want to modify from the table view.
3. Click **Open**. The policy is displayed in the Data Collection Policy form view.
4. Make the required changes to the policy.
5. Click to save changes to the policy.

Delete a Data Collection Policy

If you are deleting a policy, ensure that the underlying elements associated with the policy are associated with some other policy if you still want to continue to collect data from them.

To delete a data collection policy, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Data Collection Settings**> **Data Collection Policies**. The Data Collection Policy view is displayed.
2. Select the policy that you want to delete from the table view.

   Note: The default data collection policy cannot be deleted.

3. Do one of the following.
   - Click **Delete**. The delete confirmation message is displayed. Click **OK** to delete the policy.
   - Click **Open**. The policy is displayed in the Data Collection Policy form view. Click . The delete confirmation message is displayed. Click **OK** to delete the policy.

Create a Blackout Period

Use the Blackout Period form to define a blackout period.

To define a blackout period, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Data Collection Settings**> **Blackout Periods**.
2. Click **New** on the view toolbar. The Blackout Period form is displayed.
3. Make your configuration choices. (See the Blackout Period form attributes table below.)
4. Click one of the save options.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

BlackOut Period Attributes	Description
Name	The name of a blackout period.
Start Time End Time	The time when a blackout period starts in HH:MM format. The time when a blackout period ends in HH:MM format. **Note**: The time is in 24-hour format, for example, 1:00 AM is 0100 hours and 11:00 PM is 2300 hours.
Days of the week	The days of the week for which the blackout period is effective.

Modify a Blackout Period

Caution: Modifying blackout periods results in re-computation of scheduled data collection policies and could potentially impact system performance. Hence exercise caution if you need to modify a blackout period.

To modify a blackout period, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Data Collection Settings** > **Blackout Periods**. The Blackout Periods view is displayed.
2. Select the blackout period that you want to modify from the table view.
3. Click **Open**. The blackout period is displayed in the Blackout Period Form view.

4. Make the necessary modifications to the blackout period.
5. Click one of the save options to apply your changes.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

   The Blackout Period view is refreshed to display the modified blackout period.

Delete a Blackout Period

To delete a blackout period, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Data Collection Settings** > **Blackout Periods**. The Blackout Periods view is displayed.
2. Select the blackout period that you want to delete from the table view.

   Note: You cannot delete a blackout period that is associated with a data collection policy.

3. Do one of the following.
   - Click **Delete**. The delete confirmation message is displayed. Click **OK** to delete the blackout period.
   - Click **Open**. The blackout period is displayed in the Blackout Period Form view. Click . The delete confirmation message is displayed. Click **OK** to delete the blackout period.

Data Collection Control

Data collection settings enable you to control the subset of data that can be collected based on the device profile of the element. There are two levels of control defined for each device profile

- All
- Default

Note: By default, data collection level is set to 'Default' for all elements.

Storage Systems – Default Collection Level

Device Profile	Missing functionality	Impact
EMC Clariion/VNX Storage	Disk Drives and Storage Extents	End-to-end topology is not shown
EMC Symmetrix DMX Storage		
HPE 3PAR Storage		End-to-end topology is not shown
HP EVA 6000 Storage		
Hitachi Storage Series		
HP P9500 Storage Series		
HPE XP24000 Storage Series		
HPE XP7 Storage Series		

Hosts – Default Collection Level

Device Profile	Missing functionality
HP UX	Disk Partition, Multipath Extent, Volume Manager Volume, Raw Disk Extent, Link Partition, Port Target
Linux	Disk Partition, Multipath Extent, Volume Manager Volume, Raw Disk Extent, Link Partition, Port Target, Device Mapper Partition
Linux agentless	Disk Partition, Multipath Extent, Volume Manager Volume, Raw Disk Extent, Link Partition, Port Target, Device Mapper Partition
Windows host	Disk Partition, Multipath Extent, Volume Manager Volume, Raw Disk Extent, Link Partition, Port Target
Solaris host	Disk Partition, Multipath Extent, Volume Manager Volume, Raw Disk Extent, Link Partition, Port Target

Note: The Drive Type (Inventory > Hosts > Filesystems tab) is **Local** for hosts with the 'Default' data collection control.

Change the Data Collection Control for a Device Profile

Modifying a data collection control results in re-computation of the extent of data to be collected. A change in the collection level implies that the subsequent data collection for the selected device profile excludes or includes the data subset based on the defined level for collection. As a result, some of the tabs in the inventory form view might not have information or might show information collected from earlier collection cycles.

Caution: It is advisable that you do not modify the data collection control level while data collection is in progress. If you attempt to do so, the results of data collection cannot be accurately predicted.

To change the Data Collection Control for a device profile:

1. From the workspace navigation panel, click **Configuration** > **Data Collection Settings**> **Data Collection Control**.
2. Select the device profile that you want to view and click **Open**. The Data Collection Control is displayed for the selected device profile.
3. Modify the collection level using the drop-down list.
4. Click to save changes to the Data Collection Control. The change is effective from the subsequent data collection.

   Note: If you want the changes to the data control to take effect immediately, then you can to trigger a manual data collection from the Inventory view.

Planning Licenses

The HPE Storage Operations Manager restricts the number of elements it manages through licenses. Licensing is based on Managed Access Ports (MAP) count. Refer to the MAP Count Calculation table for details.

Key points on SOM licensing:

- SOM identifies the licensed MAP count (available capacity) limit from the installed license. SOM calculates the MAP count consumption (used capacity) based on the discovered elements in your environment. If the used capacity exceeds the available capacity, SOM will prevent discovery of further elements. In such a case if you attempt to discover an element, you will receive an error “License capacity exceeded.” However, there is no restriction on discovery for a valid temporary Instant-On license.
- Only one type of license is active at a time. You cannot have a mix of Premium and Ultimate-Perf license types. If both SOM Premium license and SOM Ultimate-Perf are installed, then Ultimate-Perf supersedes the Premium license. Available capacity is derived from the superseded license.
- You need SOM Ultimate-Perf license to collect performance metrics from devices that support performance collection. The current release of SOMallows configuring and collecting performance metrics from 25 devices simultaneously by a single instance of the management server.
- You can extend the licensed MAP count (available capacity) by procuring additional licenses. Available capacity will be aggregated and refreshed after installation of new licenses. However, the license capacity for performance is not aggregated and is fixed to 25 devices by a single instance of the management server.

License Types

There are three types of licenses available with the current release of SOM.

License Type	Validity	Supports Performance
SOM Instant-on	60 days	Yes
SOM Premium	Unlimited	No
SOM Ultimate-Perf	Unlimited	Yes

Temporary Instant-On License

When you install HPE Storage Operations Manager, it comes with a temporary Instant-On license. The temporary Instant-On license is valid for 60 days. You should obtain and install a permanent license as soon as possible to continue using SOM.

Obtain and Install New License

To request a permanent license, gather the following information:

- The Entitlement Certificate, which contains the HPE product number and order number
- The IP address of one of the SOM management servers
- Your company or organization information

You can install the permanent license using the Autopass user interface or the command line interface.

Install a Permanent License from the Command Line

To install the license at a command prompt on the SOM management server, enter the following command:

```
somlicensemanager.ovpl SOM -install <path_of_license_file>
```

where `<path_of_license_file>` is the location where the license file is stored.

Install a Permanent License using Autopass

To install a permanent license, follow these steps:

1. At the command prompt, enter the following command to open the Autopass user interface:
   `somlicensemanager.ovpl SOM -gui`
2. On the left pane of the Autopass window, click **License Management**.
3. Click **Install License Key**.
4. Click **Install/Restore License Key**.
5. Browse to the location where the license key is stored.
6. View file content.
7. Select the license and click **Install**.

Extend a Licensed Capacity

To extend the licensed capacity, purchase and install an additional SOM Premium or SOM Ultimate Perf license.

Contact your HPE Sales Representative or your Authorized Hewlett Packard Enterprise Reseller for information about the SOM licensing structure. To obtain additional license keys, go to the HPE License Key Delivery Service:

https://h30580.www3.hpe.com/poeticWeb/portalintegration/hppWelcome.htm

View License Information

1. From the SOM console, click **Help** > **System Information** > **View Licensing Information**.
2. Look for the value shown in the **Consumption** field. This is the number of MAPs that SOM is currently managing (used capacity).

Viewing Consumed MAP Count for Each Element

You can view the number of MAPs consumed by each element being managed by SOM. This information is displayed in the **MAP Count** field in the Analysis Pane of each element in the Inventory views.

About MAP Count Calculation

Element	Description	Number of MAPs	Comments
Hosts	Host with a single port HBA Host with a dual port HBA	1 MAP 2 MAPs	No additional counting for CIM Extension.
	Host without FC ports	1 MAP	
	Host with one iSCSI network card port	1 MAP	
	Host with no FC port and no iSCSI network card port with CIM Extension.	1 MAP	
	Standalone server with no FC HBA discovered through CIM Extension.	1 MAP	
	Windows server agentless discovery through Windows Management Instrumentation (WMI)	1 MAP at a minimum or 1 MAP per FC HBA port.	
	Linux server agentless discovery through SSH	1 MAP at a minimum or 1 MAP per FC HBA port.	
	AIX agentless discovery	1 MAP at a minimum or	

Element	Description	Number of MAPs	Comments
	through SSH	1 MAP per FC HBA port.	
	Solaris agentless discovery through SSH	1 MAP at a minimum or 1 MAP per FC HBA port.	
Virtual Servers	VMware ESX servers	1 MAP at a minimum or 1 MAP per FC HBA port.	Five ESX servers with two dual-ported HBAs count as 10 MAPs (5*2=10)
	Each FC port on a virtual server	1 MAP	Virtual servers are treated like physical hosts.
	A virtual server without FC ports.	1 MAP	The software assumes one MAP.
Virtual Machines	A virtual machine if it is running VMTools irrespective whether it was discovered through its virtual server or its VirtualCenter	1 MAP	
	A virtual machine with an installed CIM Extension regardless if VMTools is running.	1 MAP	
	Each VMware Virtual Machine Guest OS discovered directly through WMI (Windows) or SSH (Linux), or CIM Extension	1 MAP	A VMware Virtual Machine Guest OS discovery through VMTools, and subsequently discovered through agentless WMI or CIM Extension counts as only 1 MAP.
Switches	Each port on a switch Physical switches, all ports are counted as MAPs.	1 MAP	• All switch ports with GBICS installed are counted as MAPs. • ISL links are not counted as MAPs.

Element	Description	Number of MAPs	Comments
			• If the Switch port is not licensed then it’s not counted as MAP. • When GBIC is not there or if the port is not licensed, SOM does not discover these port numbers. Only ports that are discovered are counted as MAPs.
Isilon		No. of nodes * 5 MAPs	
HPE XP / P9500 External Storage	Each port	1 MAP	All backend ports count as MAPs.
EVA, 3PAR, EMC VNX/CLARiiON, DMX/VMAX, VPLEX, HUS/USP	Each port	1 MAP	All backend ports count as MAPs.
NetApp 7-Mode		5 MAPs	Only single node supported.
NetApp C-Mode Logic		No. of nodes * 5 MAPs	
EMC VNX Filer		5 MAPs	
EMC Data Domain Storage Systems	EMC Data Domain System	5 MAPs + No. of FC ports * 1 MAP	
	EMC Data Domain Virtual Edition	5 MAPs	
Block Storage Systems	Block Storage System without FC ports	20 MAPs	

Configure Performance Pack

You must have the SOM Ultimate-Perf license to configure performance collection for storage systems. With the Ultimate-Perf Pack, the current release of SOM supports performance collection from 25 devices simultaneously for a single instance of the management server.

To configure a performance pack, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **License** > **Perf-Pack Configuration**. The Perf-Pack Configuration dialog box is displayed.
2. Select the storage system from the list of Available Storage Arrays for which you want to collect performance data. Use the selection buttons to move your selection to the Selected Storage Arrays. (See the Attributes for details.)
3. Click **Submit**.

   Note: The performance collection does not begin until you configure a monitoring policy for the selected storage systems.

Attributes	Description
Available Storage Arrays	Lists the storage systems discovered by SOM and that are supported for performance with the current release of SOM. For storage systems that support performance, see the *SOM Device Support Matrix*.
Selected Storage Arrays	Displays your current selections. You can select as many storage systems as your license supports.
Total Performance Licenses Available	Displays the available perf-pack capacity of your license.
Total Performance Licenses Consumed	Displays the number of systems already configured for performance collection.

Monitoring Performance

You can monitor the performance of your storage environment using a monitoring policy. A monitoring policy enables you to configure the collection of a set of metrics from the elements (hosts, storage systems, and switches) of a node group.

A monitoring policy acts on a node group. SOM comes with a predefined set of collectors. You can group the collectors logically to form a monitoring group. Define the monitoring policy by associating a monitoring group to a node group and then define parameters such as priority and interval to determine the preference and schedule at which the metrics will be collected.

A monitoring policy consists of the following:

- **Node group**: Determines the target set of devices on which the metrics are to be collected. For example, a storage system node group.
- **Monitoring group**: Determines the set of metrics that will be collected. The collectors are grouped logically to form a monitoring group. For example, the collectors – 3PAR SMI-S Controller collector, 3PAR SMI-S Physical Disk collector, and 3PAR SMI-S Volume collector – can be grouped to form a 3PAR Monitoring Group.
- **Schedule**: Determines the time interval at which the metrics are collected. For example, you can schedule a collection of metrics from the elements of a node group at an interval of 15 minutes.

  Note: Collection schedules for individual elements are available in the **Collector Schedules** tab of the Analysis pane.

- **Priority**: Determines which monitoring policy applies to a given device. Lower priority value means higher preference. For example, if the system determines that multiple policies apply to a device then the policy with the least priority number will be applied.

Recommendations for Monitoring Policies

The following are important recommendations for monitoring the performance of your storage environment:

- Creating too many monitoring policies can add overheads to the system. You should create monitoring policies only for devices and the metrics on those devices that you want to monitor.
- The default interval set during creation of a policy is 15 minutes. It is recommended that you do not have intervals less than 15 minutes as this overloads the system. If you must use intervals less than 15 minutes, it is strongly recommended that you apply this to a very limited set of devices and change it to default interval as early as possible.
- When you assign priorities to policies, do not use numbers in a continuous sequence such as 0, 1, 2, 3, 4, 5, and so on. Ideally use multiples of a positive integer to set the priorities. For example, if you use multiples of 5 as the priority such as 5, 10, 15, 20, and so on. And suppose you want to modify the policy which has a priority of 10. You can change the priority to any number such as 12. This practice is helpful as you don’t have to change priorities of all policies that have priorities in immediate succession.
- Since metric collection is policy-driven, optimize your metric collection with a carefully planned approach:
  - Plan your node groups effectively by identifying high priority devices in your environment. Group collectors logically that is relevant to the node groups, for example do not associate host collectors to a storage system node group.
  - Set schedule intervals judiciously, as explained above.
  - Before configuring monitoring policies in your environment, ensure that one round of data collection is completed for the bulk of the environment. This can be verified from the collection status dashboard. As

a rule of thumb, do not configure monitoring policies when a large number of data collections are in ‘Running’ state.

Prerequisites for a Monitoring Policy

The following are the prerequisites to create a monitoring policy:

- SOM Ultimate Perf license. See "License Types" on page 185 for more information.
- Monitoring group.
- Node group. (Available by default in SOM or create a new node group)
- Successful data collection of the discovered elements.

Create a Monitoring Group

You can create a monitoring group using one of the following:

- "Monitoring Groups View" below
- " Create Monitoring Groups Dialogue" on the next page

Monitoring Groups View

To create a monitoring group, follow these steps:

1. Select **Object Groups** > **Monitoring Groups**. The Monitoring Groups view is displayed.
2. Click ✱ **New** on the view tool bar. The Monitoring Group form is displayed.
3. Enter the monitoring group details as follows.

Attribute	Description
Name	Name of the monitoring group.
Description	Description of the monitoring group.

4. On the Collector Settings tab, click ✱ **New** to associate collectors to the monitoring group. The Collector Settings form is displayed.
5. Select the **Collector** from the drop-down list.
6. Click **Save** to associate the collector to the monitoring group.

> **Note**: You can associate multiple collectors to a monitoring group. You must have at least one collector associated with a monitoring group.

7. Click one of the options to save the monitoring group.
   - **Save** – To save the form.

- **Save and New** – To save and open a new form.
- **Save and Close** – To save and close the form.

Create Monitoring Groups Dialogue

To create a monitoring group using the **Create Monitoring Group** dialog, follow these steps:

1. Select **Object Groups** > **Create Monitoring Group**. The Create Monitoring Group form is displayed.
2. Enter the monitoring group details as follows.

Attribute	Description
Name	Name of the monitoring group.
Description	Description of the monitoring group.

3. Select collectors from the list of **Available Collectors** list. Use the selection buttons to drop your choices to the **Selected Collectors** list.

   Note: Use a combination of Ctrl and Shift keys for multiple selections.

4. Click **Submit** to create the monitoring Group.

Create a Monitoring Policy

To create a Monitoring Policy, follow these steps:

1. From the workspace navigation panel, select **Configuration** > **Monitoring Settings** > **Monitoring Policies**. The Monitoring Policies view is displayed.
2. Click **New** on the view tool bar. The Monitoring Policy form is displayed.
3. Specify the monitoring policy details. (See the "Create a Monitoring Policy" above below).
4. Associate a node group to the policy with the following steps:
   a. On the Node Group Settings tab click **New**on the form tool bar. The Monitoring Policy Node Group Settings form is displayed.
   b. Select the node group from the drop-down list or click for additional options.
   c. Click one of the save options
      - **Save** – To save the form.
      - **Save and New** – To save and open a new form.
      - **Save and Close** – To save and close the form.

> **Note**: You can associate multiple node groups to the policy.

5. Associate a monitoring group to the policy with the following steps:
   a. On the Monitoring Settings tab **New** on the form tool bar. The Monitoring Policy Group Settings form is displayed.
   b. Select the **Monitoring Group** from the drop-down list. If you have not already created a monitoring group, click and **New** . See ""Create a Monitoring Group" on page 192" for information.

   > **Note**: You can associate multiple monitoring groups to a policy.

   c. Click one of the save options to associate the monitoring group to the policy.
      - **Save** – To save the form.
      - **Save and New** – To save and open a new form.
      - **Save and Close** – To save and close the form.
6. Click one of the options to save the monitoring policy.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.
   - **Save and Close** – To save and close the form.

Name	Attributes
Policy Name	Name of the performance monitoring policy.
Priority	Priority of the policy. This can be any positive integer.
Active	Enabled by default. If this is unchecked, all the elements associated with the policy will be removed from scheduling or associated with the next priority policy.
Schedule Interval (in minutes)	Time interval at which the metrics will be collected. Default interval is 15 minutes.
Description	Description of the performance monitoring policy.

View Collectors

To view the collectors provided by SOM, from the workspaces panel go to **Configuration** > **Monitoring Settings** > **Collectors**.

Double-click a collector to view metrics associated with each collector. The metric name is displayed with its unit.

Viewing Performance Data

The performance metrics are displayed on the analysis pane in the Inventory Views. Individual metrics are grouped under tabs and displayed through charts. The data points in the charts are plotted based on the schedule interval specified in the monitoring policy.

At any given time, the graphs show the data for the last 24 hours. The metrics that are displayed on the user interface are auto-refreshed every 5 minutes. You have the flexibility to refresh the data for each individual metric. For detailed historical analysis of performance data, use the OBR reports.

Modify a Monitoring Policy

You can modify the following in an existing monitoring policy:

- Schedule – Modify the schedule of a policy.
- Priority – Modify the priority of a policy.
- Monitoring Group – Associate additional monitoring groups to a policy or remove monitoring groups from a policy.
- Node Group – Associate additional node groups to a policy or remove node groups from a policy.
- Active – Deactivate a policy or activate a policy.
- Collectors – At least one collector must be associated with a monitoring group.

To modify a monitoring policy, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Monitoring Settings** > **Monitoring Policies**. The Monitoring Policies view is displayed.
2. Select the policy that you want to modify from the table view.
3. Click **Open**. The policy is displayed in the Monitoring Policy form view.
4. Make the required changes to the policy.
5. Click one of the options to save the policy.
   - **Save** – To save the form.
   - **Save and New** – To save and open a new form.

- **Save and Close** – To save and close the form.

The Monitoring Policies view is refreshed to display the changes to the policy.

Delete a Monitoring Policy

To delete a monitoring policy, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Monitoring Settings** > **Monitoring Policies**. The Monitoring Policies view is displayed.
2. Select the policy that you want to delete from the table view.
3. Do one of the following.
   - Click **Delete**. The delete confirmation message is displayed. Click **OK** to delete the policy.
   - Click **Open**. The policy is displayed in the Monitoring Policy form view. Click . The delete confirmation message is displayed. Click **OK** to delete the policy.

Note: All associated performance collection schedules are deleted.

Managing Storage Tiers

SOM provides flexible automated rules-based assignments for categorizing storage systems, volumes and pools into storage tiers. You can define storage tiers based on rules and SOM automatically assigns the elements to tiers based on the tier definitions. A rule has attributes such as type of storage, disk size, disk type, replication type, RPM, RAID levels and such others that you can use to define it. You can assign priority to each tier based on which the system runs these rules.

By default, storage tiers are enabled in SOM. If you want to disable storage tiers, follow these steps:

1. Open the `custom.properties` file located in the following directory:
   - *Windows*: `%OvDataDir%\conf\som`
   - *Linux*: `/var/opt/OV/conf/som`
2. Set the following property to `N`:
   `StorageTierEnabled=N`
3. Save the `custom.properties` file.
4. Restart the `somjboss` service.

Manual Association of Elements

Apart from rule-based associations, SOM also supports definition of manual rules in the form of manual association of elements to tiers. You can add or delete elements from storage tiers as exceptions to the defined rule.

The following points elaborates how manual associations are handled by the system:

- Manual associations of elements to tiers always override the rule-based assignments.
  Example

  Assume you created a dynamic storage tier that requires its element to have a disk size of more than 900 GB. Then, you manually add an element that has a disk size of less than 900 GB to the storage tier.

  During the next refresh of rule-based membership, elements that do not fit the criteria for being a member of the storage tier are removed, except for the elements you manually added. The elements you manually added stay members of the storage tier even if they do not meet the criteria of the storage tier.

- When you add elements manually to a tier and if the elements belong to other storage tiers because of rule-based assignments, they are removed from other tiers automatically without having to run the tier rules again.
  Example

  Assume Volume 1 is a member of Tier 1 and Tier 2 dynamically due to rule-based assignments. Assume you create Tier 3 and manually add Volume 1 to it. Volume 1 is automatically removed from Tier 1 and Tier 2 membership with immediate effect. You do not have to wait for the next refresh of the rules or run the rules manually for the changes to take effect.

- When you associate an element to a tier manually, the element is not available for selection and addition to another tier.
  Example

  If you add an element X manually to Tier 1, then element X is not available for selection for manual addition when creating other tiers.

- When you are modifying a tier, elements that are already mapped to the tier by the dynamic rules are not available for selection for manual addition.

How Do Rule-Based Assignments Work?

When new elements are discovered in the environment, the system dynamically assigns these elements to tiers based on the tier definitions.

The rules are run based on priorities. A priority determines the order in which a tier is picked up by the system for a refresh. A priority with lower numeric value has a higher priority. For example, a storage tier with priority 0 will be updated first before a storage tier with priority 5. If an element belongs to two tiers , then the element belonging to the tier with higher priority will remain during dynamic rule evaluation and the element belonging to the lower rule will be removed.

Typically tier memberships are updated in the following situations:

- At the end of successful data collection
  When data collection is completed for a storage system, the tier rules applicable to that storage system are evaluated to update tier membership.
- On saving a tier rule definition
  Any manual assignments of elements to that tier rule will be updated immediately.
- On manual execution of tier rules
  You can manually execute the tier rules using the option “Run Rule for All Tiers”. This option runs all the tier rules simultaneously in the order of their priority.
- As a rule of thumb, before doing any data export of the tiers perform the “Run Rule for All Tiers” so that the system data with respect to all the tiers is updated.

There are two important timestamp related attributes displayed in the Storage Tiers view:

- **Last Modification Time** – Denotes the last time the tier was modified.
- **Last Rule Run Time**– Denotes the last time the tier rule was run.

If the Last Rule Run Time is greater than the Last Modification Time, it indicates that the tier rule was run after the last edit of the tier rule and the changes to the tier rule are effective.

Best Practices for Creating Storage Tiers

The following are some best practices to follow while creating storage tiers:

- Create the storage tier to match the attributes of the elements that you want to monitor. Elements that match the criteria will be automatically added.
- When you assign priorities to tiers, do not use numbers in a continuous sequence such as 0, 1, 2, 3, and so on. Ideally use multiples of a positive integer to set the priorities. For example, use multiples of 5 as the priority such as Priority 5 for Tier 1, Priority 10 for Tier 2, and so on. This way, when you want to modify the priority of one tier you do not have to modify the priorities of all the other tiers that have priorities in immediate succession.
- Before you export any data of the tiers , ensure that you run the rule for all tiers so that the system data with respect to all the tiers is updated.

Create a Storage Tier

Use the Storage Tiers Wizard to create storage tiers. Launch the wizard from the Storage Tiers folder in the Configuration workspace. You can access any page of the wizard after launching it, however, you can save the tier only after you have entered all the mandatory fields for the storage tier.

To create a storage tier, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Storage Tiers** > **Storage Tier Wizard**. The Welcome to Storage Tier Wizard page is displayed on the right pane.
2. Click **Next**. The Storage Tier Properties page is displayed.
3. Enter the following information on the Storage Tier page.

Attribute	Description
Name	Name of the storage tier.
Description	Enter text that describes the storage tier.
Priority	Enter any positive integer.
Active	Enabled by default. Clear the selection to disable the rule.

4. Click **Next**. The Storage Systems page is displayed.
5. On the Storage Systems page
   a. Select one of the options for **Storage Systems**:
      - **All** – Use this option to associate all storage systems that are discovered to the tier.
      - **Selected** – Use this option to associate only selected storage systems to the tier. You can select storage systems based on **Vendors**, **Models**, or **Systems**. Use the selection buttons to make your selections.

   c. Select the **Storage System Type** from the drop-down list.

   d. Select **Offering** from the drop-down list.
6. Click **Next**. The Storage System Attributes page is displayed.
7. Define the rule for the storage tier using the following disk attributes:
   - Select the **Single Rule** or **Double Rule** option to specify the Disk Size. The drop-down provides options such as >, <, >=, or <= and MiB, GiB, or TiB. Enter a value in the text box to specify the disk size.
   - Specify disk attributes using the options – **Disk RPM**, **Disk Types**, **RAID Levels**, and **Replication Types**. The values listed here are values that are populated after successful data collection.
8. Click **Next**. The Add/Remove Elements from Tier page is displayed.
9. Click any of the tabs – **Storage Systems**, **Storage Pools**, or **Storage Volumes** – to browse for the elements that you want to add or delete from the tier.
   - To add an element, select the element from the table and click .
   - To delete an element from the tier, select the element from the table on the lower pane and click  to delete it from the tier.
10. Click **Next**. The Summary page is displayed.

11. Review your choices and click **Save & Close** to save the tier.

Modify a Storage Tier

You can modify the following attributes of a storage tier:

- Disk attributes such as disk RPM, disk type, RAID levels, and replication types, rule conditions, or priority of a storage tier.
- Activate or deactivate a storage tier.
- Add elements to a storage tier as an exception to the rule.
- Delete elements from a storage tier as an exception to the rule.

To modify a storage tier, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Storage Tiers** > **Storage Tiers**. The Storage Tiers view is displayed.
2. Select the storage tier that you want to modify from the table view.
3. Right-click and select Edit Tier Rule. The storage tier is displayed in the wizard view.
4. Make the required changes to the storage tier.
5. Click **Save & Close** to save changes to the storage tier.

Delete a Storage Tier

To delete a storage tier, follow these steps:

1. From the workspace navigation panel, click **Configuration** > **Storage Tiers** > **Storage Tiers**. The Storage Tiers view is displayed.
2. Select the storage tier that you want to delete from the table view.
3. Right-click and select . The selected storage tier is deleted.

Configuring Incidents

Incidents are information that SOM considers important to bring to your attention regarding your storage environment. SOM provides a set of incident configurations for the following:

- Traps generated from an SNMP agent (SNMPv1, SNMPv2c, or SNMPv3)
- Management incidents that are generated by SOM

SOM provides one centralized location, the incident views, where the management events and SNMP traps are visible to your team. You and your team can easily monitor the incidents and take appropriate action to preserve the health of your storage environment.

As a SOM administrator, you can configure the following in SOM:

- Communication settings to enable SOM to retrieve information required for processing SNMP traps from devices. For more information, see "Managing Communication Configurations" on page 228.
- Modify the incident configurations provided by SOM. For more information, see "Manage Incident Configurations" below.

Manage Incident Configurations

SOM provides one centralized location, the Incidents folder in the Configuration workspace, where the management events and SNMP trap incident configurations are visible to SOM administrators. These configurations enable you to control which SNMP traps are considered important enough to show up as incidents. You can also configure how incidents that are generated by SOM are displayed.

You can modify the incident configurations provided by SOM or create new incident configurations. To do so, see the following topics:

- "Enable or Disable Incidents" on the next page
- "Configure Incident Logging" on the next page
- "Configuring SNMP Traps" on page 204
- "Configuring Management Event Incidents" on page 205
- "Correlate Pairwise Incidents" on page 207
- "Basic Settings" on page 210
- "Suppression Settings" on page 212
- "Enrichment Settings" on page 213
- "Dampening Settings" on page 214
- "Deduplication Settings" on page 215
- "Rate Settings" on page 221
- "Transition Action Settings" on page 226

View Incident Configurations

SOM provides extensive configurations for SNMP traps, management events, and pairwise configurations.

To see the incident configurations provided by SOM, follow these steps:

1. Navigate to the **Incidents** folder:
    a. From the workspace navigation pane, select the **Configuration** workspace.
    b. Expand the **Incidents** folder.
2. Select the required configuration type:

- SNMP Trap Configurations
- Management Event Configurations
- Pairwise Configurations

The configurations are displayed in a table.

> **Tip:** Each row in a table displays data about one configuration. You can reduce the amount of information displayed by applying column filters. You can also export the contents of a table view for use in other applications.

3. Double-click a row to view configuration details.

Enable or Disable Incidents

You can enable or disable incidents. All incidents configurations provided by SOM are enabled by default.

You may want to selectively disable incidents when you have a scheduled downtime on devices in your storage environment.

To enable or disable incidents, follow these steps:

1. Navigate to the **Incidents** folder.
   a. From the workspace navigation panel, select the **Configuration** workspace.
   b. Expand the **Incidents** folder.
2. Select the required incident configuration: **SNMP Trap Configurations**, **Management Event Configurations**, or **Pairwise Configurations**.
3. Double-click a row.
4. In the **Basics** pane, do the following:
   - Select the **Enabled** check box to enable the incident.
   - Clear the **Enabled** check box to disable the incident.
5. Click **Save and Close** to save your changes.

Configure Incident Logging

You can configure incident logging so that SOM writes the incoming incident information into the `incident.csv` file. This feature is useful when you want to track and archive incident history.

The `incident.csv` file is located as follows on the SOM management server.

Windows

`%OvDataDir%\log\nnm\incident.csv`

Linux

`%OvDataDir%/log/nnm/incident.csv`

To configure incident logging, follow these steps:

1. Navigate to the **Incidents** folder.
   a. From the workspace navigation panel, select the **Configuration** workspace.
   b. Expand the **Incidents** folder.
2. Select **Incident Configuration**.
3. In the **Incident Logging Configuration** tab, provide the required information (see General Configuration and Log File Configuration).
4. Click **Save and Close** to save your changes.

General Configuration

Attribute	Description
Enable Incident Logging	If enabled, SOM logs incoming incident information to the `incident.csv` file. If disabled, SOM does not log the incident information.

Log File Configuration

Attribute	Description
Enable Compression	If enabled, SOM saves the `incident.csv` file in compressed (`.gz`) format. If disabled, incident information is saved in the uncompressed format.
Maximum File Size (MB)	Specify the maximum amount of disk space in megabytes that SOM should use for the `incident.csv` file. The default value is 128 megabytes. **Note:** After the maximum file size is reached, the log file is renamed to `incident.csv.<gz>.old` and a new `incident.csv` file is created. If an `incident.csv.<gz>.old` file exists, it is overwritten.
Logging Interval (ms)	Specify the time interval for SOM to log incident information. The default value is 6 seconds (6000 milliseconds). **Tip:** To optimize performance, use a longer Logging Interval with a larger Maximum Number of Incidents.

Log File Configuration, continued

Attribute	Description
	Note the following: • The minimum value is 0.01 second (10 milliseconds). • The maximum value is 1 minute (60000 milliseconds).
Maximum Number of Incidents per Logging Interval	Specify the maximum number of incidents to be logged. The default value is 1024. **Tip:** To optimize performance, use a longer Logging Interval with a larger Maximum Number of Incidents.

Configuring SNMP Traps

SOM provides default configurations to manage incidents for all supported storage devices. You can open each SNMP trap configuration to view or modify its details. You can also enable or disable an SNMP trap configuration.

Note: Make sure that the following requirements are satisfied to successfully receive and process SNMP traps from supported devices:

- Configure SOM as a trap receiver in supported devices to receive SNMP traps from those devices. For more information, see the device manufacturer's documentation.
- Configure communication settings to enable SOM to establish communication with the devices in your SAN environment. For details, see "Managing Communication Configurations" on page 228.

To configure incidents originating from SNMP traps, follow these steps:

1. Navigate to the **Incidents** folder:
   a. From the workspace navigation pane, select the **Configuration** workspace.
   b. Expand the **Incidents** folder.
2. Select **SNMP Trap Configurations**.
3. Double-click a row.
4. Make your configuration choices (see table).
5. Click **Save and Close** to save your changes and return to the previous form.

Tasks for SNMP Trap Configuration

Settings	Purpose
"Basic Settings" on page 210	To configure the **Basics** pane of the SNMP Trap Configuration

Tasks for SNMP Trap Configuration, continued

Settings	Purpose
	form.
"Suppression Settings" on page 212	To view the **Suppression** settings for the SNMP Trap Configuration.
"Enrichment Settings" on page 213	To view the **Enrichment** settings for the SNMP Trap Configuration.
"Dampening Settings" on page 214	To view the **Dampening** settings for SNMP Trap Configuration.
"Deduplication Settings" on page 215	To view the **Deduplication** settings for SNMP Trap Configuration.
"Rate Settings" on page 221	To view the **Rate** settings for the SNMP Trap Configuration.
"Transition Action Settings" on page 226	To view the **Actions** settings for SNMP Trap Configuration.

Configuring Management Event Incidents

SOM provides default configurations for management event incidents. You can open each management event configurations to view or modify its configuration details. You can also enable or disable a management event incident configuration.

To configure incidents originating from management events, follow these steps:

1. Navigate to the **Management Event Configuration** form:
   a. From the workspace navigation pane, select the **Configuration** workspace.
   b. Expand the **Incidents** folder.
   c. Select **Management Event Configurations.**
2. Make your configuration choices (see table).
   a. To add a management event configuration, click **New** , and continue.
   b. To edit a management event configuration, double-click a row, and continue.
   c. To delete a management event configuration, click **Delete**.
3. Click **Save and Close** to save your changes and return to the previous form.

Tasks for Management Event Incident Configuration

Task	How
"Basic Settings" on page 210	Use the **Basics** pane of the Management Event Configuration form.

About Pairwise Configurations

Often two incidents have a logical relationship to each other, for example, `CiscoLinkDown` followed by `CiscoLinkUp`. There is no need for both incidents to take up room in your Incident view. Nesting the two together helps you do your job quickly and efficiently.

Use the Pairwise Configuration to pair the occurrence of one incident with another subsequent incident. When the second incident in the pair occurs, the first incident becomes a correlated child incident within the parent incident.

SOM provides default pairwise configurations for devices that SOM supports. You can view the default pairwise configurations in a table when you navigate to the Pairwise Configurations form in SOM.

When using Pairwise Configurations, note the following:

- You can use Payload Filters (for example, with trap varbinds) to identify the first and second incidents in a Pairwise Configuration.
- You can specify the same incident (for example, the same trap OID) as both the first and second incident configuration for a Pairwise Configuration.
- Using the Payload Filter to distinguish the first and second incidents (the first could represent a non-normal state and the second a normal state), different instances of the same incident configuration can cancel one another.
- You can also set up the Payload Filters such that the same incident instance cancels itself.
- You can use the same incident configuration in multiple Pairwise Configurations. For example:
  - Incident configuration A cancels both incident configuration B and incident configuration C
  - Incident configuration A cancels incident configuration B and incident configuration B cancels incident configuration C.
- Single incident instance can cancel multiple incident instances (for example, one Link Up trap cancels multiple instances of a Link Down trap).

  > **Note:** If multiple Link Up/Link Down trap pairs are received within a 30 seconds, SOM investigates only once.

- Use the Duration time to specify the time in which the second incident configuration cancels the first incident configuration. This Duration is calculated from the `originOccurrenceTime` of the second incident backwards in time, canceling any number of first incidents within the Duration specified.
- You can also specify whether to delete any incidents that were canceled according to the Pairwise Configuration and that occurred within the time period specified by the Duration attribute.
- When matching incidents, SOM automatically takes into account the following values:
  - **SNMP Trap incidents**. SOM takes into account from which device the trap originated using the `cia.address` value of the source address of the trap.
  - **Management Event incidents**. SOM takes into account the name of the incident's Source Object and Source Node.

    Tip: SOM displays the Name value used to identify the Source Node and Source Object in the **Source Node** and **Source Object** attribute for each incident in the Incident form.

Tip: When configuring the Matching Criteria, you do not need to specify any of the cia names that SOM automatically takes into account. See "Matching Criteria Configuration Form (Identify Incident Pairs)" on page 537 for more information.

Related Topics:

"Pairwise Incidents Prerequisites" on page 539

"Correlate Pairwise Incidents" below

Correlate Pairwise Incidents

Use the Pairwise Configuration to pair the occurrence of one incident with another subsequent incident. See "About Pairwise Configurations" on page 524 for more information.

To configure incident pairs, follow these steps:

1. Complete the steps in "Pairwise Incidents Prerequisites" on page 539 so you know exactly which two incidents or traps belong to this logical pair.
2. Navigate to the **Pairwise Configurations** form.
   a. From the workspace navigation panel, select the **Configuration** workspace.
   b. Expand the **Incidents** folder.
   c. Select **Pairwise Configurations**.
   d. Do one of the following:

- To create a new pair configuration, click **New**, and continue.
- To edit or view an existing pair configuration, double-click a row, and continue.
- To delete a pair configuration, select a row and click **Delete**.

3. Provide the basic definition of the pair of incidents for this correlation (see table).
4. When matching incidents, SOM automatically takes into account the following values:
   - **SNMP Trap incidents:** SOM takes into account from which device the trap originated using the `cia.address` value of the trap's source address.
   - **Management Event incidents:** SOM takes into account the name of the incident's Source Object and Source Node.

     Tip: SOM displays the Name value used to identify the Source Node and Source Object in the **Source Node** and **Source Object** attribute for each incident in the Incident form.

   Some incident pairs require additional details to verify an accurate match.
5. *Optional*. Navigate to the **First Incident Payload Filter** and **Second Incident Payload Filter** tabs, and specify the payload filter to use when identifying a valid pair of incidents. See "Payload Filter Details" on page 539.
6. *Optional*. Navigate to the **Matching Criteria** tab, and provide one or more custom incident attribute sets for SOM to use as a filter when identifying a valid pair of incidents. See "Matching Criteria Configuration Form (Identify Incident Pairs)" on page 537.

   Tip: When configuring the Matching Criteria, you do not need to specify any of the ciaNames that SOM automatically takes into account . See "Matching Criteria Configuration Form (Identify Incident Pairs)" on page 537 for more information.
7. Click **Save and Close** to save your changes and return to the previous configuration form.

The next time the two incidents in this pair are generated, the first one becomes a Child Incident of the second one. See "About Pairwise Configurations" on page 524 for an example.

Pairwise Configuration Definition

Attribute	Description
Name	The name is used to identify the pairwise configuration and must be unique. Use a name that will help you to remember the purpose for this pairwise configuration. Maximum length is 64 characters. Alpha-numeric characters are permitted. No spaces are permitted.

Pairwise Configuration Definition , continued

Attribute	Description
Enabled	In the **Basics** group, verify that **Enabled** is selected.
First Incident Configuration	Identify the incident in the pair that would occur first in the logical sequence. Click the Lookup icon and select **Quick Find**. Choose the name of one of the predefined incident configurations. This first incident becomes the child incident when the second (parent) incident occurs. For example, in the CiscoLinkDownUp pairwise configuration, if a CiscoLinkUp (second incident) occurs after a CiscoLinkDown (first incident), the CiscoLinkDown is canceled and correlated as a child incident under the CiscoLinkUp incident.
Second Incident Configuration	Identify the incident in the pair that would occur second in the logical sequence. Click the Lookup icon and select **Quick Find**. Choose the name of one of the predefined incident configurations. This Second Incident becomes the Parent Incident if it occurs after the First Incident. For example, in the CiscoLinkDownUp Pairwise configuration, if a Cisco Link Up (Second Incident) occurs after a Cisco Link Down (First Incident), the Cisco Link Down is cancelled and correlated as a Child Incident under the Cisco Link Up.
Description	*Optional*. Explain the purpose of your pairwise configuration for future reference. Type a maximum of 1024 characters. Alpha-numeric, spaces, and special characters are permitted.
Author	Indicates who created or last modified the Correlation Rule. **Caution:** If the Author attribute value is **HPE Storage Operations Manager**, any changes are at risk of being overwritten in the future. • Click **Lookup** and select **Show Analysis** to display details about the currently selected Author. • Click **Quick Find** to access the list of existing Author values. • Click **New** to create an Author value.
Duration	SOM uses the value you enter to determine the duration window in which it correlates the Pairwise incidents you specify. During the timeframe specified, SOM enables a single (parent) incident to cancel multiple (child) incidents.

Pairwise Configuration Definition , continued

Attribute	Description
	The Duration is calculated from the `originOccurrenceTime` of the parent incident backwards in time, canceling any child incidents within the Duration specified. Note the following: • By default, the Duration value is 0 (zero). When the Duration value is 0, SOM finds the most recently occurring incident that matches the First Incident specified in the Pairwise configuration, regardless of time. See First Incident Configuration for more information. • The maximum duration value is 365 days.
Delete when Canceled	When enabled , after the Duration is reached, SOM deletes any incidents that were canceled according to the Pairwise configuration and that occurred within the time frame specified by the Duration attribute. When disabled, SOM cancels the pairwise incidents as configured, but does not delete them.

Basic Settings

The Basics pane in an incident specifies general information for an incident configuration, including the name, severity, and message. Table 1 and table 2 provide descriptions for attributes in the Basics pane.

Note: To configure incident logging, see "Configure Incident Logging" on page 202.

Basic Attributes for SNMP Trap and Management Event Configurations

Name	Description
Name	Displays the name for the incident configuration.
SNMP Object ID (not for Pairwise Configurations)	Displays the SNMP Object ID.
Enabled	Verify that **Enable** is selected for each configuration you want to use.

Basic Attributes for SNMP Trap and Management Event Configurations , continued

Name	Description
Root Cause (only for SNMP Trap Configurations)	Select this check box to display the SNMP trap incident as a Root Cause incident. SNMP traps normally appear as symptoms rather than root cause incidents.
Category and Family	Specify category and family attribute values for organizing your incidents.
Severity	Specify the incident severity. The incident severity represents seriousness calculated for an incident.
Message Format	Displays the default message for the incident configuration.
Description	Use the Description attribute to provide additional information you would like to store about the current incident configuration. Type a maximum of 4000 characters. Alpha-numeric and special characters are permitted.
Author	Use the Author attribute to indicate who created or last modified the incident configuration. **Note:** If the Author attribute is HPE Storage Operations Manager, upgrading SOM might overwrite your changes.

Basic Attributes for Pairwise Configurations

Name	Description
Name	Displays the name for the incident configuration.
Enabled	Verify that **Enable** is selected for each configuration you want to use.
First Incident Configuration	Identifies the incident in the pair that would occur first in the logical sequence. Click the Lookup icon and select Quick Find. Choose the name of one of the predefined incident configurations.
Second Incident Configuration	Identifies the incident in the pair that would occur second in the logical sequence. Click the Lookup icon and select Quick Find. Choose the name of one of the predefined incident configurations.
Description	Use the Description attribute to provide additional information you would like to store about the current incident configuration.

Basic Attributes for Pairwise Configurations , continued

Name	Description
Author	Use the Author attribute to indicate who created or last modified the incident configuration. **Note:** If the Author attribute is HPE Storage Operations Manager, any changes are at risk of being overwritten in the future.
Duration	SOM uses the value you enter to determine the duration window in which it correlates the Pairwise incidents you specify. During the time frame specified, SOM enables a single (parent) incident to cancel multiple (child) incidents. The Duration is calculated from the origin Occurrence Time of the parent incident backwards in time, canceling any child incidents within the Duration specified. **Note:** By default, the Duration valus is 0 (zero). When the Duration value is 0, SOM finds the most recently occurring incident that matches the First Incident specified in the Pairwise configuration, regardless of time. The maximum duration value is 365 days.
Delete When Canceled	When enabled , after the Duration is reached, SOM deletes any incidents that were canceled according to the Pairwise configuration and that occurred within the time frame specified by the Duration attribute. When disabled, SOM cancels the pairwise incidents as configured, but does not delete them.

Suppression Settings

SOM can suppress incidents based on the payload filters. SOM tries to match the details in the incoming SNMP traps with the payload filter definitions. If there is a match, SOM suppresses the incidents. For example, SNMP trap varbind names and values can be used as payload filters. You might want SOM to suppress a particular status change notification trap. Suppressed incidents are not displayed in the Incident Browsing views.

> **Note:** SOM provides default suppression settings for certain incident configurations. These settings cannot be changed.

The following table provides descriptions for attributes in the Suppression tab.

Suppression Attributes

Name	Description
Enabled	Indicates whether suppression is enabled or disabled for the incident configuration.
Payload Filter	The Payload Filter area displays the filter expression used for the incident configuration.

Enrichment Settings

SOM can fine tune and enhance incidents based on the enrichment settings. SOM provides certain incident configurations with suitable enrichment settings.

Note: SOM provides default enrichment settings for certain incident configurations. These settings cannot be changed.

The following table provides descriptions for attributes in the Enrichment tab.

Enrichment Attributes

Name	Description
Enabled	Indicates whether enrichment is enabled or disabled for the incident configuration.
Enrichments	Double-click a row to view the enrichment tab.
Category	Displays the default category for this incident configuration. See "Incident Form: General Tab" on page 529 for more information on each Category attribute.
Family	Displays the default family for this incident configuration. See "Incident Form: General Tab" on page 529 for more information on each Family attribute.
Severity	The incident severity represents the seriousness calculated for the incident. Possible values are described below: • **Normal:** Indicates there are no known problems related to the associated object. This Severity is meant to be informational. Generally, no action is needed for these incidents. • **Warning:** Indicates there might be a problem related to the associated object. • **Minor:** Indicates SOM has detected problems related to the associated object that require further investigation. • **Major:** Indicates SOM has detected problems related to the associated object to be resolved before they become critical. • **Critical:** Indicates SOM has detected problems related to the associated object that

Enrichment Attributes , continued

Name	Description
	require immediate attention.
Priority	Priority helps in communicating the urgency of resolving the selected incident. The lower the number the higher the priority.
Correlation Nature	Correlation Nature helps in customizing the Correlation Nature for this incident configuration. Possible values include: • Root Cause (or User Root Cause) • Secondary Root Cause • Symptom • Stream Correlation • None • Info • Service Impact • Dedup Stream Correlation • Rate Stream Correlation For more information, see "Incident Form: General Tab" on page 529
Message Format	Displays the default message format used in this incident configuration.
Assigned To	Used to specify the owner of any incident generated for this incident configuration. Click the Lookup icon and select Quick Find to select a valid user name. **Note:** You can also begin to type a valid user name and use the auto-complete feature to select the user name of interest.
Description	Displays the default description for the enrichment.
Payload Filter	The Payload Filter area displays the filter expression used for the incident configuration.

Dampening Settings

SOM can delay (dampen) appearance of an incident within incident views.

When using dampening configuration, note the following:

- For all incident configurations except deduplication and rate incidents, if the dampened Incident is closed before the dampen interval has passed, SOM deletes the incident. If the Incident is the Root Cause Incident, SOM also deletes any child incidents
- SOM always retains the parent deduplication or rate incident even if its child incidents are closed within the dampen interval and subsequently deleted.
- Deduplication incidents and child incidents inherit dampening settings.
- If an incident is a Root Cause Incident and a child incident's dampen Interval is less than the parent incident's dampen interval, SOM holds any child incidents until the dampen Interval for the parent incident has passed or until the parent incident is closed and subsequently deleted.
- To make sure SOM handles both Incidents in a pairwise configuration the same, configure the same dampen interval for each Incident in a pairwise incident configuration. For more information, see "Correlate Pairwise Incidents" on page 207
- After the dampen interval has passed, SOM changes the Lifecycle State to `REGISTERED`.
- You can use a payload filter to fine tune the incidents you want to dampen.

Note: SOM provides default dampening settings for certain incident configurations. These settings cannot be changed.

The following provides descriptions for attributes in the Dampening tab.

Dampening Attributes

Name	Description
Enable	Indicates whether enrichment is enabled or disabled for the incident configuration.
Hour	Specifies the number of hours to be used for the Dampen Interval.
Minutes	Specifies the number of minutes to be used for the Dampen Interval.
Seconds	Specifies the number of seconds to be used for the Dampen Interval.
Payload Filter	Payload Filter Editor enables you to create expressions that further refine the filters used to select the incidents to be suppressed, enriched, or dampened. For information about using payload filters, see "Payload Filter Details" on page 539.

Deduplication Settings

The deduplication configuration determines what values SOM should match to detect when an Incident is a duplicate.

Note the following:

- Suppression, Enrichment, and Dampening are not supported for Deduplication incidents.
- SOM applies only one deduplication configuration per incident. If SOM generates an incident using a specified deduplication configuration, SOM continues to correlate duplicate incidents using the original configuration. To use a different deduplication configuration for an incident, first delete the current deduplication incident (created using the original deduplication configuration). SOM generates the next deduplication incident according to the new deduplication configuration settings.
- SOM continues to update the duplicate count regardless of an incident's lifecycle state. For example, if an incident's **Lifecycle State** is set to **Closed**, the duplicate count continues to be incremented. This behavior helps you identify situations in which the incident is not yet fixed. Take note if the Duplicate Count is incremented after a lengthy time period has elapsed, which might indicate there is a new problem with the node, interface, or address.
- Each time you stop and restart `somjboss`, any incidents that have not yet been correlated or persisted are lost. This means that after a restart of `somjboss`, an incoming incident might not be correlated as expected. For example, after a restart of `somjboss`, a duplicate incident might not be correlated under its original parent incident. Instead, a new parent incident might be generated.
- If a Duplicate Correlation Incident is dampened, note the following:
  - Duplicate Correlation Incidents inherit the Dampening settings from its Correlated Children.
  - SOM always retains the Parent Duplicate Correlation incident, even if its Child Incidents are Closed and subsequently deleted.

> **Note:** SOM provides default deduplication settings for certain incident configurations. These settings cannot be changed.

The following table provides descriptions for attributes in the Deduplication tab.

Deduplication Attributes

Name	Description
Enabled	Use this attribute to temporarily disable an incident's deduplication configuration: **Disable** ☐ = Temporarily disable the selected configuration. **Enable** ☑ = Enable the selected configuration. **Note:** After a deduplication configuration is enabled, SOM increments the **Duplicate Count** for an associated incident regardless of the **Lifecycle State** value. For example, if an incident's **Lifecycle State** is set to **Closed**, the duplicate count continues to be incremented.

Deduplication Attributes, continued

Name	Description
Count	Specifies the number of duplicate incidents for the current configuration that SOM stores at one time. For example, if the Count is 10, after SOM receives 10 duplicate incidents, SOM deletes the first (oldest) duplicate incident and keeps the eleventh. (SOM stores ten maximum.) **Note:** By default, SOM updates the Duplicate Count every 30 seconds. This interval cannot be changed.
Hours	Used with the Minute and Second Intervals to specify the time that must elapse before a new duplicate incident is generated for this incident configuration. For example, if the Hour Interval value is 1, and no Minute or Second Intervals are specified, and the duplicate incident is not generated within one hour, SOM generates a new duplicate incident the next time it occurs.
Minutes	Used with the Hour and Second interval to specify the time that must elapse before a new duplicate incident is generated for this incident configuration. For example, if the Minute Interval is 30 and no Hour or Second Intervals are specified, and the duplicate incident is not generated within 30 minutes, SOM generates a new duplicate incident the next time it occurs.
Seconds	Used with the Hour and Minute Intervals to specify the time that must elapse before a new duplicate incident is generated for this incident configuration. For example, if the Second Interval is 120 and no Hour or Minute Intervals are specified, and the duplicate incident is not generated within 120 seconds, SOM generates a new duplicate incident the next time it occurs.
Parent Incident	Used to specify the Incident Configuration that will be the Parent Incident for the incident you are configuring. For example, you might have created a Management Event Incident Configuration that could be used as the **Parent Incident** for SNMP Trap Incidents. When specifying the **Parent Incident**, you have the following options: • When you want to use a configuration that SOM provides, use the default **Duplicate Correlation** incident configuration . If you select this option, the incident message for the Parent Incident begins as follows: `Duplicate Correlation for <incident_configuration_name>` For example if you are configuring a **Node Down** incident and select **Duplicate Correlation** as the **Parent Incident**, the Parent Incident message begins with: **Duplicate Correlation for Node Down**. Each **Node Down** incident that is a duplicate then appears correlated under the **Duplicate Correlation for Node Down**

Deduplication Attributes, continued

Name	Description
	incident. • SOM also enables you to customize the Parent Incident for a given deduplication scenario. If you have created a Management Event Incident Configuration to use for this deduplication scenario, select the Management Event Incident Configuration that you have created.

Comparison Criteria	Specify the attribute values that must match before the incident is identified as a duplicate. The possible attributes consist of the following choices.
	• **Name** - The **Name** attribute value from the Incident form: General tab.
	• **Name CIA** - Represents any of the following items configured as a Parameter Value.
	  ◦ An SNMP varbind Object ID
	  ◦ An SNMP varbind position number
	• **Name SourceNode** - The **Source Node** attribute value from the Basics attributes listed on the Incident form. The Source Node value is the IP Address or Name of the node for which the incident was generated.
	Note: The Source Node must be stored in the SOM database.
	• **Source Object** - The **Source Object** attribute value from the Basics attributes listed on the Incident form.
	Note: The Source Object must be stored in the SOM database.
	Caution: Each attribute value in the option you select must match before the incident is identified as a duplicate. For example, if you select **Name**, only the Incident Name value must match. If you select **Name SourceNode SourceObject CIA**, the Incident Name, Source Node, Source Object, and all Custom Incident Attribute values that you configure as a Parameter Value must match before SOM identifies the incident as a duplicate.
	Selecting an option that includes CIA enables you to further refine the deduplication criteria. For example, you might want to configure deduplication for incidents with CIA values that specify the same State attribute value for a particular network object.
	For a description of each Comparison Criteria option, click here.

Deduplication Attributes, continued

<table>
<tr><th>Name</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>Comparison Criteria</th><th>Description</th></tr>
<tr><td>Name</td><td>Value of the Name attribute from the Incident form: General tab must match.</td></tr>
<tr><td>Name CIA</td><td>Each of the following values must match:

Name attribute from the Incident form: General tab
CIA - Represents the Value associated with any of the following items:

Name of a Custom Incident Attribute (CIA) provided by SOM. (See the Incident form: Custom Attributes tab.)
An SNMP varbind Object ID
An SNMP varbind position number

</td></tr>
<tr><td>Name SourceNode</td><td>Note: Select this option only if the Source Node is stored in the SOM database.

Each of the following values must match:

Name attribute from the Incident form: General tab
The Source Node attribute value from the Basics attributes listed on the Incident form.
</td></tr>
<tr><td>Name SourceNode CIA</td><td>Note: Select this option only if the Source Node is stored in the SOM database.

Each of the following values must match:

Name attribute from the Incident form: General tab
The Source Node attribute value from the Basics attributes listed on the Incident form.
CIA - Represents the Value associated with any of the following items:

The Value attribute from the Incident form: Custom

</td></tr>
</table>
</td></tr>
</table>

Deduplication Attributes, continued

<table>
<tr><th>Name</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>Comparison Criteria</th><th>Description</th></tr>
<tr><td></td><td>
Attributes tab
• An SNMP varbind Object ID
• An SNMP varbind position number
</td></tr>
<tr><td>Name SourceObject</td><td>
Note: Select this option only if the Source Object is stored in the SOM database.
Each of the following values must match:
• Name attribute from the Incident form: General tab
• The Source Object attribute value from the Basics attributes listed on the Incident form.
</td></tr>
<tr><td>Name SourceObject CIA</td><td>
Note: Select this option only if the Source Object is stored in the SOM database.
Each of the following values must match:
• Name attribute from the Incident form: General tab
• The Source Object attribute value from the Basics attributes listed on the Incident form
• CIA - Represents the Value associated with any of the following items:
  ◦ The Name attribute from the Incident form: Custom Attributes tab
  ◦ An SNMP varbind Object ID
  ◦ An SNMP varbind position number
</td></tr>
<tr><td>Name SourceNode SourceObject</td><td>
Note: Select this option only if the Source Node and Source Object are stored in the SOM database.
</td></tr>
</table>
</td></tr>
</table>

Deduplication Attributes, continued

<table>
<tr><th>Name</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>Comparison Criteria</th><th>Description</th></tr>
<tr><td></td><td>Each of the following values must match:
Name attribute from the Incident form: General tab
The Source Node attribute value from the Basics attributes listed on the Incident form
The Source Object attribute value from the Basics attributes listed on the Incident form</td></tr>
<tr><td>Name SourceNode SourceObject CIA</td><td>Note: Select this option only if the Source Node and Source Object are stored in the SOM database.

Each of the following values must match:
Name attribute from the Incident form: General tab
The Source Node attribute value from the Basics attributes listed on the Incident form
The Source Object attribute value from the Basics attributes listed on the Incident form
CIA - Represents the Value associated with any of the following items configured as a Parameter Value
The Name attribute from the Incident form: Custom Attributes tab
An SNMP varbind Object ID
An SNMP varbind position number</td></tr>
</table>
</td></tr>
<tr><td>Deduplication Comparison Parameters</td><td>. Comparison Parameter values enable accurate identification of duplicate incidents. The values are pre-populated for supported deduplication configurations.</td></tr>
</table>

Rate Settings

Note: SOM currently does not support rate settings therefore the Rate tab is disabled in SOM for all incident confgurations.

Use rate configuration to track incident patterns *based on the number of incident re-occurrences within a specified time period*. After the count within the specified time period is reached, SOM emits a rate correlation incident and continues to update the Correlation Notes field with the number of occurrences within that rate.

Note: Suppression, enrichment, and dampening are not supported for rate incidents.

As long as your defined criteria (Count and Hours, Minutes, Seconds) is sustained, the following information is updated in the Correlation Notes of the Rate Correlation incident:

- The actual number of occurrences of incidents for that sustained rate (Count)
- The sustained time interval (Hours, Minutes, Seconds)

For example, you can set a Rate configuration to track when a link is intermittently down at least three times in 30 minutes. SOM shows the first occurrence of the rate incident in the incident view and uses Correlation Notes to update the number of incidents and time interval to reflect all the incremental incident occurrences and time periods. To continue the example, if the rate of three times in 30 minutes is sustained for 90 minutes, SOM updates the Correlation Notes to specify that 9 incidents occurred in 90 minutes.

SOM provides pre-configured Rate correlations. You can add new Rate correlations.

When you open the Incident form of the newest instance:

- On the General tab, two fields notify you that the Rate correlation is working:
  - **Correlation Nature:** Rate Stream Correlation
  - **Count:** Incremental incident occurrences
- On the **Correlated Children** tab, each incident is listed in the table.
- If a Rate Correlation Incident is dampened, note the following:
  - Rate Correlation Incidents inherit the Dampening configuration settings from its Correlated Children.
  - SOM always retains the Parent Rate Correlation Incident, even if its Child Incidents are Closed and subsequently deleted.

The following table provides descriptions for attributes in the Rate tab.

Rate Configuration Definition

Attribute	Description
Enable	Use this attribute to temporarily disable an incident's rate settings: ☐ Enabled = Temporarily disable the selected configuration.

Rate Configuration Definition , continued

Attribute	Description
	Enabled = Enable the selected configuration. If enabled, SOM actively tracks any reoccurrences of the designated incident within the time period you specify, and generates a Rate incident.
Count	Specify the number of re-occurrences required before your rate configuration starts working.
Hours	Used with the Minutes and Seconds attributes to specify the time duration within which the re-occurrences are measured.
Minutes	Used with the Hours and Seconds attributes to specify the time duration within which the re-occurrences are measured.
Seconds	Used with the Hours and Minutes attributes to specify the time duration within which the re-occurrences are measured.
Parent Incident	Click the icon and select Quick Find. Select **Rate Correlation** from the list.

Comparison Criteria	Specify which group of attributes must match before the incident is identified as a duplicate. The possible groups of attributes consist of the following choices.
	
	• **Name** attribute value from the Incident form: General tab.
	• **CIA** represents any of the following items configured as a Parameter Value:
	  ◦ The **Value** attribute from the Incident form: Custom Attributes tab
	  ◦ An SNMP varbind Object ID
	  ◦ An SNMP varbind position number
	• **Source Node** attribute value from the Basics attributes listed on the Incident form. The Source Node value is the IP Address or Name of the node for which the incident was generated.
	
	Note: The Source Node must be stored in the SOM database.
	
	• **Source Object** attribute value from the Basics attributes listed on the Incident form.
	
	For a description of each Comparison Criteria option, click here.

Rate Configuration Definition , continued

<table>
<tr><th>Attribute</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>Comparison Criteria</th><th>Description</th></tr>
<tr><td>Name</td><td>Value of the Name attribute from the Incident form: General tab must match.</td></tr>
<tr><td>Name CIA</td><td>Each of the following values must match:

Name attribute from the Incident form: General tab
CIA represents the Value associated with any of the following items configured as a Parameter Value:

Name of a Custom Incident Attribute (CIA) provided by SOM. (See the Incident form: Custom Attributes tab.)
An SNMP varbind Object ID
An SNMP varbind position number

</td></tr>
<tr><td>Name SourceNode</td><td>Note: Select this option only if the Source Node is stored in the SOM database.

Each of the following values must match:

Name attribute from the Incident form: General tab
Source Node attribute value from the Basics attributes listed on the Incident form.
</td></tr>
<tr><td>Name SourceNode CIA</td><td>Note: Select this option only if the Source Node is stored in the SOM database.

Each of the following values must match:

Name attribute from the Incident form: General tab
Source Node attribute value from the Basics attributes listed on the Incident form.
CIA - Represents the Value associated with any of the following items configured as a Parameter Value

The Value attribute from the Incident form: Custom

</td></tr>
</table>
</td></tr>
</table>

Rate Configuration Definition , continued

Attribute	Description
	<table><tr><th>Comparison Criteria</th><th>Description</th></tr><tr><td></td><td>Attributes tab - An SNMP varbind Object ID - An SNMP varbind position number </td></tr><tr><td>Name SourceObject</td><td>**Note:** Select this option only if the Source Object is stored in the SOM database. Each of the following values must match: **Name** attribute from the Incident form: General tab **Source Object** attribute value from the Basics attributes listed on the Incident form. </td></tr><tr><td>Name SourceObject CIA</td><td>**Note:** Select this option only if the Source Object is stored in the SOM database. Each of the following values must match: **Name** attribute from the Incident form: General tab **Source Object** attribute value from the Basics attributes listed on the Incident form. **CIA** - Represents the Value associated with any of the following items configured as a Parameter Value: **Name** attribute from the Incident form: Custom Attributes tab - An SNMP varbind Object ID - An SNMP varbind position number </td></tr><tr><td>Name SourceNode SourceObject</td><td>**Note:** Select this option only if the Source Node and Source Object are stored in the SOM database.</td></tr></table>

Rate Configuration Definition , continued

<table>
<tr><th>Attribute</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>Comparison Criteria</th><th>Description</th></tr>
<tr><td></td><td>Each of the following values must match:
Name attribute from the Incident form: General tab
The Source Node attribute value from the Basics attributes listed on the Incident form.
The Source Object attribute value from the Basics attributes listed on the Incident form.</td></tr>
<tr><td>Name SourceNode SourceObject CIA</td><td>Note: Select this option only if the Source Node and Source Object are stored in the SOM database.

Each of the following values must match:
Name attribute from the Incident form: General tab
Source Node attribute value from the Basics attributes listed on the Incident form
Source Object attribute value from the Basics attributes listed on the Incident form
CIA - Represents the Value associated with any of the following items configured as a Parameter Value.
Name attribute from the Incident form: Custom Attributes tab
An SNMP varbind Object ID
An SNMP varbind position number</td></tr>
</table>
</td></tr>
<tr><td>Rate Comparison Parameters</td><td>Optional. If you selected a Comparison Criteria that includes CIA, you must populate one or more rows in this table.</td></tr>
</table>

Transition Action Settings

Note: SOM currently does not support transition action settings therefore the Transition Action tab is disabled in SOM for all incident configurations.

You can configure actions to automatically run at any point in the incident lifecycle. For example, you might want to configure an action to occur when an incident of the type you are configuring is generated (**Registered**). When an incident is generated, you might want to automatically open a trouble ticket or send email or page your storage administrator. After the incident is **Closed**, you might want to automatically close the trouble ticket.

Note: Your actions will not be executed until you enable the Actions configuration by clicking Enable on the Actions tab.

You can configure actions for all incidents. Any time an incident configuration changes, the action directory is rescanned and any executable or script files (for example, Jython) are reloaded to the SOM database.

Tip: Copy any required executable or script files to the SOM actions directory before you configure an incident action. New or updated actions are loaded into SOM only when an incident configuration is updated or created.

When the defined Incident Action runs, output is logged to the `incidentActions.*.*.log` file.

The following table provides descriptions for attributes in the Actions tab.

Create Action Attributes

Attribute	Description
Lifecycle State	Select a Lifecycle State from the list.
Command Type	If you provided a Jython command, select **Jython** from the list. If you are using an executable or bat file, select **ScriptOrExecutable** from the list.
Command	Enter one of the following: • A Jython method with the required parameters • Executable command for the current operating system with the required parameters. When entering a **Command** value, note the following: • Left or right bracket ([ ]) and backtick (` Unicode character: 0060 hex = 96 dec) characters are not permitted in the **Command** attribute. If you need these characters in your shell script, place them in a shell script file and reference that file from the **Command** attribute. • **Windows only**: Shell commands are not permitted in the **Command** attribute. To use shell

Create Action Attributes, continued

Attribute	Description
	commands, place them in a shell script file and reference that file from the **Command** attribute. • Use absolute paths to executables instead of relying on the PATH variable as it might not be set correctly. • Verify that you do not have two Jython methods with the same name. Otherwise, SOM is not able to tell which is the correct method to load. • You can use the same Jython method for more than one incident configuration. • Jython (.py) files must reside in the following directory: **Note:** All the functions defined in the Jython files that reside in this directory are also accessible by SOM. The files are also executed by SOM on startup. Windows: `<Install_Dir>\ProgramData\HP\HP BTO Software\shared\nnm\actions` Linux: `/var/opt/OV/shared/nnm/actions` • When using executable files, specify the absolute path to the executable command or make sure the directory in which the executable file resides is in your PATH environment variable.

You can also configure Payload Filters to further define the filters to be used for selecting the incidents that should participate in an operation. For information about configuring Payload Filters, see "Payload Filter Details" on page 539.

Managing Communication Configurations

SOM needs to communicate with devices in your SAN environment to gather specific information before it can process SNMP traps originating from that device. Using the Communication Configuration form, you can specify the details that SOM requires to establish communication with the devices.

To configure communication settings, follow these steps:

1. Navigate to the Communication Configuration form:
   a. From the workspace navigation panel, select the **Configuration** workspace.
   b. Expand the **Incidents** folder and select **Communication Configuration**.
2. Make your configuration choices:
   - "Default SNMP Settings" on the next page
   - "Default SNMP V1/V2 Community String" on the next page

- "Node Specific Settings" on the next page

3. Click **Save and Close** to apply your changes.

Default SNMP Settings

Attribute	Description
SNMP Timeout	(Seconds:Milliseconds) Maximum 1 millisecond less than a minute: 59 seconds 999 milliseconds. Time that SOM waits for a response to an SNMP query before reissuing the request.
SNMP Retries Count	Maximum number of retries that SOM issues for an SNMP query before determining the query result to be "unresponsive". Zero means no retries.
SNMP Port	Default is 161. Specifies the SOM management server's port that SOM uses when generating SNMP traffic.

Default SNMP V1/V2 Community String

Attribute	Description
New	Click **New** to add a new read community.
Read Community String	The SNMPv1 or SNMPv2c "Get" (read-only) Community String that is used as the default value for each SNMP Agent (case-sensitive).
Ordering	*Optional*. A numeric value. SOM uses the first Community String that results in successful SNMP communication: • Each ordering number must be unique (no duplicate numbers). • SOM tries the provided Community Strings in the order you define (lowest number first). Consider incrementing by 10s or 100s to provide flexibility when adding new Read Community Strings over time. • If no Ordering numbers are specified, SOM tries all community strings in parallel. If some but not all the community strings have an Ordering number, SOM tries the community strings with a specified Ordering number first.

Default SNMP V1/V2 Community String, continued

Attribute	Description
	Then, SOM tries all the community strings without an Ordering number in parallel.

Node Specific Settings

Use the Specific Node Settings tab when you want to provide exceptions to the communication configurations that are generic for most nodes in your storage environment. The Specific Node Settings tab enables you to fine tune communication protocol usage and settings for a particular device within your environment.

Note: If you reconfigure the managed device (node) from SNMPv1 to operate on SNMPv3, you need to refresh the SNMP credentials using the utility `somrefreshsnmpcredentials.ovpl`. See the CLI reference page for details.

If no value is provided for an attribute in the form, SOM uses the default settings.

To configure communication protocol settings for a specific node:

1. Access the Specific Node Settings form:
   a. From the workspace navigation panel, select the **Configuration** workspace.
   b. Select **Communication Configuration**.
   c. Select **Specific Node Settings** tab.
   d. Do one of the following:
      - To establish settings for a node, click **New** and continue.
      - To edit settings for a node, double-click a row, and continue.
      - To delete settings for a node, select a row and click **Delete**.
2. Provide the communication protocol settings for the node. For details, see:
   - "Basic Settings" on the next page
   - "SNMP Settings" on the next page
3. *Optional*. Make additional configuration choices. For details, see:
   - "SNMP v1/v2 Community Strings" on page 232
   - "Configure SNMPv3 settings for a Specific Node" on page 233

4. Click **Save and Close** to return to the Communication Configuration form.
5. Click **Save and Close** to apply your changes.

Basic Settings

Attribute	Description
Target Hostname	Enter the fully-qualified host name as registered in your Domain Name System (DNS).
Preferred Management Address	Do one of the following: • Specify the address you want SOM to use for SNMP communications with this device. • Leave this attribute empty. SOM dynamically selects the management address, based on responses from the device's SNMP agent.
Description	*Optional*. Provide a description for this configuration that would be useful for communication purposes within your team.

SNMP Settings

Attribute	Description
Enable SNMP Communication	I f enabled, the Discovery Process and State Poller Service generate network traffic with SNMP protocol to discover and monitor this device. If disabled, SOM does not generate any SNMP traffic to this device.
SNMP Timeout	(Seconds:Milliseconds) Maximum 1 millisecond less than a minute: 59 seconds 999 milliseconds. Time that SOM waits for a response to an SNMP query before reissuing the request.
SNMP Retries Count	Maximum number of retries that SOM issues for an SNMP query before determining the query result to be "unresponsive". Zero means no retries.
SNMP Port	Default is 161. Specifies the management server's port that SOM uses when generating SNMP traffic.
SNMP Proxy Address	*Optional*. IP address of the your SNMP Proxy Server. To enable a proxy, you must also provide the port number of your SNMP Proxy Server. See SNMP Proxy Port (next attribute).

SNMP Settings, continued

Attribute	Description
SNMP Proxy Port	*Optional*. Port number of the SNMP Proxy Server. To enable a proxy, you must also provide the IP address of your SNMP Proxy Server. See SNMP Proxy Address (previous attribute).
SNMP Preferred Version	Specifies the SNMP version that SOM should use when communicating with a device.

SNMP v1/v2 Community Strings

Attribute	Description
New	Click **New** to add a new read community.
Read Community String	The SNMPv1 or SNMPv2c "Get" (read-only) Community String that is used as the default value for each SNMP Agent (case-sensitive).
Ordering	*Optional*. A numeric value. SOM uses the first Community String that results in successful SNMP communication: • Each ordering number must be unique (no duplicate numbers). • SOM tries the provided Community Strings in the order you define (lowest number first). Consider incrementing by 10s or 100s to provide flexibility when adding new Read Community Strings over time. • If no Ordering numbers are specified, SOM tries all community strings in parallel. If some but not all the community strings have an Ordering number, SOM tries the community strings with a specified Ordering number first. Then, SOM tries all the community strings without an Ordering number in parallel.

SNMPv3 Settings

Attribute	Description
New	Click **New** to add a new read community.

SNMPv3 Settings, continued

Attribute	Description
Read Community String	The SNMPv1 or SNMPv2c "Get" (read-only) Community String that is used as the default value for each SNMP Agent (case-sensitive).
Ordering	*Optional*. A numeric value. SOM uses the first Community String that results in successful SNMP communication: • Each ordering number must be unique (no duplicate numbers). • SOM tries the provided Community Strings in the order you define (lowest number first). Consider incrementing by 10s or 100s to provide flexibility when adding new Read Community Strings over time. • If no Ordering numbers are specified, SOM tries all community strings in parallel. If some but not all the community strings have an Ordering number, SOM tries the community strings with a specified Ordering number first. Then, SOM tries all the community strings without an Ordering number in parallel.

Configure SNMPv3 settings for a Specific Node

SOM can use SNMPv3 user-based security model (USM) settings to access devices. SOM uses the current SNMPv3 Settings provided for a node, if available.

> **Note:** If you reconfigure the managed device (node) from SNMPv1 to operate on SNMPv3, you need to refresh the SNMP credentials using the utility `somrefreshsnmpcredentials.ovpl`. See the CLI reference page for details.

To configure an SNMPv3 Settings for a specific node, follow these steps:

1. Access the Specific Node Settings form:
   a. From the workspace navigation panel, select the **Configuration** workspace.
   b. Select **Communication Configuration**.
   c. Select **Specific Node Settings** tab.
   d. Do one of the following:

- To establish settings for a node, click **New** and continue.
- To edit settings for a node, double-click a row, and continue.

2. Click the **SNMPv3 Settings** tab.
3. Click the SNMPv3 Settings Lookup icon and select one of the options from the drop-down menu:
   - Show Analysis to display Analysis Pane information for the currently configured (selected) SNMPv3 Setting name.
   - Quick Find to view and select from the list of all existing SNMPv3 Settings.
   - Open to display the details of the currently configured (selected) SNMPv3 Setting.
   - New to create a new SNMPv3 Setting (see "SNMPV3 Settings Form" below for more information).
4. Click **Save and Close** to return to the Specific Node Settings form.
5. Click **Save and Close** to return to the Communication Configuration form.
6. Click **Save and Close** to apply your changes.

SNMPV3 Settings Form

If your network environment is using the SNMPv3 user-based security model (USM), provide the information SOM needs for communication with the SNMPv3 agents in your environment. SOM uses the SNMPv3 settings to discover the SNMPv3 information about your network.

SNMPv3 Settings for the User-Based Security Model (USM)

Attribute	Description
Unique Name	Provide a unique name for this configuration record. You can reuse SNMPv3 Settings for defaults, communication regions, or specific nodes.
User Name	The SNMPv3 User Name is the text string used for SNMPv3 requests in your network environment.
Context Name	The SNMPv3 context name text string used in your network environment.
Authentication Protocol	The SNMPv3 authentication protocol. Determines whether authentication is required and indicates the type of authentication protocol used. SOM supports the following protocols: • HMAC -MD5-96 authentication protocol • HMAC-SHA-1 authentication protocol Leaving this attribute empty means SNMP Minimum Security Level = No

SNMPv3 Settings for the User-Based Security Model (USM), continued

Attribute	Description
	Authentication for this SNMPv3 configuration.
Privacy Protocol	Specify the SNMPv3 USM privacy protocol used by the SOM management server. The SNMPv3 USM privacy protocol determines whether encryption is required and indicates the type of privacy protocol used. SOM supports the following privacy protocols: • DES-CBC Symmetric Encryption Protocol • TripleDES - Triple Data Encryption Algorithm • AES128 - Advanced Encryption Standard 128 Protocol • AES192 - Advanced Encryption Standard 192 Protocol • AES256 - Advanced Encryption Standard 256 Protocol Leaving this attribute empty means SNMP Minimum Security Level = No Privacy for this SNMPv3 configuration.
Privacy Passphrase	The SNMPv3 USM privacy passphrase for the specified SNMPv3 User Name. If required for privacy, provide the appropriate encryption passphrase for use with the privacy protocol. The length limitations of the privacy passphrase depend on the privacy protocol. Leaving this attribute empty means SNMP Minimum Security Level= No Privacy for this SNMPv3 configuration.

Add SOM as a Trap Receiver

Analyzing SNMP traps helps SOM to monitor devices in real-time and acquire information regarding their health, performance, faults, and so on. Further, SOM generates incidents based on the SNMP traps. Incidents are notifications, alerts, or warnings that provide vital information about the device. For information about configuring incidents based on SNMP traps, see "Configuring SNMP Traps" on page 204. SOM also includes default pairwise configurations for the SNMP trap incidents. For information about default pairwise configurations and its benefits, see "About Pairwise Configurations" on page 524

Most SAN devices contain an SNMP agent that sends SNMP traps to the registered SNMP managers. To receive SNMP traps from the devices, register SOM as an SNMP manager. SOM supports various types of SNMP traps for each device. You must configure each device to send SNMP traps to the SOM management server. For information about adding SOM as a trap receiver for a device, see the device documentation.

SNMP Traps Supported for HPE 3PAR

An alertNotify trap contains details about an event that may affect system operations and performance. All alerts generated by the HPE 3PAR storage system, as well as alert status change events, are translated into alertNotify traps and forwarded to all registered managers.

SOM includes the HPE 3PAR MIB (`ThreeParMIB.mib`) used to decode the incoming alertNotify traps.

The incoming HPE 3PAR SNMP `alertNotify` traps are displayed in SOM in **Incident Browsing** > **SNMP Traps**.

SNMP Traps Supported for the Cisco Switch

SOM supports the following SNMP traps received from the Cisco Switch:

Sr. No.	MIB	SNMP Trap	Description
1	CISCO-FC-FE-MIB	fcTrunkIfDownNotify	This trap is generated when a trunk interface status changes to down.
2		fcTrunkIfUpNotify	This trap is generated when a trunk interface status changes to up.
3	CISCO-VSAN-MIB	vsanStatusChange	This trap is generated when a VSAN status changes to down.
4		vsanPortMembershipChange	This trap is generated when a port is added to a VSAN.
5	CISCO-FC-FE-MIB	dmNewPrincipalSwitchNotify	This trap is generated when a new principal switch is selected in a VSAN.
6	FCMGMT-MIB	fcMgmtNotification	This trap is generated when a switch port comes online.
7	CISCO-IF-EXTENSION-MIB	cieLinkUp	This trap is generated when a communication link on an FC port comes on. For example, when a connected device is switched on.
8		cieLinkDown	This trap is generated when a communication link on an FC port goes off. For example, when a connected device is switched off.

Sr. No.	MIB	SNMP Trap	Description
9	CISCO-ZS-MIB	zoneActivateNotify	This trap is generated when a zone set is activated in a VSAN.
10	CISCO-FEATURE-CONTROL-MIB	ciscoFeatureOpStatusChange	This trap is generated when a switch feature is enabled or disabled.
11		ciscoFeatureOpStatusChange2	This trap is generated when a switch feature is enabled or disabled.
12	CISCO-ENTITY-FRU-CONTROL-MIB	cefcPowerStatusChange	This trap is generated when a switch FRU is powered off, due to insufficient system power, power translation errors, temperature problems, and so on.

SNMP Traps Supported for the Brocade Switch

SOM supports the following SNMP traps received from Brocade Switch:

Sr. No.	MIB	SNMP Trap	Description
1	FCMGMT-MIB	fcMgmtNotification	This trap is generated when a switch port comes online.
2	HA-MIB	fruStatusChanged	This trap is generated when a switch FRU fails, is powered off, or removed.
3	SW-MIB	swFCPortScn	This trap is generated when an FC port changes its operational state or port type.
4	LINK-INCIDENT-MIB	linkRNIDDeviceRegistration	This trap is generated when a device is registered with a switch.
5		linkRNIDDeviceDeRegistration	This trap is generated when a device is unregistered with a switch.
6		linkRLIRFailureIncident	This trap is generated when a link failure occurs.
7	SW-MIB	swZoneConfigChangeTrap	This trap is generated when there is a change in the

Sr. No.	MIB	SNMP Trap	Description
			local zone database.
8		swPortMoveTrap	This trap is generated when the virtual ports are moved from one switch to another.
9		swStateChangeTrap	This trap is generated when a switch state changes to online or offline.
10		swPmgrEventTrap	This trap is generated when any partition manager ports are moved to or from a logical switch alpina in the fabric, Deville_Fabric.
14		swEventTrap	This trap is generated when an event occurs.
11	BD MIB	bdTrap	This trap is generated when latency and congestion bottlenecks occur.
12		bdClearTrap	This trap is generated after a latency or congestion bottleneck is cleared.
13	MAPS-MIB	mapsTrapAM	This trap is generated for MAPS threshold events.

Note: The MIBs mentioned in the table are not available in all the firmware versions. MIBs are supported in Brocade Switches based on the firmware version.

Ensure that SNMP Informs are disabled in the Brocade Switch to view SNMP traps in SOM (**Incident Browsing** > **SNMP Traps**). SOM does not support SNMP Informs.

SNMP Traps Supported for EMC VNX Celerra

SOM supports the following SNMP traps received from EMC VNX Celerra, which are part of the EMC CELERRA MIB (`EMC-CELERRA.mib`):

Sr.No.	SNMP Trap	Description
1	celMasterCtlFault	This trap message is generated when there is a serious fault in the Master Control.
2	celHWFailure	This trap message is generated when there is a failure in the hardware.
3	celSlotStale	This trap message is generated when there is a stale reason

Sr.No.	SNMP Trap	Description
		code.
4	celSlotPanicked	This trap message is generated when there is a panicked slot.
5	celIntfFailure	This trap message is generated when there is a failure in both interfaces.
6	celAAF	This trap message is generated when there is an AAF.
7	celReboot	This trap message is sent when a Data Mover Reboots. The trap message includes the facility, event id, severity, and a text description of the event.
8	celCSStart	This trap message is sent when the Control Station is started.
9	celJServer	This trap message is sent in the event of a JServer alert.
10	celWebGUI	This trap message is sent in the event of a Web GUI notification request.
11	celUFS	This trap message is sent when a UFS quota is exceeded.

SNMP Traps Supported for EMC VNX CLARiiON

SOM supports the following SNMP traps received from EMC VNX CLARiiON which are part of the CLARIION MIB (`CLARIION-MIB.mib`):

Sr.No.	SNMP Trap	Description
1	EventMonitorTrapInfo	This is an informational trap generated in response to a user-specified event.
2	EventMonitorTrapWarn	This is a warning trap generated in response to a user-specified event.
3	EventMonitorTrapError	This is an error trap generated in response to a user-specified event.
4	EventMonitorTrapFault	This is a fault trap generated in response to a user-specified event.

SNMP Traps Supported for Hitachi Data Systems

SOM receives and analyzes SNMP traps from various Hitachi Data Systems (HDS) such as, Hitachi Unified Storage (HUS), Virtual Storage Platform (VSP), and Adaptable Modular Storage (AMS). The SNMP traps received from these devices help SOM to monitor the devices in real time.

The types of SNMP traps received from various HDS devices vary depending on the device and its model. These SNMP traps belong to different MIBs (Management Information Base). The following table summarizes the information about the SNMP traps received from each HDS device and the associated MIBs:

HDS Device	MIB	Description	Number of SNMP Traps
HUS 110, HUS 130, and HUS 150	dfraid.mib	The traps provide information about the various events occurring in the device including, failures, alerts, thresholds, and warnings.	64
HUS VM	HUSVMIB.mib	The traps indicate the severity of the impact on the device because of an event.	4
VSP G1000	VSPG1000MIB.mib	The traps indicate the severity of the impact on the device because of an event.	4
VSP 700	VSPMIB.mib	The traps indicate the severity of the impact on the device because of an event.	4
AMS 2000	AMS.mib	The traps provide information about the various events occurring in the device including, failures, alerts, and thresholds.	56

To view the list of SNMP traps supported by SOM for a particular device, do the following:

1. In the SOM console, in the workspace navigation panel, select the **Configuration workspace**.
2. Select **Incidents** > **SNMP Trap Configurations**.
3. In **SNMP Trap Configurations** window, right-click the **Author** column.
4. Select **Filter** > **Create filter...**.
5. In the **Filter: Author** window, select the device value for which you want to view the list of supported SNMP traps. For example, to view the supported SNMP traps from the HDS and HPE XP devices,

select **HP SOM HDS/XP Event Manager**.

The SNMP trap Configurations window displays the supported SNMP traps for the selected device.

Note: The framework treats each SNMP trap configuration as unique even if the configuration has a wild card character in the SNMP OID. Therefore, all events that match the OID are treated as similar traps.

SNMP Traps Supported for the EMC VMAX and EMC Symmetrix Arrays

SOM supports the following SNMP traps received from the EMC VMAX and EMC Symmetrix arrays which are part of the EMCGATEWAY-MIB (`EMCGATEWAY-MIB.mib`) :

Sr. No.	SNMP Trap	Description
1	eccUnitStatusChange	The overall status of the Control Center-monitored device has changed. Recommended severity level for filtering is alert.
2	eccUnitDeletedTrap	A Control Center-monitored device has been deleted from the Console. Recommended severity level for filtering is warning.
3	eccUnitEventTrap	An event has been generated by the Control Center-monitored device. Recommended severity level for filtering is info.

These traps are dependent on FCMGMT-MIB (`FCMGMT-MIB.mib`) for trap-related descriptions.

SNMP Traps Supported for NetApp Devices

SOM receives and analyzes SNMP traps from NetApp 7-Mode and NetApp C-Mode devices. Traps received from these devices help SOM to monitor the devices in real time. The traps received from NetApp devices belong to the `netapp.mib` MIB (Management Information Base). SOM supports a total of 167 traps from NetApp devices. NetApp devices send these traps for various events such as disk or power failures, errors, critical alerts, warnings, device health status, and so on. SOM does not support user created events in NetApp devices.

To view the list of SNMP traps supported by SOM for a particular device, do the following:

1. In the SOM console, in the workspace navigation panel, select the **Configuration workspace**.
2. Select **Incidents** > **SNMP Trap Configurations**.

3. In **SNMP Trap Configurations** window, right-click the **Author** column.
4. Select **Filter** > **Create filter...**.
5. In the **Filter: Author** window, select the device value for which you want to view the list of supported SNMP traps. For example, to view the supported SNMP traps from the HDS and HPE XP devices, select **HP SOM HDS/XP Event Manager**.

   The SNMP trap Configurations window displays the supported SNMP traps for the selected device.

SNMP Traps Supported for IBM SAN Volume Controller

SOM support the following SNMP traps from IBM SVC:

Sr.No.	SNMP Trap	Description	MIB
1	IbmSvcError .1.3.6.1.4.1.2.6.190.1	This trap represents error in IBM SAN Volume Controller.	IBM-SVC-MIB
2	IbmSvcWarn .1.3.6.1.4.1.2.6.190.2	This trap represents warning in IBM SAN Volume Controller	
3	IbmSvcInfo .1.3.6.1.4.1.2.6.190.3	This trap represent informational messages from IBM SAN Volume Controller.	

SNMP Traps Supported for HP P9500 and XP Arrays

SOM receives and analyzes SNMP traps from various HP P9500 and XP array models such as HP XP, HP XP 24000, HP XP7, and HP P9500. The SNMP traps received from these devices help SOM to monitor the devices in real time.

SOM supports the following SNMP traps received from HP P9500 and XP arrays:

Sr. No.	MIB	SNMP Trap	Description
1	HUSVMIB	raideventUseracute	Indicates that the impact of this event on the subsystem is acute.
2		raideventUserserious	Indicates that the impact of this event on the subsystem is serious.
3		raideventUsermoderate	Indicates that the impact of this event on the subsystem is moderate.
4		raideventUserservice	Indicates that the impact of this event on the subsystem is low.

Sr. No.	MIB	SNMP Trap	Description
5	VSPG1000MIB	raideventUseracute	Indicates that the impact of this event on the subsystem is acute.
6		raideventUserserious	Indicates that the impact of this event on the subsystem is serious.
7		raideventUsermoderate	Indicates that the impact of this event on the subsystem is moderate.
8		raideventUserservice	Indicates that the impact of this event on the subsystem is low.

Chapter 3: Managing your Storage Environment with SOM

SOM provides the following features that enable you to manage your storage environment.

Feature	Description
"Using Inventory Views" on page 445	Provides a collection of views to access details of elements managed by SOM.
"Using Analytics and Dashboards" on page 248	Contains information panels pertaining to the entire storage environment, an element category, or an individual element.
"Using Topology Maps" on page 284	Displays the connectivity maps of the top level elements in the storage infrastructure.

Using Incident Browsing Views

Incidents are information that SOM considers important to bring to your attention regarding your storage environment.

SOM provides the following views to help you monitor incidents and take appropriate action to preserve health of your storage environment:

Open Incidents view

The Open Incidents view in the Incident Browsing workspace displays the incidents any of the following lifecycle states:

- Registered
- In Progress
- Completed

Closed Incidents view

The Closed Incidents view in the Incident Browsing workspace displays any Incident with a Life Cycle state of Closed. This view is useful for identifying the Incidents that have been resolved. This view might be particularly useful for reporting on how many incidents were closed within a given time period.

All Incidents view

The All Incidents view in the Incident Browsing workspace is useful for viewing all of the incidents generated by SOM within the specified time period. This view is useful to identify both open and closed incidents.

SNMP Traps view

The SNMP Traps view in the Incident Browsing workspace is useful for identifying all of the traps that were received from devices in your storage environment.

See the Lifecycle State information for the Incident form for more information.

In any of the Incident views, double-click a row to see the "Incident Form" on page 527 that displays details about the incident. The Incident Form contains the following tabs:

- "Incident Form: General Tab" on page 529
- "Incident Form: Correlated Parents Tab" on page 534
- "Incident Form: Correlated Children Tab" on page 534
- "Incident Form: Custom Attributes Tab" on page 535
- "Incident Form: Registration Tab" on page 535

Manage Incident Assignments

One of the first things to do with an incident is to assign it to yourself or to another operator. The following table displays the ways you can assign or un-assign an incident and the SOM user role that is required for each.

Tasks Related to Assigning Incidents

Task	How	Required Minimum SOM User Role
Own an incident	Select an incident and use **Actions > Assign > Own Incident**. See "Own One or More Incidents" on page 538 for more information.	Level 1 Operator (with limited access privileges than Level 2 Operators)
Assign an incident to someone else	There are two ways to assign an incident to someone else (see "Assign Incidents" on page 525 for more information): • From any Incident view, select one or more Incidents and use **Actions > Assign > Assign Incident**. • From an Incident form, use **Actions > Assign > Assign Incident**.	Level 1 Operator
Un-assign an incident	Select an incident and use **Actions > Assign > Unassign Incident**. See "Unassign Incidents" on page 546 for more information.	Level 1 Operator

Alternatively, you can manage incident assignments using the following procedure:

To assign or change assignment for an incident:

1. Navigate to the Incident form of interest.
   a. From the workspace navigation panel, select the **Incident Browsing** workspace.
   b. Select any Incident view.
   c. Double click the row representing the incident you want to assign.
2. In the incident form's **Basics** pane, locate the **Assigned To** field.
3. From the **Assigned To** drop-down menu, select the required operator (assignee).
4. Click **Save** to save your changes or **Save and Close** to save your changes and exit the form.

The user name you selected appears in the **Assigned To** column in any Incident views that include that incident.

Change Incident's Lifecycle State

Use an incident's Lifecycle State to track an incident's progress. Your network administrator might have additional or different guidelines for using lifecycle states. In some cases, SOM updates an incident's Lifecycle State automatically based on the configuration.

Learn the guidelines specific to your organization before updating the incident's lifecycle state.

Note: You can also change the incident's priority or severity. Additionally, you can change specific attributes in incident tabs. For more information about various views and tabs, see "Using Incident Browsing Views" on page 245.

To update an incident's lifecycle state, follow these steps:

1. In an open incident, in the **Basics** pane, select a Lifecycle State.
2. Click **Save** to save your changes or **Save and Close** to save your changes and exit the form.
3. After performing an action on a form that modifies the object being viewed, refresh the form before making additional changes.

Update Incident Notes

Use the **Notes** field to explain steps that were taken to date to troubleshoot the problem, workarounds, solutions, and ownership information.

To update an incident, follow these steps:

1. In an open incident, in the **Basics** pane of the incident form, locate the **Notes** area.
2. Expand if required and type the annotations that you want to be displayed within the **Notes** field. You can type a maximum of 1024 characters.
3. Click **Save** to save your changes or **Save and Close** to save your changes and exit the form.

To keep your incident's Lifecycle State information current, see "Change Incident's Lifecycle State" on the previous page.

Delete an Incident

After an incident is generated, it goes through different lifecycle states before it is finally closed. The closed incidents are available to you irrespective of the status of the source from where the incident originated.

You may delete the incidents that you no longer require.

To delete an incident, follow these steps:

1. Navigate to the Incident form of interest.
   a. From the workspace navigation panel, select the **Incident Browsing** workspace.
   b. Select any Incident view.
2. Do one of the following:
   - Click a row and then click **Delete**.
   - Drag across the rows that you want to select and then click **Delete**.
   - Click a row, then hold down the Ctrl key while you click other rows that you want to select, and then click **Delete**.

Using Analytics and Dashboards

SOM includes dashboards that provide the latest information about the number of discovered devices, data collection statuses, storage utilization, and performance analytics. Dashboard views help compare and isolate the details required to analyze data in the environment.

A dashboard contains multiple panels of data pertaining to the entire storage environment, an element category, or an individual element (storage system, host, switch, and so on). Dashboard panels might contain a variety of tables and pie charts.

The following dashboards are available at the environment level:

- Environment Capacity
  Information panels that illustrate the overall capacity utilization in the environment. Dashboards for element

categories (storage systems, hosts, and switches) and individual device utilization views provide additional perspectives for data analysis.

- Asset Dashboard

  Information panels that illustrate the number of discovered devices based on Device Family, Device Vendor, or the OS Type of a device.
- Collection Status Dashboard

  Information panels that illustrate device data collection status and quarantined devices in the environment (storage systems, hosts, and switches). Inventory views of devices based on the data collection status help to analyze discovered devices.
- "Storage Systems DTT Analytics" on page 256

  Information panels that forecast the number of days to threshold (DTT) values for the capacity utilization (Raw Used, Actual Allocated, and Actual Used) of storage systems.
- "Storage Pools DTT Analysis" on page 257

  Information panels that forecast the number of days to threshold (DTT) values for the capacity utilization (Actual Allocated, and Actual Used) of storage pools.
- "Analytics for Virtual Servers" on page 258

  Information panels that forecast over provisioned datastores, details of snapshots and powered-off VMs, in a VMware virtual environment.
- "Storage Systems Unused Volumes Analytics" on page 283

  Displays the storage systems with volumes that have not been accessed for the specified number of days.
- "Path Analytics" on page 260

  Information panels for analytic details about the volumes accessed by the hosts in the environment.

Environment Capacity Dashboard

The Environment Capacity dashboard displays information panels that give you an insight into the storage consumption.

The following dashboard panels are available:

- **Environment Summary**

  Displays a pie chart with the total number (count) of discovered devices in each device category.

  For a detailed view of a device category's capacity utilization, click the pie sector of a category to see the following device capacity dashboards:

  - Host Capacity
  - NAS System Capacity
  - Storage System Capacity

  - Backup System Capacity
  - Switch Capacity

- **Storage System Logical Capacity**

  Displays a pie chart that illustrates the total logical capacity visible to hosts.

  To see the aggregate **Allocated Storage** and **UnAllocated Storage** of storage systems in the environment, mouse over the pie chart sectors.

  To see the capacity metrics (**Name**, **Allocated**, and **UnAllocated**) of the storage systems, click a pie chart sector to display the **Storage System Capacity** view with storage systems sorted in the descending order by the selected capacity metric.

  For additional properties and related components of an individual storage system, double-click or **Open** a storage system from the **Storage System Capacity** view, to see its form view.

- **NAS System Capacity**

  Displays a pie chart that illustrates the total NAS system capacity.

  To see the aggregate **Free Space** and **Used Space** of NAS systems in the environment, mouse over the pie chart sectors.

  To see the capacity metrics (**Name**, **Total**, **Used**, and **Free**) of the file storage systems, click a pie chart sector to display the **NAS System Capacity** view with storage systems sorted by the selected capacity metric.

  For additional properties and related components of an individual file storage system, double-click or **Open** a storage system from the **NAS System Capacity** view, to see its form view.

- **Backup System Capacity**

  Displays a pie chart that illustrates the total backup system capacity.

  To see the aggregate **Free Space** and **Used Space** of backup systems in the environment, mouse over the pie chart sectors.

  To see the capacity metrics (**Name**, **Total**, **Used**, and **Free**) of the file storage systems, click a pie chart sector to display the **Backup System Capacity** view with storage systems sorted by the selected capacity metric.

  For additional properties and related components of an individual file storage system, double-click or **Open** a storage system from the **Backup System Capacity** view, to see its form view.

- **Host Logical Capacity**

  Displays a pie chart that illustrates the total volume capacity that is consumed by the hosts in the environment.

  The aggregate capacity at the host level excludes network filesystems such as nfs, nfs4, cifs, smbfs, and ncpfs.

To see the aggregate **Free Space** and **Used Space** of storage utilized by hosts, mouse over the pie chart sectors.

To see the capacity metrics (**Name**, **Total**, **Used**, **Free**, **%Used**, and **%Free**) of the hosts, click a pie chart sector to display the **Host Capacity** view with hosts sorted in the descending order by the selected capacity metric.

For additional properties and related components of an individual host, double-click or **Open** a host from the **Host Capacity** view, to see its form view.

- **Switch Port Utilization**

  Displays a pie chart that illustrates the utilization of all the switch ports in the environment. Only physical switches are considered and not virtual or quarantined switches.

  To see the total number of **Free Ports** and **Used Ports** of the physical switches discovered in the environment, mouse over the pie chart sectors.

  To see the capacity metrics (**Name**, **Total**, **Used**, **Free**, **%Used**, and **%Free**) of the physical switches, click a pie chart sector to display the **Switch Capacity** view with switches sorted in the descending order by the selected capacity metric.

  For additional properties and related components of an individual switch, double-click or **Open** a switch from the **Switch Capacity** view to see its form view.

- **Stale Elements**

  Displays a table view of all stale elements.

Host Capacity Dashboard

The Host Capacity dashboard displays the overall storage capacity utilized by the hosts that are discovered and managed in the environment.

Note: By default, raw partitions are not used in capacity calculations of a host.

Note: The capacity of all VMs—including that of devices (RDMs, LUNs assigned via passthrough HBAs, and NPIV configurations) connected directly to the VMs—at the environment level is excluded from the capacity calculations of ESX servers.

To view the dashboard of an individual host, click a host name in a panel.

The dashboard of an individual host displays its capacity and performance (at the top level only) information. These details are also available in the Analysis pane.

The Host Capacity dashboard panels give you the following insights into the storage consumed by the host elements in the environment:

- **Top 10 by Used Space**
  Highlights the top 10 hosts by the storage volume capacity that is utilized.
- **Top 10 by Free Space**
  Highlights the top 10 hosts by the storage volume capacity that is not utilized.
- **Top 10 by Unused Volume Group Capacity**
  Highlights the top 10 hosts by unused capacity in the volume groups. This is the space available in a volume group to create more volumes.
- **Top 10 by Unused Storage**
  Highlights the top 10 hosts by the unused disk space, which is defined as any one of the following:
  - The disk is not part of a volume group.
  - The disk has no storage volume.
- **Host Capacity**
  Displays the storage volume utilization and percentage (**Name**, **Total (GiB)**, **Used (GiB)**, **Free (GiB)**, **% Used**, **% Free**, **VG Available/Grey Space (GiB)**, **VG Used Space (GiB)**) for all the discovered hosts. The percentage values are based on the total space.

  For the properties and related components of a host, double-click or **Open** a host to see its Host Form.

NAS Capacity Dashboard

The NAS System Capacity dashboard displays the overall file storage capacity utilization of the NAS systems that are discovered and managed in the environment.

To view the dashboard of an individual NAS system, click a NAS system name in a panel.

NAS System Capacity

Displays the file storage capacity (**Name**, **Total**, **Used**, and **Free**) information of the NAS systems discovered in the environment.

For the properties and related components of a NAS system, double-click or **Open** a NAS system to see its Storage System Form.

Storage System Capacity Dashboard

The Storage System Capacity dashboard displays the overall storage capacity utilization of the storage systems that are discovered and managed in the environment.

Capacity utilization for a storage system is based on the total overall capacity versus the free space, that is calculated from the unused, unmapped, and reclaimable space. The capacity utilization rate takes into account pool capacity that is not provisioned, provisioned but unmapped capacity, and unused space reserved for snapshot and thin provisioned volumes.

To view the dashboard of an individual storage system, click a storage system name in a panel.

The dashboard of an individual storage system displays its capacity and performance information. These details are also available in the Analysis pane.

The Storage System Capacity dashboard panels are designed keeping in mind the following perspectives of analyzing storage consumption in an environment:

- **Top 10 by Raw Used Capacity**
  Highlights the top 10 arrays in the environment by the raw space that is used. This dashboard highlights devices that may potentially require additional disks to be provisioned due to a high raw space utilization.
- **Top 10 by Raw Available**
  Highlights the top 10 arrays in the environment by the raw space that is available. This dashboard highlights devices with raw storage that can be provisioned.
- **Top 10 by Logical Unallocated Storage**
  Highlights the top 10 arrays with space available for creation of storage volumes. You can use this dashboard, in conjunction with the Top 10 by Raw Used dashboard to plan procurement of additional storage. A device that shows up in the Raw Used dashboard and not in Unallocated dashboard indicates that the device is reaching the limits of storage configuration.
- **Top 10 by Logical Unmapped Storage**
  Highlights the top 10 arrays with volumes created but not exposed to hosts. This dashboard can be used to identify devices from which storage can be reclaimed by either exporting the volumes to hosts or by deleting the unmapped volumes.
- **Storage System Capacity**
  Displays the storage capacity utilization (**Name**, **Raw Total (GiB)**, **Raw Used (%)**, **Raw Used (GiB)**, **Raw Available (%)**, **Raw Available (GiB)**, **Unallocated (%)**, **Unallocated Storage (GiB)**, **Unmapped (%)**, **Unmapped Storage (GiB)**) for the discovered storage systems.
  The percentage values are calculated based on the total raw space that is available on the storage system.
  For the properties and related components of a storage system, double-click or **Open** a storage system to see its Storage System Form.

Backup System Capacity Dashboard

The Backup System Capacity dashboard displays the overall file storage capacity utilization of the backup systems that are discovered and managed in the environment.

To view the dashboard of an individual backup system, click a backup system name in a panel.

Backup System Capacity

Displays the file storage capacity (**Name**, **Total**, **Used**, and **Free**) information of the backup systems discovered in the environment.

For the properties and related components of a backup system, double-click or **Open** a backup system to see its Storage System Form.

Switch Capacity Dashboard

The Switch Capacity dashboard displays the port utilization of the physical switches that are discovered and managed in the environment.

To view the dashboard of an individual physical switch, click a switch name in a panel.

The dashboard of an individual switch displays its capacity and performance information. These details are also available in the Analysis pane.

The Switch Capacity dashboard panels are designed to give you the following utilization perspectives of the ports of the physical switches in the environment:

- **Top 10 by Used Ports**
  Highlights the top 10 physical switches by the number of used ports.
- **Top 10 by Free Ports**
  Highlights the top 10 physical switches by the number of free ports.
- **Switch Capacity**
  Displays the percentage of port utilization (**Name**, **Used Ports**, **Free Ports**, **Total Ports**, **% Used Ports** and **% Free Ports**) for all the discovered physical switches in the storage network.
  The percentage values are based on the total number of switch ports.

  For the properties and related components of a physical switch, double-click or **Open** a switch to see its Switch Form.

Asset Dashboard

The Asset Dashboard displays pie chart views of discovered devices based on a device attribute.

The following panels are available:

- Hosts - based on the OS Type
- Storage Systems - based on the Device Family
- (Physical) Switches - based on Device Vendor
- Virtual Machines - based on the OS Type

In each panel, the number of discovered devices are available on mouse rollover of a pie sector.

For the inventory view of a set of devices, click the corresponding pie sector.

For additional properties and related components of an individual discovered device, double-click or **Open** a selected device from the inventory view to see its form view.

Collection Status Dashboard

The Collection Status Dashboard gives an overview of the data collection status for discovered elements across the entire storage infrastructure.

The following information panels display the different data collection statuses and the percentage of devices in a particular collection state:

- Elements Collection Status - for all the discovered elements in the environment
- Hosts Collection Status
- Storage Systems Collection Status
- Switches Collection Status

Quarantined Status

The Quarantined Status panel displays the number of elements (fabrics, switches, hosts, and storage systems) that are quarantined.

An element is quarantined if the following are true:

- Data collection fails for three schedules (implying non-transient data collection errors)
- An element is under maintenance (for a firmware/hardware/software upgrade) and the administrator excludes the element from data collection.

To quarantine an element, select **Actions** > **Quarantine** or use the context menu in the Inventory and Topology workspaces.

The administrator must include an element for data collection after maintenance/data collection errors are resolved.

> **Note:** By default, SOM automatically quarantines an element if data collection fails. If you do not want SOM to quarantine elements automatically, you can disable the auto-quarantine feature. For more information about disabling the auto-quarantine feature, see "Quarantine/Unquarantine Elements" on page 434.

Data Collection Status

Data may or may not be collected from devices for various reasons. The collection status of a device can be any of the following:

- Success
- Running
- Queued
- Failed to start
- Provider problem
- Remote agent unavailable

- IP unreachabe
- Bad username
- Bad password
- Device busy
- Lost connection
- No result
- Canceled
- Remote agent unavailable
- Agent problem
- Timeout
- Unknown
- Internal error

To see the inventory details of devices with a particular collection status, click the corresponding pie chart sector for an inventory view filtered by the selected collection status.

For example, in the **Hosts Collection Status** panel, click the **Success** sector, to see the inventory details of hosts with the **Collection Status** as **Success**.

Storage Systems DTT Analytics

The Storage Systems Days to Threshold (DTT) Analytics dashboard forecasts the number of storage systems (Y-axis) that are expected to reach the capacity utilization (Raw Used, Actual Allocated, and Actual Used) threshold percentage values of 80, 90, and 100, in the number of days plotted along the X-axis in each panel.

Note: The data in this dashboard comes from the SOM reporting server. For information about setting up the connection to the SOM reporting server, see "Establish the Connection to the SOM Reporting Server for Gathering Analytics Data" on page 283.

The following dashboard panels are available:

- **Raw Used Percentage**
  The Raw Used Percentage is the percentage of raw disk capacity that is used.
  This panel displays a bar chart of the number of storage systems for which the Raw Used percentage will reach the threshold percentage values in the specified number of days.
- **Actual Allocated Percentage**
  The Actual Allocated Percentage is the percentage of capacity that is actually allocated at the storage storage system level to the Total (Logical) Capacity of the storage system.

This panel displays a bar chart of the number of storage systems for which the Actual Allocated percentage will reach the threshold percentage values in the specified number of days.

- **Actual Used Percentage**
  The Actual Used percentage is the percentage of capacity that is actually used by the volumes in a storage pool.
  This panel displays a bar chart of the number of storage systems for which the Actual Used percentage will reach the threshold percentage values in the specified number of days.
- **Storage Systems DTT Analytics Table**
  This panel displays all the storage systems and the number of days in which the capacity utilization (Raw Used, Actual Allocated, and Actual Used) of each storage system will reach the respective threshold percentage values.
  To see the properties of a storage system and details of its related components, double-click to open the "Block Storage Systems View" on page 450.

  > **Note:** A zero value indicates that a storage system has exceeded a threshold value. A value of 99999 indicates that a storage system will reach a threshold value in the future.

Hover over a bar in a chart to see the threshold percentage and the number of storage systems for a particular capacity utilization metric in a panel.

To see the storage systems and the actual number of days to the selected threshold value, click a bar in a chart, to view the Analytics table sorted in the ascending order by the selected capacity utilization metric. For a selected threshold value, the Analytics table displays the number of days for all the capacity utilization metrics (Raw Used, Actual Allocated, and Actual Used) of a storage system.

Storage Pools DTT Analysis

The Storage Pools Days to Threshold (DTT) Analytics dashboard forecasts the number of storage pools (Y-axis) that are expected to reach the capacity utilization (Actual Allocated, and Actual Used) threshold percentage values of 80, 90, and 100, in the number of days plotted along the X-axis in each panel.

> **Note:** The data in this dashboard comes from the SOM reporting server. For information about setting up the connection to the SOM reporting server, see "Establish the Connection to the SOM Reporting Server for Gathering Analytics Data" on page 283.

The following dashboard panels are available:

- **Actual Allocated Percentage**
  The Actual Allocated Percentage is the percentage of capacity that is actually allocated to the storage pools to the Total Capacity of the storage pools.
  This panel displays a bar chart of the number of storage pools for which the Actual Allocated percentage will reach the threshold percentage values in the specified number of days.

- **Actual Used Percentage**
  The Actual Used percentage is the percentage of capacity that is actually used by the volumes in a storage pool.
  This panel displays a bar chart of the number of storage pools for which the Actual Used percentage will reach the threshold percentage values in the specified number of days.
- **Storage Pools DTT Analytics Table**
  This panel displays all the storage pools and the number of days in which the capacity utilization (Actual Allocated, and Actual Used) of each storage pool will reach the respective threshold percentage values.
  To see the properties of a storage pool and details of its related components, double-click to open the "Storage Pool Form" on page 464.

  > **Note:** A zero value indicates that a storage pool has exceeded a threshold value. A value of 99999 indicates that a storage system will reach a threshold value in the future.

Hover over a bar in a chart to see the threshold percentage and the number of storage pools for a particular capacity utilization metric in a panel.

To see the storage pools and the actual number of days to the selected threshold value, click a bar in a chart, to view the Analytics table sorted in the ascending order by the selected capacity utilization metric. For a selected threshold value, the Analytics table displays the number of days for both the capacity utilization metrics (Actual Allocated, and Actual Used) of a storage pool.

Analytics for Virtual Servers

The analytics dashboard for virtual servers—discovered through a virtual center—provides the following information:

- Datastores and virtual servers that are over provisioned
- The number of snapshots and the size of all snapshots on a VM
- The number of powered-off VMs and the total size of such VMs on a virtual server

SOM must discover the virtual servers and storage devices to populate this dashboard. If SOM is unable to determine whether a storage volume or datastore is thinly provisioned, it assumes that the storage volume or datastore is thick.

The following dashboard panels provide information to analyze the over allocation of provisioned storage at the environment level:

- **Top 10 Datastores by Over Allocation**
  The top 10 datastores with high physical disk space utilization.
  - **Over Allocation (%)**
    The percentage of over provisioned storage out of the storage that is actually consumed by a datastore.

  - **Space Available for Expansion**
    The physical storage in the pool that is available for a datastore to grow. This value is applicable only for datastores that are composed of thin provisioned LUNs.

- **Top 10 Virtual Servers by Over Allocated Datastores Count**
  The top 10 ESX servers with the largest number of over allocated datastores with high physical disk space utilization.
- **Top 10 Virtual Machines by Snapshot Count**
  The top 10 virtual machines (VMs) with the highest number of snapshots.
- **Top 10 Virtual Machines by Snapshot Size**
  The top 10 VMs with the largest snapshot size.
- **Top 10 Virtual Servers by Powered-off VM Size**
  The top 10 virtual servers with the largest size of powered-off VMs.
  - **Size**
    The total size of VMs that are powered off on a virtual server.

    > **Note:** If the time when a VM was last powered off cannot be determined, the VM is not considered in this calculation.

  - **VM Count**
    The total number of powered-off VMs on a virtual server.

  By default, SOM considers VMs that have been powered off for three days. You can customize the default value, in the custom.properties file. For more information about customizing such properties, see "Configure Analytics" on page 284.
- **All Datastores by Over Allocation**
  All the datastores that are susceptible to an outage of physical disk space.

  To see the properties and related components of a datastore (filesystem), double-click a datastore for its inventory view.
- **All Virtual Servers by Over Allocated Datastores Count**
  All the virtual servers with the number of over allocated datastores that are susceptible to resource outage.

  To see the properties and related components of a host (virtual server), double-click a host for its inventory view.
- **All Virtual Machines by Snapshot Size and Count**
  All the VMs with the number of snapshots and snapshot size of each VM.

  To see the properties and related components of a host (virtual server), double-click a host for its inventory view.
- **All Virtual Servers by Powered-off VM Size**
  All the virtual servers with the number of VMs powered-off for the specified number of days and the size of the powered-off VMs on each server.

By default, SOM considers VMs that have been powered off for three days. You can customize the default value, in the custom.properties file. For more information about customizing such properties, see "Configure Analytics" on page 284.

To see the properties and related components of a VM, double-click a VM for its inventory view.

To view the dashboard of an individual datastore (filesystem) or virtual server (host), click the name of a datastore or virtual server from any of the top 10 panels. Individual dashboards display the summary and analysis information of a selected datastore or host.

Path Analytics

The Path Analytics dashboard provides the following information about hosts. This dashboard also considers inferred and created hosts.

- Storage volumes that are available through a single port
- Storage volumes with missing paths
- Storage volumes that are shared with multiple hosts

Storage paths are obtained from the host security groups (HSG) configured for a storage system after successful data collection of the storage system. Therefore, when a new host is discovered, data collection must be rerun for the connected storage systems to see information in this panel. Likewise, for any configuration changes for a host, data collection must be rerun for the connected storage systems. When storage system ports are unavailable or removed from the HSGs, the path information does not include these ports after the next data collection.

The path analytics information is available in the following panels:

- **Top 10 Hosts by Storage Volumes with Single Point of Failure Ports**
  The top 10 hosts with the total number of storage volumes at risk due to single port connectivity.

  If the storage volume is connected to a single storage system port or HBA port, the storage volume is at risk and considered in the total count.

  For additional information about the underlying storage volumes and the connected ports, click a host to see the Single Point of Failure Ports table. The information in this table is also available in the **Storage Volumes with Single Point of Failure Ports** tab in the Analysis pane for a selected host.

- **Top 10 Hosts by Missing LUN Paths due to Zoning Misconfigurations**
  The top 10 hosts with the total number of paths to storage volumes that are not visible.

  Zones can be configured with either node WWNs or port WWNs or both. Host security groups are configured using port WWNs. Since SOM uses the HSGs to obtain the path information, SOM considers only port WWNs and displays paths with node WWNs as missing paths due to zoning misconfigurations.

  For additional information about the missing LUN paths as well as the connected storage volumes (storage systems, ports) click a host to see the Zoning Misconfiguration table. The information in this table is also available in the **Missing LUN Paths due to Zoning Misconfigurations** tab in the Analysis pane for a selected host.

- **Top 10 Hosts by Shared Storage Volumes Count**
  The top 10 hosts with the total number of storage volumes (LUNs) shared with multiple hosts.

  > **Note:** These volumes are not at risk if the hosts belong to the same cluster.

  For additional information about the top 10 hosts, shared volumes, and the additional hosts and clusters, click a host to see the Shared Storage Volumes dashboard. The information in this dashboard is also available in the **Shared Storage Volumes** tab in the Analysis pane for a selected host.
- **Top 10 Storage Systems by Count of Volumes exposed via Single Storage Port**

  The top 10 storage systems with the total number of storage volumes (LUNs) exposed from single storage system port.

  For additional information about the storage volumes, storage system ports, and the storage systems, click a storage system to see the Storage Volumes exposed via Single Storage Port dashboard. The information in this dashboard is also available in the **Storage Volumes exposed via Single Storage Port** tab in the Analysis pane for a selected storage system.
- **All Hosts by Storage Volumes with Single Point of Failure Ports**
  All hosts with the total number of storage volumes at risk due to single port connectivity.

  For the inventory view of a host, double-click a host to see its Form view.
- **All Hosts by Missing LUN Paths due to Zoning Misconfigurations**
  All hosts with the total number of paths to storage volumes that are not visible.

  For the inventory view of a host, double-click a host to see its Form view.
- **All Hosts by Shared Storage Volumes Count**
  All hosts with the number of volumes that are shared with multiple hosts.

  For the inventory view of a host, double-click a host to see its Form view.
- **All Storage Systems by Count of Volumes exposed via Single Storage Port**

  All storage systems with the number of storage volumes (LUNs) that are exposed from single storage system port.

  For additional properties and related components of an individual file storage system, double-click or  **Open** a storage system to see its form view. The Analysis pane of the Storage System form page displays a table view of Storage Volumes exposed via Single Storage Port for a selected storage system.

Performance Workbench

The Performance Workbench enables you to analyze performance metrics in customizable dashboards. By default, the Performance Workbench comprises predesigned dashboard templates for the top level elements. You can launch the Performance Workbench and view a dashboard template corresponding to the selected element (see "Launch a Saved Dashboard" on page 278). Additionally, as an SOM Administrator, you can also create and customize dashboards for the elements or device components that you monitor. Performance

Workbench can display performance metrics for the last 7 days. For more information about the definitions of the performance metrics, see "Viewing Device Performance in the Analysis Pane" on page 302.

The Performance Workbench helps you perform the following tasks:

- Explore and visualize performance metrics of an element (host, switch, or, storage device) or device components on dashboards
- Visualize the performance of important elements by using personalized favorites
- Embed charts such as Graphs, Single Value charts, Text charts, Pie charts, and Tables into the dashboard or copy the URL of any chart and embed it into your dashboard
- Visualize time-specific data by using Date Range Panel
- Assign dashboards to elements

Performance Workbench can be used by Performance Experts to diagnose specific problems and to troubleshoot device performance.

Launch Performance Workbench

You can launch the Performance Workbench from the workspace navigation panel using one of the following methods:

- From the Analytics and Dashboard view: Click **Analytics and Dashboards** > **Performance Workbench**. The Performance Workbench opens in a new window.
- From the Inventory view:
  a. Click **Inventory** and then navigate to the respective element folder.
  b. Select the element of interest from the table view.
  c. Right-click and then select **Launch Performance Workbench**. The Performance Workbench opens in a new window.

Components of the Performance Workbench

The components that help you to visualize the performance metrics in the Performance Workbench are as follows:

- " Element Explorer Tree" on the next page
- " Dashboard " on the next page

Performance Workbench - the Components

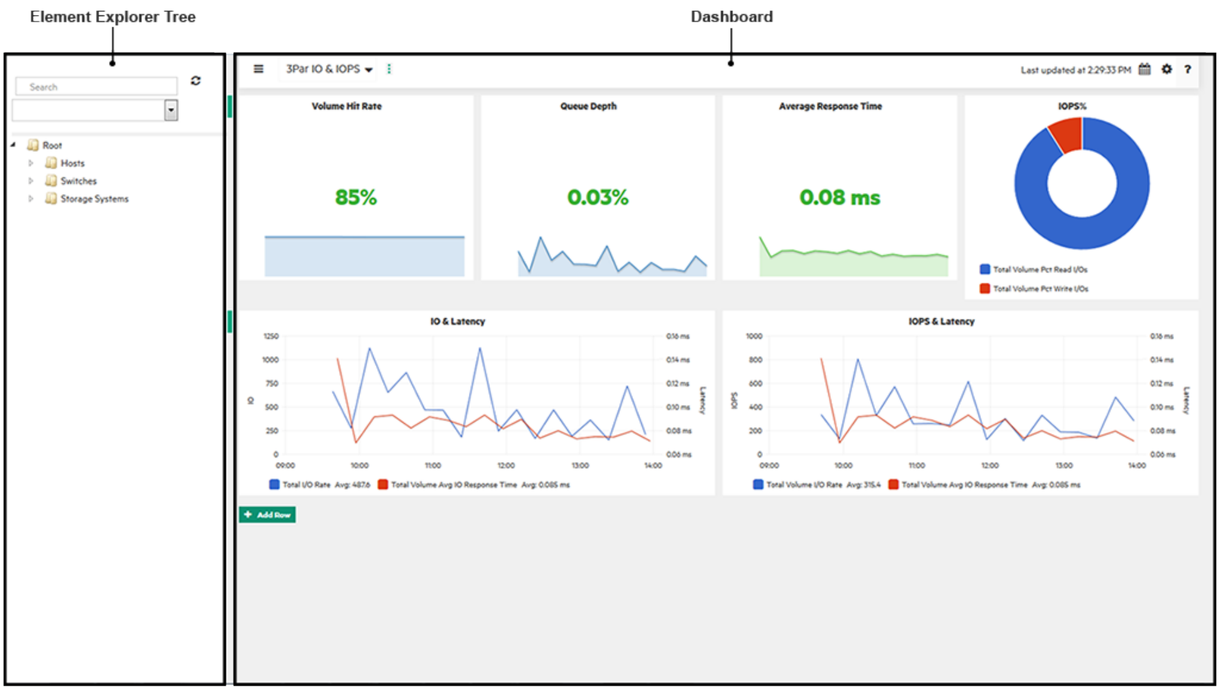

Element Explorer Tree

The element explorer tree displays a list of all discovered elements and their device components. You can select the required device component for which you want to view an existing dashboard or create a new dashboard. You can select only one device component at a time from the element explorer tree to create dashboards.

> **Note:** To find an element or a device component, you can also use the search box at the top of the element explorer tree. However, the search option displays results only from those element category folders that are expanded in the element explorer tree.

Dashboard

The dashboard enables you to create, launch, customize, and visualize charts. The following table explains the options available in the dashboard:

UI Element	Description
≡	Hides/unhides the element explorer tree
▼	For the current element category, it lists the saved dashboards and favorites
⋮	Displays the following options:

UI Element	Description
	• **New Dashboard** () - Creates a new dashboard • **Save** () - Saves the dashboard • **Refresh** () - Refreshes the data when connection is lost and retrieved • **Save as** - Saves a copy of the dashboard • **Clear Cache** - Clears the cache and reloads meta data • **Mark as Default** - Marks a dashboard as the default dashboard for the current element category • **Dashboard JSON** - Displays the contents of the dashboard JSON file • **Export Dashboard** - Exports the dashboard in JSON file format
Last Updated	Shows the last updated time for the dashboard in hh:mm:ss format
	Displays the date range panel
	Displays the dashboard configuration settings
	Displays options to add charts, set the height of a chart, move a row up or down, edit a row (row editor), collapse a row, and delete a row
+ Add Row	Adds a new empty row to the dashboard

Designing a Dashboard

You can design and create any number of dashboards for a given element category or a device component and customize the dashboards as needed.

Prerequisites

Before you design a dashboard, configure the following:

1. "Create a Monitoring Group" on page 192.
2. "Create a Monitoring Policy" on page 193.
3. Allow SOM to collect performance data for at least two consecutive data collection cycles.

Design a Dashboard

To design a dashboard, complete the following tasks:

1. "Create a New Dashboard" on the next page.
2. "Configure Rows and Charts" on page 266.

3. "Edit Charts" on page 268.
4. "Save a Dashboard" on page 276.

Create a New Dashboard

You can create a new dashboard using one of the following methods:

- "From the Analytics and Dashboards View" below
- "From the Inventory View" below

From the Analytics and Dashboards View

To create a new dashboard from the Analytics and Dashboards view, follow these steps:

1. On the workspace navigation panel, click **Analytics and Dashboards** > **Performance Workbench**. The Performance Workbench opens in a new window.

   Note: By default, the element explorer tree is visible. If the element explorer tree is not visible, on the dashboard toolbar, click ☰ to view it.

2. On the element explorer tree, navigate to the desired element category folder, and then select an element or device component. A new dashboard appears.

   If you have already set a default dashboard for the selected element category, performing this step launches the default dashboard. To create a new dashboard from a default dashboard, click **Menu** (⋮) > **New Dashboard** ().

   Note: To find an element or device component, you can also use the search box at the top of the element explorer tree. However, the search option displays results only from those element category folders that are expanded in the element explorer tree.

3. Configure rows and charts on the dashboard (see "Configure Rows and Charts" on the next page).

From the Inventory View

To create a new dashboard from the Inventory view, follow these steps:

1. On the workspace navigation panel, click **Inventory**, and then navigate to the desired element category folder.
2. On the table view, right-click the element of interest, and then select **Launch Performance Workbench**. The Performance Workbench, with a new dashboard, opens in a new window.

   If you have already set a default dashboard for the selected element category, performing this step launches the default dashboard. To create a new dashboard from the default dashboard, click **Menu** (⋮)

> **New Dashboard** ( ).

3. Configure rows and charts on the dashboard (see "Configure Rows and Charts" below).

Configure Rows and Charts

By default, a new dashboard contains an empty graph. You can choose to populate the graph with metrics (see "Edit Charts" on page 268) or you can add more rows and charts to the dashboard, and then configure them as follows:

- "Add a Row" below
- "Add Chart" below
- "Configure Rows" on the next page
- "Configure Charts " on the next page

Note: To replace the default graph with a different chart, delete the graph (click the title of the graph and then click), and then add the required chart. For more information about adding charts, see "Configure Rows and Charts" above.

Add a Row

To add a row to the dashboard, click **+ Add Row**. A new row is added.

Add Chart

To add a chart to the dashboard, follow these steps:

1. On the desired row, point to , and then click > **Add Chart**.
2. Select the required chart type. The available options are:
   - Graph
   - Single Value
   - Text
   - Pie
   - Table

   The selected chart type is added to the row.

   Note: By default, a Single Value chart displays the average metric value.

Configure Rows

To configure a row, point to , click , and then make your selections from the following options:

- Set Height - Select the height of the chart in pixels from the list
- Move - Use this option to move the row up or down in the dashboard
- Row Editor - Gives you the following options:

General

UI Element	Description
Title	Type the title of the row.
Height	Enter the height of the row in pixels
Show Title	If checked, displays the title of the row

Charts

UI Element	Description
Title	Displays the title of the row
Type	Displays the type of chart
Span	Select the span of the chart
↑ or ↓	Moves charts up or down within the row
×	Deletes the chart

- Collapse Row - Collapses the row
- Delete Row - Deletes the row

Configure Charts

To configure a chart, click the title of the chart, and then make your selections from the following options:

UI Element	Description
−	Decreases the width of the chart
+	Increases the width of the chart
×	Deletes the chart

UI Element	Description
≡	Provides the following options: **Chart JSON** - Displays the JSON file for the chart **Export CSV** - Exports the chart in CSV format **Toggle Legend** - Shows or hides the legends
Maximize	Maximizes the chart
Edit	Opens the chart for editing. For more information about editing charts, see "Edit Charts" below
Duplicate	Duplicates the chart
Share	Provides the following share options: • **Current Time Range** - Uses the current time range when the dashboard is shared using the URL. If this option is not selected, you can share the dashboard with the time-range that is already saved. • **Current Panel Only** - Displays only the current chart when the dashboard is shared using the URL. If you do not select this option, then the dashboard will contain all charts. • **Include Parameters** - Displays the existing parameters when the dashboard is shared using the URL. If you do not select this option, the parameters are not included. If the chart does not contain any parameters, then this setting is not applicable.
Add to Favorites	Adds the chart to favorites

Edit Charts

This section provides information about populating the charts with metrics and customizing the charts.

The following table describes the tabs available for each chart:

Tab	Graph	Single Value	Text	Pie	Table
Metrics Tab	✓	✓		✓	✓
Axis & Grid Tab	✓				
Display Styles Tab	✓				

Tab	Graph	Single Value	Text	Pie	Table
Options Tab		✓			
Styles Tab					✓
Link Tab	✓			✓	✓
Edit Text Tab			✓		

To edit a chart, follow these steps:

1. Click the title of the chart, and then click **Edit**.
2. On the tabs that appear, supply the required information. The tabs that appear are as follows:
   - "Metrics Tab" below
   - "Axis & Grid Tab" on page 271
   - "Display Styles Tab" on page 272
   - "Options Tab" on page 273
   - "Styles Tab" on page 274
   - "Link Tab" on page 275
   - "Edit Text Tab" on page 275

Note: To populate a Graph, Single Value chart, Pie chart, and Table, it is necessary to supply the required information in the Metrics tab. Similarly, to populate a Text chart, it is necessary to supply the required information in the Edit Text tab. The information in the other tabs are optional.

Metrics Tab

The Metrics tab is available for the following chart types:

- Graph
- Single Value
- Pie
- Table

The following table describes the fields that appear on this tab::

UI Elements	Display Styles
Required Attributes	
Title	Type the title for the chart.
Data Source	Select a data source from the list. Currently, SOM Performance Workbench supports only SOM Database as the data source.
Class Name	Select a class name from the list. The class names are listed depending on the data source you selected. A class is a group of related metrics.
Metric Name	Select a metric name from the list. The metric names are listed depending on the class name you selected. To view the Metric Description of the selected metric, click . For more information about the definitions of the performance metrics, see "Viewing Device Performance in the Analysis Pane" on page 302.
Optional Attributes	
Label	Type a label. The label is displayed in the chart for the corresponding metric name. The label identifies this metric in the chart.
Actions	Select the action to perform: • **Hide Metric** () - Hides the metric. • **Duplicate Metric** () - Duplicates the metric name and its attributes. This is useful while defining multiple instances for the same metric. • **Remove Metric** () - Deletes the metric.
Add Metric	Click to add another metric to the chart.
Import Graph Template	Import graph templates from a saved file. For more information about importing graph templates, see "Import Graph Templates" on page 280. **Note:** If you are importing graph templates, do not type the title of the chart.

Axis & Grid Tab

This tab is available only for Graphs. The following table describes the fields that appear on this tab:

UI Elements	Description
Left Y and Right Y	These options help you specify the Units, Grid (Max and Min), and labels for the Y-axes. These also help you create dual Y-axis Graphs (see "Create Dual Y-Axis Graphs" on page 276). Based on the selections for Left Y and Right Y, you can make your selections for the following: • **Units** - The units of measurement of the metrics. • **Grid Max** and **Min** - The maximum and minimum values for the grids. All values outside these values are trimmed off from the chart. • **Label** - The labels appear in the chart along the axes.
Show Axis	Displays the selected axis (X-Axis, Y- Axis, or both).
Range	Use this option to make your selections for the following: • **Level 1** and **Level 2** - If the metric values exceed the entered values, the exceeded areas of the chart are applied with the selected colors or a line. •  - Color selection for areas of the chart that exceed Level 1 or Level 2. • **Line Mode** - A horizontal line appears at Level 1 and Level 2 when the metric values exceed them.
Legend	Select the desired legend settings: • **Show** - Displays the legend for the chart. • **Table** - Displays the legend as a table. • **Right Side** - Displays the legend to the right of the chart. • **Hide Empty** - Hides empty legends.
Legend Values	Select the values of the legends that should appear for the given chart. The available options are: • **Min**

UI Elements	Description
	• **Max** • **Avg** • **Current** • **Total**

Display Styles Tab

This tab is available only for Graphs. The following table describes the fields that appear on this tab:

UI Element	Description
Chart Options	Visualize the metrics on the graph as bars, lines, or points.
Line Options	Use this option to make your selections for the following: • **Line Fill** - Fills the area of the graph below the line with the selected fill color intensity. • **Line Width** - Increases or decreases the line widths. • **Null Point Mode** - Displays the null points as Connected, Null, or Null as Zero. • **Staircase Line** - Visualizes the lines on the graph as staircases.
Multiple Metrics	Use this option to make your selections for the following: • **Stack** - Stacks the metric values in the chart. You can also choose to display the stack values as cumulative or individual. • **Percent** - Displays the metric values as percentage.
Tooltip	**All Metrics** - Displays all metrics on the same tooltip with a shared crosshair. This helps to follow all the metrics on the same tooltip.
Metric Specific Overrides	Use this option to add custom styles for each metric using metric override rule as follows: a. Click **Add Metric Override Rule**. b. In the **Label** field, select a metric label from the list. c. Click and choose a custom style from the list. The graph reflects the customized style for the selected metric label.

UI Element	Description
	To delete a metric override rule, click .

Options Tab

This tab is available only for Single Value charts. The following table describes the fields that appear on this tab:

UI Element	Description
Big Value	The Big Value comprises the following options: • **Prefix** - Prefixes the value is with the specified prefix name. For example, if you specify Sample as the Prefix. The chart is displayed as Sample <Value>. • **Value** - Values that need to appear in the chart. You can choose Min, Max, Avg, Current, or Total. • **Postfix** - Postfixes the value with the specified postfix name. For example, if you specify MB as the Postfix. The chart is displayed as <Value>MB.
Font Size	Font size for the Prefix, Value, and Postfix data.
Unit	Units of measurement for the metrics.
Coloring	Select the range options and the coloring options: • **Background** - Displays background color. • **Value** - Displays values in colors. • **Range** - If the metric values exceed the entered values, the exceeded values are applied with the selected colors. • **Colors** - Color selection for the values of the chart that exceed Level 1, Level 2 or Level 3. • **Invert Order** - Inverts the order of the range values (Level 1, Level 2, and Level 3).
Spark Lines	Displays the historical data that provides valuable context at a glance. It does not include X-Axis or Y-Axis. You can specify the following options for spark lines:

UI Element	Description
	• **Show** - Displays the spark lines. • **Background Mode** - Displays the background. • **Line Color** - Line color for spark lines. • **Fill Color** - Fill color for the spark lines.
Value to Text Mapping	This option enables you to map values to text. For example, you can map value such as 12 to appear as N/A. You can do this by using the following options: • ✖ - Deletes the value to text mapping. • **Value** - Value that needs to be mapped to text. • **Text** - The text to which the value is to be mapped. • ✚ - Adds the value to text mapping.

Styles Tab

This tab is available only for Tables. The following table describes the fields that appear on this tab:

UI Elements	Description
Table Type	Choose the table type: • **TimeSeries** - Displays the metric values are for a specific duration. • **Structured** - Displays the latest metric values.
Styling Options	Defines the column width.
Sorting	Enable this option to choose table sorting. By enabling sorting, you can sort the column headers in ascending or descending order. It allows multi-level sorting. For example, you can sort the data of one column and then sort the data of another column.
Formatting	Defines the formatting by choosing the decimal places. The values are updated accordingly.
Time as Date	Displays the time as date. The UNIX timestamp is displayed in a regular date format.

UI Elements	Description
Column Coloring	Select the column coloring options from the following: • **Background** - Displays background color. • **Value** - Displays values in colors. • **Range** - If the metric values exceed the entered values, the exceeded areas of the chart are applied with the selected colors. • **Colors** - Color selection for areas of the chart that exceed Level 1, Level 2 or Level 3. • **Invert Order** - Inverts the order of the range values (Level 1, Level 2, and Level 3).

Link Tab

This tab is available only for the following chart types:

- Graph
- Pie
- Table

Use this tab to provide a link to another dashboard from a chart. This helps to quickly launch the linked dashboard from the current chart. The following table describes the fields that appear on this tab:

UI Elements	Description
Link Label	Type a name for the link label
Dashboard	Select a dashboard from the list. The label with a link to the selected dashboard appears below the chart.
URL	Enter the URL of the dashboard that is to be linked. The label with a link to the provided URL appears below the chart.
	Deletes the link
Add Link	Adds another link

Edit Text Tab

This tab is available only for Texts. To add text as chart, follow these steps:

   a. Type a Title for the chart.
   b. Select the Mode as Markdown, HTML or Text.
   c. In the Content area, type the text.
3. To return to the dashboard, click **Back to Dashboard**. For Text charts, click **Close**.
4. Save the dashboard (see "Save a Dashboard" below).

Create Dual Y-Axis Graphs

Dual Y-Axis Graphs help you visualize on the same Graph, metrics with two different units, or similar metrics with a large difference in the data points between them. To create a dual Y-axis Graphs, follow these steps:

1. On the "Axis & Grid Tab" on page 271, assign appropriate **Units** to the Left Y and Right Y axes to match the units reflected in the Metric Description (to see the Metric Description, on the "Metrics Tab" on page 269, click).
2. On the "Display Styles Tab" on page 272, do the following:
   a. Click **Add Metric Override Rule**.
   b. In the **Label** field, select the metric from the list that you assigned to the Right Y-axis in Step 1 above.
   c. Click > **Y-Axis** > **2**. A Right Y-axis is added to the Graph.

Save a Dashboard

This section provides information that helps you perform the following tasks:

- "Save a New Dashboard" below
- "Save a Copy of a Dashboard" on the next page
- "Mark a Dashboard as Default" on the next page

Save a New Dashboard

To save a new dashboard, follow these steps:

1. On the dashboard toolbar, click **Menu** ( ) > **Save** ( ). The **Save Dashboard** dialog box appears.
2. Type a name for the dashboard, and then click **Save**. A message appears stating that the dashboard is saved.

> **Note:** You can view a list of all saved dashboards for the current element category by clicking . Clicking next to a dashboard name deletes the dashboard. However, the save and the delete options are not available for the predesigned dashboard templates, as they are read-only. If required, you can edit a predesigned dashboard template and save a copy of it.

Save a Copy of a Dashboard

To save a copy of an existing dashboard, follow these steps:

1. On the dashboard toolbar, click ☰. The element explorer tree appears.
2. On the element explorer tree, navigate to the desired element category folder, and then select an element or device component.
3. On the dashboard toolbar, click ▾, and then select a dashboard from the list. The selected dashboard appears.
4. Click **Menu** (⋮) > **Save as**. The **Save a copy of the Dashboard** dialog box appears.
5. Type a name for the dashboard, and then click **Save**. A message appears stating that the dashboard is saved.

Mark a Dashboard as Default

When you mark a dashboard as default for an element or a device component, it becomes the default dashboard for the entire element category. You cannot set different dashboards as defaults for individual elements or device components. For example, if you mark a dashboard as default for one of the storage arrays, the dashboard becomes the default dashboard for the entire Storage Systems element category. Subsequently, selecting any storage element launches the default dashboard automatically.

To mark a dashboard as default for an element category, follow these steps:

1. On the dashboard toolbar, click ☰. The element explorer tree appears.
2. On the element explorer tree, navigate to the desired element category folder, and then select an element or device component.
3. On the dashboard toolbar, click ▾, and then select a dashboard from the list. The selected dashboard appears.
4. Click **Menu** (⋮) > **Mark as default**. A message appears stating that the dashboard is saved as the default dashboard.

Managing a Dashboard

This section provides information about the tasks related to managing dashboards. To manage a dashboard, you can perform the following tasks:

- "Launch a Saved Dashboard" on the next page
- "Add and Use Favorites" on page 279
- "Export a Dashboard" on page 279
- "Import a Dashboard" on page 280
- "Import Graph Templates" on page 280
- "Duplicate a Chart" on page 280

- "Share a Dashboard" on page 280
- "Use the Date Range Panel" on page 281
- "Configure Dashboard Settings" on page 281

Launch a Saved Dashboard

By default, the Performance Workbench comprises predesigned dashboard templates for the top level elements. You can also readily access the dashboards that you have designed or imported, and then saved. Before you launch a dashboard template or a dashboard that you have imported and saved, make sure that the "Prerequisites" on page 264 are met.

You can launch a dashboard template or a dashboard that you have saved, using one of the following methods:

- "From the Analytics and Dashboards View" below
- "From the Inventory View" below

From the Analytics and Dashboards View

To launch a saved dashboard from the Analytics and Dashboards view, follow these steps:

1. On the workspace navigation panel, click **Analytics and Dashboards** > **Performance Workbench**. The performance Workbench opens in a new window.

   > **Note:** By default, the element explorer tree is visible. If the element explorer tree is not visible, on the dashboard toolbar, click ☰ to view it.

2. On the element explorer tree, navigate to the desired element category folder, and then select the element of interest.

   > **Note:** To find an element or device component, you can also use the search box at the top of the element explorer tree. However, the search option displays results only from those element category folders that are expanded in the element explorer tree.

3. On the dashboard toolbar, click ▼, and then select a dashboard from the list. The selected dashboard appears.

From the Inventory View

To launch a saved dashboard from the Inventory view, follow these steps:

1. From the workspace navigation panel, click **Inventory**, and then navigate to the desired element category folder.
2. Select the element of interest from the table view, right-click, and then select **Launch Performance Workbench**. The Performance Workbench opens in a new window.

3. On the dashboard toolbar, click ▼, and then select a dashboard from the list. The selected dashboard appears.

Add and Use Favorites

The Favorites feature helps you to quickly visualize the performance metrics of frequently monitored elements. The Favorites comprises group of charts from different dashboards for faster visualization of performance data.

You can add any number of Favorites. Each user can see the charts that they have added.

You can manage the Favorites from a dashboard using the following steps:

- To add a chart to Favorites:
  a. On a dashboard, click the title of the chart that you want to add to Favorites, and then click **Add to Favorite** > ✚.
  b. Type a name to create a new Favorite or select an existing Favorite from the list.
  c. Click **Save**.
- To view a Favorite:
  a. On the dashboard toolbar, click ▼.
  b. Select a favorite from the list.
- To delete a Favorite:
  a. On the dashboard toolbar, click ▼.
  b. Click 🗑 next to the favorite that you want to delete.

Export a Dashboard

You can export a dashboard as a JSON file to edit or review the dashboard content or to import it later.

To export a dashboard, follow these steps:

1. On the dashboard toolbar, click ☰. The element explorer tree appears.
2. On the element explorer tree, navigate to the desired element category folder, and then select an element or device component.
3. On the dashboard toolbar, click ▼, and then select a dashboard from the list. The selected dashboard appears.
4. On the dashboard toolbar, click **Menu** (⋮) > **Export Dashboard**. The **Export Dashboard** dialog box appears.
5. Click **Export**. A file explorer window appears.
6. Navigate to the location where you want to save the JSON file and then click **Save**.

Import a Dashboard

To import a dashboard from a JSON file, follow these steps:

1. On the dashboard toolbar, click ≡. The element explorer tree appears.
2. On the element explorer tree, navigate to the desired element category folder, and then select an element or device component.
3. On the dashboard toolbar, click ▾, and then click **Import**.
4. Click **Browse**, and select the JSON file that you want to import.
5. Click **Open**. The dashboard is imported.
6. Save the dashboard (see "Save a Dashboard" on page 276).

Import Graph Templates

You can import graph templates to a new dashboard or to an existing dashboard.

To import graph templates to a dashboard, follow these steps:

1. On a dashboard, click the title of the destination graph, and then click **Edit**.
2. On the Metrics tab, click **Import Graph Template**.
3. Navigate to the appropriate graph template.
4. Click **Import**. The graph template is imported.

Duplicate a Chart

You can duplicate the charts within the dashboard. This feature enables you to create a copy of the chart that you can edit and customize easily.

1. On a dashboard, click the title of the chart that you want to duplicate.
2. Click **Duplicate**. A copy of the chart is created and added to the subsequent row in the dashboard.

Share a Dashboard

You can share the URL of a dashboard or charts (graph, single value, text, pie, and table). This feature helps you to quickly and easily share dashboards or charts.

To share a dashboard, follow these steps:

1. On a dashboard, click the title of a chart, and then click **Share**. The following options appear:
   - Current Time Range - Uses the current time range when the dashboard is shared using the URL. If this option is not selected, you can share the dashboard with the time-range that was already saved.
   - Current Chart Only - Displays only the current chart when the dashboard is shared using the URL. If you do not select this option, then the dashboard will contain all charts.

- Include Parameters - Displays the existing parameters when the dashboard is shared using the URL. If you do not select this option, the parameters are not included. If the chart does not contain any parameters, then this setting is not applicable.

2. Copy the URL of the dashboard, and then click **Close**.
3. Share the URL as needed.

Use the Date Range Panel

Use the Date Range Panel to visualize performance data for specific date and time range. To use the Date Range Panel, follow these steps:

1. On the dashboard toolbar, click **Date Range Panel** (). The Date Range Panel appears.
2. On the Date Range Panel, click the date field, and do one of the following:
   - To view data for a previous duration relative to the current time, click one of the predefined time durations. For example, **Last 5 mins** or **Last 1 Day**.
   - To view data for a specific date and time range:
     i. Click **Custom Range**.
     ii. Enter the **Start Date** and time, and the **End Date** and time.

     > **Note:** You can also select a date range by clicking the required dates on the calendar while pressing the Ctrl key on the keyboard.

     iii. Click **Apply**.
3. Drag the along the timeline to view data for a specific time range. You can also expand or contract the time range by dragging the by its edges.

Configure Dashboard Settings

This option enables you to configure and define general dashboard settings. To configure the **Dashboard Settings**, follow these steps:

1. On the dashboard toolbar, click **Configure Dashboard** (). The **Dashboard Settings** panel appears.
2. Make your selections for the dashboard settings from the following tabs:
   - "General Tab" on the next page
   - "Features Tab" on the next page

General Tab

You can use this tab to edit the following properties:

UI Element	Description
Title	The title of the dashboard.
Time Correction	Select Coordinated Universal Time (UTC) or Browser time.
Auto-Refresh	Select the timeline for auto-refresh by choosing 5 minutes, 15 minutes, 30 minutes, 1 hour, 2 hour, or 1 day.
Hide Controls	Check this option to hide controls when viewing the dashboard.

Features Tab

You can use this tab to edit the following properties:

UI Element	Description
Shared Crosshair	Compare the data within a dashboard.

Dashboard JSON

You can use the Dashboard JSON to view and edit the metadata or attributes of the dashboard. The dashboard JSON contains the following attributes:

Element	Attribute
Row	Title
	Height
Charts	Title
	Span
Targets	Data Source
	Class Name
	Metric
	Instance

To view the contents of the Dashboard JSON, click **Menu** ( ) > **Dashboard JSON**.

Establish the Connection to the SOM Reporting Server for Gathering Analytics Data

Some analytics dashboards display information obtained from the SOM reporting server database.

To support this data gathering:

- Port `5433` must be open on the SOM reporting server.
- SOM must be connected to the SOM reporting server database that processes the capacity utilization data exported from the SOM management server.

  To establish this connection, run the following command:

  `somdatatransfercertconfig.ovpl -shrdbconfig` ***<OBR_database_hostname> <OBR_database_port number> <OBR_database_username> <OBR_database_password>***

  Replace *<OBR_database_hostname>* with the IP address or fully qualified domain name of the database server used by the SOM reporting server.

  Replace *<OBR_database_port number>* with the port for connecting to the database used by the SOM reporting server. The default port number is 5433.

  Replace *<OBR_database_username>* with the user name for accessing the database used by the SOMreporting server.

  Replace *<OBR_database_password>* with the password for the specified user name as configured post installation of the SOM reporting server.

  This command is located in the following directory:

  - *Windows*: `%OvInstallDir%\bin`
  - *Linux*: `/opt/OV/bin`

  For more information about `somdatatransfercertconfig.ovpl`, see the *SOM CLI Reference Pages*.

For information about customizing the SOM analytics, see "Configure Analytics" on the next page.

Storage Systems Unused Volumes Analytics

The Storage Systems Unused Volumes Analytics dashboard displays the list of storage systems that have not been accessed for the last 30 days based on the volume performance data.

> **Note:** The data in this dashboard comes from the SOM reporting server. For information about setting up the connection to the SOM reporting server, see "Establish the Connection to the SOM Reporting Server for Gathering Analytics Data" above.

To see the volumes of a storage system that have not encountered I/O operations for the specified number of days, click a storage system name to display the **Analytics - Storage Systems by Unused Volumes Capacity** table.

To see a storage volume's properties and related components, double-click a storage volume to see the "Storage Volume Form" on page 465.

Configure Analytics

Some analytics dashboards display information obtained from the SOM reporting server database. For information about setting up the connection to the SOM reporting server database, see "Establish the Connection to the SOM Reporting Server for Gathering Analytics Data" on the previous page.

The SOM reporting server forecasting algorithm uses a 42-day baseline to compute capacity metric values received from the SOM management server.

You can customize much of the SOM analytics configuration in the following file:

- *Windows*: `%OvDataDir%\conf\som\custom.properties`
- *Linux*: `/var/opt/OV/conf/som/custom.properties`

Available customizations in the `custom.properties` file include:

- Powered-off VM monitoring (`som.vm.analytics.PoweredOffDurationInDays=3`)

  By default, SOM shows data for VMs that have been powered off for more than three days.
- Definition of an unused storage system volume (`analytics.shr.volume.perf.duration.days`)

  By default, SOM defines an unused storage system volume as one that has not been accessed for 30 days.

Additional customization is available in the following file:

- *Windows*: `%OvDataDir%\shared\nnm\conf\props\ovjboss.jvmargs`
- *Linux*: `/var/opt/OV/shared/nnm/conf/props/ovjboss.jvmargs`

The available customization in the `ovjboss.jvmargs` file is:

- Frequency (`analytics.shr.data.pull.interval.millis`)

  By default, the SOM queries the database used by the SOM reporting server every 12 hours (43200000 milliseconds).

Using Topology Maps

The Topology Maps workspace displays the connectivity maps of the top-level elements in the storage infrastructure that your user role is authorized to see.

In a map view, storage systems, hosts, and physical switches are represented pictorially on the map. The connectivity lines between the storage objects represent the connection or relationship between these objects.

To view the topology map of your entire storage infrastructure, from the workspace navigation panel, click **Topology** > **System Topology**. The System Topology map view displays the physical connectivity of the storage elements in your network. You can access storage element nodes and filter the view by fabrics and element types.

You can select a device to view its capacity and performance details in the analysis pane. To see additional properties and related components, either select a device and click **Open** or double-click a device to display its form view.

The System Topology view uses the information gathered from the elements in your storage environment to generate a topology map of the environment. The topology shows the fabric and network connections among the discovered devices. The map view changes dynamically as new devices are discovered in the environment.

For Fabric topology, the port connectivity information is gathered from the fabric name server and then correlated to the devices containing the ports.

For Network Attached Storage (NAS) system topology, the connectivity is established between all NAS devices and the connected hosts in SOM. To view NAS topology alone, select **IP Fabric** in the **Fabric** list.

Only one link is shown between the connected elements in the topology map, even if there are multiple physical connections. The Port Connector form displays details of the connected nodes and the physical port connections between the nodes. You can double-click the connectivity line or path between two nodes to see the Port Connector Form.

The following filter options on the System Topology toolbar enable you to modify the system topology view:

- **Fabric**: Displays the list of discovered fabrics. The system topology map displays the topology of the topmost fabric.
  - Select a particular fabric to see the connectivity among the storage elements within that fabric.
  - Select **Show All** to view the connectivity among all discovered elements in the network.

    Note: The maximum for the **Show All** option is 1000 elements.

  - Select **IP Fabric** to view the connectivity among all NAS devices and the connected hosts discovered in SOM.
- **Show Devices**: Displays the connectivity among the discovered devices as selected. The following options are available:

- **Show All**
- **Hosts + Switches**
- **Arrays + Switches**

For example, if you select **Hosts + Switches**, the storage systems are not displayed in the map. The map shows only the connectivity between the hosts and switches.

Click  **Apply** after you select a filter.

Right-click a device to perform the following tasks:

Task	Description
Open Dashboard	Displays the element dashboard pane.
Start Collection	Triggers data collection for the selected device.
Data Collection Logs	Displays the recent data collection log messages of the selected device in the **Data Collection Logs** window. The log messages can be filtered by **Start Date**, **Start Time**, **End Date**, **End Time**, and **Log Severity** (Info/Severe). You can select **Recent Only** to display the most recent log messages.
Launch Topology	Navigates the topology map of the selected device.
Delete	Deletes the device, its components, and nodes. All historical capacity and performance data is also deleted during this process. To monitor and manage the device again, rediscover the device using the Configuration workspace.

System Topology uses the following icons to depict storage elements on a map:

Icon	Description
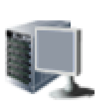	Host. If the host has a question mark and the word “inferred” after its name, the host was discovered through rule-based inference.
	Storage system or subsystem.

Icon	Description
	Switch.

Port Connector Form

The Port Connector form is displayed when you double-click the connecting line or path between two nodes on a map. It displays the details of the connected nodes and the physical port connections between the nodes.

The following details of the port connections between the two nodes are displayed in the table view of the Port Connections tab:

Attribute	Description
Switch Name	The name of the connected switch.
Port Name	The port number on the connected switch.
Port Type	Indicates the type of switch port. For example, F, E, FL, and so on.
Port WWN	The unique 64-bit World Wide Name identifier of the switch port.
Connected Device Name	The name of the other device that is connected.
Connected Port Name	The port number on the connected device.
Connected Port Type	The type of port on the connected device.
Connected Port WWN	The unique 64-bit World Wide Name identifier of the port on the connected device.
Port Link Speed in GBPs	The port link speed in GBPs.

The **Properties** tab displays the following details of the connected nodes:

Attribute	Description
Switch Name	The name of the connected switch.
Connected Device	The name of the other device that is connected.
Connected Device Type	The type of the device that is connected. For example, storage system, host, switch.

Storage System Topology

The Storage System Topology map displays an overview of the hosts and switches connected to a selected storage system.

To navigate to the Storage System Topology map, select **Launch Topology** from the context menu of the storage system.

The **Host Options** on the Storage System Topology toolbar provides the following view filters:

- **Top 25 by Presented Storage**
  Displays the top 25 hosts to which the selected storage system presents storage. This is the default view.
- **Top 25 by Unused Disks**
  Displays the top 25 hosts that are not using the storage presented to them by the selected storage system. This view highlights the hosts from which storage can be reclaimed.

Click **Apply** after you specify a filter to display the connectivity for the selected hosts. You can navigate to the Host Topology map by selecting the **Launch Topology** from the context menu of the host to further analyze the storage configuration.

The **Analysis** pane displays capacity and performance information of the selected storage system.

Storage Virtualizer and Backend Storage Topology

A storage virtualizer is a storage system that uses external, or a combination of local and external storage. For storage virtualizers, the Storage System Topology map displays the connection path between a storage virtualizer and backend storage system including components such as storage extents.

> **Note:** In the System Topology map, the relation between a virtualizer and backend storage system is indicated by a dotted line. To navigate to the Storage System Topology map of the virtualizer, select **Launch Topology** from the context menu of the virtualizer.

For storage virtualizers, the Storage System Topology toolbar displays an additonal menu, **Storage Extent** options. In the **Storage Extent** options menu, select a storage extent to view the connection path between the virtualizer and backend storage system containing the selected storage extent, and then click **Apply**.

The map displays various components within the virtualizer and backend storage system connection path. The connection path may contain various components such as represented below:

Virtualizer < -- > Extents < -- > Virtualizer initiator ports < -- >Backend system port < -- >Backend system volume < -- > Backend system.

To switch back to the default view, select the blank option in the **Storage Extent** menu and click **Apply**.

You can apply either **Storage Extents** or **Host Options** menu options, but not both together.

Host Topology

The Host Topology map displays an overview of the storage systems and switches connected to a selected host (discovered or inferred hosts).

To open the Host Topology map, select **Launch Topology** from the context menu of the Host.

From the Host Topology toolbar, you can either select a volume from **Host Volume** or any one of the options from **Host Volumes Options** for the connectivity map. Click Apply after you specify a filter.

- **Top 25 by Size**
  Displays the top 25 host volumes by size.
- **Top 25 by % Used**
  Displays the top 25 host volumes by the percentage used.
- **Top 25 by % Free**
  Displays the top 25 host volumes by the percentage of free capacity.
- **Show All**
  Displays all the volumes that are visible to the selected host.

On selecting a volume in the map, the storage path details for the volume are shown in the **Analysis** pane.

Note: For Inferred Hosts, the **Host Volume** and **Host Volume Options** are not available in the Host Topology toolbar.

Switch Topology

The Switch Topology map displays the connectivity between a selected physical switch and its logical switches. If you select a logical switch, the switch topology displays the physical switch, storage devices and hosts that are connected to the logical switch.

To navigate the Switch Topology map, select **Launch Topology** from the context menu of the switch. You can launch the topology of a virtual (logical) switch from the context menu of its physical switch.

The **Analysis** pane displays the port utilization details of the selected switch.

Fabric Topology

The Fabric topology map displays the connectivity between the switches, storage systems, and hosts within the selected fabric.

To navigate a Fabric Topology map, use one of the following:

- **The Fabric filter**
  From the Fabric list (System Topology toolbar), select a fabric to see the connectivity among the elements within the fabric.

  or

  Select **Show All** to view the connectivity among all the discovered elements.
- **Launch Topology**
  Select Launch Topology from a Fabric context menu in any view.

By default, System Topology displays the topology of the topmost Fabric.

Only one link is shown between the connected devices even if there are multiple physical connections. For details of the connected devices and the physical port connections between them, double-click the connectivity line to see the Port Connector Form.

Using Inventory Views

The Inventory workspace is a collection of views to access details of storage infrastructure objects (elements) that are discovered by Storage Operations Manager.

Inventory views are categorized into element groups. Each view displays a pre-determined subset of properties of the elements in a group. Inventory form views display additional properties and sub-components of individual elements.

The information in a view is refreshed whenever data collection is triggered based on the freshness threshold that is specified for a data collection policy. The Collection Status indicates the status of data collection for an element.

Use the following inventory views to gain an in-depth understanding of a particular element's properties, and related components:

- "Hosts Views" on page 446
- "Switches View" on page 292
- "Storage Systems Views" on page 293
- "Fabrics View" on page 294
- "Nodes View" on page 294
- "Node Groups View" on page 296
- "FC HBA View" on page 297
- "HBA Ports View" on page 297
- "Switch Ports View" on page 463
- "Storage System Ports View" on page 298

Using the Analysis Pane

Use the Analysis pane to view the following information about a selected device:

- **Summary**
  Key information about a selected element.
  For example:
  - Map Count - The number of Managed Access Points (MAP) for a discovered element.
  - Last Data Collection Time - The time when a storage system, host, switch, or fabric was last contacted for data collection. This value is populated only if data can be collected from an element.
  - Access Point - The IP address that was used to discover and collect data from an element.
  - Data Collection Policy - The policy used to collect data from a discovered element.
  - Node Groups - The node groups to which a discovered element belongs.
- **Capacity**
  Overall capacity utilization of a selected element. For more information, see "Viewing Device Capacity in the Analysis Pane" on page 298.
- **Analytics**
  Analytics information of a selected host. For more information, see "Viewing Host Analytics in the Analysis Pane" on page 301
- **Performance**
  Performance information about a selected element. For more information, see "Viewing Device Performance in the Analysis Pane" on page 302.

Hosts Views

Hosts are categorized into the following views:

- Discovered Hosts
  Includes the list of hosts discovered by Storage Operations Manager. This includes hosts, virtual servers and member nodes that belong to host clusters but not inferred or created hosts. For more information about the properties and components of a selected host, see Viewing Details of Discovered Hosts.
- Virtual Servers
  Includes the list of discovered virtual servers.
  For more information about the properties and components of a selected virtual server, see Viewing Details of Virtual Servers.
- Virtual Machines
  Includes the list of virtual machines hosted on the discovered virtual servers. For more information about the properties and components of a selected host, see Viewing Details of Virtual Machines.

- Inferred Hosts
  Includes hosts inferred based on host security groups, zones, and zone aliases configured in the environment. These hosts are managed without installing a CIM extension. For more information about the properties and components of a selected host, see Viewing Details of Inferred Hosts.
- Created Hosts
  Includes hosts that are created by the administrator using the CLI `somagentlesshostcreator.ovpl`. An administrator can group WWNs and create hosts that contain these WWNs. Host details such as, hostname, IP, DNS, Version, and OS can be specified along with the port WWNs (to be added or deleted) to create such hosts. For more information about the properties and components of a selected host, see Viewing Details of Created Hosts.
- Host Clusters
  Includes host clusters that are discovered through their cluster member nodes. Cluster members are also displayed in the Discovered Hosts inventory view. Use the Host Cluster column in the Discovered Hosts inventory view to link to the host cluster. Information about cluster member nodes and shared resources such as filesystems, disk drives, and volume manager volumes is available in the form view of the host cluster. For more information about the properties and components of a cluster, see Viewing Details of Host Clusters.

The **Analysis** pane displays the Host Capacity and Host Performance Metrics of a selected host.

Switches View

Switches are categorized into the following views:

- **Physical Switches**
  Includes the list of physical switches. Certain vendors such as Cisco, configure the physical switch to be discovered as the top level element. See the Collection Status column in this view to determine if SOM has successfully collected information from a switch after discovery.
- **Virtual Switches**
  Includes the list of virtual switches created on the physical switches. Certain vendors such as Brocade, configure the virtual switch to be discovered as the top level element. Therefore, SOM collects information from the virtual switch listed in this view.

Double-click or **Open** a selected switch in either of the inventory views to see its properties and related components, in the following tab views:

- **Ports**
  Lists the FC ports of a selected switch. Double-click or **Open** a selected switch port to see its properties and connectivity details in "Switch Ports View" on page 463.
- **Virtual Switches**
  Lists the virtual switches created on a selected physical switch. This tab appears only for physical switches with configured virtual switches. Double-click or **Open** a selected virtual switch to see its properties and ports in the Switch Form.

- **Connected Storage Systems**
  Lists the storage systems that are connected to a selected virtual switch. This tab provides details about each pair of ports connected on a storage system and the switch.
- **Asset Record**
  Provides general asset information about a Fabric switch if the switch is an asset, that is, an asset record is created for the switch. For more information about the details that are specified in an asset record, see the "Asset Record Tab" on page 479.
- **Connected Hosts**
  Lists the hosts that are connected to a selected virtual switch. This tab provides details about each pair of ports connected on a host (HBA Port) and the switch.
- **Connected Switches**
  Lists the switches that are connected to a selected switch. This tab provides details about each pair of ports connected on both the switches.
- **Incidents**

  Lists all incidents associated with the selected object. For more information, see "Open Incidents Tab" on page 499

The **Properties** pane displays the properties of a selected switch.

The **Analysis** pane displays the "Switch Capacity" on page 300 and "Performance Collectors for Switches" on page 312 of a selected switch.

Storage Systems Views

Storage systems are categorized into the following inventory views:

- **Top Level Storage Systems**

  Includes top level physical storage systems that can be any of the following:
  - Standalone Storage Systems
    The functionality of standalone storage systems can be broadly categorized as the following:
    - Block Storage
    - File Storage
  - Cluster Storage Systems
    - Cluster storage systems that comprise internal nodes.
    - Distributed storage systems that are logical clusters of multiple storage systems.
    - Hybrid storage systems that are clusters of block and file storage systems.
- **All Storage Systems**

  Includes top level physical storage systems and their underlying internal nodes, and so on. Standalone block and file storage systems are also listed in this view.

For example, a VNX storage system dispalys the top level physical system, the block component, and the filer component. A NetApp cluster displays the top level cluster, nodes, and vservers.

- **Backup Systems**

  Includes physical storage devices which are used as backup devices. For more information about backup systems, see Backup Systems.

The **Analysis** pane displays the Storage System Capacity and Storage System Performance Metrics of a selected storage system.

Fabrics View

The **Fabrics** view displays the list of Fabrics associated with the switches that SOM discovers and manages. A Fibre Channel (FC) Fabric consists of one or more switches that provide optimized interconnections between communicating devices.

Use this view to see the properties of a fabric, the switches, device aliases, zone sets, zones, and zone aliases associated with a fabric.

To see the properties of a fabric and its components, double-click or **Open** a selected fabric for the following tab views:

- "Fabrics View: Switches Tab" on page 517
- "Fabrics View: Device Aliases Tab" on page 517
- "Fabrics View: Zone Aliases Tab" on page 518
- "Fabrics View: Zone Sets Tab" on page 518
- "Fabrics View: Zones Tab" on page 518

The following shows an overview of a Fabric zoning structure.

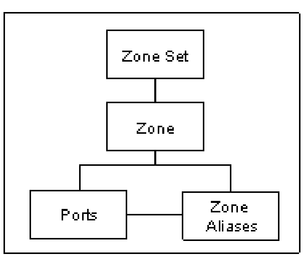

The **Properties** pane displays the properties of a fabric.

Nodes View

The Nodes view displays the list of nodes that are automatically created for each element after an element is successfully discovered.

SOM creates nodes that are associated with the following predefined device categories:

- FC Fabric
- FC Switch (Physical and Virtual)
- Host
- Storage System

Based on its device category, a node is automatically assigned to a node group and consequently scheduled for data collection and performance monitoring.

Double-click or **Open** a node to see its details in the following tab views of the Node Form view:

- Capabilities
- Node Groups
- Registration

If your role permits, you can use the node form to add a node to additional node groups or add notes to communicate information about a node to the team.

The **Basics** pane displays the following properties of a node:

Attribute	Description
Name	The dynamically generated name assigned to this device.
Hostname	The fully-qualified hostname currently stored in the SOM database for this device (according to any hostname resolution strategy currently in use in your network environment; for example, DNS).
Device Profile	Name of the device profile. The device profile comprises the device model, family, vendor, category, and author. The device profile determines how devices of this type are managed, including the icon displayed in topology maps. For more information about the attributes that comprise a device profile, click **Lookup** and select **Open** to display the Device Profile Form .
Notes	Additional information about a node. For example, the location of the node, serial number, if applicable, to which customer, department, or service the node is related, and so on. You could also track maintenance history in this attribute if your role permits you to add the information. A maximum of 1024 characters, alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _+ -) are permitted. **Note**: You can sort the nodes view based on this value. Therefore, you might want to include keywords for this attribute value.

The Analysis pane displays node information in the following tabs:

- **Node Summary**
  Includes the Hostname, Tenant, Security Group, and the number of incidents.
- **Details**
  Includes the Node Management Mode (whether the node is currently managed), Device Profile, Device Category (the nodes view shows an icon for this column), Capabilities (predefined by SOM), and the Status Last Modified (the date and time when the node information was refreshed).
- **Security**
  Includes the security groups (determine the level of security) to which the node belongs and the access privilege.

Node Groups View

The Node Groups view displays the list of node groups that are provided by SOM and those that are created by an administrator.

A node group is a collection of element nodes or child node groups with the same device filters. Element nodes are categorized into node groups to facilitate administration, monitoring, and security to a specific set of nodes.

Node group definitions specify membership using combinations of device filters, such as, device category, vendor, family, and profile. If you provide more than one filter specification for a particular node group, the node group includes nodes that fulfill any one of the device filters.

Note: Additional nodes if specified are included in the node group, regardless of any filters.

SOM uses the Device Category filter to provide the following predefined node groups for discovered elements within the storage infrastructure:

- **All Elements** - Comprises the predefined SOM node groups: FC Fabrics, FC Switches, Hosts, and Storage Systems.
- **FC Fabrics** - All FC fabrics
- **FC Switches** - Physical and virtual FC switches.
- **Hosts** - Physical hosts and virtual servers.
- **Storage Systems** - All storage systems (block, file, and clusters).

Note: Only administrators can create node groups. Default node groups cannot be deleted.

Double-click or **Open** a node group to see its details in the following tab views of the Node Group form view:

- Device Filters
- Additional Filters
- Additional Nodes
- Child Node Groups
- Custom Properties

The **Analysis** pane displays information about a selected node group in the following tabs:

- Node Group Summary
- Node Information – Lists the number of nodes in a selected node group.
- Details – Displays the child node groups within a selected node group.

FC HBA View

The **FC HBA** inventory view displays the total list of host bus adapter cards that are discovered and managed by SOM in the environment.

Double-click a port in the **Ports** tab view to see the properties and ports connected to a selected HBA port in its form view.

For additional properties and the ports of an HBA card, double-click or **Open** a selected card to see the HBA Card Form.

HBA Ports View

The **HBA Ports** inventory view displays the entire list of host bus adapter ports that are discovered and managed by SOM in the environment.

Use this view to see the switch ports or (target) storage system ports that an HBA port is connected to. These ports are visible only if the connected switches and storage systems are discovered by SOM.

For additional properties and the connected ports of a selected HBA port, double-click or **Open** a selected HBA port to see the HBA Port Form.

Switch Ports View

The **Switch Ports** inventory view displays the entire list of switch ports in the environment that are discovered and managed by SOM. Use this view to see the host initiator ports, storage system target ports or other FC switch ports that a switch port is connected to. These ports are visible only if the connected switches, hosts, inferred hosts, or storage systems are discovered by SOM.

To see additional properties and ports connected to a switch port, double-click or **Open** a switch port to see the Switch Port Form.

Double-click a port in the following tabs to see its form view:

- **Connected Switch Ports**
- **Connected Host Ports**
- **Connected Storage System Ports**

The **Properties** pane displays the properties of a selected switch port.

The **Analysis** pane displays the summary details and performance information of a selected switch port.

Storage System Ports View

The **Storage System Ports** inventory view displays the entire list of storage system FC ports in the environment that are discovered and managed by SOM.

To see additional properties and ports connected to a selected storage system port, double-click or **Open** a selected port to see the Storage System Port Form.

In the Storage System Port Form view, the **Properties** pane displays the properties of a selected storage system port.

> **Note:** In the **Properties** pane, the NAS System Node property is not relevant for block storage systems and the Storage System Processor property is not relevant for NAS systems.

Use the form view to see the connected switch ports and host initiator ports in the following tab views:

- Connected Switch Ports
- Connected Host Ports
- Mapped Volume Details - The list of volumes that are mapped to a selected storage system port
- Masked Volume Details - Displays the list of volumes that are masked to the host that is connected to the selected storage system port

Viewing Device Capacity in the Analysis Pane

The following sections provide device-level capacity utilization that is captured by SOM for supported devices.

- Capacity of Hosts
- Capacity of Switches

- Capacity of Storage Systems
- Capacity of Backup Storage Systems

Host Capacity

The overall capacity information of a selected host is available in the **Analysis** pane.

The property `filesystemTypesExcludedForAggregation` is set to Network in the `custom.properties` file and therefore the aggregate capacity at the host level excludes network filesystems such as nfs, nfs4, cifs, smbfs, and ncpfs. The `custom.properties` file is available in the following location:

- *Windows*: `%OvDataDir%\conf\som\custom.properties`
- *Linux*: `/var/opt/OV/conf/som/custom.properties`

The following tabs in the **Analysis** pane contain customizable charts. To customize a chart, see Customize Charts. Mouse over **Refresh** to see the last time the details were updated.

Tab	Description
Host Capacity	Customizable bar chart that illustrates the usage of the following metrics: • **Used Space** - The storage space that is used by the host. • **Total Space** - The total storage space of the host which is the sum of used space and available space. **Note:** Network shares will not be reported in the Host Capacity.
Presented Storage Summary	The total LUN size that each storage system presents to a host. The table displays the **Storage Systems** and the total **Size** of LUNs from each storage system with a bar chart that reflects the same.
Unused Storage	Unused disks on the host. This is the set of storage volumes (LUNs) presented to the host, but not used by Volume Manager, or file systems. These storage volumes can be potentially unmapped on the storage system to reclaim space. The table displays the following: • **Storage Volume** • **Size (GiB)** • **Storage System**
Unused Volume Group	Topmost volume groups by unused capacity. This is the reclaimable space available in a volume group to create more volumes. The table displays the following:

Tab	Description
Capacity	• **Volume Group** • **Available/Grey Space** - a chart that reflects the available space in a volume group. • **Used Space** - aggregated space that is used across all the volumes of a group. **Note:** This tab is not available for virtual servers.
Volumes by % Used	Topmost host volumes used by the selected host element. The table displays **Host Volume**, percentage of used volume capacity (**Used %**), and a bar chart that reflects the percentage of used capacity.

Switch Capacity

The real-time aggregated port statistics of an FC switch are available in the **Port Utilization** tab of the Analysis pane. Mouse over (the Refresh icon) to see the time when the details were last updated.

Tab	Description
Port Utilization	Displays a pie chart to highlight the utilization of FC switch ports using the following metrics: • **Used** - The total number of used ports. • **Free** - The number of free ports that are available for use.

Storage System Capacity

The overall capacity information of a storage system is available in the **Analysis** pane. The tabs in the Analysis pane contain customizable charts that present capacity usage for the last seven days and are dynamically refreshed according to the freshness schedule of the data collection policy.

The capacity information depends on the functionality of storage systems as listed here:

- Capacity of Block Storage Systems
- Capacity of File Storage Systems

The capacity information of the individual components (disk drives, storage pools, volumes, and so on) of a storage system is available in the Properties pane of the component form view.

Capacity Information of Backup Storage Systems

The overall capacity utilization of a backup storage system is available in the following tabs of the **Analysis** pane:

- **Backup System Capacity**

  Displays the capacity utilization of the file systems on a selected backup device using the following metrics:

  - **Used Size**: The space used by compressed data.
  - **Available Size**: The space available for data storage.
  - **Total Size**: The sum total of the Used and Available space on a backup system.

- **Compression**

  The amount of space utilized for data storage before and after compression.
  - **Pre-Compression Size**: Size of data written before compression.
  - **Post-Compression Size**: Space consumed after compression.

- **Total Compression Factor**

  It is the ratio of the pre-compression size to the post-compression size.

> **Note:** For individual file systems, MTrees, and DD Boost units, the capacity information is available in the Analysis pane when selected from the **File Systems**, **MTree**, and **DD Boost** tab views of the backup system.

Viewing Host Analytics in the Analysis Pane

The Analysis pane includes the following tabs with analytics information about a selected host:

Tab	Description
Datastores by % Over Allocation	All the datastores with the percentage of over provisioned storage out of the storage that is actually consumed by a datastore. **Note:** This tab is available only for virtual servers.
Missing LUN Paths due to Zoning Misconfigurations	All the storage volumes that are not visible as the HBA and storage system ports are not configured for the same zone.

Tab	Description
Powered-off VMs by Size	All the powered-off virtual machines with the largest size. **Note:** A powered-off VM is not considered if the time when the VM was last powered off cannot be determined. **Note:** This tab is available only for virtual servers.
Shared Storage Volumes	All the storage volumes that are shared with either standalone hosts or hosts that do not belong to the same cluster as the selected host.
Storage Volumes with Single Point of Failure Ports	All the storage volumes at risk due to single port connectivity.

Viewing Device Performance in the Analysis Pane

You need the following to view performance information from devices that support performance collection:

- SOM Ultimate Perf license
- Monitoring policy associated with the device

The **Analysis** pane displays the performance information of an element. The tabs in the Analysis pane contain customizable charts that display the performance information and metrics used. These charts display data collected at intervals of 15 minutes.

Note: To see data in the charts, the performance collectors must run for a minimum of thirty minutes. However running the collectors for two hours results in a meaningful pattern.

Note: If metric data is not available, the corresponding legend does not appear in the chart.

The **Collector Schedules** tab in the analysis pane displays the monitoring policies configured for a selected element.

Expand to see the properties displayed in the Collector Schedules tab.

- **Monitoring Policy** - The name of the monitoring policy.
- **Collector Name** - The name of the collector group of performance metrics.

- **Device Family** - The family of devices that the device belongs to.
- **Next Run Time** - The time when the next collection is scheduled.
- **Schedule Interval (Minutes)** - The time interval between two subsequent collections.

The following sections provide details about the performance collectors and metrics used for each element.

- "Performance Collectors for Hosts" on the next page
- "Performance Collectors for Switches" on page 312
- "Performance Collectors for HPE 3PAR Arrays" on page 317
- "Performance Collectors for HP StorageWorks EVA Arrays" on page 329
- "Performance Collectors for EMC Symmetrix DMX/VMAX Arrays" on page 341
- "Performance Collectors for CLARiiON and VNX Arrays" on page 354
- "Performance Collectors for NetApp C-mode Clusters" on page 366
- "Performance Collectors for NetApp 7-mode" on page 375
- "Performance Collectors for HDS and HPE XP Arrays" on page 386

Performance Collectors for Hosts

The following performance metrics are available for hosts that are discovered with (agent) and without (agentless) a CIM extension installed on a host:

Discovery Method	Hosts	Disk I/O Rate[2]	Disk Utilization[2]	Disk Latency[2]	Disk Queue Depth[2]	Communication[3]	Data Rate[3]	CPU/ Memory
VMWare API	ESX servers[1]	Y	N	Y	Y	N	Y	Y
Agentless	Windows	Y	Y	Y	N	Y	Y	N
	Linux	Y	Y	Y	Y	Y	Y	N
CIM Agent	Windows	Y	Y	Y	Y	Y	Y	N
	Linux (RHEL, SUSE)	Y	Y	Y	Y	Y	Y	N
	HP-UX	P	Y	N	N	N	N	N
	Solaris	Y	Y	N	N	N	N	N

[1]"ESX Server Performance Collectors" on the next page

[2]"Physical Disk Collectors" on page 307

[3]"HBA Port Performance Collectors" on page 310

Y - All the metrics are available

N - None of the metrics are available

P - Partial metrics are available

ESX Server Performance Collectors

The following collectors (Configuration > Monitoring Settings > Collectors), must be scheduled to measure the performance of disk drives, HBAs, memory and CPU utilization of ESX servers:

- ESX Host CPU Collector
- ESX Host HBA Port Collector
- ESX Host Memory Collector
- ESX Host Physical Disk Collector

The performance metrics are grouped into the following tabs of the **Analysis** pane:

Memory and CPU Utilization

Tab	Metrics	Description	Common Use
CPU	CPU Utilization (%)	The percentage of total CPU utilization for all processes running on a host. Use this metric to identify CPU bottlenecks.	Raw value from host
Memory	Free Physical Memory (KBytes)	Amount of physical memory available. Use this metric to measure available main memory for additional processes and threads.	Total * (% Free Memory)
	Used Physical Memory (%)	Percentage of physical memory being consumed by all processes. Indicates physical memory optimization and availability over a period of time.	(Total - Free / Total) * 100

Disk Drive Metrics

Tab	Metrics	Description	Common Use
Disk IO Rate	Disk Write (KBytes/Sec)	Rate of writing data on the storage path. Use this metric to compare write times for disks on an ESX Server.	Disk Writes for a time period / number of samples
	Disk Read (KBytes/Sec)	Rate of reading data on the storage path Use this metric to compare read times for disks on an ESX Server.	Disk Reads for a time period / number of samples
	Disk Total (KBytes/Sec)	Total read and write requests on the storage path in seconds. Use this metric to test maximum throughput.	Total Disk Reads and Writes for a time period / number of samples
Disk Latency	Disk Write Latency (ms)	Average amount of time for a write issued on the storage path.	Disk Write Latency for a time period / number of samples
	Disk Read Latency (ms)	Average amount of time for a read issued on the storage path.	Disk Read Latency for a time period / number of samples
Disk Queue	Disk Max Queue Depth	Maximum queue depth for a disk drive.	Disk Queue Lengths for a time period / number of samples

HBA Port Metrics

Tab	Metrics	Description	Common Use
Data Rate	Average Write	Average number of write commands issued per second by the storage adapter during the collection interval.	Average Write for a time period / number of samples

Tab	Metrics	Description	Common Use
	Average Read	Average number of read commands issued per second by the storage adapter during the collection interval.	Average Read for a time period / number of samples
	Average Commands	Average number of commands issued per second by the storage adapter during the collection interval.	Average Commands for a time period / number of samples
IO Rate	Write Rate	Rate of writing data by the storage adapter.	Write Rate for a time period / number of samples
	Read Rate	Rate of reading data by the storage adapter.	Read Rate for a time period / number of samples
Latency	Write Latency	Average amount of time for a write operation by the storage adapter.	Write Latency for a time period / number of samples
	Read Latency	Average amount of time for a read operation by the storage adapter.	Read Latency for a time period / number of samples

Physical Disk Collectors

The physical disk collectors for hosts, comprise metrics to measure the performance of the disk drives on a host.

The following available collectors (Configuration > Monitoring Settings > Collectors) must be scheduled to collect performance information:

- Windows Host Physical Disk Collector
- Windows Agentless Disk Observer
- Linux Agent Physical Disk Collector
- Linux Agentless Physical Disk Collector

- HPUX Physical Disk Collector
- Solaris Host Physical Disk Collector

Note: For Linux hosts, install the `sysstat` package to ensure collection of disk performance information.

Note: Performance information is not available for external disk drives on Windows hosts irrespective of the discovery mechanism (that is with or without a CIM extension).

The performance metrics are grouped into the following tabs of the **Analysis** pane:

Disk IO Rate Tab

Metric	Description	Formula
Disk Write (KBytes/Sec)	Rate at which data is transferred to a disk during write operations. Use this metric to compare write times for a given application (for example, writes compared to reads). **Note:** This metric is not available for HP-UX hosts discovered without a CIM agent.	Δ Disk Writes / Δ Time
Disk Read (KBytes/Sec)	Rate at which data is transferred from a disk during read operations. Use this metric to compare read times for a given application (for example, read compared to writes). **Note:** This metric is not available for HP-UX hosts discovered without a CIM agent.	Δ Disk Reads / Δ Time

Metric	Description	Formula
Disk Total (KBytes/Sec)	Total read and write requests in seconds. Use this metric to test maximum throughput.	(Δ Disk Reads + Δ Disk Writes) / Δ Time

Queue Depth Tab

Metric	Description	Formula
Disk Queue Length	Number of outstanding requests on a disk at the time the performance data is collected. It includes requests in service at the time of the snapshot. **Note:** This metric is not available for the following hosts: • Windows hosts discovered without a CIM agent • HP-UX hosts discovered with a CIM agent • Solaris hosts discovered with a CIM agent	Raw value from host

Latency Tab

Metric	Description	Formula
Disk Write Latency (ms)	Average time for a data write operation to a disk.	Raw value from host
Disk Read Latency (ms)	Average time for a data read operation from a disk.	Raw value from host
Total Latency (ms)	Average time required to respond to request, including queue time and service time.	Raw value from host

Disk Utilization Percent Tab

Metric	Description	Formula
Disk Utilization (%)	Based on the IRP (I/O request packets) round trip times the Average Disk Sec/Transfer. Determine the average disk utilization for a given application or known number of processes. Utilization indicates how busy a disk is.	Δ Use (Time for which Disk is in use) / Δ Time

HBA Port Performance Collectors

The host bus adapter (HBA) performance metrics are collected from fabric switch ports. These are inverted and mapped against SAN facing host/array ports and stored along with the hosts/arrays.

The metrics for the bytes that are transmitted and received help determine transmission loads for load balancing, multi-pathing optimization and identifying the need for increased bandwidth. The goal is to determine total IOPS over an interval. If the total IOPS approaches the limitations of an HBA over a sustained period of time, then it might be beneficial to provide an additional path or upgrade the bandwidth.

The following available collectors (Configuration > Monitoring Settings > Collectors) must be scheduled to collect performance information:

- Windows Agent HBA Port Collector
- Windows Agentless HBA Port Collector
- Linux Agent Port Collector
- Linux Agentless Port Collector

Best Practices

Follow these recommended practices for optimal HBA performance:

- Leave the FC ports at auto-negotiate.
- Do not increase the HBAs on a server beyond the limit of the bus throughput.
- Turn on the I/O coalesce parameter in high-performance environments.

- Install the latest HBA firmware and driver.
- Use multipathing software that supports both load balancing and path failover.

HBA Performance Metrics

The performance metrics of an HBA are grouped into the following tabs of the **Analysis** pane:

Communication Tab

Metric	Description	Formulas
Invalid CRC Error Count	The count of the number frames with invalid cyclic redundancy checksums.	Raw value from host
Invalid TX Words Count	The count of the number of invalid transmissions.	Raw value from host
Primitive Sequence Protocol Error Count	The count of the primitive sequence protocol errors.	Raw value from host
Loss of Signal Count	The count of the loss of signals.	Raw value from host
Loss of Sync Count	The count of the loss of synchronizations.	Raw value from host
Link Failure Count	The count of the link failures.	Raw value from host
NOS Count	The number of non-operational state (NOS) primitive sequence events that have occurred on the switched fabric.	Raw value from host
LIP Count	The number of loop initialization primitive (LIP) sequence events that have occurred on a arbitrated loop.	Raw value from host

Data Rate Tab

Metric	Description	Formulas
Dumped Frames Rate (Frames/Sec)	The rate of the number of frames lost due to the lack of available host buffers.	Δ Dumped_Frames / Δ Time
Error Frames Rate (Frames/Sec)	The rate of the number of frames received in error.	Δ Error_Frames / Δ Time
RX_Words Rate (Words/Sec)	The rate of the number of Fibre Channel words received across all protocols and classes.	Δ RX_Words / Δ Time
TX_Words Rate (Words/Sec)	The rate of the total number of Fibre Channel words transmitted across all protocols and classes.	Δ TX_Words / Δ Time
RX_Frames Rate (Frames/Sec)	The rate of the number of Fibre Channel frames received across all protocols and classes.	Δ RX_Frames / Δ Time
TX_Frames Rate (Frames/Sec)	The rate of the total number of Fibre Channel frames transmitted across all protocols and classes.	Δ TX_Frames / Δ Time

Other

Metric	Description	Formulas
Seconds Since Last Reset (Sec)	The number of seconds since the statistics were last reset.	Raw value from host

Performance Collectors for Switches

The following performance collectors (Configuration > Monitoring Settings > Collectors) are available for discovered switches:

- Brocade SMI-S Switch Port Collector
- Cisco Switch Port Collector

Note: For Brocade switches, performance information is available for the linked virtual switches.

The Aggregated Port metrics provide port performance information at a switch level. Whereas, port level metrics (CRC Errors, Link Failures, and so on) collect data only at a port level.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Switch level metrics - Data Rate Tab		
Aggregated Port Bytes Received (MBytes/Sec)	The rate of aggregated bytes received over time. Use this metric to measure inbound traffic for all ports on the switch.	sum of all bytes received (Δ bytes_received_cnt / Δ Time) Aggregated port bytes received = sum of receive rates for all switch ports Switch port receive rate (bytes received) = Δ bytes_received_cnt / Δ Time
Aggregated Port Bytes Transmitted (MBytes/Sec)	Rate of aggregated bytes transmitted over time. Use this metric to measure outbound traffic for all ports on the switch.	sum of all bytes transmitted (Δ bytes_transmitted_cnt / Δ Time) Aggregated port bytes transmitted = sum of transmission rates for all switch ports Switch port transmission rate (bytes transmitted) = Δ bytes_transmitted_cnt / Δ Time
Port Metrics		

Metric	Description	Formula
BB Credit Tab		
Transmit BB Credit Zero	Amount of time that frame transmission is blocked by a transmit credit of zero.	Δ bb_credit_zero_cnt / Δ Time
Communication Tab		
CRC Errors	Number of Cyclic Redundancy Check errors over a period of time. Use this metric to isolate CRC errors on a specific initiator or between devices.	Δ crc_errors / Δ Time
Link Failures	Number of link Failures over a period of time. Use this metric to isolate connection failures and the effect on performance.	Δ link_failures / Δ Time
Data Rate Tab		
Bytes Received (MBytes/Sec)	Number of bytes received over a given interval. Use this metric to measure inbound traffic for specific ports on a switch.	Δ bytes_received_cnt / Δ Time
Bytes Transmitted (MBytes/Sec)	Number of bytes transmitted over a given interval. Use this metric to measure outbound traffic for specific ports on a switch.	Δ bytes_transmitted_cnt / Δ Time
Data Utilization (%) Tab		
Receive Utilization (%)	Utilization percent of the number of received bytes.	((Δ bytes_received_cnt * 100 /(port_speed * Δ Time))

Metric	Description	Formula
Transmit Utilization (%)	Utilization percent of the number of transmitted bytes.	((Δ bytes_transmitted_cnt * 100 / (port_speed * Δ Time))

Switch performance metrics might not be available for all switch vendors.

Best Practices

Switch performance best practices should focus on the establishment of baselines. Use Aggregate Port and Port I/O metrics to establish typical IOPS rates and throughput rates as well as common error rates, average queue depths, and response times. Monitoring SAN switch and overall SAN performance primarily involves three metrics: IOPS (I/O operations per seconds), bandwidth, and latency.

Measuring IOPS and bandwidth can tell you how much work or activity is taking place in the SAN. Measuring latency tells you how effectively the SAN is doing its work, as well as whether the SAN is meeting its service objectives. By using switches and HBAs to view error rates, you can pinpoint the source of SAN performance problems. Error rates can include loss of signal or synchronization, re-transmissions, link failure, or invalid CRC.

Follow these best practices to optimize switch performance:

- Keep the highest performing directors at the core of the SAN.
- Connect storage devices and the highest performing applications to the core.
- Benchmark the performance on oversubscribed ports.
- Leave the Fibre Channel (FC) ports at auto-negotiate for host and storage connections.

Switch Performance Issues

Fibre channel (FC) performance issues can be identified by performing a Cyclic Redundancy Check (CRC). CRC is a method of data integrity assurance across a transmission link. On the transmitting end, a mathematical computation is performed on the bitstream, and the result is added to the data frame. The process is reversed on the receiving end. If the two results do not match, a CRC error is generated, resulting in retransmission of the frame to maintain data integrity.

FC Errors

CRC errors are not the only cause for FC errors. Other types of FC errors could also potentially occur. The following FC errors may be observed due to CRC errors:

- During high I/O traffic
  - FC-attached storage path goes down.
  - CRC errors in conjunction with Microsoft Windows error message: device not accessible.
- After an HBA link reset, no response from ProLiant BL20p G3 server blades and "Link failure," "loss of sync" or "loss of signal" errors logged at the switch.
- Multiple path failures in multi-path environments.

CRC Errors

Brief CRC errors in SOM are a normal occurrence when an HBA is first powered on or off, or when cables are attached or detached. Excessive CRC errors during data transfers can cause performance degradation but do not compromise data integrity.

Link Failure

Link failure is the result of a loss of signal, loss of synchronization, or NOS primitive received. A link failure indicates that a link is actually "broken" for a period of time. It can possibly be due to a faulty connector, media interface adapter (MIA), or cable. The recovery for this type of an error is disruptive. This error is surfaced to the application using the SAN device that encountered this link failure. This causes the system to run degraded until the link recovery is complete. These errors should be monitored closely as they typically affect multiple SAN devices.

I/O Traffic

I/O traffic results have different implications in different operating systems. The Linux and UNIX operating systems bundle small block I/O into large 128 KB block requests, and performance at the upper end of the I/O block spectrum is an important concern. Microsoft Windows, on the other hand, defaults to a maximum I/O block of 64 KB and does not bundle small requests into larger ones.

Performance Collectors for HPE 3PAR Arrays

The following performance collectors (Configuration > Monitoring Settings > Collectors) are available for 3PAR arrays:

- "3PAR SMI-S Storage System Collector" below
- "3PAR SMI-S Controller Collector" on page 320
- "3PAR SMI-S Volume Collector" on page 322
- "3PAR SMI-S Physical Disk Collector" on page 326
- "3PAR SMI-S Fiber Channel Port Collector" on page 328

3PAR SMI-S Storage System Collector

The Storage System Collector metrics are aggregated from the underlying volume statistics.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	**Description**	**Formula**
Data Rate		
Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
I/O Rate		

Metric	Description	Formula
Total I/O Rate (Req/Sec)	Average number of read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time
Queue Depth		
Total Volume Average Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
Total Volume Avg Write IO Response Time (ms)	Average time to complete a write I/O operation.	(Δ WriteIOTimeCounter / 1000) / Δ TotalWriteIOs
Total Volume Avg Read IO Response Time (ms)	Average time to complete a read I/O operation.	(Δ ReadIOTimeCounter / 1000) / Δ TotalReadIOs
Total Volume Avg IO Response Time (ms)	Average time to complete an I/O operation.	(Δ IOTimeCounter / 1000) / Δ TotalIOs
Volume Data Rate		
Total Volume Write Data Rate (Bytes/Sec)	Write throughput rate.	(Δ KBytesWritten * 1024) / Δ Time
Total Volume Read Data Rate (Bytes/Sec)	Read throughput rate.	(Δ KBytesRead * 1024) / Δ Time
Total Volume Data Rate (Bytes/Sec)	Rate data can be transmitted between devices for all volumes.	(Δ KBytesTransferred x 1024) / Δ Time
Volume Data Size		

Metric	Description	Formula
Total Volume Avg Write Size (Bytes)	Average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Total Volume Avg Read Size (Bytes)	Average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
Volume I/O Percent		
Total Volume Percent Hit Rate (%)	Ratio of read and write cache hit rate to total number of I/O operations.	100 * ((Δ ReadHitIOs + Δ WriteHitIOs) / Δ TotalIOs)
Total Volume Average Percent Busy (%)	Average time the storage system was busy.	(Δ PercentBusy) / time
Total Volume Pct Write I/Os (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Total Volume Pct Read I/Os (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
Volume I/O Rate		
Total Volume Read Hit Rate (Req/Sec)	Read cache hit requests per second.	Δ ReadHitIOs / Δ Time
Total Volume Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time
Total Volume Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time

Metric	Description	Formula
Total Volume I/O Rate (Req/Sec)	Average number of I/O operations per second for both sequential and non-sequential read and write operations for all volumes.	Δ TotalIOs / Δ Time

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the HPE 3PAR storage system.

3PAR SMI-S Controller Collector

The controller performance metrics are collected by the SMI-S provider from the underlying port metrics.

The controller performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Write throughput rate.	(Δ KBytesWritten * 1024) / Δ Time
Read Data Rate (Bytes/Sec)	Read throughput rate.	(Δ KBytesRead * 1024) / Δ Time
Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
Average Write Size	Average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs

Metric	Description	Formula
(Bytes)		
Average Read Size (Bytes)	Average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		
Utilization (%)	Utilization rate of the storage system processes.	100 * (Δ Time – (Δ IdleTimeCounter / 1000)) / Δ Time
Percent Hits (%)	Percentage of read and write cache hit rate to total number of I/O operations.	100 * ((Δ ReadHitIOs + Δ WriteHitIOs) / Δ TotalIOs)
Percent Writes (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Percent Reads (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Write Hits (Req/Sec)	The cumulative count of Write Cache Hits (Writes that went directly to Cache).	Δ WriteHitIOs / Δ Time
Read Hits (Req/Sec)	Read cache hit rate.	ReadHitRate = deltaReadHitIOsTotal / duration
Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time
Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time

Metric	Description	Formula
Queue Depth		
Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
Service Time (ms)	The service time since the system start time, for all read and write I/O operations.	Utilization / Total I/O Rate
I/O Response Time (ms)	Time to complete an I/O operation.	(Δ IOTimeCounter / 1000) / Δ TotalIOs

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the HPE 3PAR storage system.

3PAR SMI-S Volume Collector

The volume performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Write throughput rate.	(Δ KBytesWritten * 1024) / Δ Time
Read Data Rate (Bytes/Sec)	Read throughput rate.	(Δ KBytesRead * 1024) / Δ Time

Metric	Description	Formula
Total Data Rate	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
Average Write Size (Bytes)	Average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Average Read Size (Bytes)	Average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		
Volume Percent Hit Rate (%)	Ratio of read and write cache hit rate to total number of I/O operations.	100 * ((Δ ReadHitIOs + Δ WriteHitIOs) / Δ TotalIOs)
Volume Average Percent Busy (%)	Average time the storage volume was busy.	(Δ PercentBusy) / time
Percent Writes (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Percent Reads (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Read Hits (Req/Sec)	Number of read requests (per second) completed from the array cache memory.	Δ ReadHitIOs / Δ Time
Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time
Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time

Metric	Description	Formula
Total I/O Rate (Req/Sec)	Average number of read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time
Queue Depth		
Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
Avg Write IO Response Time (ms)	Average time to complete a write I/O operation.	(Δ WriteIOTimeCounter / 1000) / Δ TotalWriteIOs
Avg Read IO Response Time (ms)	Average time to complete a read I/O operation.	(Δ ReadIOTimeCounter / 1000) / Δ TotalReadIOs
IO Response Time (ms)	Time to complete an I/O operation.	(Δ IOTimeCounter / 1000) / Δ TotalIOs
Aggregate Pools		
Total Pool Average Percent Busy (%)	Average time the pool was busy.	(Δ PercentBusy) / time
Total Pool Average Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Total Pool Avg IO Response Time (ms)	Average time to complete an I/O operation.	(Δ IOTimeCounter / 1000) / Δ TotalIOs
Total Pool Avg Read IO	Average time to complete a read I/O operation.	(Δ ReadIOTimeCounter / 1000) /

Metric	Description	Formula
Response Time (ms)		Δ TotalReadIOs
Total Pool Avg Read Size (Bytes)	Average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
Total Pool Avg Write IO Response Time (ms)	Average time to complete a write I/O operation.	(Δ WriteIOTimeCounter / 1000) / Δ TotalWriteIOs
Total Pool Avg Write Size (Bytes)	Average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Total Pool Data Rate (Bytes/sec)	Rate data can be transmitted between devices for all pools.	(Δ KBytesTransferred x 1024) / Δ Time
Total Pool I/O Rate (Req/Sec)	Average number of I/O operations per second for both sequential and non-sequential read and write operations for all pools.	Δ TotalIOs / Δ Time
Total Pool Pct Read I/Os (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
Total Pool Pct Write I/Os (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Total Pool Percent Hit Rate (%)	Ratio of read and write cache hit rate to total number of I/O operations.	100 * ((Δ ReadHitIOs + Δ WriteHitIOs) / Δ TotalIOs)
Total Pool Read Data Rate (Bytes/sec)	Read throughput rate.	(Δ KBytesRead * 1024) / Δ Time
Total Pool Read Hit Rate	Read cache hit requests per second.	Δ ReadHitIOs / Δ Time

Metric	Description	Formula
(Req/Sec)		
Total Pool Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time
Total Pool Write Data Rate (Bytes/sec)	Write throughput rate.	(Δ KBytesWritten * 1024) / Δ Time
Total Pool Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the HPE 3PAR storage system.

3PAR SMI-S Physical Disk Collector

The Disk Collector metrics are used to understand the performance of the physical disks on the storage system.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Write throughput rate (Bytes per second).	(Δ KBytesWritten * 1024) / Δ Time
Read Data Rate	Read throughput rate (Bytes per second).	(Δ KBytesRead * 1024) / Δ Time

Metric	Description	Formula
(Bytes/Sec)		
Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
Average Write Size (Bytes)	Average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Average Read Size (Bytes)	Average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		
Avg Percent Busy (%)	Time required to complete I/O in seconds	(Δ PercentBusy) / time
Percent Writes (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Percent Reads (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time
Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of read and write I/O operations in requests per second.	Δ TotalIOs / Δ Time
Queue Depth		

Metric	Description	Formula
Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
Avg Write IO Response (ms)	Average time to complete a write I/O operation.	(Δ WriteIOTimeCounter / 1000) / Δ TotalWriteIOs
IO Response Time (ms)	Time to complete an I/O operation.	(Δ IOTimeCounter / 1000) / Δ TotalIOs

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the HPE 3PAR storage system.

3PAR SMI-S Fiber Channel Port Collector

The Port Collector metrics are used to monitor the performance of the FC ports in the array.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate Tab		
Total Data Rate (Bytes/Sec)	The rate that data is transmitted through the selected FC port.	(Δ KBytesTransferred * 1024) / Δ Time
I/O Rate Tab		
Total I/O Rate (Req/Sec)	Average number of read and write I/O operations in requests per second.	Δ TotalIOs / Δ Time

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the HPE 3PAR storage system.

Performance Collectors for HP StorageWorks EVA Arrays

SOM provides the following performance collectors (Configuration > Monitoring Settings > Collectors) for the components of EVA storage arrays:

- EVA SMI-S Storage System Collector
- EVA SMI-S Controller Collector
- EVA SMI-S Volume Collector
- EVA SMI-S Physical Disk Collector
- EVA SMI-S Fiber Channel Port Collector

EVA SMI-S Storage System Collector

The storage system collector provides performance information for HP StorageWorks Enterprise Virtual Arrays (EVA) at the top level.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Total Data Rate (Bytes/Sec)	The rate that data can be transmitted between devices for the storage system.	(Δ KBytesTransferred x 1024) / Δ Time
I/O Rate		

Metric	Description	Formula
Total I/O Rate(Req/Sec)	Average number of I/O operations in requests per second for both sequential and non-sequential reads and writes for the storage system.	Δ TotalIOs / Δ Time
Volume Data Rate		
Total Volume Prefetch Data Rate (Bytes/Sec)	The rate that data is read from the physical disk to cache in anticipation of subsequent reads when a sequential stream is detected.	(Δ PrefetchKBytes x 1024) / Δ Time)
Total Volume Mirror Data Rate(Bytes/Sec)	Rate at which data travels across the mirror port to complete read and write requests to all virtual disks.	(Δ MirrorKBytes x 1024) / Δ Time
Total Volume Flush Data Rate(Bytes/Sec)	Rate at which data is written to physical disks in array.	(Δ FlushKBytes x 1024) / Δ Time)
Total Volume Read Miss Data Rate (Bytes/Sec)	Rate at which data is read from physical disks because the data was not present in the array cache memory.	(Δ ReadMissKBytes x 1024) / Δ Time
Total Volume Read Hit Data Rate (Bytes/Sec)	Rate at which data is read from the array cache memory because of read hit requests.	(Δ ReadHitKBytes x 1024) / Δ Time)
Total Volume Read Data Rate (Bytes/Sec)	Rate data is read from the virtual disk by all hosts and includes transfers from the source array to the destination array.	(Δ KBytesRead x 1024) / Δ Time
Total Volume Write Data Rate (Bytes/Sec)	Rate at which data is written to the virtual disk by all hosts, including transfers from the source array to the destination array.	Δ KBytesWritten x 1024) / Δ Time
Total Volume Data Rate (Bytes/Sec)	Rate data can be transmitted between devices for all volumes.	(Δ KBytesTransferred x 1024) / Δ Time
Volume Data Size		

Metric	Description	Formula
Total Volume Avg Write Size (Bytes)	Average write size for all volumes.	(Δ KBytesWritten x 1024) / Δ WriteIOs
Total Volume Avg Read Size (Bytes)	Average data read size for all volumes.	(Δ KBytesRead x 1024) / Δ ReadIOs
Volume I/O Percent		
Total Volume Pct Write I/Os (%)	Percentage of write I/O operations per second for both sequential and non-sequential writes for all volumes.	100 x (Δ WriteIOs / Δ TotalIOs)
Total Volume Pct Read I/Os (%)	Percentage (%) of read I/O operations per second for both sequential and non-sequential reads for all volumes	100 x (Δ ReadIOs / Δ TotalIOs)
Volume I/O Rate		
Total Volume Read Miss Rate (Req/Sec)	Number of read requests (per second) that were not available from cache memory and therefore were completed from the physical disks instead.	Δ ReadMissRequests / Δ Time
Total Volume Read Hit Rate (Req/Sec)	Number of read requests per second completed from the array cache memory	Δ ReadHitIOs / Δ Time
Total Volume Flush Rate (Req/Sec)	Aggregate of all flush counters: mirror flush, cache flush, host writes to snapshots and snapclones	Δ FlushRequests / Δ Time
Total Volume Write Rate (Req/Sec)	Number of write requests per second completed to a virtual disk that were received from all hosts.	Δ WriteIOs / Δ Time
Total Volume Read Rate (Req/Sec)	Number of read requests per second completed from a virtual disk that were sent to all hosts.	Δ ReadIOs / Δ StatisticTime

Metric	Description	Formula
Total Volume I/O Rate (Req/Sec)	Average number of I/O operations per second for both sequential and non-sequential read and write operations for all volumes.	Δ TotalIOs / Δ Time
Volume Latency		
Total Volume Avg Write Latency (Sec)	Average time to complete a write request (from initiation to receipt of write completion) for all volumes.	(Δ KBytesTransferred x 1024) / Δ Time
Total Volume Avg Read Miss Latency (Sec)	Average time to complete a read request (from initiation to information receipt) from the physical disks for all volumes.	(Δ ReadMissLatency / 1000) / Δ ReadMissIOs
Total Volume Avg Read Hit Latency (Sec)	Average time to complete a read request (from initiation to information receipt) from the array cache memory for all volumes in the array.	(Δ ReadHitLatency / 1000) / Δ ReadHitIOs

EVA SMI-S Controller Collector

SOM monitors the following performance metrics for EVA controllers.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Rate at which data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array	(Δ KBytesWritten x 1024) / Δ Time
Read Data Rate	Rate at which data is read from the controller by all disks	(Δ KBytesRead x 1024) / Δ Time

Metric	Description	Formula
(Bytes/Sec)		
Total Data Rate (Bytes/Sec)	Rate at which data can be transmitted between devices for the controller	(Δ KBytesTransferred x 1024) / Δ Time
Data Size		
Average Write Size (Bytes)	Amount of data written (per second) to physical disks	(Δ KBytesWritten x 1024) / Δ WriteIOs
Average Read Size (Bytes)	Amount of data read (per second) from physical disk	(Δ KBytesRead x 1024) / Δ ReadIOs
I/O Percent		
Percent Writes (%)	Percentage (%) of CPU time dedicated to writes	100 x (Δ WriteIOs / Δ TotalIOs)
Percent Reads (%)	Percentage (%) of CPU time dedicated to reads	100 x (Δ ReadIOs / Δ TotalIOs)
Data Transfer Percent (%)	Similar to % Processor Time except that it does not include time for internal processes not related to host-initiated data transfers	100 x (Δ DataTxCounter / Δ StatisticsTime)
CPU Utilization (%)	Percentage of time that the central processing unit on the controller is active. A completely idle controller shows 0%. A controller saturated with activity shows 100%.	100 x (Δ CpuBusyCounter / Δ StatisticsTime)
I/O Rate		
Write Rate	Number of write requests per second completed to a virtual disk that were received from all	Δ WriteIOs / Δ Time

Metric	Description	Formula
(Req/Sec)	hosts	
Read Rate (Req/Sec)	Rate at which data is read from each host port	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of I/O operations as requests per second for both sequential and non-sequential reads and writes for the controller	Δ TotalIOs / Δ Time
Latency		
Write Latency (Sec)	Average time it takes to complete a write request (from initiation to receipt of write completion)	(Δ WriteLatency / 1000) / Δ WriteIOs
Read Latency (Sec)	Average time it takes to complete a read request (from initiation to receipt of write completion) through the controller	(Δ ReadLatency / 1000) / Δ ReadIOs

EVA SMI-S Volume Collector

SOM monitors performance metrics of HP EVA volumes.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Prefetch Data Rate (Bytes/Sec)	Rate at which data is read from the physical disk to cache in anticipation of subsequent reads when a sequential stream is detected.	(Δ PrefetchKBytes x 1024) / Δ Time)
Mirror Data Rate (Bytes/Sec)	Rate at which data travels across the mirror port to complete read and write requests for the associated virtual disk	(Δ MirrorKBytes x 1024) / Δ Time

Metric	Description	Formula
Flush Data Rate (Bytes/Sec)	Rate at which data is written to a physical disk for the associated virtual disk	(Δ FlushKBytes x 1024) / Δ Time)
Read Miss Data Rate (Bytes/Sec)	Rate at which data is read from physical disks because the data was not present in the array cache memory	(Δ ReadMissKBytes x 1024) / Δ Time
Read Hit Data Rate (Bytes/Sec)	Rate at which data is read from the array cache memory because of read hit requests.	(Δ ReadHitKBytes x 1024) / Δ Time)
Read Data Rate (Bytes/Sec)	Rate at which data is read from the virtual disk by all hosts, including transfers from the source array to the destination array.	(Δ KBytesRead x 1024) / Δ Time
Write Data Rate (Bytes/Sec)	Rate at which data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array.	(Δ KBytesWritten x 1024) / Δ Time
Total Data Rate (Bytes/Sec)	Rate at which data can be transmitted between devices for the host port	(Δ KBytesTransferred x 1024) / Δ Time
Data Size		
Average Write Size (Bytes)	Amount of data written (per second) to physical disks	(Δ KBytesWritten x 1024) / Δ WriteIOs
Average Read Size (Bytes)	Amount of data read (per second) from physical disks	(Δ KBytesRead x 1024) / Δ ReadIOs
I/O Percent		
Percent Writes (%)	Percentage of CPU time dedicated to writes.	100 x (Δ WriteIOs / Δ TotalIOs)

Metric	Description	Formula
Percent Reads (%)	Percentage of CPU time dedicated to reads.	100 x (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Read Miss Rate (Req/Sec)	Number of read requests (per second) that were not available from cache memory and therefore were completed from the physical disks instead.	Δ ReadMissRequests / Δ Time
Read Hits (Req/Sec)	Number of read requests per second completed from the array cache memory.	Δ ReadHitIOs / Δ Time
Flush Rate (Req/Sec)	Aggregate of all flush counters: mirror flush, cache flush, host writes to snapshots and snapclones.	Δ FlushRequests / Δ Time
Write Rate (Req/Sec)	Number of write requests received from all hosts and completed to a virtual disk per second.	Δ WriteIOs / Δ Time
Read Rate (Req/Sec)	Number of read requests that were sent to all hosts from a virtual disk per second.	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of I/O operations in requests per second for both sequential and non-sequential reads and writes for the hostport	Δ TotalIOs / Δ Time
Latency		
Write Latency (Sec)	Average time to complete a write request (from initiation to receipt of write completion)	(Δ WriteLatency / 1000) / Δ WriteIOs
Read Miss Latency (Sec)	Average time it takes to complete a read request (from initiation to information receipt) from the physical disks for all volumes	(Δ ReadMissLatency / 1000) / Δ ReadMissIOs
Read Hit Latency (Sec)	Average time to complete a read request (from initiation to information receipt) from the array volume	(Δ ReadHitLatency / 1000) / Δ ReadHitIOs

EVA SMI-S Physical Disk Collector

The Physical Disk Collector provides performance statistics of EVA physical disks.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Rate at which data is written to the virtual disk by all hosts, including transfers from the source array to the destination array	(Δ KBytesWritten x 1024) / Δ Time
Read Data Rate (Bytes/Sec)	Rate at which data is read from the virtual disk by all hosts, including transfers from the source array to the destination array	(Δ KBytesRead x 1024) / Δ Time
Total Data Rate (Bytes/Sec)	Rate at which data can be transmitted between devices for the host port	(Δ KBytesTransferred x 1024) / Δ Time
Data Size		
Average Write Size (Bytes)	Amount of data written to physical disk	(Δ KBytesWritten x 1024) / Δ WriteIOs
Average Read Size (Bytes)	Amount of data read from physical disk	(Δ KBytesRead x1024) / Δ ReadIOs
I/O Percent		
Percent Writes (%)	Percentage (%) of CPU time dedicated to writes	100 x (Δ WriteIOs / Δ TotalIOs)
Percent Reads (%)	Percentage (%) of CPU time dedicated to reads	100 x (Δ ReadIOs / Δ TotalIOs)

Metric	Description	Formula
I/O Rate		
Write Rate (Req/Sec)	Number of write requests per second completed to a virtual disk that were received from all hosts	(Δ KBytesRead x 1024) / Δ Time
Read Rate (Req/Sec)	Rate at which data is read from each host port	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of I/O operations (requests per second) for both sequential and non-sequential reads and writes for the host port	Δ TotalIOs / Δ Time
Latency		
Write Latency (Sec)	Average time to complete a write request (from initiation to receipt of write completion)	(Δ WriteLatency / 1000) / Δ WriteIOs
Read Latency (Sec)	Average time to complete a read request (from initiation to information receipt) from the array volume	(Δ ReadLatency / 1000) / Δ ReadIOs
Drive Latency (Sec)	Average time to complete read/write requests from the physical disk drive	(Δ DriveLatency / 1000) / Δ TotalIOs
Queue Depth		
Queue Depth	Average number of outstanding requests against the physical disk	Δ DriveQueueDepth / Δ Statistic Time

EVA SMI-S Fibre Channel Port Collector

SOM monitors the performance metrics for EVA FC Ports.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Communication		
Receive Abnormal End of Frame (count)	Number of times a bad frame was detected during data transmission.	–
Protocol Error (count)	Number of errors in the protocol between the channel and the control unit. Use to differentiate between protocol errors and link errors.	–
Loss of Sync (count)	Number of times the receiver logic reports loss of sync has timed-out. Use to determine the number of times an intermittent loss of synchronization in communication signals was received by an enclosure connected to a Fibre Channel (FC) loop.	–
Loss of Signal (count)	Number of times the receiver reports loss of signal. Indicates that fiber optic signal no longer exists. Use to assist in troubleshooting signal loss.	–
Link Fail (count)	Number of link failures. Use to find issues with the fiber optic cable or transceiver or the SAN infrastructure.	–
Discard Frames (count)	Number of frames discarded due to Bad CRCs. Frames are the basic unit of communication between two N_ports, and are composed of a starting delimiter, header, payload, CRC, and end delimiter.	–
Bad Receive	Number of bad receive characters in the bit stream. Use to determine the number of bad frames associated with the Bad CRC metric above.	–

Metric	Description	Formula
Characters (count)		
Bad CRC (count)	Number of bad CRC errors. Indicates that the Cyclic Redundancy Check (CRC) which compares a data stream against a stored checksum, has found the data stream changed and therefore no longer reliable. Use to help the transmitter detect errors in the frame that are caused by bad writes, bad media, damaged links/hardware, excessive link errors, and transfer rates.	–
Queue Depth (count)	Average number of outstanding host requests against all virtual disks accessed through this host port	Δ QDepth / Δ Time
Data Rate		
Write Data Rate (Bytes/Sec)	Rate at which data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array	(Δ KBytesWritten x 1024) / Δ Time
Read Data Rate (Bytes/Sec)	Rate at which data is read from the controller by all disks.	(Δ KBytesRead x 1024) / Δ Time
Total Data Rate (Bytes/Sec)	Rate at which data can be transmitted between devices for the host port.	(Δ KBytesTransferred x 1024) / Δ Time
I/O Rate		
Write Rate (Req/Sec)	Number of write requests per second completed to a virtual disk that were received from all hosts.	(Δ KBytesRead x 1024) / Δ Time

Metric	Description	Formula
Read Rate (Req/Sec)	Rate at which data is read from each host port.	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of I/O operations as requests per second for both sequential and non-sequential reads and writes for the host port.	Δ TotalIOs / Δ Time
Latency		
Write Latency (Sec)	Average time to complete a write request (from initiation to receipt of write completion)	(Δ WriteLatency / 1000) / Δ WriteIOs
Read Latency (Sec)	Average time to complete a read request (from initiation to receipt of write completion) through the controller	(Δ ReadLatency / 1000) / Δ ReadIOs

Performance Collectors for EMC Symmetrix DMX/VMAX Arrays

The following performance collectors (Configuration > Monitoring Settings > Collectors) are available for an EMC Symmetrix array:

- "EMC Symmetrix DMX SMI-S Storage System Collector" below
- "EMC Symmetrix DMX SMI-S Controller Collector" on page 346
- "EMC Symmetrix DMX SMI-S Volume Collector" on page 348
- "EMC Symmetrix DMX SMI-S Fibre Channel Port Collector" on page 353

EMC Symmetrix DMX SMI-S Storage System Collector

The EMC Symmetrix storage system collector includes metrics used to collect and display performance information at the storage system level.

The following table lists the performance metrics of the storage system collector grouped by the tabs in the Analysis pane:

Metric	Description	Formula
Data Rate		
Delayed DFW Rate (Bytes/Sec)	Delayed DFW request rate. A delayed deferred fast write (DFW) is a write-miss. A delayed DFW occurs when the I/O write operations are delayed because the system or device write-pending limit was reached and the cache had to de-stage slots to the disks before the writes could be written to cache.	DelayedDfwRate = deltaEMCDelayedDFWIOs / duration
Deferred Write Rate (Bytes/Sec)	Rate of deferred write request. A deferred write is a write hit. A deferred write occurs when the I/O write operations are staged in cache and will be written to disk at a later time.	DeferredWriteRate = deltaEMCDeferredWriteIOs / duration
Write Flush Data Rate (Bytes/Sec)	Number of tracks written per second from cache to disks.	WriteFlushRate = (deltaEMCWriteKBytesFlushed x 1024) / duration
Write Data Rate (Bytes/Sec)	Write throughput rate.	WriteDataRate = (deltaKBytesWritten x 1024) / duration
Prefetch Data Rate (Bytes/Sec)	Rate of pre-fetched bytes per second.	PrefetchRate = (deltaEMCKBPrefetched * 1024) / duration
Read Data Rate (Bytes/Sec)	Read throughput rate.	ReadDataRate = (deltaKBytesRead x 1024) / duration
Total Data	Total bytes read and written per second.	TotalDataRate =

Metric	Description	Formula
Rate (Bytes/Sec)		(deltaKBytesTransferred x 1024) / duration
Data Size		
Average Write Size (Bytes)	Average write size.	AvgWriteSize = (deltaKBytesWritten x 1024) / deltaTotalWriteIOsRandomAndSeq
Average Read Size (Bytes)	Average read size.	AvgReadSize = (deltaKBytesRead x 1024) / deltaTotalReadIOsRandomAndSeq
I/O Percent		
Percent Write Hits (%)	Percentage of cache write hit I/O operations performed by the Symmetrix device.	PctWriteHitIOs = 100 x (deltaWriteHitIOsTotalRandomAndSeq / deltaTotalWriteIOsRandomAndSeq)
Percent Read Hits (%)	Read cache hit ratio (percentage of read hits).	PctReadHitIOs = 100 x (deltaReadHitIOsTotal / deltaTotalReadIOsRandomAndSeq)
Percent Hits (%)	Ratio of total hits (random and sequential) to total I/Os (random and sequential).	PctHitIOs = 100 x (deltaTotalHitIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)
Percent Writes (%)	Ratio of write I/Os to total I/Os.	PctWriteIOs = 100 x (deltaTotalWriteIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)

Metric	Description	Formula
Percent Reads Seq (%)	Sequential read rate (percentage of sequential reads to Total IOs including Sequential Reads).	PctSeqReadIOs = 100 x (deltaReadIOsSeq / deltaTotalReadIOsRandomAndSeq)
Percent Reads (%)	Ratio of read I/Os to total I/Os.	PctReadIOs = 100 x (deltaTotalReadIOsRandomAndSeq / deltaTotalIOsRandomAndSeq
I/O Rate		
Write Hits (Req/Sec)	Write cache hit rate.	WriteHitRate = deltaWriteHitIOsTotal / duration
Read Data Rate (Req/Sec)	Read throughput rate (Bytes per second).	ReadDataRate = (deltaKBytesRead x 1024) / duration
Write Rate (Req/Sec)	Number of write operations performed each second by the Symmetrix disk.	Req/s Δ WriteIOs / Δ Time
Read Rate Total (Req/Sec)	Read request rate that includes both random and sequential reads.	ReadRateTotal = deltaTotalReadIOsRandomAndSeq / duration
Read Rate Random (Req/Sec)	Random read cache request rate (requests per second).	ReadRate = deltaReadIOs / duration
Total I/O Rate	I/O rate which includes random and sequential reads and writes.	TotalIORate = deltaTotalIOsRandomAndSeq /

Metric	Description	Formula
(Req/Sec)		duration
Pending Count		
Pending Format	Number of format pending tracks. This count can be less than the last-taken statistic; it is a point-in-time value captured at the time the statistics are taken.	PendingFormat = EMCKBPendingFormat x 1024
Pending Flush	Number of tracks in cache that are waiting to be de-staged to disk and cannot be overwritten. This is a point-in-time value captured at the time the statistics are taken.	PendingFlush = EMCKBPendingFlush x 1024
Max Pending Flush Limit	Maximum number of write-pending slots for the entire Symmetrix. System write-pending limit is equal to 80% of the available cache slots. Symmetrix write-pending limit is not simply a sum of all Symmetrix device write-pending slots. It depends on other factors such as cache size and the Symmetrix configuration. System property. This is a point-in-time value captured at the time the statistics are taken.	MaxPendingFlushLimit = EMCMaxKBPendingFlush x 1024
Sequential I/O Rate		
Write Hits Seq (Req/Sec)	Rate of write cache hits per second (sequential hits only).	SeqWriteHitRate = deltaWriteHitIOsSeq / duration
Write Rate Seq (Req/Sec)	Write cache request rate (requests per second) and includes only sequential writes.	SeqWriteRate = deltaWriteIOsSeq / duration
Read Hits Seq (Req/Sec)	Rate of read cache hits per second (sequential hits only).	SeqReadHitRate = deltaReadHitIOsSeq / duration

Metric	Description	Formula
Read Rate Seq (Req/Sec)	Sequential read rate.	SeqReadRate = deltaReadIOsSeq / duration

EMC Symmetrix DMX SMI-S Controller Collector

The Symmetrix controller metrics are used to monitor performance of the front-end controllers in the array.

The following table lists the performance metrics of the front-end controller collector, grouped by the tabs in the Analysis pane:

Metric	Description	Formula
Communication		
System Write Pending Event Rate (Events/Sec)	Number of times in each second when the write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.	SystemWritePendingEventRate = deltaEMCSystemFlushPendingEvents / duration
Device Write Pending Event Rate (Events/Sec)	Number of times in each second when the write-pending limit for a specific Symmetrix device was reached. When the limit is reached, additional write I/O operations are deferred while waiting for data in cache to be destaged to the disk.	DeviceWritePendingEventRate = deltaEMCDeviceFlushPendingEvents / duration
Slot Collision Rate (Slot Collisons/Sec)	Number of slot collisions each second. A slot collision occurs when two or more directors try to access the same cache slot and the slot happens to be locked for an update operation by one of the directors.	SlotCollisionRate = deltaEMCSlotCollisions / duration
Data Rate		

Metric	Description	Formula
Total Data Rate (Bytes/Sec)	Number of Bytes transferred through the Symmetrix Director each second.	TotalDataRate = (deltaKBytesTransferred x 1024) / duration
I/O Percent		
Utilization (%)	Percentage of time that the disks in the array group are busy.	100 * (Δ Time – (Δ IdleTimeCounter / 1000)) / Δ Time
Percent Hits (%)	Percentage of requests performed by the host director and immediately satisfied by cache.	PctHitIOs = 100 x (deltaEmcTotalHitIOs / deltaTotalIOs)
Percent Writes (%)	Percentage of write requests performed by the host director over the sample interval.	PctWriteIOs = 100 x (deltaWriteIOs / deltaTotalIOs)
Percent Reads (%)	Percentage of read requests performed by the host director.	PctReadIOs = 100 x (deltaReadIOs / deltaTotalIOs)
I/O Rate		
Total Hit Rate (Req/Sec)	Number of read and write requests performed each second by the host director that was immediately satisfied by cache.	TotalHitRate = deltaEMCTotalHitIOs / duration
Write Rate (Req/Sec)	Number of write requests performed each second by the host directors.	WriteRate = deltaWriteIOs / duration
Read Rate (Req/Sec)	Number of random read requests performed each second by Symmetrix host director.	ReadRate = deltaReadIOs / duration
Total I/O Rate (Req/Sec)	Number of I/O operations performed each second by the Symmetrix host director. This metric represents activity between the Symmetrix device and the host or SAN device.	TotalIORate = deltaTotalIOs / duration

EMC Symmetrix DMX SMI-S Volume Collector

The Symmetrix volume metrics are used to monitor the performance of the volumes in the array.

The following table lists the performance metrics of the volume collector grouped by the tabs in the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Rate Seq (Bytes/Sec)	Number of sequential write I/O operations performed each second by the Symmetrix device.	SeqWriteRate = deltaWriteIOsSeq / duration
Read Rate Seq (Bytes/Sec)	Number of sequential read I/O operations performed each second by the Symmetrix device.	SeqReadRate = deltaReadIOsSeq / duration
Write Data Rate (Bytes/Sec)	Number of Bytes written by the Symmetrix device each second.	WriteDataRate = (deltaKBytesWritten * 1024) / duration
Read Data Rate (Bytes/Sec)	Number of Bytes read by the Symmetrix device each second.	ReadDataRate = (deltaKBytesRead x 1024) / duration
Total Data Rate (Bytes/Sec)	Total Bytes read and written per second.	TotalDataRate = deltaKBytesTransferred x 1024) / duration
Data Size		
Average I/O	Average size of an I/O operation performed by the Symmetrix device.	AvgIOSize = (deltaKBytesTransferred x

Metric	Description	Formula
Size (Bytes)		1024) / deltaTotalIOsRandomAndSeq
Average Write Size (Bytes)	Average size of a write I/O operation performed by the Symmetrix device.	AvgWriteSize = (deltaKBytesWritten x 1024) / deltaTotalWriteIOsRandomAndSeq
Average Read Size (Bytes)	Average size of a read I/O operation performed by the Symmetrix device.	AvgReadSize = (deltaKBytesRead x 1024) / deltaTotalReadIOsRandomAndSeq
I/O Percent		
Percent Write Miss (%)	Percentage of write I/O operations performed by the Symmetrix device that were write misses.	PctWriteMissIOs = 100 x (deltaWriteMissIOsTotalRandomAndSeq / deltaTotalWriteIOsRandomAndSeq)
Percent Write Hits (%)	Percentage of cache write hit I/O operations performed by the Symmetrix device.	PctWriteHitIOs = 100 x (deltaWriteHitIOsTotalRandomAndSeq / deltaTotalWriteIOsRandomAndSeq)
Percent Read Miss I/Os Total (%)	Percentage of read miss I/O operations performed by the Symmetrix device.	PctReadMissIOsTotal (%) = 100 * (delta Total ReadMissIOs / delta ReadIOsTotal)
Percent Read Hit I/Os Total	Percentage of read hit I/Os (including both random and sequential) operations performed by the Symmetrix device.	PctReadHitIOsTotal (%) = 100 * (delta ReadHitIOsTotal) / delta ReadIOsTotal)

Metric	Description	Formula
(%)		
Percent Read Miss I/Os Random (%)	Ratio of read miss I/Os to Total I/Os.	PctReadMissIOsRandom (%) = 100 * (deltaReadMissIOsRandom / delta IOsTotal)
Percent Read Hit I/Os Random (%)	Ratio of read hit I/Os to Total I/Os.	PctReadHitIOsRandom (%) = 100 * (delta ReadHitIOsRandom / delta IOsTotal)
Percent Miss (%)	Percentage of read and write operations performed by the Symmetrix device that were misses.	PctMissIOs = 100 - PctHitIOs
Percent Hits (%)	Percentage of I/O cache hit operations performed by the Symmetrix device that were immediately satisfied by cache.	PctHitIOs = 100 x (deltaTotalHitIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)
Percent Writes (%)	Percentage of total write I/O operations performed by the Symmetrix device.	PctWriteIOs = 100 x (deltaTotalWriteIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)
Percent Reads (%)	Percentage of read I/O operations performed by the Symmetrix device.	PctReadIOs= 100 x (deltaTotalReadIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)

Metric	Description	Formula
I/O Rate		
Write Hits Seq (Req/Sec)	Rate of write cache hits per second (sequential hits only).	SeqWriteHitRate = deltaWriteHitIOsSeq / duration
Read Hits Seq (Req/Sec)	Rate of read cache hits per second (sequential hits only).	SeqReadHitRate = deltaReadHitIOsSeq / duration
Total Miss Rate (Req/Sec)	Total number of I/O operations (random and sequential) performed each second by the Symmetrix device that were NOT immediately satisfied by cache.	TotalMissRate = TotalIORate - TotalHitRate
Total Hit Rate (Req/Sec)	Total number of I/O operations (random and sequential) performed each second by the Symmetrix device that were immediately satisfied by cache.	TotalHitRate = readHitRateTotalRandomAndSeq + writeHitRateTotalRandomAndSeq
Write Miss Rate (Req/Sec)	Number of write misses that occurred for the Symmetrix device each second.	WriteMissRate = deltaWriteMissIOsTotalRandomAndSeq / duration
Write Hits (Req/Sec)	Write cache hit rate.	WriteHitRate = deltaWriteHitIOsTotal / duration
Read Hit Rate Total (Req/Sec)	Total number of read hit operations (random and sequential) performed each second by the Symmetrix device.	ReadHitRateTotal = deltaReadHitIOsTotalRandomAndSeq / duration
Read Hit	Number of random read hit I/O operations performed each second by the Symmetrix	ReadHitRateRandom = deltaReadHitIOs

Metric	Description	Formula
Rate Random (Req/Sec)	device. The read hits per sec metric for the Symmetrix device statistic does not include sequential read hits. In contrast, the Read Hit Rate Total metric includes random and sequential read hits per second.	/ duration
Write Rate Total (Req/Sec)	Write cache request rate (requests per second) including both random and sequential I/Os performed for the Symmetrix device.	WriteRateTotal = deltaTotalWriteIOsRandomAndSeq / duration
Write Rate (Req/Sec)	Number of write requests performed each second by the host directors.	WriteRate = deltaWriteIOs / duration
Read Rate Total (Req/Sec)	Read request rate including both random and sequential read operations performed each second by the Symmetrix device.	ReadRateTotal = deltaTotalReadIOsRandomAndSeq / duration
Read Rate Random (Req/Sec)	Number of I/O operations performed each second by the Symmetrix device that were random reads. This Random Reads per sec metric for the Symmetrix device statistic does not include sequential reads. In contrast, the Read Rate Total metric includes random and sequential read hits per second.	ReadRateRandom = deltaReadIOs / duration
Total I/O Rate Random (Req/Sec)	Number of I/O operations performed each second by the Symmetrix device, including writes and random reads. In contrast, the Total IO Rate metric includes writes, random reads, and sequential reads.	TotalIORateRandom = deltaTotalIOsRandom / duration
Total I/O Rate (Req/Sec)	Total number of read I/O and write I/O operations (random and sequential) performed each second by the Symmetrix device.	TotalIORate = readRateTotalRandomAndSeq + writeRateTotalRandomAndSeq

Metric	Description	Formula
I/O Time		
Sampled Average Write Time (ms)	Completion time of a write as measured by the host director. Measurements are taken for a sample set of approximately 30% of the I/Os.	SampledAvgWriteTimeMs = current_ EMCSampledWritesTime / current_ EMCSampledWrites
Sampled Average Read Time (ms)	Completion time of a read as measured by the host director. Measurements are taken for a sample set of approximately 30% of the I/Os.	SampledAvgReadTimeMs = curr.getEMCSampledReadsTime(), curr.getEMCSampledReads(), null
Pending Count		
Max Write Pending Threshold	Maximum number of write-pending slots available (expressed in Bytes) for the Symmetrix device.	MaxWritePendingThreshold = current_ EMCMaxKBPendingFlush x 1024
Pending Flush	Number of cache slots (expressed in Bytes) that were write pending for the logical volume at a point in time. This number changes according to the cache de-stage activity rate and the number of writes. A write is pending when it has been written to cache but has not yet been written to the disk.	PendingFlush = current_ EMCKBPendingFlush x 1024

EMC Symmetrix DMX SMI-S Fibre Channel Port Collector

The Symmetrix FC port metrics are used to monitor the performance of the FC ports of the array.

The following table lists the performance metrics collected for Symmetrix FC ports, grouped by the tabs in the Analysis pane:

Metric	Description	Formula
Data Rate		
Total Data Rate (Bytes/Sec)	Number of Bytes transferred through the Symmetrix host port each second.	TotalDataRate = (deltaKBytesTransferred x 1024) / duration
I/O Rate		
Total I/O Rate (Req/Sec)	Number of I/O operations performed each second by the Symmetrix host port. This metric represents activity between the Symmetrix device and the host or SAN device.	TotalIORate = deltaTotalIOs / duration
I/O Size		
Average I/O Size (Bytes)	Average number of Bytes transferred through the Symmetrix host port per I/O operation.	AvgIOSize = (deltaKBytesTransferred x 1024) / deltaTotalIOs

Performance Collectors for CLARiiON and VNX Arrays

The following performance collectors (Configuration > Monitoring Settings > Collectors) are available for CLARiiON and VNX arrays:

- "EMC CLARiiON and VNX SMI-S Storage System Collector" on the next page
- "EMC CLARiiON and VNX SMI-S FrontEnd Controller Collector" on page 356
- "EMC CLARiiON and VNX SMI-S Volume Collector" on page 358
- "CLARiiON and VNX SMI-S Physical Disk Collector" on page 361
- "EMC CLARiiON and VNX SMI-S FrontEnd Port Collector" on page 363

EMC CLARiiON and VNX SMI-S Storage System Collector

The EMC CLARiiON and VNX SMI-S storage system collector includes metrics used to collect and display performance information at the storage system level.

The storage system metrics are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Write throughput rate (Bytes per second).	WriteDataRate = (deltaKBytesWritten x 1024) / duration
Read Data Rate (Bytes/Sec)	Read throughput rate (Bytes per second).	ReadDataRate = (deltaKBytesRead x 1024) / duration
Total Data Rate (Bytes/Sec)	Total bytes read and written per second.	TotalDataRate = (deltaKBytesTransferred x 1024) / duration
Data Size		
Average Write Size (Bytes)	Average write size.	AvgWriteSize = (deltaKBytesWritten x 1024) / deltaTotalWriteIOsRandomAndSeq
Average Read Size (Bytes)	Average read size.	AvgReadSize = (deltaKBytesRead x 1024) / deltaTotalReadIOsRandomAndSeq
I/O Percent		
Percent Hits (%)	Ratio of total hits (random and sequential) to total I/Os (random and sequential).	PctHitIOs = 100 x (deltaTotalHitIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)
Percent Writes (%)	Ratio of write I/Os to total I/Os.	PctWriteIOs = 100 x (deltaTotalWriteIOsRandomAndSeq /

Metric	Description	Formula
		deltaTotalIOsRandomAndSeq)
Percent Reads (%)	Ratio of read I/Os to total I/Os.	PctReadIOs = 100 x (deltaTotalReadIOsRandomAndSeq / deltaTotalIOsRandomAndSeq
I/O Rate		
Write Hits (Req/Sec)	Write cache hit rate.	WriteHitRate = deltaWriteHitIOsTotal / duration
Read Hits (Req/Sec)	Read cache hit rate.	ReadHitRate = deltaReadHitIOsTotal / duration
Write Rate (Req/Sec)	Number of write operations performed each second.	Req/s Δ WriteIOs / Δ Time
Read Rate (Req/Sec)	Number of random read requests performed each second.	ReadRate = deltaReadIOs / duration
Total I/O Rate (Req/Sec)	I/O rate which includes random and sequential reads and writes.	TotalIORate = deltaTotalIOsRandomAndSeq / duration

EMC CLARiiON and VNX SMI-S FrontEnd Controller Collector

The CLARiiON and VNX front-end controller metrics are used to monitor performance of the front-end controllers in the array.

The front-end controller performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Rate at which data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array.	WriteDataRate = (deltaKBytesWritten * 1024) / duration
Read Data Rate (Bytes/Sec)	Rate at which data is read from the virtual disk by all hosts, including transfers from the source array to the destination array.	ReadDataRate = (deltaKBytesRead x 1024) / duration
Total Data Rate (Bytes/Sec)	Host port rate at which data is transmitted between devices.	TotalDataRate = (deltaKBytesTransferred x 1024) / duration
Data Size		
Average Write Size (Bytes)	Amount of data written (per second) to physical disks.	AvgWriteSize = (deltaKBytesWritten x 1024) / deltaTotalWriteIOsRandomAndSeq
Average Read Size (Bytes)	Amount of data read (per second) from physical disks.	AvgReadSize = (deltaKBytesRead x 1024) / deltaTotalReadIOsRandomAndSeq
I/O Percent		
Utilization (%)	Percentage of time that disks in the array group are busy.	100 * (Δ Time – (Δ IdleTimeCounter / 1000)) / Δ Time
Percent Writes (%)	Percentage of CPU time dedicated to writes.	PctWriteIOs = 100 x (deltaWriteIOs / deltaTotalIOs)
Percent	Percentage of CPU time dedicated to reads.	PctReadIOs = 100 x (deltaReadIOs /

Metric	Description	Formula
Reads (%)		deltaTotalIOs)
I/O Rate		
Write Rate (Req/Sec)	Number of completed write requests received per second from all hosts to a virtual disk.	WriteRate = deltaWriteIOs / duration
Read Rate (Req/Sec)	Rate at which data is read from each host port.	ReadRate = deltaReadIOs / duration
Total I/O Rate (Req/Sec)	Average number of I/O operations for both sequential and non-sequential reads and writes for a host port. This metric represents activity between the CLARiiON/VNX device and the host or SAN device.	TotalIORate = deltaTotalIOs / duration
Queue Depth		
Queue Depth	List of tasks in queue.	Total I/O Rate * I/O Response Time
Response Time		
Service Time (ms)	Time taken while controller is in use.	Utilization / Total I/O Rate
I/O Response Time (ms)	Time required to complete a read or write I/O in seconds.	(Δ IOTimeCounter / 1000) / Δ TotalIOs

EMC CLARiiON and VNX SMI-S Volume Collector

The CLARiiON and VNX volume collector provides performance information of the volumes in the array.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Number of Bytes written by the CLARiiON/VNX device each second.	WriteDataRate = (deltaKBytesWritten * 1024) / duration
Read Data Rate (Bytes/Sec)	Number of Bytes read by the CLARiiON/VNX device each second.	ReadDataRate = (deltaKBytesRead x 1024) / duration
Total Data Rate (Bytes/Sec)	Total Bytes read and written per second.	TotalDataRate = deltaKBytesTransferred x 1024) / duration
Data Size		
Average Write Size (Bytes)	Average size of a write I/O operation performed by the CLARiiON/VNX device.	AvgWriteSize = (deltaKBytesWritten x 1024) / deltaTotalWriteIOsRandomAndSeq
Average Read Size (Bytes)	Average size of a read I/O operation performed by the CLARiiON/VNX device.	AvgReadSize = (deltaKBytesRead x 1024) / deltaTotalReadIOsRandomAndSeq
I/O Percent		
Utilization (%)	Percentage of time that disks in the array group are busy.	100 * (Δ Time – (Δ IdleTimeCounter / 1000)) / Δ Time
Percent Hits (%)	Percentage of CPU time dedicated to hits.	PctHitIOs = 100 x (deltaTotalHitIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)

Metric	Description	Formula
Percent Writes (%)	Percentage of CPU time dedicated to writes.	PctWriteIOs = 100 x (deltaTotalWriteIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)
Percent Reads (%)	Percentage (%) of CPU time dedicated to reads .	PctReadIOs= 100 x (deltaTotalReadIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)
I/O Rate		
Write Hits (Req/Sec)	Number of completed write hits requests received per second from all hosts to a virtual disk.	WriteHitRate = deltaWriteHitIOsTotal / duration
Read Hits (Req/Sec)	Number of completed read hits requests received per second from all hosts to a virtual disk.	ReadHitRate = deltaReadHitIOsTotal / duration
Write Rate (Req/Sec)	Number of write requests performed each second by the host directors.	WriteRate = deltaWriteIOs / duration
Read Rate (Req/Sec)	Number of random read requests performed each second by CLARiiON/VNX host director.	ReadRate = deltaReadIOs / duration
Total I/O Rate (Req/Sec)	Total number of read I/O and write I/O operations (random and sequential) performed each second by the CLARiiON/VNX device.	TotalIORate = readRateTotalRandomAndSeq + writeRateTotalRandomAndSeq
Queue Depth		
Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		

Metric	Description	Formula
Service Time (ms)	The service time since the system start time, for all read and write I/O operations.	Utilization / Total I/O Rate
I/O Response Time (ms)	Time to complete an I/O operation.	(Δ IOTimeCounter / 1000) / Δ TotalIOs

CLARiiON and VNX SMI-S Physical Disk Collector

The CLARiiON and VNX physical disk collector metrics are used to monitor performance of the physical disk drives in the array.

The disk performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Number of Bytes written by the CLARiiON/VNX array each second.	WriteDataRate = (deltaKBytesWritten * 1024) / duration
Read Data Rate (Bytes/Sec)	Number of Bytes read by the CLARiiON/VNX device each second.	ReadDataRate = (deltaKBytesRead x 1024) / duration
Total Data Rate (Bytes/Sec)	Number of Bytes transferred through the CLARiiON/VNX Director each second.	TotalDataRate = (deltaKBytesTransferred x 1024) / duration
Data Size		

Metric	Description	Formula
Average Write Size (Bytes)	Average size of a write I/O operation performed by the CLARiiON/VNX device.	AvgWriteSize = (deltaKBytesWritten x 1024) / deltaTotalWriteIOsRandomAndSeq
Average Read Size (Bytes)	Average size of a read I/O operation performed by the CLARiiON/VNX device.	AvgReadSize = (deltaKBytesRead x 1024) / deltaTotalReadIOsRandomAndSeq
I/O Percent		
Utilization (%)	Percentage of time that the disks in the array group are busy.	100 * (Δ Time – (Δ IdleTimeCounter / 1000)) / Δ Time
Percent Writes (%)	Percentage of write requests performed by the host director over the sample interval.	PctWriteIOs = 100 x (deltaWriteIOs / deltaTotalIOs)
Percent Reads (%)	Percentage of read requests performed by the host director.	PctReadIOs = 100 x (deltaReadIOs / deltaTotalIOs)
I/O Rate		
Write Rate (Req/Sec)	Number of write requests performed each second by the host directors.	WriteRate = deltaWriteIOs / duration
Read Rate (Req/Sec)	Number of random read requests performed each second by CLARiiON/VNX host director.	ReadRate = deltaReadIOs / duration
Total I/O Rate (Req/Sec)	Number of I/O operations performed each second by the CLARiiON/VNX host director. This metric represents activity between the CLARiiON/VNX device and the host or SAN device.	TotalIORate = deltaTotalIOs / duration

Metric	Description	Formula
Queue Depth		
Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
Service Time (ms)	The service time since the system start time, for all read and write I/O operations.	Utilization / Total I/O Rate
I/O Response Time (ms)	Time to complete an I/O operation.	(Δ IOTimeCounter / 1000) / Δ TotalIOs

EMC CLARiiON and VNX SMI-S FrontEnd Port Collector

The CLARiiON and VNX SMI-S FrontEnd port metrics are used to monitor the performance of the FC ports of the array.

The performance metrics for ports are grouped into the following tabs in the Analysis pane:

Metric	Description	Formula
Data Rate		
Total Data Rate (Bytes/Sec)	Number of Bytes transferred through the CLARiiON/VNX host port each second.	TotalDataRate = (deltaKBytesTransferred x 1024) / duration
I/O Rate		
Total I/O Rate (Req/Sec)	Number of I/O operations performed each second by the CLARiiON/VNX host port. This metric represents activity between the CLARiiON/VNX device and the host or SAN device.	TotalIORate = deltaTotalIOs / duration

Performance Collectors for EMC Celerra and VNX Arrays

The following performance collectors (**Configuration** > **Monitoring Settings** > **Collectors**) are available for Celerra and VNX arrays:

"EMC Celerra/VNX Data Mover Collector" below

"EMC Celerra Volume Collector" on the next page

EMC Celerra/VNX Data Mover Collector

The Celerra and VNX Data Mover metrics are used to monitor performance of the Data Movers in the EMC VNX array. The Celerra Data Mover performance metrics are grouped into the **Data Mover Resource Usage**, **NFS Usage**, and **CIFS Usage** tabs of the Analysis pane:

Metric	Description	Formula
Data Mover Resource Usage		
CPU Utilization (%)	The percentage of total CPU utilization for all processes running on a Data Mover. Use this metric to identify CPU bottlenecks.	Raw value from device.
Memory Utilization (%)	Percentage of memory being consumed by all processes running on a Data Mover. Indicates memory utilization and availability over a period of time.	Raw value from device.
NFS Usage		
NFS Avg Call Time (micro seconds)	The average call time of the NFS protocol.	Raw value from device.
NFS Write Operations (ops/sec)	The number of NFS write operations processed per second.	Raw value from device.

Metric	Description	Formula
NFS Read Operations (ops/Sec)	The number of NFS read operations processed per second.	Raw value from device
NFS Operations(ops/sec)	The total number of NFS operations processed per second.	NFS Write Operations + NFS Read Operations
CIFS Usage		
CIFS Avg Call Time (sec)	The average call time of the CIFS protocol.	Raw value from device
CIFS Open Connections	The number of CIFS connections for the Data Movers.	Raw value from device
CIFS Open Files	The number of open CIFS files for the Data Movers.	Raw value from device
CIFS Operations (ops/sec)	The number of CIFS operations processed per second.	Raw value from device

EMC Celerra Volume Collector

The Celerra and VNX Volume metrics are used to monitor the usage of volumes in the array.

The Celerra Share performance metrics are grouped into the **NAS Extents** tab of the Analysis pane:

Metric	Description	Formula
NAS Extents		
Read Request (reqs/sec)	The number of read requests processed per second.	Rate (Read requests)
Write Request (reqs/sec)	The number of write requests processed per second.	Rate (Write requests)

Metric	Description	Formula
Read Rate (bytes/sec)	Data read from the volumes per second.	Rate (Bytes read)
Write Rate (bytes/sec)	Data written to the volumes per second.	Rate (Bytes written)

Performance Collectors for NetApp C-mode Clusters

The following performance collectors (Configuration > Monitoring Settings > Collectors) are available for the nodes and vservers in a NetApp C-mode cluster:

- "NetApp Cluster Node Collector" below
- "NetApp Cluster Aggregate Collector" on page 369
- "NetApp Cluster Disk Collector" on page 369
- "NetApp Cluster Vserver Collector" on page 371
- "NetApp Cluster FileSystem Collector" on page 372
- "NetApp Cluster LUN Collector" on page 374

NetApp Cluster Node Collector

The NetApp Cluster Node Collector includes metrics used to collect and display performance information of a cluster node.

The performance metrics for a node are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
CIFS/NFS Operations		
NFS Operation	Number of Network File System (NFS) operations per second.	Δ NFS_Ops / Δ Time

Metric	Description	Formula
(Req/Sec)		
CIFS Operation (Req/Sec)	Number of Common Internet File System (CIFS) operations per second.	Δ CIFS_Ops / Δ Time
Cache Operations		
Name Cache Misses Count (Req/Sec)	Number of name cache misses per second. A cache miss is data that is not in the cache. This results in the system fetching the data from the disk. The name cache behaves in the same way. Use it to determine if the name cache needs to be increased.	Δ Name_Cache_Miss / Δ Time
Name Cache Hits Count (Req/Sec)	Number of name cache hits per second. Use to determine the frequency of the name cache hits. The name cache improves file lookup in a file system.	Δ Name_Cache_Hit / Δ Time
iNode Cache Misses Count (Req/Sec)	Number of inode cache misses per second. A cache miss is data that is not in the cache. This results in the system fetching the data from the disk. The inode cache behaves in the same way. Use it to determine if the inode cache needs to be increased.	Δ Inode_Cache_Miss / Δ Time
iNode Cache Hits Count (Req/Sec)	Number of hits that are cached and subsequently accessed for inodes read from a disk. Use it to determine the increase in file system performance.	Δ Inode_Cache_Hit / Δ Time
Buffer Cache Misses Count (Req/Sec)	Buffer cache miss count per second. A cache miss is data that is not in the cache. This results in the system fetching the data from the disk.	Δ Buff_Miss_Cnt / Δ Time
Buffer Cache Hits Count	Buffer Cache or system memory read cache hits per second. Use it to determine if the access latency is contributing to performance issues.	Δ Buf_Load_Cnt / Δ Time

Metric	Description	Formula
(Req/Sec)		
Latency		
iSCSI Write Latency (ms)	Average latency of write operations observed for all LUNs in the system accessed over iSCSI.	Δ ISCSI_Write_ Latency / Δ ISCI_ Write_Ops
ISCSI Read Latency (ms)	Average latency of read operations observed for all LUNs in the system accessed over iSCSI.	Δ ISCSI_Read _ Latency / Δ ISCI_ Read_Ops
CIFS Latency (ms)	Average latency for Common Internet File System (CIFS) operations in milliseconds.	Δ CIFS_Latency / Δ CIFS_Latency_Base
Processor Utilization		
Average Processor Utilization (%)	Total CPU utilization (%) by all the processes running on a node. Indicates the percentage (%) of time that the processor is active. A completely idle processor shows 0%. A processor saturated with activity shows 100%. Use to identify CPU bottlenecks.	100 x (Δ Processor_ Busy / Δ Processor_ Elapsed_Time)
iSCSI Operations		
iSCSI Write Operation (Req/Sec)	Total number of write operations per second observed for all the LUNs in the system accessed by iSCSI.	Δ ISCSI_Write_Ops / Δ Time
iSCSI Read Operation (Req/Sec)	Total number of read operations per second observed for all the LUNs in the system accessed by iSCSI.	Δ ISCSI_Read_Ops / Δ Time

Metric	Description	Formula
iSCSI Operation (Req/Sec)	Number of Internet Small Computer System Interface (iSCSI) operations per second.	Δ ISCSI_Ops / Δ Time

NetApp Cluster Aggregate Collector

The NetApp Cluster Aggregate Collector includes metrics used to collect and display performance information of a cluster node extent.

The performance metrics for an extent are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Block Rate		
User Write Blocks (Blocks/Sec)	Number of blocks written per second to the aggregate.	Δ User_Write_Blocks / Δ Time
User Read Blocks (Blocks/Sec)	Number of blocks read per second from the aggregate.	Δ User_Read_Blocks / Δ Time
Data Rate		
Total Transfers (Req/Sec)	Total number of transfers per second serviced by the aggregate.	Δ Total_Transfers / Δ Time
User Writes (Req/Sec)	Number of user writes per second to the aggregate.	Δ User_Writes / Δ Time
User Reads (Req/Sec)	Number of user reads per second from the aggregate.	Δ User_Reads / Δ Time

NetApp Cluster Disk Collector

The NetApp Cluster Disk Collector includes metrics used to collect and display performance information of a cluster node disk drive.

The performance metrics for a disk drive are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Block Rate		
User Write Blocks (Blocks/Sec)	Number of blocks transferred for user write operations per second.	Δ User_Write_Blocks / Δ Time
User Read Blocks (Blocks/Sec)	Number of blocks transferred for user read operations per second.	Δ User_Read_Blocks / Δ Time
Data Rate		
Total Transfers (Req/Sec)	Total number of disk operations involving data transfer initiated per second.	Δ Total_Transfers / Δ Time
User Writes (Req/Sec)	Number of disk write operations initiated each second for retrieving data or metadata associated with user requests	Δ User_Writes / Δ Time
User Reads (Req/Sec)	Number of disk read operations initiated each second for retrieving data or metadata associated with user request	Δ User_Reads / Δ Time
Disk Busy		
Disk Busy (%)	Percentage of time there was at least one outstanding request to the disk	100 x (Δ Disk_Busy / Δ Disk_Busy_Base)
Latency		
User Write Latency (ms)	Average latency per block in milliseconds for user write operations.	Δ User_Write_Latency / Δ User_Write_Blocks
User Read Latency (ms)	Average latency per block in milliseconds for user read operations.	Δ User_Read_Latency / Δ User_Read_Blocks

NetApp Cluster Vserver Collector

The NetApp Cluster Vserver Collector includes metrics used to collect and display performance information of a vserver in a cluster.

The performance metrics for a vserver are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
CIFS/NFS Operations		
NFS Operation (Req/Sec)	Number of Network File System (NFS) operations per second.	Δ NFS_Ops / Δ Time
CIFS Operation (Req/Sec)	Number of Common Internet File System (CIFS) operations per second.	Δ CIFS_Ops / Δ Time
Data Latency		
iSCSI Write Latency (ms)	Average latency of write operations observed for all LUNs in the system accessed over iSCSI.	Δ ISCSI_Write_Latency / Δ ISCI_ Write_Ops
ISCSI Read Latency (ms)	Average latency of read operations observed for all LUNs in the system accessed over iSCSI.	Δ ISCSI_Read _Latency / Δ ISCI_Read_Ops
CIFS Latency (ms)	Average latency for Common Internet File System (CIFS) operations in milliseconds.	Δ CIFS_Latency / Δ CIFS_ Latency_Base
iSCSI Operations		
iSCSI Write Operation (Req/Sec)	Total number of write operations per second observed for all the LUNs in the system accessed by iSCSI.	Δ ISCSI_Write_Ops / Δ Time
iSCSI Read Operation (Req/Sec)	Total number of read operations per second observed for all the LUNs in the system accessed by iSCSI.	Δ ISCSI_Read_Ops / Δ Time

Metric	Description	Formula
Bytes Transmission		
Bytes Received (Bytes/Sec)	Number of bytes received per second.	Δ Bytes_Received / Δ Time
Bytes Transmitted (Bytes/Sec)	Number of bytes sent per second.	Δ Bytes_Transmitted / Δ Time

NetApp Cluster FileSystem Collector

The NetApp Cluster FileSystem Collector includes metrics used to collect and display performance information of a vserver file system.

The performance metrics for a file system are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Data Operations		
Total Operations (Req/Sec)	Total number of operations per second serviced by a volume.	Δ Total_Ops / Δ Time
Other Operations (Req/Sec)	Number of other operations per second to a volume.	Δ Other_Ops / Δ Time
Write Operations (Req/Sec)	Number of writes per second to a volume.	Δ Write_Ops / Δ Time

Metric	Description	Formula
Read Operations (Req/Sec)	Number of reads per second to a volume.	Δ Read_Ops / Δ Time
Data Rate		
Write Data (Bytes/Sec)	Number of bytes written per second to a volume.	Δ Write_Data / Δ Time
Read Data (Bytes/Sec)	Number of bytes read per second from a volume.	Δ Read_Data / Δ Time
Latency		
Other Latency (ms)	Average latency time for other writes to the volume in milliseconds.	Δ Other_ Latency / Δ Other_Ops
Read Latency (ms)	Average latency time for reads to the volume in milliseconds.	Δ Read_ Latency / Δ Read_Ops
Write Latency (ms)	Average latency time for writes to the volume in milliseconds.	Δ Write_ Latency / Δ Write_Ops
Average Latency (ms)	Average latency in milliseconds for all operations on a volume.	Δ Avg_Latency / Δ Total_Ops
iNode		

Metric	Description	Formula
Used Inodes	Total number of inodes that are currently used. Can be used to alert an admin when the inode utilization approaches the total number of available nodes.	NA
Total Inodes	Total number of inodes. Inodes are file system data structures or metadata used to store basic file data like ownership and file permissions. Can be used to view the inode limit and determine if more are needed.	NA
Reserved Inodes	Provides a count of reserved inodes in a file system. The first ten (10) inodes of a file system are special inodes. Inodes 7-10 are reserved and usually not used.	NA

NetApp Cluster LUN Collector

The NetApp Cluster LUN Collector includes metrics used to collect and display performance information of a vserver volume.

The performance metrics for a volume are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Data Operations		
Queue Full Rate (Req/Sec)	Number of queue full responses per second.	Δ Queue_Full / Δ Time
Total Operations (Req/Sec)	Total number of operations on the LUN per second.	Δ Total_Ops / Δ Time
Other Operations (Req/Sec)	Number of other operations per second.	Δ Other_Ops / Δ Time
Write Operations (Req/Sec)	Number of write operations per second.	Δ Write_Ops / Δ Time
Read Operations (Req/Sec)	Number of read operations per second.	Δ Read_Ops / Δ Time
Data Rate		

Metric	Description	Formula
Write Data (Bytes/Sec)	Number of bytes written per second.	Δ Write_Data / Δ Time
Read Data (Bytes/Sec)	Number of bytes read per second.	Δ Read_Data / Δ Time
Latency		
Average Latency (ms)	Average latency in milliseconds for all operations on the LUN.	Δ Avg_Latency / Δ Total_Ops

Performance Collectors for NetApp 7-mode

The following performance collectors (Configuration > Monitoring Settings > Collectors) are available for a NetApp 7-mode NAS device:

- "NetApp System Collector" below
- "NetApp Filesystem Collector" on page 379
- "NetApp Qtree Collector" on page 381
- "NetApp Aggregate Collector" on page 381
- "NetApp LUN Collector" on page 382
- "NetApp Disk Drive Collector" on page 383
- "NetApp IPPort Collector" on page 384

NetApp System Collector

The NetApp System Collector includes metrics used to collect and display performance information of a NetApp 7-mode NAS device.

The performance metrics for a NetApp 7-mode NAS device are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
CIFS/NFS Operations		
NFS Operations (Req/Sec)	Number of Network File System (NFS) operations per second.	Δ NFS_Ops / Δ Time
CIFS Operations (Req/Sec)	Number of Common Internet File System (CIFS) operations per second.	Δ CIFS_Ops / Δ Time
Cache Operations		
Name Cache Misses Count (Req/Sec)	Number of name cache misses per second. A cache miss is data that is not in the cache. This results in the system fetching the data from the disk. The name cache behaves in the same way. Use it to determine if the name cache needs to be increased.	Δ Name_Cache_Miss / Δ Time
Name Cache Hits Count (Req/Sec)	Number of name cache hits per second. Use to determine the frequency of the name cache hits. The name cache improves file lookup in a file system.	Δ Name_Cache_Hit / Δ Time
iNode Cache Misses Count (Req/Sec)	Number of inode cache misses per second. A cache miss is data that is not in the cache. This results in the system fetching the data from the disk. The inode cache behaves in the same way. Use it to determine if the inode cache needs to be increased.	Δ Inode_Cache_Miss / Δ Time
iNode Cache Hits Count (Req/Sec)	Number of hits that are cached and subsequently accessed for inodes read from a disk. Use it to determine the increase in file system performance.	Δ Inode_Cache_Hit / Δ Time
Buffer Cache Misses Count	Buffer cache miss count per second. A cache miss is data that is not in the cache. This results in the system fetching the data from the disk.	Δ Buff_Miss_Cnt / Δ Time

Metric	Description	Formula
(Req/Sec)		
Buffer Cache Hits Count (Req/Sec)	Buffer Cache or system memory read cache hits per second. Use it to determine if the access latency is contributing to performance issues.	Δ Buf_Load_Cnt / Δ Time
Latency		
iSCSI Write Latency (ms)	Average latency of write operations observed over all LUNs in the system accessed over iSCSI in milliseconds.	Δ ISCSI_Write_Latency / Δ ISCI_Write_Ops
iSCSI Read Latency (ms)	Average latency of read operations observed over all LUNs in the system accessed over iSCSI in milliseconds.	Δ ISCSI_Read _ Latency / Δ ISCI_ Read_Ops
NFSv3 Average Operations Latency (ms)	Average latency of the NFS v3 operations in milliseconds.	Δ NFSv3_Avg_Op _ Latency / Δ NFSv3_ Avg_Op_Latency_Base
NFSv3 Write Latency (ms)	Average latency for NFS v3 write operations in milliseconds.	Δ NFSv3_Write _ Latency / Δ NFSv3_ Avg_Write_Latency_ Base
NFSv3 Read Latency (ms)	Average latency for NFS v3 read operations in milliseconds.	Δ NFSv3_Read _ Latency / Δ NFSv3_ Avg_Read_Latency_ Base

Metric	Description	Formula
CIFS Latency (ms)	Average latency for Common Internet File System (CIFS) operations in milliseconds.	Δ CIFS_Latency / Δ CIFS_Latency_Base
Processor Utilization		
Processor Utilization (%)	Total CPU utilization (%) by all the processes running on a node. Indicates the percentage (%) of time that the processor is active. A completely idle processor shows 0%. A processor saturated with activity shows 100%. Use to identify CPU bottlenecks.	100 x (Δ Processor_ Busy / Δ Processor_ Elapsed_Time)
iSCSI Operations		
iSCSI Write Operations (Req/Sec)	Total number of write operations per second observed for all the LUNs in the system accessed by iSCSI.	Δ ISCSI_Write_Ops / Δ Time
iSCSI Read Operations (Req/Sec)	Total number of read operations per second observed for all the LUNs in the system accessed by iSCSI.	Δ ISCSI_Read_Ops / Δ Time
iSCSI Operations (Req/Sec)	Number of Internet Small Computer System Interface (iSCSI) operations per second.	Δ ISCSI_Ops / Δ Time
iSCSI Rate		
iSCSI Write Data (Bytes/Sec)	iSCSI bytes written per second.	Δ ISCSI_Write_Data / Δ Time
iSCSI Read Data	iSCSI bytes read per second.	Δ ISCSI_Read_Data / Δ Time

Metric	Description	Formula
(Bytes/Sec)		

NetApp Filesystem Collector

The NetApp IPPort Collector includes metrics used to collect and display performance information of a selected NetApp 7-mode NAS filesystem.

The performance metrics for a NetApp 7-mode NAS filesystem are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Data Operations		
Total Operations (Req/Sec)	Total number of operations per second serviced by a volume.	Δ Total_Ops / Δ Time
Other Operations (Req/Sec)	Number of other operations per second to a volume.	Δ Other_Ops / Δ Time
Write Operations (Req/Sec)	Number of writes per second to a volume.	Δ Write_Ops / Δ Time
Read Operations (Req/Sec)	Number of reads per second to a volume.	Δ Read_Ops / Δ Time
Data Rate		
Write Data (Bytes/Sec)	Bytes written per second to a volume.	Δ Write_Data / Δ Time
Read Data (Bytes/Sec)	Bytes read per second from a volume.	Δ Read_Data / Δ Time

Metric	Description	Formula
INodes		
Total INodes	Count of total number of inodes. Inodes are file system data structures or metadata used to store basic file data like ownership and file permissions. Can be used to view the inode limit and determine if more are needed.	NA
Reserved INodes	Count of reserved inodes in a file system. The first 10 inodes on a file system are special inodes. Inodes 7-10 are reserved and usually not used.	NA
Used INodes	Count of total number of inodes that are currently used. Can be used to highlight the utilization of inodes if it approaches the total available inodes.	NA
Latency		
Average Latency (ms)	Average latency in milliseconds for all operations on a volume.	Δ Avg_Latency / Δ Total_Ops
Other Latency (ms)	Average latency time for other writes to a volume in milliseconds.	Δ Other_Latency / Δ Other_Ops
Write Latency(ms)	Average latency time for writes to a volume in milliseconds.	Δ Write_Latency / Δ Write_Ops
Read Latency (ms)	Average latency time for reads to a volume in milliseconds.	Δ Read_Latency / Δ Read_Ops

NetApp Qtree Collector

The NetApp Qtree Collector includes metrics used to collect and display performance information of a selected NetApp 7-mode NAS device qtree.

The performance metrics for a NetApp 7-mode NAS device qtree are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Internal Operations (Req/Sec)	Number of internal operations generated by activities such as snapmirror and backup per second to a qtree.	Δ Internal_Ops / Δ Time
CIFS Operations (Req/Sec)	Number of Common Internet File System (CIFS) operations per second to a qtree.	Δ CIFS_Ops / Δ Time
NFS Operations (Req/Sec)	Number of NFS operations per second to a qtree.	Δ NFS_Ops / Δ Time

NetApp Aggregate Collector

The NetApp Aggregate Collector includes metrics used to collect and display performance information of a selected NetApp 7-mode NAS extent.

The performance metrics for a NetApp 7-mode NAS extent are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Block Rate		
User Write Blocks (Blocks/Sec)	Number of blocks written per second to an extent.	Δ User_Write_Blocks / Δ Time
User Read Blocks (Blocks/Sec)	Number of blocks read per second from an extent.	Δ User_Read_Blocks / Δ Time
Data Rate		

Metric	Description	Formula
Total Transfers (Req/Sec)	Total number of transfers per second serviced by an extent.	Δ Total_Transfers / Δ Time
User Writes (Req/Sec)	Number of user writes per second to an extent.	Δ User_Writes / Δ Time
User Reads (Req/Sec)	Number of user reads per second from an extent.	Δ User_Reads / Δ Time

NetApp LUN Collector

The NetApp LUN Collector includes metrics used to collect and display performance information of a selected NetApp 7-mode NAS volume.

The performance metrics for a NetApp 7-mode NAS volume are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Data Operations		
Queue Full Responses (Req/Sec)	Number of queue full responses per second.	Δ Queue_Full / Δ Time
Total Operations (Req/Sec)	Total number of operations on the LUN per second.	Δ Total_Ops / Δ Time
Other Operations (Req/Sec)	Number of other operations per second.	Δ Other_Ops / Δ Time
Write Operations (Req/Sec)	Number of write operations per second.	Δ Write_Ops / Δ Time
Read Operations (Req/Sec)	Number of read operations per second.	Δ Read_Ops / Δ Time
Data Rate		
Write Data (Bytes/Sec)	Number of bytes written per second	Δ Write_Data / Δ Time
Read Data (Bytes/Sec)	Number of bytes read per second	Δ Read_Data / Δ Time

Metric	Description	Formula
Latency		
Average Latency (ms)	Average latency in milliseconds for all operations on the LUN	Δ Avg_Latency / Δ Total_Ops

NetApp Disk Drive Collector

The NetApp Disk Drive Collector includes metrics used to collect and display performance information of a selected NetApp 7-mode NAS disk drive.

The performance metrics for a NetApp 7-mode NAS disk drive are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Block Rate		
User Write Blocks (Blocks/Sec)	Number of blocks transferred for user write operations per second	Δ User_Write_Blocks / Δ Time
User Read Blocks (Blocks/Sec)	Number of blocks transferred for user read operations per second	Δ User_Read_Blocks / Δ Time
Data Rate		
Total Transfers (Req/Sec)	Total number of disk operations involving data transfer initiated per second	Δ Total_Transfers / Δ Time
User Writes(Req/Sec)	Number of disk write operations initiated each second for retrieving data or metadata associated with user requests	Δ User_Writes / Δ Time
User Reads(Req/Sec)	Number of disk read operations initiated each second for retrieving data or metadata associated with user request	Δ User_Reads / Δ Time

Metric	Description	Formula
Disk Busy		
Disk Busy (%)	Percentage of time there was at least one outstanding request to the disk	100 x (Δ Disk_Busy / Δ Disk_Busy_Base)
Latency		
User Write Latency (ms)	Average latency per block in milliseconds for user write operations	Δ User_Write_Latency / Δ User_Write_Blocks
User Read Latency (ms)	Average latency per block in milliseconds for user read operations	Δ User_Read_Latency / Δ User_Read_Blocks

NetApp IPPort Collector

The NetApp IPPort Collector includes metrics used to collect and display performance information of a selected NetApp 7-mode NAS network interface.

The performance metrics for a NetApp 7-mode NAS network interface are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Data Rate		
Bytes Received (MBytes/Sec)	Inbound traffic in megabytes per second through the filer network interface controller (NIC). This can be used to assess network traffic for load balancing, multi-path optimization, and network performance.	(Δ Bytes_Received / Δ Time) / (1000.0 * 1000.0)
Bytes Transmitted	Outbound traffic in megabytes per second through the filer NIC.	(Δ Bytes_Transmitted / Δ

Metric	Description	Formula
(MBytes/Sec)	This can be used to assess network traffic for load balancing, multi-path optimization, and network performance.	Time) /(1000.0 * 1000.0)
Error Rate		
Received Errors (Req/Sec)	Errors per second while receiving packets.	Δ Recv_Errors / Δ Time
Send Errors (Req/Sec)	Errors per second while sending packets.	Δ Send_Errors / Δ Time
Packet Rate		
Packets Received (Packets/Sec)	Inbound traffic in packets per second through the filer NIC. Network packets contain data headers, address source and destination, payload and CRC fields. This metric can be used to measure network traffic for load balancing, multi-path optimization, and network performance. Contrary to testing the bytes received, packet testing is preferred as it is easier to test whether or not inbound packets have arrived. Byte testing does not indicate whether packet transmission is complete.	Δ Packets_ Received / Δ Time
Packets Transmitted (Packets/Sec)	Outbound traffic in packets per second through the filer NIC. Network packets contain data headers, address source and destination, payload and CRC fields. This metric can be used to measure network traffic for load balancing, multi-pathing optimization and network performance. Contrary to testing the bytes transmitted, packet testing is preferred as it is easier to test whether or not outbound packets are sent. Byte testing does not indicate whether packet transmission is complete.	Δ Packets_ Transmitted / Δ Time

Performance Collectors for HDS and HPE XP Arrays

The following performance collectors (Configuration > Monitoring Settings > Collectors) are available for HDS and HPE XP Storage arrays:

- "HDS/XP Storage System Collector" below
- "HDS/XP Front-End Controller Collector" on page 388
- "HDS/XP Storage Volume Collector" on page 388
- "HDS/XP Fibre Channel Port Collector" on page 391
- "HDS/XP Array Group Collector" on page 392
- "HDS/XP Back-End Controller Collector" on page 395
- "HDS/XP MPB Controller Collector" on page 396

HDS/XP Storage System Collector

The HDS/XP Storage System Collector includes metrics used to collect and display the cache performance information of a selected XP array.

The performance metrics are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Memory Rate		
CHA Cache Memory Busy Rate (%)	Rate at which the CHA cache is busy de-staging data to the DKA.	Raw value collected from the array.
CHA Shared Memory Busy Rate (%)	Rate at which the CHA shared memory is busy.	Raw value collected from the array.
DKA Cache Memory Busy Rate (%)	Rate at which the DKA cache is busy de-staging data to the disk.	Raw value collected from the array.

Metric	Description	Formula
DKA Shared Memory Busy Rate (%)	Rate at which the DKA shared memory is busy.	Raw value collected from the array.
Utilization		
Percent Cache Usage (%)	Cache utilization percent.	CacheUsage/CacheSize
Percent Sidefile Usage (%)	Percent utilization of the sidefile. A sidefile is an internal buffer that saves a copy of the data to be transmitted to a remote XP array. Use to track continuous access (CA) sidefile cache utilization and the potential impact of DR activities.	SidefileUsage/CacheSize
Percent Write Pending Data (%)	Percentage of pending writes based on the percentage of cache being used to buffer writes on the selected controller. Use to determine if a CLPR is needed or if attention needs to be directed toward journal parity groups.	WritePendingData/CacheSize
IO Rate		
Read Hits (Req/Sec)	Read I/O requests per second satisfied from cache.	Δ ReadHitIOs/ Δ StatisticTime
Usage		
Cache Usage (MBytes)	Cache utilization in megabytes.	Raw value collected from the array.
Sidefile Usage (MBytes)	Sidefile cache utilization in megabytes.	Raw value collected from the array.
Write Pending Data (MBytes)	Indicator of pending writes based on cache in megabytes used to buffer writes on the selected controller.	Raw value collected from the array.

HDS/XP Front-End Controller Collector

The HDS/XP Front-End Controller Collector includes metrics used to collect and display the utilization of the processors (MP) of channel adapters (CHA) of XP arrays.

The Storage System Processors tab view (Inventory > Storage Systems) of an XP array displays all the front-end channel adapters (CHA). Selecting a CHA allows you to view the utilization of its processors in the Analysis pane.

The performance metrics are grouped into the **Utilization (%)** tab of the **Analysis** pane:

Metric	Description	Formula
MP (%)	Utilization rate of a processor on a selected CHA controller.	Δ BusyTimeCounter / Δ ElapsedTimeCounter

Note: XP24000 arrays have multiple processors per CHA.

The MPB (Micro Processor Blade) controllers for XP P9500 storage arrays have a separate set of metrics to measure performance. These front-end controller metrics are therefore not used to measure the performance for XP P9500 arrays.

HDS/XP Storage Volume Collector

The HDS/XP Storage Volume Collector includes metrics used to collect and display the performance information of a selected volume of an XP array.

The volume performance metrics are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formulas
Data Rate		
Write Data Rate	Rate at which data is written to the volume by all hosts. Includes transfers from the source array to the destination	(Δ RandomWriteData + Δ SequentialWriteData) / Δ StatisticTime

Metric	Description	Formulas
(Bytes/Sec)	array.	
Data Write Rate Seq (Bytes/Sec)	Rate at which sequential data is written to the volume by all hosts. Includes transfers from the source array to the destination array.	Δ SequentialWriteData/ Δ StatisticTime
Data Write Rate Random (Bytes/Sec)	Rate at which random data is written to the volume by all hosts.	Δ RandomWriteData / Δ StatisticTime
Read Data Rate (Bytes/Sec)	Rate at which data is read from the volume by all hosts. Includes transfers from the source array to the destination array.	(Δ RandomReadData + Δ SequentialReadData) / Δ StatisticTime
Data Read Rate Seq (Bytes/Sec)	Rate at which sequential data is read from the volume by all hosts. Includes transfers from the source array to the destination array.	Δ SequentialReadData/ Δ StatisticTime
Data Read Rate Random (Bytes/Sec)	Rate at which random data is read from the volume by all hosts.	Δ RandomReadData / Δ StatisticTime
Total Data Rate (Bytes/Sec)	Rate in which data can be transmitted between devices for the selected volume.	(Δ RandomReadData + Δ RandomWriteData + Δ SequentialReadData + Δ SequentialWriteData) / Δ StatisticTime
Data Size		
Average Write Size (Bytes)	Average number of writes in bytes to the volume.	(Δ RandomWriteData + Δ SequentialWriteData) / (Δ RandomWriteIOs + Δ SequentialWriteIOs)

Metric	Description	Formulas
Average Read Size (Bytes)	Average number of reads in bytes to the volume.	(Δ RandomReadData + Δ SequentialReadData) / (Δ RandomReadIOs + Δ SequentialReadIOs)
I/O Percent		
Percent Read Hit Random IO (%)	Percentage of random read I/Os that were de-staged from cache.	100 x Δ RandomReadHitIOs / Δ RandomReadIOs
Percent Read Hit Seq IO (%)	Percentage of sequential read I/Os that were de-staged from cache.	100 x Δ SequentialReadHitIOs / Δ SequentialReadIOs
Percent Read Hit IO (%)	Percentage of cache reads for the volume.	100 x (Δ RandomReadHitIOs + Δ SequentialReadHitIOs) / (Δ RandomReadIOs + Δ SequentialReadIOs)
Percent Write IO (%)	Percentage of writes from cache.	100 x (Δ RandomWriteIOs + Δ SequentialWriteIOs)/ Δ RandomReadIOs + Δ RandomWriteIOs + Δ SequentialReadIOs + Δ SequentialWriteIOs)
Percent Read IO (%)	Percentage of reads from cache.	100 x (Δ RandomReadIOs + Δ SequentialReadIOs)/ (Δ RandomReadIOs + Δ RandomWriteIOs + Δ SequentialReadIOs + Δ SequentialWriteIOs)
I/O Rate		
Write Rate (Req/Sec)	Number of write I/Os per second.	(Δ RandomWriteIOs + Δ SequentialWriteIOs) / Δ StatisticTime
Read Hits (Req/Sec)	Cache read hits per second.	(ΔRandomReadHitIOs + Δ SequentialReadHitIOs) / Δ StatisticTime

Metric	Description	Formulas
Read Rate (Req/Sec)	Number of read I/Os per second.	(Δ RandomReadIOs + Δ SequentialReadIOs) / Δ StatisticTime
Total I/O Rate (Req/Sec)	Total number of read or write operations taking place per second for a selected volume.	(Δ RandomReadIOs + Δ RandomWriteIOs + Δ SequentialReadIOs + Δ SequentialWriteIOs) / Δ StatisticTime
I/O Response Time		
Write Response Time (Sec)	Time required to complete a write I/O in seconds.	Δ WriteResponseTimeCounter / (Δ RandomWriteIOs + Δ SequentialWriteIOs)
Read Response Time (Sec)	Time required to complete a read I/O in seconds.	Δ ReadResponseTimeCounter / (Δ RandomReadIOs + Δ SequentialReadIOs)
Utilization (%)		
Percent Utilization (%)	Percentage of time that disks in an volume are busy.	100 x Δ ActiveTime fields / Δ StatisticTime

HDS/XP Fibre Channel Port Collector

The HDS/XP Fibre Channel Port Collector includes metrics used to collect and display the performance information of a selected host Fibre Channel (FC) port (Inventory > Storage Systems > Ports) of an XP array.

The performance metrics are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Data Rate		
Total Data Rate (Bytes/Sec)	Processor utilization for the mainframe external initiator processes as a percentage of total processor time.	(Δ KBytesTransferred * 1024) / Δ StatisticTime
I/O Rate		
Total IO Rate (Req/Sec)	Difference of TotalIOs / difference of statisticTime	Δ TotalIOs / Δ StatisticTime

HDS/XP Array Group Collector

The HDS/XP Array Group Collector includes metrics used to collect and display the performance information of a selected array group (Inventory > Storage Systems > Storage Extents) of an XP array.

The performance metrics are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formulas
Data Rate		
Write Data Rate (Bytes/Sec)	Rate at which data is written to the array group by all hosts. Includes transfers from the source array to the destination array.	(Δ RandomWriteData + Δ SequentialWriteData) / Δ StatisticTime
Data Write Rate Seq (Bytes/Sec)	Rate at which sequential data is written to the array group by all hosts. Includes transfers from the source array to the destination array.	Δ SequentialWriteData/ Δ StatisticTime
Data Write Rate Random	Rate at which random data is written to the array group by all hosts.	Δ RandomWriteData / Δ StatisticTime

Metric	Description	Formulas
(Bytes/Sec)		
Read Data Rate (Bytes/Sec)	Rate at which data is read from the array group by all hosts. Includes transfers from the source array to the destination array.	(Δ RandomReadData + Δ SequentialReadData) / Δ StatisticTime
Data Read Rate Seq (Bytes/Sec)	Rate at which sequential data is read from the array group by all hosts. Includes transfers from the source array to the destination array.	Δ SequentialReadData/ Δ StatisticTime
Data Read Rate Random (Bytes/Sec)	Rate at which random data is read from the array group by all hosts.	Δ RandomReadData / Δ StatisticTime
Total Data Rate (Bytes/Sec)	Rate in which data can be transmitted between devices for the selected array group.	(Δ RandomReadData + Δ RandomWriteData + Δ SequentialReadData + Δ SequentialWriteData) / Δ StatisticTime
Data Size		
Average Write Size (Bytes)	Average number of writes in bytes to the array group.	(Δ RandomWriteData + Δ SequentialWriteData) / (Δ RandomWriteIOs + Δ SequentialWriteIOs)
Average Read Size (Bytes)	Average number of reads in bytes to the array group.	(Δ RandomReadData + Δ SequentialReadData) / (Δ RandomReadIOs + Δ SequentialReadIOs)
I/O Percent		
Percent Read Hit Random IO	Percentage of random read I/Os that were de-staged from cache.	100 x Δ RandomReadHitIOs / Δ RandomReadIOs

Metric	Description	Formulas
(%)		
Percent Read Hit Seq IO (%)	Percentage of sequential read I/Os that were de-staged from cache.	100 x Δ SequentialReadHitIOs / Δ SequentialReadIOs
Percent Read Hit IO (%)	Percentage of cache reads for the array group.	100 x (Δ RandomReadHitIOs + Δ SequentialReadHitIOs) / (Δ RandomReadIOs + Δ SequentialReadIOs)
Percent Write IO (%)	Percentage of writes from cache.	100 x (Δ RandomWriteIOs + Δ SequentialWriteIOs)/ Δ RandomReadIOs + Δ RandomWriteIOs + Δ SequentialReadIOs + Δ SequentialWriteIOs)
Percent Read IO (%)	Percentage of reads from cache.	100 x (Δ RandomReadIOs + Δ SequentialReadIOs)/ (Δ RandomReadIOs + Δ RandomWriteIOs + Δ SequentialReadIOs + Δ SequentialWriteIOs)
I/O Rate		
Write Rate (Req/Sec)	Number of write I/Os per second.	(Δ RandomWriteIOs + Δ SequentialWriteIOs) / Δ StatisticTime
Read Hits (Req/Sec)	Cache read hits per second.	(ΔRandomReadHitIOs + Δ SequentialReadHitIOs) / Δ StatisticTime
Read Rate (Req/Sec)	Number of read I/Os per second.	(Δ RandomReadIOs + Δ SequentialReadIOs) / Δ StatisticTime
Total I/O Rate (Req/Sec)	Total number of read or write operations taking place per second for a selected array group.	(Δ RandomReadIOs + Δ RandomWriteIOs + Δ SequentialReadIOs + Δ SequentialWriteIOs) / Δ StatisticTime

Metric	Description	Formulas
I/O Response Time		
Write Response Time (Sec)	Time required to complete a write I/O in seconds.	Δ WriteResponseTimeCounter / (Δ RandomWriteIOs + Δ SequentialWriteIOs)
Read Response Time (Sec)	Time required to complete a read I/O in seconds.	Δ ReadResponseTimeCounter / (Δ RandomReadIOs + Δ SequentialReadIOs)
Utilization (%)		
Percent Utilization (%)	Percentage of time that disks in an array group are busy.	100 x Δ ActiveTime fields / Δ StatisticTime

HDS/XP Back-End Controller Collector

The HDS/XP Back-End Controller Collector includes metrics used to collect and display the utilization of the processors (MP) of disk controller adapters (DKA) of XP arrays (such as XP24000) that expose back-end controllers.

The SCSI Controllers tab view (Inventory > Storage Systems) of an XP array displays all the disk controller adapters (DKA). Selecting a DKA allows you to view the utilization of its processors in the Analysis pane.

The performance metrics are grouped into the **Utilization (%)** tab of the **Analysis** pane:

Metric	Description	Formula
MP (%)	Utilization rate of a processor on a selected DKA controller.	Δ BusyTimeCounter / Δ ElapsedTimeCounter

HDS/XP MPB Controller Collector

The HDS/ XP MPB Controller Collector includes metrics used to collect and display the performance information of a selected Micro Processor Blade (MPB) controller (Inventory > Storage Systems > SCSI Controllers) of XP7 and P9500 arrays. MBP processor statistics are obtained using a CIM extension to discover an XP array.

The performance metrics are grouped into the following tabs of the **Analysis** pane:

Metric	Description	Formula
Backend Utilization (%)	Processor utilization for the back-end processes (back-end activities for target I/O requests) as a percentage of total processor time.	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Mainframe Ext Initiator Utilization (%)	Processor utilization for the mainframe external initiator processes as a percentage of total processor time.	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Mainframe Target Utilization (%)	Processor utilization rate for the mainframe target processes (front-end activities for processing mainframe I/O requests) as a percentage of total processor time.	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Open Ext Initiator Utilization (%)	Processor utilization rate for the open external initiator processes (external storage access activities) as a percentage of total processor time.	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Open Initiator Utilization (%)	Processor utilization rate for the open initiator processes (continuous access replication activities) as a percentage of total processor time.	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Open Target Utilization (%)	Processor utilization rate for open target processes (front-end activities) as a percentage of total processor time.	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Processor Utilization (%)	Processor utilization rate on the selected MPB controller. This rate is the sum of the other seven metrics listed in this table.	Δ BusyTimeCounter / Δ ElapsedTimeCounter

Metric	Description	Formula
System Utilization (%)	Processor utilization rate of the array system processes as a percentage of total processor time.	Δ BusyTimeCounter / Δ ElapsedTimeCounter

Enabling Performance Data Collection for HPE XP Storage and HP P9000 Arrays

You must perform configuration tasks on the SOM to enable performance data collection from HPE XP Storage and HP P9000 arrays.

SOM collects performance statistics in-band from a command device for HPE XP Storage and HP P9000 arrays. Therefore, you must set up at least one proxy host in your SAN with a command device from the array and install a CIM extension on the host. The CIM extension automatically detects the command device.

HPE XP and HP P9000 Array-Related Software Requirements

To enable performance data collection for HPE XP and HP P9000 arrays in SOM, RAID Manager Library and built-in provider must be installed on the proxy host.

The software requirements for proxy host are as follows:

- CIM extension – A CIM extension must be installed and running on the proxy host. For information about operating systems supported by the proxy host server, see *SOM Device Support Matrix*.
- RAID Manager Lib XP (RMLIBXP) – This library is not included in the CIM extension installation package. It must be installed separately on the proxy host server. RAID Manager Lib XP (RMLIBXP) must be 32-bit version. For more information about locating and installing RAID Manager Lib XP, see the HPE XP Storage arrays documentation.

Performance Data Collection Architecture for HPE XP Storage Arrays

The following figure shows the collection path for performance data of HPE XP Storage arrays. As the figure shows, SOM uses the RMI-API to collect data from the HPE XP Storage array. However, to collect performance data, SOM uses a proxy host server, which is connected to the

HPE XP Storage array. You can use a single proxy host server to collect performance data from multiple arrays. The RAID Manager Library (2) collects the performance data from the array using a special command device LUN (3). The command device LUN sends the statistical data to SOM, which stores it in the database.

Command device LUN generates additional input/output traffic while collecting performance data, which causes increase in input/output traffic of the array group.

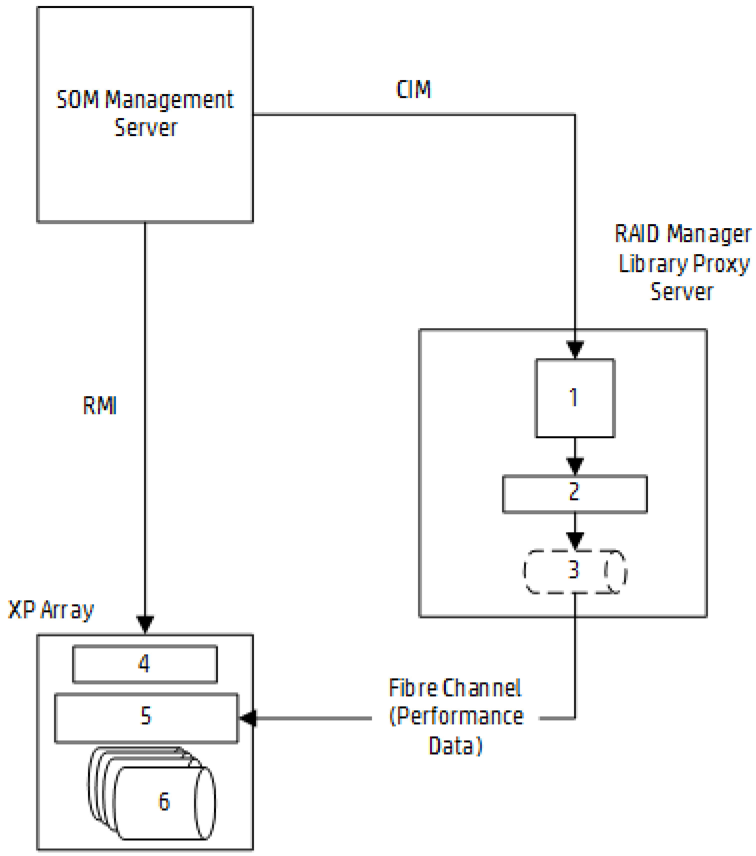

1	HPE SOM Host Agent/CIM Extension
2	RAID Manager Library
3	Command Device LUN
4	Service Processor (SVP)
5	Controllers
6	Data

Configuring the SOM Management Server

To enable performance data collection in SOM for HPE XP and HP P9000 arrays, follow these steps:

1. Discover the HPE XP or HP P9000 array as described in "Prerequisites to Discover HPE XP/P9500 Arrays" on page 144.
2. Discover the proxy host by entering the DNS or IP information and appropriate credentials of the CIM extension running on the proxy host.
3. Enable the license for the performance pack (SOM Ultimate-Perf license), as described in the "Configure Performance Pack" on page 189.
4. In the **All Storage Systems** inventory view, right-click the HPE XP Storage array, and then click **Add/Edit Proxy**.
5. In the **Assign Proxy** dialog box, specify the proxy host name that was set up, as described in "Create the Proxy Host " below (also displayed in the **Discovered Hosts** inventory view).
6. Create a monitoring group and associate a set of HPE XP collectors to it. For more information, see "Create a Monitoring Group" on page 192.

Create the Proxy Host

SOM requires a proxy host to access both historical and real-time performance data for HPE XP and HP P9000 arrays.

If you use the HPE XP Performance Advisor software, which collects performance data for the HPE XP arrays, it is recommended that you use the same proxy host for SOM as used for the HP Performance Advisor software. Both the SOM and HP XP Performance Advisor use a similar proxy host configuration, RAID Manager Library (RMLIB API), and a command device LUN.

Operating System Requirements

For HPE XP arrays, you can set up a proxy host server on any supported operating system. For information on operating systems supported by proxy host, see the *SOM Device Support Matrix*.

Setting Up the Proxy Host

To set up the proxy host, follow these steps:

1. Verify that the command device LUN is accessible to the host bus adaptor (HBA) on the proxy host by using the native HBA tool set.

   If you have HPE XP Performance Advisor and you already installed the RAID Manager Library (RMLIB API), go to Step 3. If you are not sure where RMLIB API is installed, look at the HPE XP Performance Advisor configuration to see where the agents for Performance Advisor are installed.
2. If not installed already, install the RAID Manager Library (RMLIB API). Obtain the RAID Manager Library (RMLIB API) from the array firmware CD. If you do not have the RAID Manager Library software, contact HPE services for assistance.
3. If not created already, create a command device LUN (LUN:0). For information about creating a command device, see "Create Command Device" below.
4. Install the CIM Extension on the same proxy host server that has RMLIB API and LUN:0.
   This is the same CIM Extension that SOM uses to manage and discover the proxy host server.

Create Command Device

Create a command device LUN on the SLPR0 partition using the HP StorageWorks XP Remote Web Console (RWC). Present the command device LUN to the port that is accessible to the proxy host server. A command device LUN can be any LUN that is accessible to the proxy host server.

For RAID600-based or RAID500-based XP array models (which support SLPR partitioning), the command device LUN should be from the first SLPR0 partition of the XP array. The SLPR0 command device LUN provides visibility to the array regardless of its partitioning.

When creating a command device LUN, ensure the following:

- No data exists on the volume that you select as the command device. Any data that exists on the volume is unavailable to the proxy host.
- The volume that you designate as the command device is used only by the disk array and is blocked from the user.
- No file system has been mounted and no data is stored in the volume.

Creating a Command Device LUN

> **Caution:** After you create the command device, do not mount any file systems on the command device.

To create a command device LUN:

1. Launch the Remote Web Console (RWC) interface using administrator privileges.
2. From the RWC, click **GO** > **Lun Manager** > **LUN Path and Security**.
3. Right-click the logical device in the **LDEV** column, and then click **Enable/Disable** to convert the device into a command device.
4. Click **Apply**.

Additional Related Tasks

Creating a command device LUN might also require you to perform the following tasks:

- Zone the SAN switches between the proxy host and the HPE XP array port to open the path.
- Expose the command device LUN on the HPE XP array port to the HBA WWN on the RMLIB proxy server to create a host security group.

HPE XP and HP P9000 Array Performance Data Collection

Performance metrics for HPE XP and HP P9000 arrays include global resources as described in "Performance Collectors for HDS and HPE XP Arrays" on page 386.

When collecting performance statistics for HPE XP and HP P9000 logical array groups, be aware that a parity group is divided into several logical RAID groups. Thus, for example, parity group X-X can contain logical array groups X-X-1, X-X-2, and so on. The collection of performance statistics for logical array groups is only done at the parity group level. Since SOM does not show the parity group, all the statistics are gathered at the first logical RAID group (in our example, X-X-1), and no statistics are gathered at X-X-2 and the other subsequent logical RAID groups.

SOM requires two points to plot the first data point. Therefore, depending on the collector setup, interval, and other factors, it can take a while for the data to begin to display, in some cases a couple of hours.

The minimum collection interval that can be set for the HPE XP performance data collectors is 15 minutes. It is recommended to set an interval between 15 minutes and 1 hour. Optionally, set the collection interval lower while analyzing a problem, and then restore the collector interval to the default.

Verify that the RAID Manager Library Returns Data

Use the arrayScan tool on the proxy host to verify that the RAID Manager Library (RMLIB API) is returning data through the command device LUN. This tool is located in the `<CIM_extension_installation_directory>\tools` directory on the proxy host server.

> **Note:** Use the `./` prefix for the arrayScan tool on non-Windows operating systems.

For information about the arrayScan tool, type `arrayScan -?` or `arrayScan -help`.

Output Example

When the arrayScan tool is used with no parameters, it returns the selected command LUN used to get statistics. Here is an example of the output from the arrayScan tool:

```
arrayScan build date: May 21 2009:16:24:19
Return string...
\\.\PHYSICALDRIVE4 :"HP ","OPEN-V-CM ", Rev"5001"
( Serial# 10118, RAID600or500,LDKC0, SLPR0, CLPR0, RG1-1, LDEV 00:1E,
CU 0, RAID5 , Port1A, PortWWN:10000000C95C763F,
NodeWWN:20000000C95C763F )
...1 Array Cmd Dev Lun device paths found including any SLPR0 ones
just shown.
...Return string.
```

```
Return string length: 293 (0 percent of current max 14680064 bytes).
Largest line length: 116
```

Performance Collectors for IBM SAN Volume Controller and IBM V7000

The following performance collectors (**Configuration** > **Monitoring Settings** > **Collectors**) are available for IBM SVC and IBM V7000:

- "IBM SVC SMI-S Node Controller Collector" below
- "IBM SVC SMI-S MDisk Collector" on the next page
- "IBM SVC SMI-S Physical Disk Collector" on page 407
- " IBM SVC SMI-S Volume Collector" on page 409

IBM SVC SMI-S Node Controller Collector

The Node Controller Collector metrics are used to monitor the performance of the nodes on the storage system.

The node controller performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
I/O Percent		
Percent Writes (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Percent Reads (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Write Hits (Req/Sec)	The cumulative count of Write Cache Hits (Writes that went directly to Cache).	Δ WriteHitIOs / Δ Time

Metric	Description	Formula
Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time
Read Hits (Req/Sec)	Read cache hit rate.	Δ ReadHitIOs/ Δ Time
Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the IBM SAN Volume Controller and IBM V7000.

IBM SVC SMI-S MDisk Collector

The MDisk Collector metrics are used to monitor the performance of storage extents created under concrete pools in the **Analysis** pane.

To view performance information for concrete pool extents, follow these steps:

1. In the **Storage System** view, click the **Pools** tab.
2. In the **Pools** tab, select a concrete storage pool and click **Open**.
3. In the **Storage Pool** view, click the **Storage Extents** tab, and then select a storage extent.
   The performance information for the extent is displayed in the Analysis pane.

Note: The underlying storage extents are aggregated at **Pool** level.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Write throughput rate.	(Δ KBytesWritten * 1024) / Δ Time
Read Data Rate (Bytes/Sec)	Read throughput rate.	(Δ KBytesRead * 1024) / Δ Time
Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
Average Write Size (Bytes)	Average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Average Read Size (Bytes)	Average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		
Percent Writes (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Percent Reads (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time
Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time
Queue Depth		

Metric	Description	Formula
Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
I/O Response Time (Sec)	Time to complete an I/O operation.	(Δ IOTimeCounter/ Δ TotalIOs) * 0.001
Aggregate of Backend Volume at Pool Level		
Data Rate		
Pool Write Data Rate (Bytes/Sec)	Pool write throughput rate.	(Δ KBytesWritten * 1024) / Δ Time
Pool Read Data Rate (Bytes/Sec)	Pool read throughput rate.	(Δ KBytesRead * 1024) / Δ Time
Pool Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
Pool Average Write Size (Bytes)	Pool average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Pool Average Read Size (Bytes)	Pool average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		
Pool Percent Writes (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)

Metric	Description	Formula
Pool Percent Reads (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Pool Write Rate (Req/Sec)	Number of pool write requests per second.	Δ WriteIOs / Δ Time
Pool Read Rate (Req/Sec)	Number of pool read requests per second.	Δ ReadIOs / Δ Time
Pool Total I/O Rate (Req/Sec)	Average number of pool read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time
Queue Depth		
Pool Queue Depth	Average number of pending pool read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
Pool I/O Response Time (Sec)	Time to complete a pool I/O operation.	(Δ IOTimeCounter/ Δ TotalIOs) * 0.001

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the IBM SAN Volume Controller and IBM V7000.

IBM SVC SMI-S Physical Disk Collector

The Disk Collector metrics are used to monitor the performance of the physical disks on the storage system.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Write throughput rate (Bytes per second).	(Δ KBytesWritten * 1024) / Δ Time
Read Data Rate (Bytes/Sec)	Read throughput rate (Bytes per second).	(Δ KBytesRead * 1024) / Δ Time
Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
Average Write Size (Bytes)	Average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Average Read Size (Bytes)	Average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		
Percent Writes (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Percent Reads (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time
Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of read and write I/O operations in requests per second.	Δ TotalIOs / Δ Time

Metric	Description	Formula
Queue Depth		
Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
IO Response Time (Sec)	Time to complete an I/O operation.	(Δ IOTimeCounter/ Δ TotalIOs) * 0.001

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the IBM SAN Volume Controller and IBM V7000.

IBM SVC SMI-S Volume Collector

Note: The underlying volume statistics are aggregated at **Storage System** and **Storage System Processor** level under I/O Groups.

The volume performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Write Data Rate (Bytes/Sec)	Write throughput rate.	(Δ KBytesWritten * 1024) / Δ Time
Read Data Rate (Bytes/Sec)	Read throughput rate.	(Δ KBytesRead * 1024) / Δ Time

Metric	Description	Formula
Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
Average Write Size (Bytes)	Average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Average Read Size (Bytes)	Average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		
Percent Writes (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Percent Reads (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Write Hits (Req/Sec)	The cumulative count of Write Cache Hits (Writes that went directly to Cache).	Δ WriteHitIOs / Δ Time
Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time
Read Hits (Req/Sec)	Number of read requests (per second) completed from the array cache memory.	Δ ReadHitIOs / Δ Time
Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time
Total I/O Rate (Req/Sec)	Average number of read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time

Metric	Description	Formula
Queue Depth		
Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
IO Response Time (Sec)	Time to complete an I/O operation.	(Δ IOTimeCounter/ Δ TotalIOs) * 0.001
Aggregate of Storage Volumes at Storage System Level (Cluster)		
Data Rate		
Cluster Write Data Rate (Bytes/Sec)	Cluster write throughput rate.	(Δ KBytesWritten * 1024) / Δ Time
Cluster Read Data Rate (Bytes/Sec)	Cluster read throughput rate.	(Δ KBytesRead * 1024) / Δ Time
Cluster Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
Cluster Average Write Size (Bytes)	Cluster average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Cluster average Read Size (Bytes)	Cluster average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		

Metric	Description	Formula
Cluster Percent Writes (%)	Ratio of cluster write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Cluster Percent Reads (%)	Ratio of cluster read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Cluster Write Hits (Req/Sec)	The cumulative count of cluster Write Cache Hits (Writes that went directly to Cache).	Δ WriteHitIOs / Δ Time
Cluster Read Hits (Req/Sec)	Cluster read cache hit rate.	Δ ReadHitIOs/ Δ Time
Cluster Write Rate (Req/Sec)	Number of cluster write requests per second.	Δ WriteIOs / Δ Time
Cluster Read Rate (Req/Sec)	Number of cluster read requests per second.	Δ ReadIOs / Δ Time
Cluster Total I/O Rate (Req/Sec)	Average number of cluster read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time
Queue Depth		
Cluster Queue Depth	Average number of cluster pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
Cluster I/O Response Time (Sec)	Time to complete a cluster I/O operation.	(Δ IOTimeCounter / 1000) / Δ TotalIOs
Aggregate of Storage Volumes at Storage System Processor Level		
Data Rate		
IOGrp Write Data Rate	IOGrp write throughput rate.	(Δ KBytesWritten * 1024) / Δ

Metric	Description	Formula
(Bytes/Sec)		Time
IOGrp Read Data Rate (Bytes/Sec)	IOGrp read throughput rate.	(Δ KBytesRead * 1024) / Δ Time
IOGrp Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
IOGrp Average Write Size (Bytes)	IOGrp average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
IOGrp Average Read Size (Bytes)	IOGrp average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		
IOGrp Percent Writes (%)	Ratio of IOGrp write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
IOGrp Percent Reads (%)	Ratio of IOGrp read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
IOGrp Write Hits (Req/Sec)	The cumulative count of IOGrp Write Cache Hits (Writes that went directly to Cache).	Δ WriteHitIOs / Δ Time
IOGrp Write Rate (Req/Sec)	Number of IOGrp write requests per second.	Δ WriteIOs / Δ Time
IOGrp Read Hits (Req/Sec)	Number of IOGrp read requests (per second) completed from the array cache memory.	Δ ReadHitIOs / Δ Time

Metric	Description	Formula
IOGrp Read Rate (Req/Sec)	Number of IOGrp read requests per second.	Δ ReadIOs / Δ Time
IOGrp Total I/O Rate (Req/Sec)	Average number of IOGrp read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time
Queue Depth		
IOGrp Queue Depth	Average number of pending IOGrp read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
IOGrp IO Response Time (Sec)	Time to complete an IOGrp I/O operation.	(Δ IOTimeCounter/ Δ TotalIOs) * 0.001

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the IBM SAN Volume Controller and IBM V7000.

IBM SVC Cluster Performance

The IBM SVC cluster performance metrics are aggregated from the underlying volume statistics.

The cluster performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Formula
Data Rate		
Cluster Write Data Rate	Cluster write throughput rate.	(Δ KBytesWritten * 1024) / Δ Time

Metric	Description	Formula
(Bytes/Sec)		
Cluster Read Data Rate (Bytes/Sec)	Cluster read throughput rate.	(Δ KBytesRead * 1024) / Δ Time
Cluster Total Data Rate (Bytes/Sec)	Rate data is transmitted between devices.	(Δ KBytesTransferred * 1024) / Δ Time
Data Size		
Cluster Average Write Size (Bytes)	Cluster average write size of I/Os written.	(Δ KBytesWritten * 1024) / Δ WriteIOs
Cluster average Read Size (Bytes)	Cluster average read size of I/Os read.	(Δ KBytesRead *1024) / Δ ReadIOs
I/O Percent		
Cluster Percent Writes (%)	Ratio of write I/Os to total I/Os.	100 * (Δ WriteIOs / Δ TotalIOs)
Cluster Percent Reads (%)	Ratio of read I/Os to total I/Os.	100 * (Δ ReadIOs / Δ TotalIOs)
I/O Rate		
Cluster Write Hits (Req/Sec)	The cumulative count of Write Cache Hits (Writes that went directly to Cache).	Δ WriteHitIOs / Δ Time
Cluster Read Hits (Req/Sec)	Read cache hit rate.	Δ ReadHitIOs/ Δ Time
Cluster Write Rate (Req/Sec)	Number of write requests per second.	Δ WriteIOs / Δ Time

Metric	Description	Formula
Cluster Read Rate (Req/Sec)	Number of read requests per second.	Δ ReadIOs / Δ Time
Cluster Total I/O Rate (Req/Sec)	Average number of read and write I/O operations given in requests per second.	Δ TotalIOs / Δ Time
Queue Depth		
Cluster Queue Depth	Average number of pending read and write I/O operations.	Total I/O Rate * I/O Response Time
Response Time		
Cluster I/O Response Time (Sec)	Time to complete an I/O operation.	(Δ IOTimeCounter / 1000) / Δ TotalIOs

Note: In the formulas shown above, the value Δ Time represents the difference in seconds between the most recent two `StatisticTime` values returned by the SMI-S provider. `StatisticTime` is a date/time raw statistic collected by the SMI-S provider for the HPE 3PAR storage system.

Device-Specific Exceptions

The following sections capture the disparities between SOM and the device consoles. Additionally, it also explains how SOM handles the collected data.

- Hosts
- Switches
- Storage Systems

Hosts

This topic captures the disparities between SOM and the device consoles of the listed hosts. Additionally, it also explains how SOM handles the collected data.

Hosts Discovered Using the Agentless Method

Hosts discovered using the agentless method have the following limitations based on the operating system:

- **Windows Hosts**
  - Public folders and mailbox information is not available.
  - Limited information related to disk partitions and disk drives is available, when the native volume manager volumes are used to obtain data.
  - The grey space is calculated for both basic and dynamic disks.
- **Linux Hosts**
  - The Analysis pane does not display the CPU and Memory performance information.
  - The following performance metrics are available:
    - Disk Read
    - Disk Total
    - Disk Utilization
    - Disk Write
  - The number of target mappings may be less than the number of target mappings returned by the CIM extension. This difference is because some entries with a SCSI LUN value of zero are not shown.
  - The following issues are observed for Linux hosts of certain vendors:

i. HBA information about the following is not available:

ii. For dual port HBAs, each port is displayed as an individual adapter in the Cards tab view with each adapter mapped to its port in the Ports tab view.

- Host capacity, disk partition, or file system sizes shown in SOM may not match with those shown on the host. This difference is because SOM calculates the total capacity using the blocks (total and available) provided by the output of the command `df -h`. This may vary when compared to the CLI output because SOM may not consider the size of the metadata during the calculations.

- While you create a node group for Linux hosts, always use Device Family as the filter criteria. Device filters based on OS Type are not supported.

- **Solaris Hosts**
  - The **Hardware Version** of the HBA card (Properties pane of the Form view) is blank for Solaris hosts.
  - For SES devices on Solaris hosts, SOM displays the LunID, SCSI Bus and SCSI Target value as blank in the Target Mappings tab.

- **AIX Hosts**

  In some cases, for AIX Hosts discovered using Powerpath with LVM setup, the **Drive Type** column in the **Filesystems** tab might display the files systems as **Local**, because of incorrect mappings between the disk drive and target ports. Incorrect output returned by the `hbatest` command causes the incorrect mappings, even though the underlying disks for the file systems are visible in SOM. For information about `hbatest` command, see "Commands for an AIX Host as a Root User" on page 131 and "Commands for an AIX Host as a Non-Root User" on page 134.

Windows Hosts

- For Windows hosts with HDLM multipathing and native Volume Manager, the **Size** (Disk Drives tab) and the **Max Media Size** (Host Disk Drive form) are blank.
- For Windows hosts with HDLM multipathing and Volume Manager, the link between Multipath Disks tab and Volume Management tab is not available.
- To view file shares (CIFS, NFS) on a Windows host with the CIME agent running, log on as Administrator to the `AppStorWin32Agent` service (Properties > Log On > This account) and run **Start Collection** from the context menu of the host inventory view.
- The grey space is calculated for both basic and dynamic disks.
- Data collection fails for hosts with disks listed as 'Unknown' in the Disk Management console.
  Do the following in the Device Manager to enable data collection for such hosts:
  - Scan for hardware changes on the host.
  - Manually disable 'Unknown' disks.

- Data collection times out for hosts with LUNs > 256.
  Change the following property in the file, `custom.properties`, in the location, `%OvDataDir%\Conf\som`, to an appropriate value to enable data collection:
  `cxws.agency.queue.operationTimeout=1800000 ms`

Virtual Hosts

- SOM discovers templates as powered off virtual machines. Templates are only discovered when you discover virtual machines through the vCenter. If you discover individual ESX servers directly, the templates are not found.
- A known third-party issue related to ESX Servers causes SOM to present incomplete or erroneous information. The issue occurs when a LUN is shared by more than one ESX Servers.

  The following exceptions are a result of this issue:

  - Some shared external storage volumes for a virtual machine are reported with drive types of Local instead of external.
  - A virtual machine's element topology appears as having only local (to the ESX Server) storage instead of external storage.
  - The Volumes property in the Multipathing tab for a virtual machine is blank instead of containing the name of the external storage volume.
  - In the End to End Connectivity Report, ESX Servers reporting back as not connected display "Not connected to external storage" in the Storage System column.
- **Details of Virtual Hosts After Discovery**
  After discovery, the following details are available in the hosts inventory views:

Discovered Virtual Host	Inventory Details
vCenter	The vCenter's access point in the Summary tab with the associated virtual servers. • **Access Point** (of the vCenter) • **Hosts Virtual Servers view** – Details and sub-components of the virtual servers managed by the vCenter.
Virtual server	The virtual server's access point in the Summary tab. • **Access Point** (of the virtual server) • **Hosts Virtual Servers view** – Details and sub-components of the Virtual servers.
Virtual machine with VMTools	The virtual server's or vCenter's access point in the Summary tab.

Discovered Virtual Host	Inventory Details
	• **Access Point** (of the virtual server or vCenter) • **Hosts Virtual Machines view** – Details and sub-components of the virtual machines.
Virtual machine with VMTools and a CIM Extension	The virtual machine's access point in the Summary tab. • **Access Point** (of the virtual machine) • **Hosts Virtual Machines view** – Details and sub-components of the virtual machines. **Note**: There is no access point for a virtual machine unless it has a CIM extension installed.

Host Clusters

- VMWare DRS clusters
  Local disks are displayed as 'Shared Disk Drives' in a cluster if the API returns a disk UUID that is not unique.
- Host Cluster Dashboard
  - Although the title of the Host Cluster Summary pane is Host Summary, the pane displays the details of the host cluster.
  - Although the title of the Host Cluster Capacity pane is Host Capacity, the pane displays the capacity utilization of the host cluster.

HP-UX Hosts

The following exceptions are noticed with HP-UX hosts:

- The Host Unused Capacity for HP-UX hosts includes the capacity from any DVD device on the host.
- The HP-UX CIM Extension does not report capacity for a VxFS file system on HP-UX if the file system's size exceeds 2TB.
- The presence of special agile devices on HP-UX causes the local disk to appear in the Multipathing tab for the host.
- The model number for the AH403A HBA is not shown when installed on HP-UX 11.31 hosts due to an issue in the SNIA HBAAPI library.
- The Link Failure counter does not report data for most HBAs supported on HP-UX. The A5158A HBA does report values correctly.

Switches

This topic captures the disparities between SOM and the device consoles of the listed switches. Additionally, it also explains how SOM handles the collected data.

Cisco

- Duplicate E Ports are shown for CISCO Multi-VSAN ISL:
  Duplicate E ports are shown in the port list for all fabrics for multi-VSAN ISLs on CISCO switches. Logical ports are shown instead of physical ports.
- Port speed is not available for CISCO switch ports with port speed greater than 4 GB/s.
- Some inactive zone aliases do not appear in the Associated Zones on Cisco SNMP switches:
  On Cisco switches managed through SNMP, some inactive zone aliases are not shown in the zones to which they belong.

Brocade

- Switches discovered through BNA appear differently in SOM.
  For information about changing the required settings, see the documentation of the switch.
  - If the physical name of the switch has not been set, it might display a default name, such as the switch model.
    You can set the physical name of the switch, by providing a value for the `chassisname` property of the switch.
  - The logical/virtual switch might display the same name as the physical switch.
    You can set the name of the virtual switch by using the `switchname` command on the switch. This differentiates the virtual switch from its corresponding physical switch.
  - The fabric name might display a World Wide Name.
    The fabric name of a switch is the name set in the BNA discovery tool.The World Wide Name of the primary switch is usually used as the name of the fabric.

Storage Systems

This topic captures the disparities between SOM and the device consoles of the listed storage systems. Additionally, it also explains how SOM handles the collected data.

HPE XP

- **Data Collection**
  Data collection of XP arrays often results in a transaction timeout. As a workaround the transaction timeout can be increased to 10 mins. To increase the transaction timeout, follow these steps:
  a. Edit the `drive:\Program Files (x86)\HP\HP BTO Software\nmsas\common\deploy\transaction-jboss-beans.xml` file and set the property

"defaultTimeout" of "CoordinatorEnvironmentBean" to 600.

b. Restart the somjboss service.

- **Volume Representation**
  SOM suffixes the following letters/symbols to volume names based on the types of array groups to identify volumes:
  - # - external volumes
  - V - snapshot volumes
  - X - THP volumes
  - D - other volumes
- **Array Replication**
  SOM maps HPE XP terminology as follows:

Property	Continuous Access	HP Continuous Access Journal	HP Business Copy	HPE XP Snapshot
Locality	Remote pair	Remote pair	Local pair	Local pair
Replica type	Full copy	Full copy	Full copy	After delta
Copy type	Sync/async depending on cache journaling in use	Async	Sync	UnSyncAssoc
Sync state	Paired, idle, failed, suspended	Active, halted, stopped	Copy, pair, psus	Idle, pair

Note: The values listed in the table are observed in the product test environment for replication pair attributes of different types of XP volume replication. You might observe additional values based on your environment.

Whenever the locality is a remote pair, the remote system serial number and volume ID are displayed. Volume ID is the devNum (CU:LDEV converted to decimal). If the remote system is also discovered by SOM, the replication table links directly to that volume on the remote system.

For Universal Replicator and Continuous Access Journal, SOM displays the individual journal groups containing the journal LDEVs and categorizes their storage capacity separately so that it is accounted for but not considered as available capacity.

- **Pool Information**
  - For this array family, LUSEable storage pools are based on emulation and RAID levels. A pool based on RAID5 can include LDEVs from any RAID5 array group, such as, RAID5(3D+1P) as well as RAID5 (7D+1P).
  - There is some free space in each of the array groups, which is reported by the storage pools "Free Space…" and is added to the aggregate ThP Total Capacity.
- **Understanding Capacity Information of XP7 and P9500 Arrays**
  The capacity information available in the Thin Provisioning Data tab (form view of the storage system) can be compared to the native Remote Web Console (RWC) as shown here:

  Internal Allocation Summary

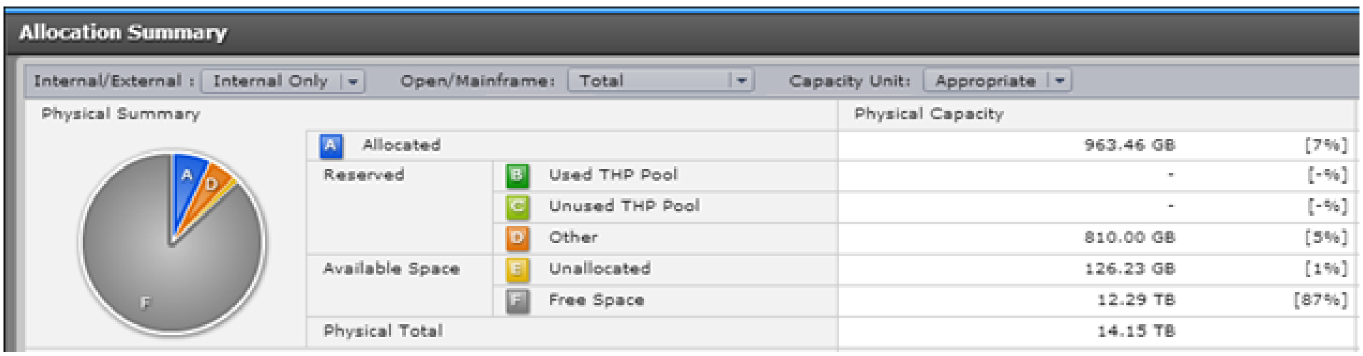

  The internal allocation metrics in RWC can be mapped to the following metrics in the Thin Provisioning Data tab:

  - **Allocated** - The sum of the values in the **Actual Mapped** column of pools that have names in the *<Emulation Type> <RAID Level>* format. For example, OPEN-V RAID 5.
  - **Other** - The sum of the values in the **Total Capacity** column of pools shown in RWC.
  - **Unallocated** - The sum of values in the **Actual Used Unmapped** column of pools that have names in the *<Emulation Type> <RAID Level>* format. For example, OPEN-V RAID 5.
  - **Free Space** - The sum of the values in the **Total Capacity** column of pools with names, such as, Free Space on Array Group 1-5-1 RAID5(3D+1P), and so on.

External Allocation Summary

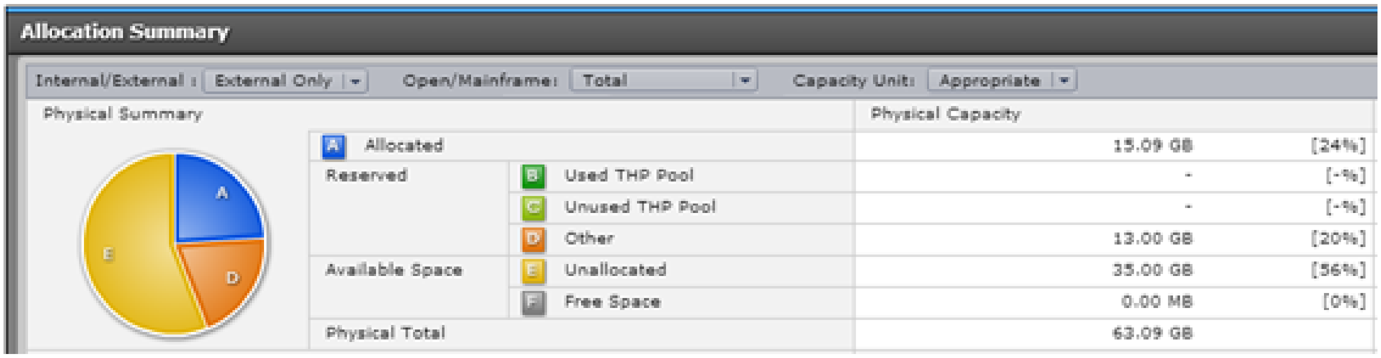

The external allocation metrics in RWC can be mapped to the following metrics in the Thin Provisioning Data tab:

- **Allocated** - The sum of values in the **Actual Mapped** column of external pools, such as, External (HP/XP7/10035).

- **Other + Unallocated** - The sum of values in the **Actual Unmapped** column of external pools, such as, External (HP/XP7/10035).

> **Note**: RWC does not include the size of **Journal Groups** in the Physical Summary Total. However, SOM includes the size of Journal Groups in capacity calculations. Therefore to compare the Physical Total value of logical devices shown by RWC, exclude the size of Journal Groups from the Total size shown in the **ThP Allocation** tab (Analysis pane) in SOM.

HDS/HUS

- When data collection runs concurrently for multiple arrays, there is a possibility that data collection might fail for one or more arrays. This is resolved in subsequent automatic data collections.
- **Volume Representation**
  SOM suffixes the following letters/symbols to volume names based on the types of array groups to identify volumes:
  - # - external volumes
  - V - snapshot volumes
  - X - THP volumes
  - D - other volumes
- **Replication Pairs**
  This table describes the HDS terminology and how SOM maps these terms:

Property	TrueCopy (Sync & Async)	Universal Replicator	Shadow Image	C.O.W. Snapshot
Locality	Remote pair	Remote pair	Local pair	Local pair
Replica Type	Full copy	Full copy	Full copy	After delta
Copy type	Sync/Async depending on cache journaling in use	Async	Sync	UnSyncAssoc
Sync State	Paired, idle, failed, suspended	Active, halted, stopped	Copy, pair, PSUS	Idle or pair

The functionality of replication pairs has not been tested due to device unavailability.

For the error message, `HdsModifier Exception No replica pairs will be returned: CIM_ERR_FAILED`, retry data collection.

- **Backend Storage**
  Data is not populated in the Backend storage tab for HDS arrays that do not support the back-end capability.
- **Storage Pools and Extents of HUS Arrays**
  The storage pools and extents of an HDS storage system are based on the HUS array groups and do not match those shown by the HUS device manager.
- **Pools Logical Data Tab**

  In the **Pools Logical Data** tab, SOM uses disk count and disk size to calculate **Total Space** of the DP pool. SOM uses both DP-VOL and V-VOL volumes from open-unallocated volumes to calculate the **Available Space** of the DP pool. However, the Device Manager uses only DP VOL type of volumes to calculate the free capacity. Therefore, the Available Space value might be inconsistent between the Device Manager and SOM.

HPE 3PAR

Replication errors

Data collection of 3PAR arrays with 3PAR InForm OS 3.1.2 (MU3) results in replication errors if there are hosts on the array that do not have any LUNs assigned to them. This behavior is due to a bug in the 3PAR InForm OS 3.1.2 (MU3) SMI-S provider.

As a workaround, remove the hosts that do not have any LUNs presented to them through the 3PAR InForm Management Console, and then restart data collection for the 3PAR array.

Replication Pairs

For HPE 3PAR 3.2.1 and later, if the remote replication configuration is broken, the **Replication Pairs** tab might not display remote replication pair information.

To view the remote replication pair information, run the CIM server in legacy mode. To run the CIM server in legacy mode, follow these steps:

1. Log in to the HPE 3PAR through the CLI.
2. Run the following commands:
```
setcim -pol no_replica_entity
stopcim
startcim
```
3. Restart data collection for the HPE 3PAR array after approximately five minutes.

HP EVA

- When the EVA firmware and the Command View EVA support RAID6, SOM creates RAID6 (enhanced) capable storage pools (disk groups) that are capable of RAID 0, 1, 5, and 6 volumes. Basic disk groups continue to be created for configurations that are not RAID6 capable, such as RAID 0, 1, and 5.

EMC Isilon

- **Raw Capacity Tab**
  Isilon clusters have a single file system and hence the entire file system is mapped to the physical space. Hence the raw capacity tab is not available in the **Analysis** pane.
- **System Nodes**
  Data is not populated in the Shares and NAS System Ports tab views of a system node as these sub-components are not relevant to Isilon.
- **CheckPoints**
  Data is not populated for this tab as it is not relevant to Isilon.
- **NAS Replication Pairs Tab**
  The **Source File System** and **Target File System** columns display the policy name of the replication instead of file system name.

EMC VNX Unified Storage

The VNX Unified storage is listed in the Top Level Storage Systems inventory view only if its underlying block and filer storage systems are discovered and collected.

Discover and collect the block storage system before the Filer to see the VNX Unified storage. If the Filer is discovered and collected before the block storage is discovered, you must rerun data collection for the Filer.

EMC Celerra/VNX Filer

The following exceptions are encountered in the inventory tab views:

- **Volumes Tab**
  Does not display NMFS volume types for VNX Filer.
- **System Nodes Tab**
  Displays information about the Data Movers.
- **NAS Extent Tab**
  Displays extent details. For example, meta volumes, slice volumes, and so on. The Description property in

the Properties pane identifies the extent type.

The tab does not display the Used Size and Available Space properties for extents, except for pool.

EMC Clariion/VNX Block

The following exceptions are encountered in the inventory tab views:

- **Sub-LUN Tiers Tab**
  The **Sub-LUN tiers** tab in the **Sub-LUN Tier Policy** form does not display the Sub-LUN tiers and related information. To view the Sub-LUN tier information, double-click or **Open** the Volume in the **Volumes** tab and then select **Sub-LUN Tiers** tab in the **Storage Volume** form.

EMC Symmetrix VMAX Arrays

- **Storage Extents Tab**
  SOM does not collect disk extents data and related associations for EMC VMAX 8.x or higher versions. Therefore, the storage extents data is not available.
- **Backend Storage Tab**
  SOM does not collect disk extents data and related associations for EMC VMAX 8.x or higher versions. Therefore, the Storage Extent column data in the Backend Storage Tab is not available.
- **Storage System Topology View**
  The Storage System Topology view of the EMC VMAX storage system displays the connectivity between the initiator ports and the backend storage system. However, it does not display elements such as storage extents.
- **Existing Access Points**
  Before you discover EMC Symmetrix VMAX array using Solution Enabler 8.x, ensure that you delete the existing access points discovered using SMI-S provider 4.x.

EMC Symmetrix DMX Arrays

- **Performance Data of Symmetrix Arrays**
  The performance data shown for Symmetrix arrays does not match with the values obtained using the CLI tool. The difference in performance values is observed because of the difference in time when the data is collected by SOM and the CLI tool.
- **Raw Used Capacity of DMX Arrays**
  Since the value of the property `RemainingRawCapacity` returned by the SMI-S is always zero, SOM derives the Raw Used capacity from the remaining extents.
- **DMX savedevs**
  SMI-S v4.5 does not return all 'savedevs'—a limitation with the SMI-S provider. However, SMI-S v4.6 returns all 'savedevs'.
- **Capacity values of DMX Arrays**
  SOM derives the capacity values (Used Raw, Actual Mapped, and Actual Used Mapped) by iterating through all the volumes of the array and the HSG information as this is not available from SMI-S pool property. Therefore, the values do not match with those obtained from the device vendor tool.

EMC VPLEX Clusters

Pool Types

The virtual volumes are logically grouped into the following pools during data collection and displayed in the Pools tab (Inventory > Storage Systems > All Storage Systems) as follows:

- Primordial Pool (Claimed Storage Volumes) – represents all the storage volumes presented to the VPLEX cluster that are claimed.
- Primordial Pool (Storage Volumes used for Logging) – represents all the storage volumes presented to the VPLEX cluster that are used for logging (Metro or Geo configuration only).
- Primordial Pool (Storage Volumes used for Meta-data) – represents all the storage volumes presented to the VPLEX cluster that are used for meta-data.
- Primordial Pool (Used Storage Volumes) – represents all the storage volumes presented to the VPLEX cluster that are used.
- Distributed Device Pool - represents the capacity of all the globally visible devices on that cluster that are used for Distributed Devices.
- Local Device Pool raid-0 - represents all Local Devices of similar RAID type.
- Unused Extents - represents all unused extents.

> **Note:** SOM does not list the primordial pools, except for Used Storage Volumes, as parent pools in the **Pools Logical Data** tab because these pools do not have any child pools. SOM lists these pools in the underlying **Storage Extents** form.

NetApp 7-Mode NAS Device

The Quota tab in the inventory view, displays the quotas seen in the Quota Report tab on the device console using the NetApp OnCommand System Manager interface.

SOM displays a space for values that are displayed as 'Unlimited' by the device console.

For NetApp devices with version 8.1, the API returns only the `size-total` property and does not return the `filesystem-size` property (through the ONTAPI query). Therefore, the `size-total` property is used to derive the total size of the file system.

The NAS Replication Pairs tab displays snapmirrors for the discovered NetApp devices. However, this feature requires that you enable the license on the applicable NetApp device.

The NAS Dependent Hosts tab displays only one share if the NIFS and CIFS shares have the same name.

NetApp C-Mode NAS Device

For NetApp Cluster Mode (C-Mode) devices, SOM displays capacity values for the selected clusters, nodes, Vservers, and individual file systems.

For Clusters:

NAS System Capacity tab

- **Used Capacity** is the sum of the used capacity of all the file systems on all the Vservers.
- **Total Capacity** is the sum of the total capacity of all the file systems on all the Vservers.

Aggregate Capacity tab

- **Used Capacity** is the sum of the used capacity of all the aggregates of all the nodes.
- **Total Capacity** is the sum of the total capacity of all the aggregates of all the nodes.

For Vservers:

NAS System Capacity tab

- **Used Capacity** is the sum of the used capacity of all the file systems of the selected Vserver.
- **Total Capacity** is the sum of the total capacity of all the file systems of the selected Vserver.

For Nodes:

Aggregate Capacity tab

- **Used Capacity** is the sum of the used capacity of all the aggregates of the selected node.
- **Total Capacity** is the sum of the total capacity of all the aggregates of the selected node.

The NAS Replication Pairs tab displays snapmirrors for the discovered NetApp Cluster Mode (C-Mode) devices. However, this feature requires that you enable the license on the applicable NetApp Cluster Mode (C-Mode) device.

In NetApp Cluster Mode (C-Mode) devices, if the NIFS and CIFS shares have the same name, the NAS Dependency Host tab displays only one share.

HPE StoreEasy Storage

- **NAS System Capacity tab**
  - **Used Capacity** is the total used file system capacity.
  - **Total Capacity** is the total file system capacity.
- **Raw Capacity tab**
  - **Used Capacity** is the total capacity of file systems.
  - **Total Capacity** is the total capacity of the device.
- **NAS Replication Pairs Tab**

  SOM supports only Distributed File System (DFS) replication for StoreEasy.
  - **Source File System** is the local path of the source folder on the primary member.
  - **Target File System** is not populated due to a third party issue.
  - **When Synced** is the synchronization completion time. Until the replication process is complete, this column displays a default value, Jan 01 00:00:00 IST 9999.

- **Sync State** is the status of replication.
- **Remote System Identifier** is the target member of the replicated folder.

Chapter 4: Common Tasks

This section describes procedures that are common to many HPE Storage Operations Manager configuration and maintenance tasks. It includes the following topics:

- "Start or Stop SOM Services" below
- "Change the Display Name of Elements" on the next page
- "Delete Elements" on page 433
- "View Stale Elements" on page 434
- "Quarantine/Unquarantine Elements" on page 434
- "Launch Topology" on page 436
- "Create an Asset Record" on page 437

Start or Stop SOM Services

Stopping the SOM services before changing the SOM configuration prevents conflicting data from being stored in the SOM database. Some procedures require restarting the SOM services to read the updated configuration.

The following SOM services are running on the management server when SOM is installed:

- OVsPMD
- somtrapreceivermd
- somdbmgr (If you have installed SOM with embedded database)
- somjboss
- somaction

To start SOM services

- **Windows**: Run the following commands:
  - %OvInstallDir%\bin\ovstart (Starts all SOM services)
  - %OvInstallDir%\bin\ovstart -c <*service_name*> (Starts the specified SOM service)
- **Linux**: Run the following commands:
  - /opt/OV/bin/ovstart (Starts all SOM services)
  - /opt/OV/bin/ovstart -c <*service_name*> (Starts the specified SOM service)

To stop SOM services

- **Windows**: Run the following commands:
  - %OvInstallDir%\bin\ovstop (Stops all SOM services)
  - %OvInstallDir%\bin\ovstop -c <*service_name*> (Stops the specified SOM service)
- **Linux**: Run the following commands:
  - /opt/OV/bin/ovstop (Stops all SOM services)
  - /opt/OV/bin/ovstop -c <*service_name*> (Stops the specified SOM service)

Change the Display Name of Elements

SOM obtains the provider name of an element (host, storage system, switch, or fabric) during discovery and displays it as the name of the element.

You must be logged in as an administrator to change an element's display name. SOM does not replace the provider name in the database but stores both the names separately.

The display name is useful to track an element as an asset and to easily identify a generic element instance across the storage network. SOM reflects the new name in all the workspace views.

However, the node that SOM automatically generates for a discovered element retains its node name irrespective of this change in the display name.

To change an element's display name:

1. From the inventory or topology view, select an element.
2. Right-click the element, and then select **Edit Display Name**.
3. Do one of the following:
   - Type a new display name, and then click **Submit**.
     The display name must adhere to the following:
     - The name should not be a null value.
     - The name must not exceed 2000 characters.
     - The name can contain the following characters and symbols: letters, numerals (0 to 9), ~, @, *, _, -, +, ., < >, (), [ ], { }, |.
     - The name is case-sensitive; for example, “Element1” and “element1” are different element names.

     Or
   - Click **Reset To Default**, to change an element's display name to its provider name (default value).

Find Elements

You can find a discovered element directly from the SOM console by providing a string of characters (*not case-sensitive*) such as an element name or part of an element name. SOM finds the associated element. If multiple elements match, SOM displays a list of all elements that match the search criteria. SOM checks the following element attributes for a match:

- Name
- Hostname (*fully-qualified*)
- System Name
- IP Address

To find an element, follow these steps:

1. On the SOM console main menu, click or click **Tools** > **Find Element**. A search window appears.
2. Type the search criteria. For example, part of the element name that you want to find.
3. Click **Find**. SOM displays a list of all matching elements along with the description, if available.
4. To view the details of an element, click the element name.

Delete Elements

You must be logged in as an administrator to delete elements.

You can delete a discovered top level element such as a host, storage system or switch. When you delete an element, all its associations are also deleted.

Key points about deleting elements:

- You can trigger delete for the collectible elements. Deleting the collectible element will also delete the other elements collected along with them.
  For example,
  - For a cisco switch, you can delete the virtual switch. This will in turn delete the physical switch.
  - For brocade switch, you can delete the physical switch. This will in turn delete the virtual switch.
  - For host cluster, you can delete the member nodes. Deleting the last member node will delete the cluster.
  - For ESX server, deleting the ESX server will delete the VMs belonging to that ESX.
- You can delete only one element at a time from the SOM web console.
- When you reset the database or delete an element, it is recommended that you delete the contents of the

`repository` folder manually if you plan to use a different user for discovery the next time. The folder is located at the location:

- **Windows**: <*Install_Dir*>\HP\HP BTO Software\se\repository
- **Linux**: <*Install_Dir*>/var/opt/OV/se/repository/root/cimv2

To delete an element, use one of the following:

The Inventory workspace

1. Navigate to the Inventory workspace. Choose an element to delete from the Hosts, Storage Systems or Switches folder.
2. Select a row from the table view, right click and select . The delete confirmation message is displayed. Click **OK** to delete the element.

The Topology workspace

- Go to System Topology, select an element, right-click and select . The delete confirmation message is displayed. Click **OK** to delete the element.

View Stale Elements

Elements managed by SOM are considered stale if data is not collected even once during the freshness criteria defined in the data collection policy. You may want to look at the data collection policy and consider increasing the freshness interval. It is recommended that you periodically view stale elements to ensure that SOM can successfully collect data on devices in your environment.

To view stale elements, follow these steps:

1. Log on to SOM as an administrator.
2. Open a command prompt and navigate to the following directory:
   Windows: INSTALL_DIR\bin

   Linux: /opt/OV/bin
3. Run the following command:
   `somstaleelements –all`

The command displays all the stale elements in SOM. The list of elements is written to a file in CSV format in the directory from where you run the command.

Quarantine/Unquarantine Elements

Quarantine an Element

The **Quarantine** option is useful to exclude elements from data collection in the following situations:

- Data collection fails for three consecutive cycles.
- An element needs maintenance, that is, a firmware, hardware, or software upgrade.

> **Note:** Only those elements for which data collection is possible can be quarantined.

For quarantined elements, the **Is Quarantined** property is set to true. This property must be manually reset when an element can be reconsidered for data collection.

However, the quarantine option does not stop data collection if it is already in progress.

To quarantine an element:

1. From the inventory or topology view, select an element (host, storage system, switch, or fabric).
2. Right-click the element, and then select **Quarantine**.
   The quarantine confirmation message appears.
3. Click **OK**.

Disable Auto-quarantine

By default, SOM automatically quarantines an element if data collection fails (the property `disable.quarantine.on.datacollection.failure` is set to false in the `custom.properties` file). If you do not want SOM to quarantine elements automatically, follow these steps:

1. Stop the `somjboss` service.
2. Open the `custom.properties` file located in the following directory:
   - *Windows*: `%OvDataDir%\conf\som`
   - *Linux*: `/var/opt/OV/conf/som`
3. Set the following property to `true`:
   `disable.quarantine.on.datacollection.failure=true`
4. Save the `custom.properties` file.
5. Start the `somjboss` service.

Unquarantine an Element

You must unquarantine an element so that SOM can resume data collection for the element.

> **Note:** Data collection is scheduled as per the data collection policy that is associated with the element. To initiate an immediate collection, from the context menu, select **Start Collection** for a selected element.

To unquarantine an element:

1. From the inventory or topology view, select an element (host, storage system, switch, or fabric).
2. Right-click the element, and then select **Un-Quarantine**.
3. Click **OK**.

Enable Auto-unquarantine

You can configure SOM to automatically unquarantine elements during "Start Discovery" on page 165. To enable the auto-unquarantine feature, follow these steps:

1. Stop the `somjboss` service.
2. Go to the following directory:
   - *Windows*: `%OvDataDir%\conf\som`
   - *Linux*: `/var/opt/OV/conf/som`
3. Add the following property from the `custom.properties.sample` file to the `custom.properties` file:
   `discovery.auto.unquarantine=true`
4. Save the `custom.properties` file.
5. Start the `somjboss` service.

Launch Topology

Use the Topology Maps feature to view the System Topology and individual element topology.

System Topology

System topology displays the physical connectivity of all the storage elements in your network. You can access storage element nodes, and filter the view by fabrics and element types.

To view the topology map of your entire storage infrastructure, from the workspace navigation panel, click **Topology** > **System Topology**.

Element Topology

The following element topologies are available.

Storage System Topology	Displays an overview of the hosts and switches connected to a selected storage system.
Host Topology	Displays an overview of the storage systems and switches connected to a selected host.
Switch Topology	Displays the connectivity between a selected physical switch and its logical switches.

Fabric Topology	Displays the connectivity between the switches, storage systems, and hosts within the selected fabric.

To launch any of the element topologies, use one of the following:

Actions Menu

1. Navigate to the Inventory workspace view and the respective element folder (Hosts, Switches, Storage Systems, and Fabrics).
2. Select the element of interest in the table view.
3. Click **Actions** > **Launch Topology**. The selected element topology is displayed.

Inventory View

1. Navigate to the Inventory workspace view and the respective element folder (**Hosts**, **Switches**, **Storage Systems**, and **Fabrics**).
2. Select the element of interest from the table view, right-click and select **Launch Topology**.

Create an Asset Record

SOM enables you to keep track of your asset information for an element.

To create an asset record, follow these steps:

1. Navigate to the Inventory workspace. Expand any folder of your choice (Hosts, Switches, or Storage Systems) and select the relevant view. Fore example, if you wish to create an asset record for a storage system, click the storage system view under the Storage System folder.
2. Select the element from the table view, right-click and select Create/Edit Asset Record. The Asset Record form is displayed.
3. Enter the information for the asset . (See "Attributes" below table for details.)
4. Click **Save** to create the asset record.

Attributes	Description
Record Name	A name that identifies the asset.
Record description	Description of the asset.
Status	Any text that identifies the status of the asset such as for example, New, In Use, or Under Maintenance.
Type	Type of the element such as a storage system.

Attributes	Description
Offering	Type of offering.
Vendor	The company that supplied the element.
Model	The model of the element.
Serial Number	The serial number of the element.
Bar Code No	The barcode on the device.
Asset Code	The asset code assigned to the element.
Asset type	The asset type assigned to the element.
Asset tag	The asset tag assigned to the element.
Asset Category	The asset category assigned to the element.
Location	The location of the element; for example, Boston, Massachusetts.

Troubleshooting

This section provides information that might help you to troubleshoot scenarios that you might encounter while using SOM to monitor your storage environment. It includes the following:

- "Troubleshooting VMware ESX Servers and Virtual Machines" below
- "Troubleshooting Windows Agentless Discovery" on the next page

Troubleshooting VMware ESX Servers and Virtual Machines

SOM might encounter the following situations while discovery or data collection from VMware ESX Servers and VMs if the "Prerequisites to Discover VMware ESX Servers and Virtual Machines" on page 140 are not met:

Subsequent data collection does not reflect the changes made to ESX servers and VM configurations

Possible reason: SOM maintains an internal cache of ESX servers and VM configurations for two hours. Therefore, it may take up to two hours to reflect any changes made to the configurations.

Solution

Retry the data collection later.

SOM reports the external storage volumes as local drives or does not show the volume topology or element topology correctly

Possible reason: The VMware APIs do not enable the associating of external volumes to VMs correctly.

Solution

Do the following:

1. Set the data collection level to 'All' for the corresponding device profile to which the VM belongs.
2. Rediscover the host either with an agent or through agentless discovery.

Duplicate entries for VMs

Possible reason:

- If you recently re-imaged or hosted the VM on a different ESX server while retaining the IP address and DNS name, SOM treats this VM as new.

Or

- The VMtools was disabled on the ESX server while discovering the VMs either with an agent or through agentless discovery

Solution

Do one of the following:

- Re-run the data collection on the old and the new ESX servers. This will delete the older entry of the VM from SOM
- Ensure that the VMtools is running on the ESX server for proper reconciliation of the VMs

Datastore size on ESX server shows 0 GB

Possible reason: The user credential used to discover the ESX server might not have "Datastore Browse" privileges.

Solution

Ensure that the user credential has "Datastore Browse" privileges.

vCenter discovery fails with new or updated vCenter credentials

Possible reason: This is a known issue where connection details persist in SOM repository and the changes are not updated.

Solution

To resolve this issue, follow these steps:

1. Delete all ESX servers for the vCenter in SOM.
2. Delete the configuration value for the particular vCenter from the VMware repository located at `<DATA_DIR>\se\repository\root\cimv2\APPIQ_VCProviderConfig`.
3. Restart the somjboss service.
4. Enter the new vCenter credentials in SOM.
5. Rediscover the vCenter server in SOM.

Troubleshooting Windows Agentless Discovery

SOM might encounter the following situations while attempting to discover a remote Windows host if the prerequisites are not met:

Discovery Fails with error code: 0x800706BA

The following error message occurs in the `Discovery Log` file (accessible from the context menu):

```
A single exception was thrown during discovery. It is probably the reason we failed to discover anything. BAD_CREDENTIALS
```

Solution

Use the following steps on a remote Windows host to enable DCOM connections on TCP port 135.

1. On the remote Windows server, open the **Control Panel** > **Windows Firewall**.
2. In the left pane, click **Advanced Settings**.
3. In the left pane, right-click **Inbound Rules**, and select **New Rule**.
   The **New Inbound Rule** wizard appears.
4. Select **Custom**, and then click **Next**.
   The **Program** page appears.
5. Select **All Programs**, and click **Next**.
   The **Protocol and Ports** page appears.
6. Do the following:
   a. From the **Protocol Type** list, select **TCP**.
   b. From the **Local Port** list, select **RPC Dynamic Ports**.
   c. From the **Remote Port** list, select **Specific Ports** and type **1024-65535**.
7. Click **Next**.
   The **Scope** page appears.
8. Do the following:
   a. From the **Which local IP addresses does this rule apply to?** list, select **Any IP Address**.
   b. From the **Which remote IP addresses does this rule apply to?** list,
      i. To allow all remote connections, select **Any IP Address**
         Or
      ii. To enter specific IP addresses, select **These IP addresses**.
9. Click **Next**.
   The **Action** plan appears.
10. Select **Allow the connection**, and then click **Next**.
    The **Profile** page appears.
11. Select **Domain**, and then click **Next**.
    The **Name** page appears.
12. Type a name to identify the rule. For example, **ArchiveOne incoming DCOM connections**.
13. Click **Finish**.
14. Verify the rule is enabled.

Discovery Fails with error code: 0x00000005

The following error message occurs in the `Discovery Log` file (accessible from the context menu):

```
A single exception was thrown during discovery. It is probably the reason we failed to
discover anything. BAD_CREDENTIALS
```

```
Windows host had a failure. Skipping it for now. Error verifying device connection:
com.hp.se.disco.common.task.exceptions.DataCollectorException: Error initializing Windows
Agentless Bundle Connection: Message not found for errorCode:0x00000005
```

Solution

1. Check for the following files in the [*drive*:]`\Windows\System32` folder on the remote Windows host.
   - `DcomConfigurator_x64.exe`
   - `DcomConfigurator_x86.exe`
   - `hbatest.exe`
2. If the files do not exist, delete the following status file:

   On Windows SOM

   [*drive*:]`\Windows\Temp\<hostname>.hbacopied`

   On Linux SOM

   `/tmp/<hostname>.hbacopied`
3. Run the script.

   On Windows SOM

   [*drive*:]`\Program Files (x86)\HP\HP BTO Software\bin\somwindowsAgentlessDiscovery.ovpl –i`

   On Linux SOM

   `/opt/OV/bin/somwindowsAgentlessDiscovery.ovpl –i`

   The option "`-i`" ensures that subsequent discoveries run automatically after the script is completed.

   If there are more than 150 hosts to be discovered, it is recommended that you do not use the "`-i`" option. Instead, select hosts (less than 50 at a time) manually, and trigger discovery in time intervals of 10 minutes.

   For more information about the `somwindowsAgentlessDiscovery.ovpl` script, see the HPE Storage Operations Manager CLI Reference Page.

If the host is not discovered, follow these steps:

1. Copy the following files from the SOM management server to the Windows remote host.

   On Windows SOM

   From [*drive*:]`\Program Files (x86)\HP\HP BTO Software\newconfig\HPOvSEi\tools\thirdparty` to [*drive*:]`\Windows\System32` on the Windows remote host.

   On Linux SOM

   From `/opt/OV/newconfig/HPOvSEi/tools/thirdparty` to `C:\Windows\Temp` on the Windows remote host.
   - `DcomConfigurator_x64.exe`
   - `DcomConfigurator_x86.exe`
   - `hbatest.exe`

2. Run `DcomConfigurator_x86.exe` and `DcomConfigurator_x64.exe`.
3. Run discovery.

Remote Agent is Unavailable

Possible reason for failure:

Either the IP is not reachable, no agents are running, or the agent is on a non-default port or a firewall on the device proxy server or SOM management server is blocking the port.

Solution

- **On Windows SOM**

  Verify that port 135 is open and not blocked by the firewall on the remote Windows host.
- **On Linux SOM**

  Verify that ports 135 and 139 are open and not blocked by the firewall on the remote Windows host.

BAD_CREDENTIALS

Possible error: Either the credentials are incorrect or the WMI service is disabled.

`Message not found for errorCode: 0x80020009`

Solution

Enable the WMI service.

HBA card details are not available

After successful discovery and data collection of a remote Windows host, the details of its HBA cards are not populated in its form view (Inventory > Hosts > Discovered hosts or Inventory > HBA Ports).

Solution

1. Copy the `hbatest.exe` file from the SOM management server to the Windows remote host.
   On Windows SOM

   [*drive*:]`\Program Files (x86)\HP\HP BTO Software\newconfig\HPOvSEi\tools\thirdparty` to [*drive*:]`\Windows\System32` on the Windows remote host.

   On Linux SOM

   `/opt/OV./newconfig/HPOvSEi/tools/thirdparty` to `C:\Windows\Temp` on the Windows remote host.
2. Run the data collection.

Discovery fails if SOM is running on Linux

If a Windows host does not get discovered by the SOM management server running on Linux, check for the following:

1. The `nmsproc` user is running as the root user.
   Run the following command: `id nmsproc`
   The uid should be zero. For example, `uid=0(root) gid=0(root) groups=0(root)`

2. The `winexe` tool is not installed.
   Run the following command: `rpm -qa | grep winexe`

   If winexe is installed, the version of the installed package is displayed. For example, `winexe-1.00-3.3.x86_64`

Appendix A: Forms and Tabs

This appendix contains the tabs and forms referenced in this guide.

Using Inventory Views

The Inventory workspace is a collection of views to access details of storage infrastructure objects (elements) that are discovered by Storage Operations Manager.

Inventory views are categorized into element groups. Each view displays a pre-determined subset of properties of the elements in a group. Inventory form views display additional properties and sub-components of individual elements.

The information in a view is refreshed whenever data collection is triggered based on the freshness threshold that is specified for a data collection policy. The Collection Status indicates the status of data collection for an element.

Use the following inventory views to gain an in-depth understanding of a particular element's properties, and related components:

- "Hosts Views" on the next page
- "Switches View" on page 292
- "Storage Systems Views" on page 293
- "Fabrics View" on page 294
- "Nodes View" on page 294
- "Node Groups View" on page 296
- "FC HBA View" on page 297
- "HBA Ports View" on page 297
- "Switch Ports View" on page 463
- "Storage System Ports View" on page 298

Using the Analysis Pane

Use the Analysis pane to view the following information about a selected device:

- **Summary**
  Key information about a selected element.
  For example:

- Map Count - The number of Managed Access Points (MAP) for a discovered element.
- Last Data Collection Time - The time when a storage system, host, switch, or fabric was last contacted for data collection. This value is populated only if data can be collected from an element.
- Access Point - The IP address that was used to discover and collect data from an element.
- Data Collection Policy - The policy used to collect data from a discovered element.
- Node Groups - The node groups to which a discovered element belongs.

- **Capacity**
  Overall capacity utilization of a selected element. For more information, see "Viewing Device Capacity in the Analysis Pane" on page 298.
- **Analytics**
  Analytics information of a selected host. For more information, see "Viewing Host Analytics in the Analysis Pane" on page 301
- **Performance**
  Performance information about a selected element. For more information, see "Viewing Device Performance in the Analysis Pane" on page 302.

Hosts Views

Hosts are categorized into the following views:

- Discovered Hosts
  Includes the list of hosts discovered by Storage Operations Manager. This includes hosts, virtual servers and member nodes that belong to host clusters but not inferred or created hosts. For more information about the properties and components of a selected host, see Viewing Details of Discovered Hosts.
- Virtual Servers
  Includes the list of discovered virtual servers.
  For more information about the properties and components of a selected virtual server, see Viewing Details of Virtual Servers.
- Virtual Machines
  Includes the list of virtual machines hosted on the discovered virtual servers. For more information about the properties and components of a selected host, see Viewing Details of Virtual Machines.
- Inferred Hosts
  Includes hosts inferred based on host security groups, zones, and zone aliases configured in the environment. These hosts are managed without installing a CIM extension. For more information about the properties and components of a selected host, see Viewing Details of Inferred Hosts.

- Created Hosts
  Includes hosts that are created by the administrator using the CLI `somagentlesshostcreator.ovpl`. An administrator can group WWNs and create hosts that contain these WWNs. Host details such as, hostname, IP, DNS, Version, and OS can be specified along with the port WWNs (to be added or deleted) to create such hosts. For more information about the properties and components of a selected host, see Viewing Details of Created Hosts.
- Host Clusters
  Includes host clusters that are discovered through their cluster member nodes. Cluster members are also displayed in the Discovered Hosts inventory view. Use the Host Cluster column in the Discovered Hosts inventory view to link to the host cluster. Information about cluster member nodes and shared resources such as filesystems, disk drives, and volume manager volumes is available in the form view of the host cluster. For more information about the properties and components of a cluster, see Viewing Details of Host Clusters.

The **Analysis** pane displays the Host Capacity and Host Performance Metrics of a selected host.

Viewing Details of Discovered Hosts

To view the properties and components of a standalone host, from the hosts inventory view, double-click or **Open** a selected host to see the Host form.

Information about a related host component is available in the following tab views:

- "Hosts View: File Systems Tab" on page 480
- "Hosts View: Cards Tab" on page 481
- "Hosts View: Ports Tab" on page 481
- "Hosts View: Target Mappings Tab" on page 482
- "Hosts View: Multipathing Tab" on page 482
- "Hosts View: Volume Management Tab" on page 483
- "Hosts View: Disk Partitions Tab" on page 483
- "Hosts View: Disk Drives Tab" on page 483
- "Asset Record Tab" on page 479
- "Hosts View: NAS System Dependencies Tab" on page 484
- "Open Incidents Tab" on page 499
- "Hosts View: Storage System Dependencies Tab" on page 484
- "Hosts View: Switch Dependencies Tab" on page 485
- "Hosts View: Zone Membership Tab" on page 486
- "Hosts View: Presented Storage Details Tab" on page 486

The Host form displays the properties of a selected host in the **Properties** pane.

Viewing Details of Virtual Servers

The Virtual Servers view displays virtual servers discovered through the vCenter. The list of virtual servers is a subset of the Discovered Hosts list.

For additional properties and related components, double-click or **Open** a virtual server to see the following tabs in the Host Form.

- Virtual Machines
- Filesystems
- Cards
- Ports
- Target Mappings
- Multipathing
- Volume Management
- Disk Partitions
- Disk Drives
- Asset Record
- "Open Incidents Tab" on page 499

The **Properties** pane displays the properties of a selected virtual server.

Viewing Details of Virtual Machines

The Virtual Machines view displays the entire list of virtual machines that are hosted on the virtual servers discovered in the environment.

For additional properties and related components, double-click or **Open** a virtual machine to see the following tabs in the Host Form.

- Filesystems
- Disk Drives
- "Open Incidents Tab" on page 499

The **Properties** pane displays the properties of a selected virtual machine.

Viewing Details of Inferred Hosts

The Inferred Hosts view displays hosts that are specified using inference rules based on host security groups, zones, zone aliases configured on storage systems, and fabrics in the SAN. You can use regular expressions or customize predefined expressions to create inference rules. For more information about creating inferred hosts, see Inferring Agentless Hosts Based on Rules.

User credentials can be assigned to inferred hosts so as to discover and collect host information. A host becomes a managed host after successful discovery and data collection and is displayed in the Discovered Hosts inventory view.

> **Note**: A rule must be executed at least once for the hosts associated with the rule to be displayed. The view is refreshed whenever a rule is run and changes to the host topology are recalculated .

To view the properties and components of an inferred host, double-click or **Open** a selected host from the Inferred Hosts inventory view for the following tab views:

- Cards
- Ports
- Presented Storage
- "Open Incidents Tab" on page 499

The **Properties** pane displays the properties of a selected inferred host.

Viewing Details of Created Hosts

An administrator can group FC switch port WWNs (that are discovered) and specify hosts to be discovered without connecting to the hosts. Such hosts can be created using the CLI `somagentlesshostcreator.ovpl` and are displayed in the Created Hosts inventory view. User credentials can be assigned to such hosts so as to discover and collect host information. A host becomes a managed host after successful discovery and data collection and is then displayed in the Discovered Hosts inventory view.

To view the properties and components of a created host, double-click or **Open** a selected host from the Created Hosts inventory view for the following tab views:

- Cards
- Ports
- Presented Storage
- "Open Incidents Tab" on page 499

The **Properties** pane displays the properties of a selected created host.

Viewing Details of Host Clusters

The Host Cluster view displays host clusters that are discovered through their member nodes.

To view the properties, shared resources, and member nodes of a host cluster, double-click or **Open** a selected host cluster to see its Host form.

Information about each related component is available in the following tabs of the Host form:

- Member Nodes
- Shared Filesystems
- Shared Volume Manager Volumes
- Shared Disk Drives
- "Open Incidents Tab" on page 499

The **Properties** pane displays the properties of a selected host cluster.

Block Storage Systems View

Block storage systems display the following tabs:

- "Storage Systems View: Storage System Processors Tab" on page 486
- "Storage Systems View: Volumes Tab" on page 487
- "Storage Systems View: Disk Drives Tab" on page 498
- "Storage Systems View: Ports Tab" on page 507
- "Storage Systems View: Host Security Groups Tab" on page 488
- "Storage Systems View: Storage Extents Tab" on page 491
- "Storage Systems View: Replication Pairs Tab" on page 492
- "Storage Systems View: Backend Storage Tab" on page 494
- "Storage Systems View: Backend Switches Tab" on page 492
- "Storage Systems View: SCSI Controller Tab" on page 497
- Sub-LUN Tier Policies
- "Storage Systems View: Connected Switches Tab" on page 492
- "Asset Record Tab" on page 479
- "Storage Systems View: Pools Logical Data Tab" on page 499
- "Storage Systems View: Thin Provisioning Data Tab" on page 500
- "Storage Systems View: Volume Presentation Details Tab" on page 492
- "Open Incidents Tab" on page 499
- "Storage Systems View: Masked Hosts Tab" on page 498
- "Storage Systems View: Dependent Hosts Tab" on page 491
- "Storage Systems View: Dependent Storage Virtualizers Tab" on page 491

Details about a storage system's components (storage pools, volumes, extents, disks, and so on) are available in the **Properties** pane of an individual component form view.

The **Properties** pane displays the properties of a selected block storage system.

The overall capacity utilization and performance information of a selected storage system is available in the tabs of the **Analysis** pane. For details about the capacity metrics that are collected at the array level, see "Capacity Information of Block Storage Systems" below.

Performance information is specific to a device and depends on the device metrics that can be collected. For details about the performance collectors of a device, see the performance information of a device in "Viewing Device Performance in the Analysis Pane" on page 302.

Capacity Information of Block Storage Systems

The overall capacity information of a block storage system is based on the capabilities of a storage system.

The following tabs display the capacity utilization in the **Analysis** pane:

Raw Capacity

Displays a customizable chart that illustrates the raw capacity usage for the last seven days with the following metrics:

- **Used Raw** - Raw disk capacity consumed by RAID groups or other such disk groups on the array. Disks configured for use in provisioning volumes, regardless of whether volumes were allocated from those disk groups.
- **Total Raw** - The sum of all raw disk capacity (Used and Unused) of a storage system.

The raw capacity values come directly from the SMI instrumentation of storage arrays, where raw capacity is modeled as primordial storage pools.

Note: For block storage systems without storage pools, the raw capacity metrics are calculated based on storage volumes.

The list of pools and details of the used and unused raw space are available in the "Storage Systems View: Pools Tab" on page 488.

Double-click a pool to see its details and associated volumes and storage extents in the "Storage Pool Form" on page 464.

Example:

The chart shows 2.09 TiB of space used from the total 2.09 TiB of raw (unused + used) space over the last seven days.

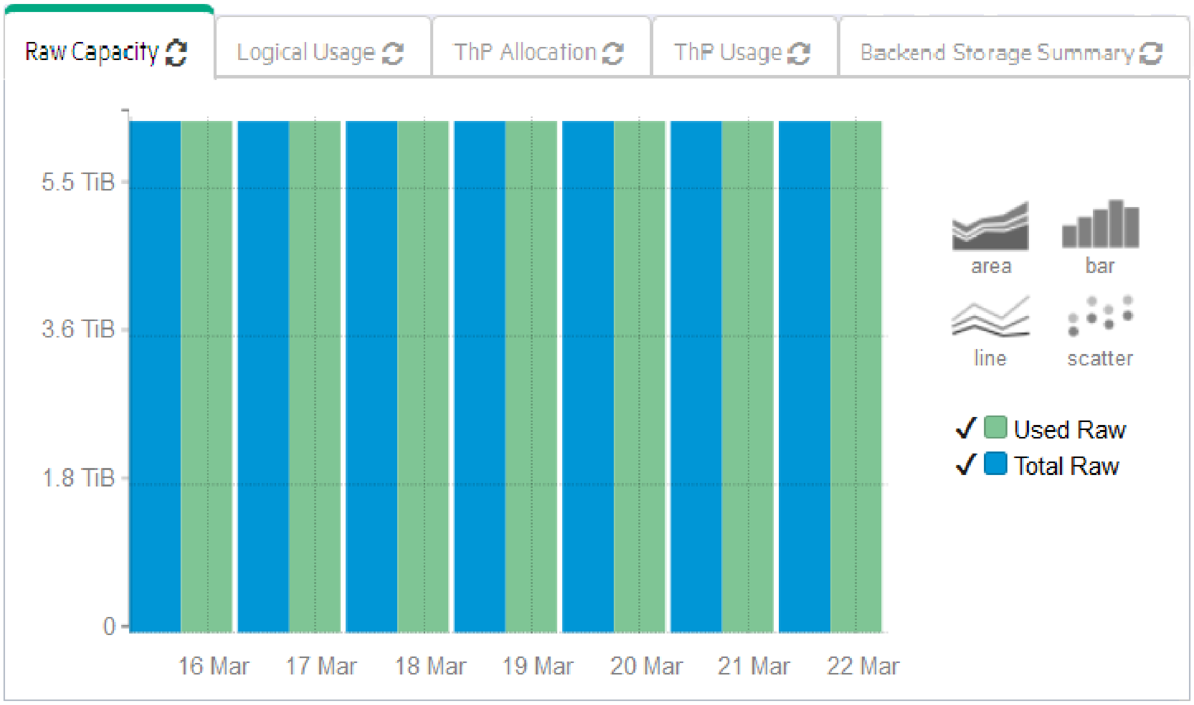

Logical Usage

Displays the aggregate capacity seen by the host. The customizable chart depicts the usage for the last seven days with the following metrics:

- **Mapped** - Sum of the volumes visible to hosts. For a volume to be mapped, it must have a logical mapping to at least one host initiator.
- **Allocated** - Sum of Mapped and Unmapped logical volumes allocated from the storage pools. Unmapped is the sum of volumes not visible to hosts. An unmapped volume is the storage committed as a single volume but not visible or potentially visible to any host initiator.

Details of Mapped and Allocated space from individual pools are available in the "Storage Systems View: Pools Logical Data Tab" on page 499.

Note: For block storage systems without storage pools, the logical usage metrics are calculated based on storage volumes.

Double-click a storage pool to see its details and associated volumes and storage extents in the "Storage Pool Form" on page 464.

Example

The total (mapped) space visible to hosts is 294.45 GiB from the 1021.08 GiB space allocated (mapped + unmapped) post RAID.

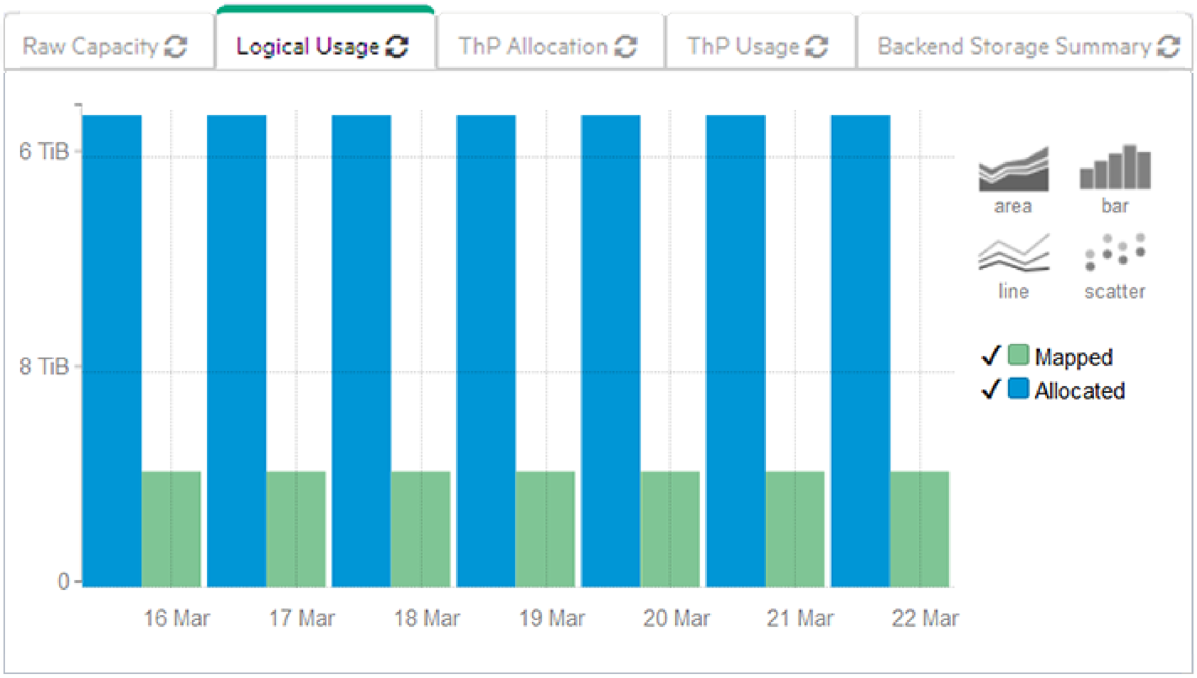

ThP Allocation

The ThP Allocation tab appears only if the selected storage system supports thin provisioning and is capable of extending volumes to a host until a volume reaches the configured maximum size.

This tab displays the aggregate physical capacity allocation of all configured storage pools. The customizable chart depicts the usage for the last seven days with the following metrics:

- **Actual Mapped** - Sum of the physical capacity that is allocated across all storage pools and visible to hosts.
- **Actual Allocated** - Sum of the physical capacity that is allocated across all storage pools. Physical capacity that is allocated cannot be used for creating volumes.
- **Total** - Sum of the physical capacity of all the configured storage pools in the array.

Details of the physical capacity allocated to individual storage pools are available in the "Storage Systems View: Thin Provisioning Data Tab" on page 500.

Double-click a storage pool to see its details and associated volumes and storage extents in the "Storage Pool Form" on page 464.

> **Note**: If you see an empty chart, it implies that the storage system has not been configured for thin provisioning although it has the capability.

Example

The chart shows 294.45 GiB of space mapped for usage from the 1021.08 GiB space that is allocated out of the total 1.57 TiB of raw space.

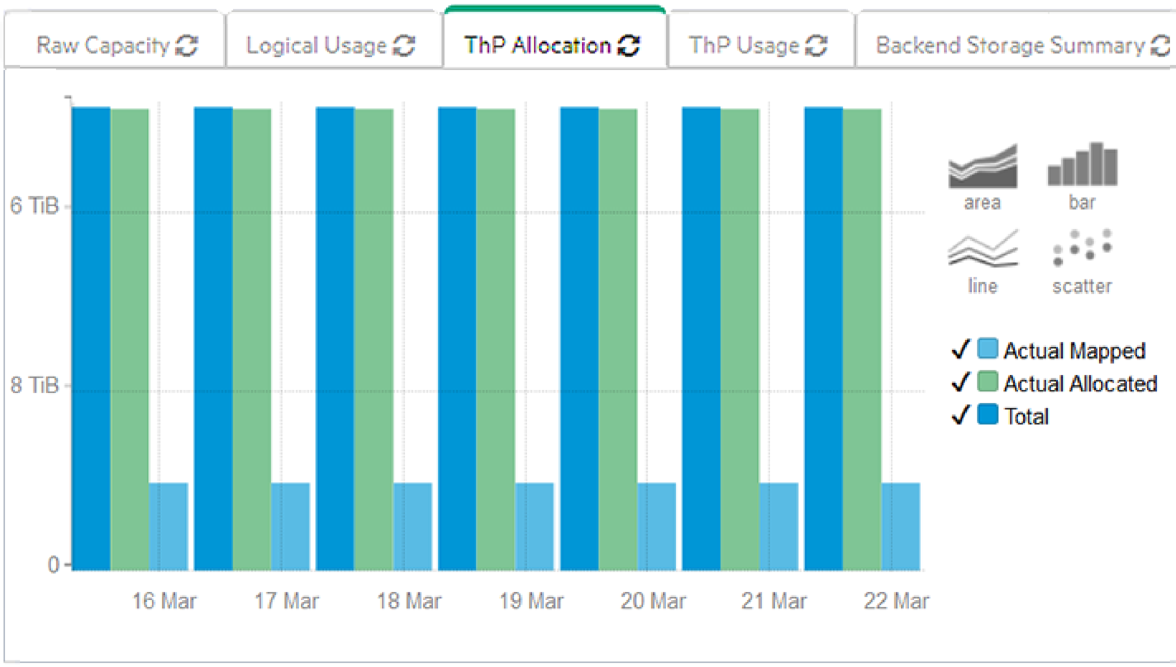

ThP Usage

The ThP Usage tab displays the usage summary of the physical capacity that is allocated. The customizable chart illustrates the usage for the last seven days with the following metrics:

- **Actual Used Mapped** - Sum of the physical capacity that is actually used by all the storage pools and is visible to the hosts.
- **Actual Used** - Sum of the capacity that is actually used by the volumes.
- **Actual Allocated** - Sum of the physical capacity that is allocated across all storage pools. Physical capacity that is allocated cannot be used for creating volumes.

Note: For block storage systems without storage pools, the ThP usage metrics are calculated based on storage volumes.

Details of the capacity utilization of the individual physical pools are available in the "Storage Systems View: Thin Provisioning Data Tab" on page 500.

Double-click a storage pool to see its details and associated volumes and storage extents in the "Storage Pool Form" on page 464.

Note: If you see an empty chart, it implies that the storage system has not been configured for thin provisioning although it has the capability.

Example

The following chart shows 294.45 GiB of mapped capacity that is actually used from 1021.08 GiB of the physical capacity that is actually allocated from 1021.08 of actually used raw capacity.

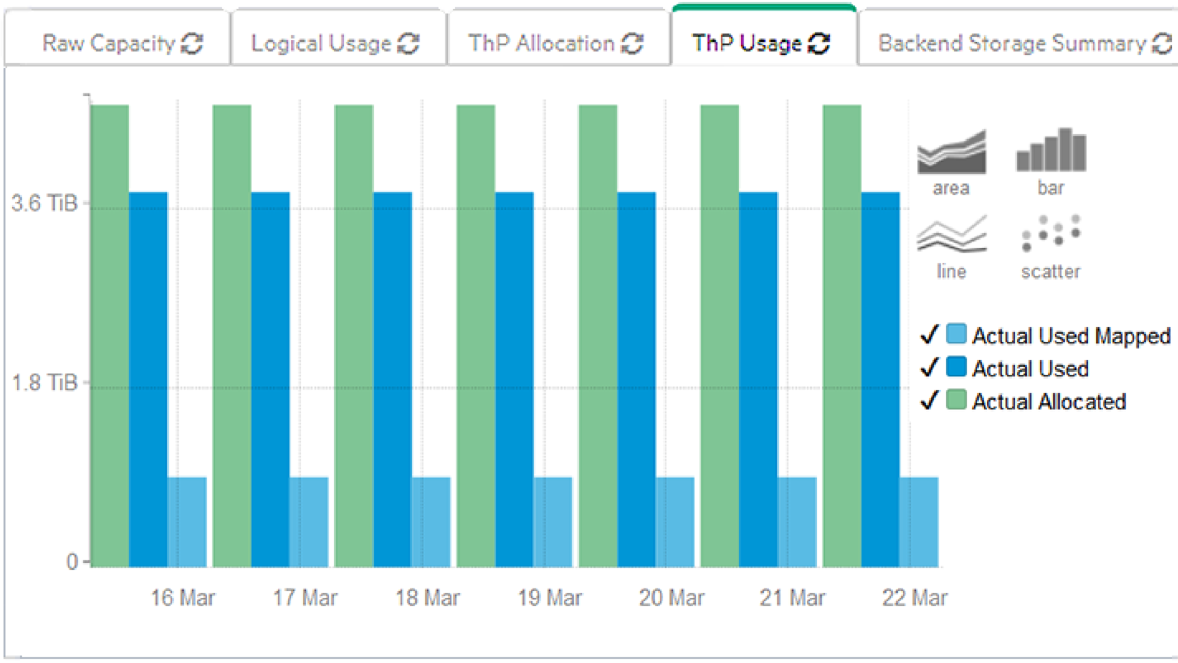

Backend Storage Summary

Displays a pie chart of the total volume capacity exposed to the selected front-end storage system. Each segment in the chart denotes the capacity of the associated backend array.

External Logical Usage

The external logical usage tab displays the capacity that is visible to a host from the external allocated capacity of the selected front-end storage system (the logical usage from backend devices).

The following metrics are used:

- **Mapped** - Sum of the volumes visible to hosts. For a volume to be mapped, it must have a logical mapping to at least one host initiator.
- **Allocated** - Sum of Mapped and Unmapped logical volumes allocated from the storage pools. Unmapped is the sum of volumes not visible to hosts. An unmapped volume is the storage committed as a single volume but not visible or potentially visible to any host initiator.

Capacity Details of a Storage Pool

The following details of a storage pool and its associated volumes and storage extents are available in the "Storage Pool Form" on page 464.

- Pool Type
- Total Space(GiB)
- Available Space(GiB)
- Used Space(GiB)

Capacity Details of a Storage Volume

The following details of a storage volume and its associated storage extents, disk drives, and target ports are available in the "Storage Volume Form" on page 465.

- LUN WWN
- Raid Type
- Volume Type
- Block Size
- Number of Blocks
- Actual Blocks
- Consumable Blocks
- Used Blocks
- Size (GiB)
- Raw Space
- Storage Pool

File Storage Systems View

File storage systems display the following tabs in the form view. Some of these tabs are visible only if data is collected for the related component.

- System Nodes
- File Systems
- Dedup and Compression (NetApp devices only)
- Snapshots
- Shares
- Qtrees
- Quotas
- NAS Extents
- NAS Replication Pairs
- Volumes
- Disk Drives

- Initiator Groups
- NAS Network Interface
- Backend Storage
- Ports
- CheckPoints
- NAS Dependent Hosts
- Open Incidents
- Asset Record

Certain File Storage Systems also support Block storage functionality. Therefore, the inventory view might include the following tabs:

- Ports
- Host Security Groups
- Volume Presentation Details
- Dependent Hosts
- Masked Hosts

Details about a storage system's related components are available in the **Properties** pane of the individual component's form view.

The **Properties** pane displays the properties of a selected file storage system.

The overall capacity utilization and performance information of a selected file storage system is available in the tabs of the **Analysis** pane. For details about the capacity metrics that are collected at the system level, see "Capacity of File Storage Systems" below.

Performance information is specific to a device and depends on the device metrics that can be collected. For details about the performance collectors of file storage systems, see the performance information in "Viewing Device Performance in the Analysis Pane" on page 302.

Capacity of File Storage Systems

The overall capacity utilization of a file storage system (NAS device) is available in the following tabs of the **Analysis** pane:

- **NAS System Capacity**

  The aggregate utilization of all the file systems on a selected NAS device using the following metrics:

  - **Used Capacity**
  - **Total Capacity**
- **Raw Capacity**

The physical capacity of a NAS device using the following metrics:

- **Used Capacity**
- **Total Capacity**

- **Backend Storage Summary**

  Displays a pie chart of the total volume capacity exposed to the selected front-end storage system. Each segment in the chart denotes the capacity of the associated backend array.

> **Note:** For individual file systems, the capacity information is available in the Analysis pane when selected from the **File Systems** tab view of a NAS device.

Cluster Storage Systems View

Cluster storage systems display the following tabs:

- Component Storage Systems

  Lists the component storage systems of a cluster, such as, nodes, vservers, block, and file storage systems.

- Asset Record

  Displays general asset information if specified for a storage cluster. The information in this tab appears only if an asset record is created for a cluster.

The **Properties** pane displays the properties of a selected storage cluster.

Backup Storage Systems View

Backup storage systems display the following tabs in the form view. Some of these tabs are visible only if data is collected for the related component.

- File Systems
- MTree
- Replication Pairs
- DD Boost
- Shares
- Quotas
- Snapshots
- Virtual Tape Library
- Network Interface
- Disk Drives

- Initiator Groups
- Ports
- Open Incidents
- Asset Record
- Dependent Hosts

Details about a backup storage system's related components are available in the **Properties** pane of the individual component's form view.

The **Properties** pane displays the properties of a selected backup storage system.

The overall capacity utilization and performance information of a selected backup storage system is available in the tabs of the **Analysis** pane.

Capacity Information of Backup Storage Systems

The overall capacity utilization of a backup storage system is available in the following tabs of the **Analysis** pane:

- **Backup System Capacity**

  Displays the capacity utilization of the file systems on a selected backup device using the following metrics:

  - **Used Size**: The space used by compressed data.
  - **Available Size**: The space available for data storage.
  - **Total Size**: The sum total of the Used and Available space on a backup system.

- **Compression**

  The amount of space utilized for data storage before and after compression.

  - **Pre-Compression Size**: Size of data written before compression.
  - **Post-Compression Size**: Space consumed after compression.

- **Total Compression Factor**

  It is the ratio of the pre-compression size to the post-compression size.

Note: For individual file systems, MTrees, and DD Boost units, the capacity information is available in the Analysis pane when selected from the **File Systems**, **MTree**, and **DD Boost** tab views of the backup system.

Forms

The SOM console includes forms for the following categories:

- "Host Forms" below
- "Switch Forms" on page 462
- "Storage System Forms" on page 463
- "Fabric Forms" on page 469
- "Node Forms" on page 471
- "Node Group Forms" on page 472

Host Forms

Filesystem Form

The Filesystem form displays the properties and related components of a filesystem that is mounted on a host.

Double-click a component in the following tab views to see its details in its form view:

- Disk Drives
- VM Volumes
- Disk Partitions

The **Properties** pane displays the properties of a shared filesystem.

HBA Card Form

The HBA Card form displays the properties of a selected Host Bus Adapter (HBA) card and its ports.

Double-click or **Open** a port in the **Ports** tab view to see its details in the HBA Port form.

The **Properties** pane displays the properties of a selected HBA card.

HBA Port Form

The HBA Port form displays the properties of a selected HBA port and the switch and storage ports that it might be connected to. Connected switch ports and target ports are visible only if the connected switches and storage systems are discovered by SOM.

Double-click or **Open** a port in the following tab views to see its details in its form view:

- Connected Switch Ports
- Target Ports

The **Properties** pane displays the properties of a selected HBA port.

Host Disk Drive Form

The **Host Disk Drive** form displays the properties and related components of a selected host drive.

Double-click a component in the following tab views to see its details in its form view:

- Filesystems
- Disk Partitions

The **Properties** pane displays the properties of a selected host drive.

Multipath Disk Form

The Multipath Disk form displays the properties of a selected multipath disk and its related components.

Double-click a component in the following tab views to see its details in its form view:

- Volume Management
- Disk Drives

The **Properties** pane displays the properties of a selected multipath disk.

Volume Manager Volume Form

The Volume Manager Volume form displays the properties of a selected logical volume manager configured on a host and its related components.

Double-click a component in the following tab views to see its details in its form view:

- Disk Partitions
- File Systems
- Multipath Disks
- Disk Drives

The **Properties** pane displays the properties of a selected logical volume manager.

Disk Partition Form

The Disk Partition form displays the properties of a selected partition and its related components on a host .

Double-click a component in the following tab views to see its details in the corresponding form view:

- Disk Drives
- Filesystems

The **Properties** pane displays the properties of a host disk partition.

Switch Forms

Switch Form

The Switch form displays the properties of a selected switch, and details about its ports.

Double-click or **Open** a port in the Ports tab to see its properties in the Switch Port form view.

The **Properties** pane displays the properties of a selected switch.

The **Analysis** pane displays the summary details and performance information of a selected switch port.

Fibre Channel Port Types

Understanding FC port types can help to identify ports along a storage path. The following table describes the different types of Fibre Channel ports:

Node Ports	Description
N_ port	Port on the node (such as, host or storage device) used with both FC-P2P or FC-SW topologies; also known as Node port.
NL_ port	Port on the node used with an FC-AL topology; also known as Node Loop port.
F_ port	Port on the switch that connects to a node point-to-point (for example, connects to an N_port); also known as Fabric port. An F_port is not loop capable
FL_ port	Port on the switch that connects to an FC-AL loop (such as, to NL_ports); also known as Fabric Loop port.
E_ port	Connection between two fibre channel switches. Also known as an Expansion port. When E_ports between two switches form a link, that link is referred to as an inter-switch link (ISL).
EX_ port	Connection between a fibre channel router and a fibre channel switch. On the side of the switch it looks like a normal E_port, but on the side of the router it is an EX_port.
TE_ port	Cisco addition to Fibre Channel, now adopted as a standard. It is an extended ISL or EISL. The TE_port provides not only standard E_port functions but allows for routing of multiple VSANs (Virtual SANs). This is accomplished by modifying the standard Fibre Channel frame (vsan tagging) upon ingress/egress of the VSAN environment. The TE_port is also known as Trunking E_ port.

General Ports	Description
Auto	Auto or auto-sensing port found in Cisco switches, can automatically become an E_, TE_, F_, or FL_port as needed.
Fx_port	Generic port that can become an F_port (when connected to a N_port) or a FL_port (when connected to an NL_port). Found only on Cisco devices where over-subscription is a factor.
G_port	G_port or generic port on a switch that can operate as an E_port or F_port. The G_port is found on Brocade and McData switches.
L_port	Loose term used for any arbitrated loop port, NL_port or FL_port. L_port is also known as Loop port.
U_port	Loose term used for any arbitrated port. U_port is also known as Universal port and is found only on Brocade switches.

Switch Ports View

The **Switch Ports** inventory view displays the entire list of switch ports in the environment that are discovered and managed by SOM. Use this view to see the host initiator ports, storage system target ports or other FC switch ports that a switch port is connected to. These ports are visible only if the connected switches, hosts, inferred hosts, or storage systems are discovered by SOM.

To see additional properties and ports connected to a switch port, double-click or **Open** a switch port to see the Switch Port Form.

Double-click a port in the following tabs to see its form view:

- **Connected Switch Ports**
- **Connected Host Ports**
- **Connected Storage System Ports**

The **Properties** pane displays the properties of a selected switch port.

The **Analysis** pane displays the summary details and performance information of a selected switch port.

Storage System Forms

Storage System Processor Form

The **Storage System Processor** form is useful to view the properties of a selected storage system (front-end) processor and its component details.

Double-click a port from the Ports tab to see its properties in the Storage System Port form.

The **Properties** pane displays the properties of a storage system processor.

Storage Pool Form

The **Storage Pool** form displays the properties of a selected storage pool and the volumes and storage extents (a contiguous array of real or virtual bytes) that are configured in a pool.

The Storage Pool Form is displayed when you open a storage pool from the following tab views of the Storage System Form:

- Pools
- Pools Logical Usage
- Thin Provisioning Data

Double-click a component in the following tab views to see its details in its form view:

- Volumes
- Storage Extents
- Pool Settings

  The RAID level configured for a storage pool. For additional properties of the RAID level of a storage pool, double-click or **Open** the pool setting to see the Pool Capabilities Form.

The **Properties** pane displays the properties of a storage pool.

The **Analysis** pane displays the summary (Name, Description, and Pool Type), capacity (Used and Available space), and performance information of a selected storage pool.

Pool Capabilities Form

The Pool Capabilities form displays the data redundancy properties that comprise the RAID level used in a selected storage pool.

The **Properties** pane displays the following properties:

- Name
- Default Spindle Redundancy
- Minimum Spindle Redundancy
- Maximum Spindle Redundancy
- Default Data Redundancy
- Minimum Data Redundancy
- Maximum Data Redundancy
- Minimum Delta Reservation

- Maximum Delta Reservation
- Default Delta Reservation
- No Single Point Of Failure
- Record Created
- Description

Storage Volume Form

The **Storage Volume** form displays the properties of the selected storage volume/ LUN and details about the ports, extents, disk drives and replication pairs associated with the selected storage volume/LUN.

Double-click a component in the following tab views to see its details in its form view:

- Storage Extents
- Disk Drives
- Sub-LUN Tiers (only for EMC VNX devices)
- Storage System Ports
- Replication Pairs (block storage volume/NAS file system)

The **Properties** pane displays the properties of a storage volume.

The **Analysis** pane displays the summary and performance information of a selected volume.

Storage Extent Form

The Storage Extent form displays the properties of a selected storage extent and details of the disk drives, volumes, pools, and source and target storage extents associated with the storage extent.

Double-click a component in the following tab views to see its details in its form view:

- Disk Drives
- Source Storage Extents
- Target Storage Extents
- Volumes
- Pools

The **Properties** pane displays the properties of a selected storage extent.

> **Note:** The Properties pane displays the **Provisioning Type** property. The **Provisioning Type** property displays the provisioning type of the storage disk. The Provisioning Type property can display the following values:

- **tdvv:** Thinly Provisioning Deduplicated Virtual Volumes
- **tpvv:** Thinly Provisioned Virtual Volumes
- **cpvv:** Copy Provisioned Virtual Volumes
- **full:** Fully provisioned virtual volumes

SCSI Card Form

The SCSI Card form is useful to see the properties of a selected internal SCSI controller card and the disk drives connected to the card.

For additional properties and related components of a disk drive connected to a SCSI controller, double-click or **Open** a selected disk drive to see the Disk Drive Form.

The **Properties** pane displays the following properties of a SCSI card:

- Name
- Controller Number
- Description
- Cluster Id
- Storage System
- Record Created

Sub-LUN Tier Policy Form

The Sub-LUN Tier Policy form displays the storage groups or volumes and tiers that are associated with a selected FAST policy.

Double-click a component in the following tab views to see its details:

- Storage Groups
  To see the volumes in a group, double-click a group to open the Storage Groups form. For the properties and components of a volume, double-click a volume from the Volumes tab to see the "Storage Volume Form" on the previous page.

  Note: Storage groups are applicable only for storage systems that support Fully Automated Storage Tiering (FAST) policies.

- Volumes (only for EMC VNX devices)
  The Volumes tab displays the list of volumes and associated storage pools for a selected Sub-LUN Tier policy. For the properties and components of a volume, double-click a volume from the Volumes tab to see the "Storage Volume Form" on the previous page.

> **Note:** In some situations, SOM might not immediately reflect the changes made to the volumes and associated storage pools for a selected Sub-LUN Tier policy. This is because of the SMI-S provider sending cached information. To resolve this issue, wait for some time and then re-run the data collection.

- Sub-LUN Tiers
  The Sub-LUN Tiers tab displays the following properties of a Sub-LUN tier:

Attribute	Description
Name	The name of the Sub-LUN tier.
Threshold Percentage	The specified maximum percentage of the associated storage group's logical capacity allocated to the tier. **Note:** This attribute is applicable only for storage systems that support FAST policies.
Threshold Size(GiB)	The specified maximum storage limit. **Note:** This attribute is applicable only for storage systems that support Adaptive Optimization.
Tier Technology	The type of storage disk drives.

To see the pools from which a tier is created, double-click a tier to open the Sub-LUN Tier form.

For the properties and components of a pool, double-click a pool from the Pools tab to see the "Storage Pool Form" on page 464.

Storage Disk Drive Form

The **Storage Disk Drive** form displays the properties and the following related components of a selected storage system disk drive:

- For block storage systems - storage extents, and volumes
- For file storage systems (NAS) - volumes, file systems, and NAS extents

The **Properties** pane displays the following properties:

- Name
- Description
- Model
- Vendor

- Architecture
- Hardware Version
- Serial Number
- Enabled Status
- Status
- RPM
- Maximum Access Time
- Compression Methodology
- Maximum Media Size (GiB)
- Default Block Size
- Maximum Block Size
- Minimum Block Size
- Uncompressed Data Rate
- Node WWN
- SCSI Port
- SCSI Target ID
- SCSI Bus
- OS LUN
- Storage System
- Record Created
- Disk Type

File Systems Form

The File Systems form displays the properties and components of a selected file system on a NAS device.

The form displays the disk drives and extents on which a file system is created and shares, snapshots or checkpoints belonging to a file system. Shares, and snapshots or checkpoints appear only if these exist for a selected file system.

Double-click a component in the following tab views to see its details in its form view:

- Shares
- Disk Drives
- NAS Extents
- Snapshots/Checkpoints

The Properties pane displays the properties of a selected file system.

NAS Extent Form

The NAS Extent form displays the properties of a selected NAS extent, the disk drives from which a NAS extent is created, and the file systems created on a NAS extent.

Double-click a component in the following tab views to see its details in its form view:

- "Storage Systems View: Disk Drives Tab" on page 498
- "Storage Systems View: File Systems Tab" on page 502

The **Properties** pane displays the following properties of a selected NAS extent:

- Name
- Description
- Block Size
- Number of Blocks
- Consumable Blocks
- Total Size (GiB)
- Used Size (GiB)
- Available Size (GiB)
- Storage System
- Record Created
- Status

Fabric Forms

Zone Alias Form

The Zone Alias form displays the list of ports associated with a selected zone alias and its properties.

For details of the Fabric to which a port belongs, double-click or **Open** a selected port to see the Port form.

The **Properties** pane displays the following properties of a zone alias:

Attribute	Description
Name	The name of the zone alias.
Description	A description of the zone alias
Record Created	The time when the zone alias was first contacted.

Attribute	Description
Fabric	The Fabric to which the zone alias belongs. For analysis information, or a detailed view of the Fabric's properties and related components, click **Lookup**.

Zone Set Form

The Zone Set form displays the properties of a selected zone set and the list of zones within a zone set. A zone can exist in more than one zone set. Zones sets are usually created for a particular task.

To see the properties of a zone and details of the aliases and ports in a zone, double-click or **Open** a selected zone to view the Zone Form.

The **Properties** pane displays the following properties for a zone set:

Attribute	Description
Name	The name of the zone set.
Description	A description of the zone set.
Record Created	The time when the zone set was first contacted.
Active	True or False. Indicates whether the zone set is active within the fabric. A switch fabric can have multiple zone sets, but only one zone set can be active.
Fabric	The Fabric to which the zone set belongs. For analysis information, or a detailed view of the properties and components of the fabric, click **Lookup**.

Zone Form

The Zone form displays the properties of a selected zone, the zone aliases, and FC switch ports within a zone.

Double-click a component in the following tab views to see its details in its form view:

- "Fabrics View: Zone Aliases Tab" on page 518
- **Ports** - For more information about a selected fabric port, double-click or **Open** a selected port to see the Port Form.

The **Properties** pane displays the following properties of a zone:

Attribute	Description
Name	The name of the zone.
Description	A description of the zone.
Record Created	The time when the zone was first contacted.
Active	True or False. Indicates whether the zone is active.
Protocol Type	
Zone Type	Specifies the type of zoning method that is implemented for the zone.
Fabric	The Fabric to which the zone belongs. For analysis information, or a detailed view of the properties and components of the Fabric, click **Lookup**.

Node Forms

Node Device Filter Form

The Node Device Filter form displays the device filters that can be used to determine the membership of a node group. Each Node Device Filter specifies a criteria that nodes must meet to qualify for inclusion in the node group. If you select more than one filter, nodes must fulfill all the criteria to be associated with the node group.

Node Device Filters

Filter	Description
Device Category	*Optional:* A particular category of devices. The drop-down list box displays the available categories. SOM provides four predefined categories – FC Fabric, FC Switch, Host, and Storage System.
Device Vendor	*Optional:* A particular device vendor. The drop-down list box displays the available device vendors.
Device Family	*Optional:* A particular family of devices. The drop-down list box displays the available device families.

Node Device Filters, continued

Filter	Description
Device Profile	*Optional*: A text string for Device Vendor and Device Family. The drop-down list box displays the available device profiles. If you are an administrator, click **Lookup** for additional options. • **Show Analysis** - To view analysis information of a selected device profile. • **Quick Find** - To select an existing device profile. • **Open** - To edit an existing device profile.

Node Group Forms

Device Category Form

The Device Category attribute indicates the pre-defined category of a device and is represented by an icon. It is displayed in the Nodes View of the Inventory workspace.

After discovery, an element is automatically associated with a pre-defined Node Group (Hosts, Storage Systems, FC Switches, and FC Fabrics) based on its device category. SOM manages an element based on its Node Group.

The Device Category attribute helps with the following:

- To determine the icon that SOM uses in map views to represent devices of a particular category.
- To determine the membership in Node Groups.

This form can be accessed from the Device Profile Form and displays the following properties:

Attribute	Description
Label	The device family name. For example, Cisco Catalyst 6500 Series Switches or HP AdvanceStack Routers. Maximum length is 255 characters. Alpha-numeric, spaces, and underline characters are permitted.
Unique Key	The required unique identifier that is important when exporting and importing device profile information within SOM. Maximum length is 80 characters. Alpha-numeric characters and periods are permitted. Spaces are not permitted.

Attribute	Description
Icon	Displays the icon that is associated with the Device Category. If you are an administrator, you can customize the icon.

Device Vendor Form

The Device Vendor attribute indicates the name of the manufacturer of a device; for example, HPE, Cisco, and so on.

This form can be accessed from the Device Profile Form and helps with the following:

- Configuring SOM monitoring behavior differently for each device vendor.
- Determining membership in a Node Group by device vendor.

The **Basics** pane displays the following properties of a Device Vendor:

Attribute	Description
Label	The device vendor name. Maximum length is 255 characters. Alpha-numeric, spaces, and underline characters are permitted.
Unique Key	The required unique identifier that is important when exporting and importing device profile information within SOM. This value must be unique. One possible strategy is to use the Java name space convention. For example: com.<your_company_name>.nnm.device_ profile.family.<family_label> Maximum length is 80 characters. Alpha-numeric characters and periods are permitted. Spaces are not permitted.
Icon	Displays the icon that is associated with the Device Category. If you are an SOM administrator, you can customize the icon.

Device Family Form

The Device Family property indicates the family name assigned by the vendor when a device is manufactured and helps with the following:

- Configuring SOM monitoring behavior differently for each device family.
- Determining membership in a Node Group by device family.

This form can be accessed from the Device Profile Form and lists the basic properties that are displayed for the Device Family:

Attribute	Description
Label	The device family name. For example, Cisco Catalyst 6500 Series Switches or HP AdvanceStack Routers. Maximum length is 255 characters. Alpha-numeric, spaces, and underline characters are permitted.
Unique Key	The required unique identifier that is important when exporting and importing device profile information within SOM. This value must be unique. One possible strategy is to use the Java name space convention. For example: com.<your_company_name>.nnm.device_profile.family.<family_label> Maximum length is 80 characters. Alpha-numeric characters and periods are permitted. Spaces are not permitted.
Management URL	Optional. The URL to the device's management page (provided by the vendor). This page is used to provide configuration information for the device and is usually organized by device family.
Icon	Displays the icon that is associated with the Device Category. If you are an administrator, you can customize the icon.

Device Profile Form

Every storage element that is discovered by the system is assigned a device profile based on the device vendor and device family provided by the vendor. The device profile is visible in the Nodes View of the Inventory workspace and determines how devices of this type are managed, including the icon and background shape displayed on maps.

The Basics pane displays the following properties of a Device Profile:

Attribute	Description
Device Model	The device model name or number designator, determined by the vendor.
Description	The description provided by the vendor. Maximum length is 255 characters: alpha-numeric, spaces, and special characters (~ ! @ # $ % ^ & * ( ) _ + -)
Device Family	Device family name provided by the vendor; for example Cisco Catalyst 6500 Series Switches or HP AdvanceStack Routers.

Attribute	Description
	Click the **Lookup** to access the Device Family Form for more information.
Device Vendor	Name of the vendor that manufactures the device. Click the **Lookup** to access the Device Vendor Form for more information.
Device Category	The value of this attribute determines which background shape NNMi uses for the map icon representing devices of this type. See About Map Symbols for more information about the possible values. Click the **Lookup** to access the Device Category Form for more information.
Author	Indicates who created or last modified the device profile. Click the **Lookup** to access the Author Form for more information.

Author Form

The Author attribute identifies who provided that instance of an object. Create a value for the Author attribute to represent you or your organization. The value you create then appears in the Author selection list in any appropriate form. A value of **HPE SOM Manager** implies that SOM created the object.

Caution: Each time a SOM upgrade is installed, objects with an Author attribute value of HPE SOM Manager are overwritten with the latest settings. When you modify an object provided by SOM, you must change the Author attribute value to ensure that your changes are not overwritten.

The Author attribute value is also useful for filtering objects in certain views and when using the SOM Export/Import feature.

To change an object's Author attribute value:

1. Open the form for the object.
2. Locate the Author attribute and click **Lookup**.
3. Do one of the following:
   - To create a new Author configuration, select **New**.
   - To select a previously defined Author attribute value, select **Quick Find**.

- To edit an existing Author configuration, select **Open**.

4. Type the text string that represents the new author.
5. Click **Save and Close** to save your changes and return to the previous form.

Tip: An administrator can set any author value as the default.

Attribute	Description
Label	The author name. The maximum length is 255 characters. Alpha-numeric, punctuation, spaces, and underline characters allowed.
Unique Key	Used as a unique identifier when exporting and importing configuration definitions. To ensure that the value you enter is unique, it is recommended that you use the Java name space convention when providing this value. It is also useful to include a part of the label value in the unique key, for example, com.<your_company_name>.author.<author_label>. **Caution**: After you click **Save and Close**, this value cannot be changed. The maximum length allowed is 80 alpha-numeric characters with periods but without spaces. **Note**: Do not begin the Unique Key value with com.hp.som. This prefix is reserved for use by HPE.

Additional Node Form

Administrators can add additional member nodes to node groups by specifying the case-sensitive Hostname or IP Address of the nodes. Such nodes belong to the node group regardless of any filters.

To add a node hostname, specify the fully-qualified, case-sensitive node Hostname attribute as it appears on the Node form.

Tip: To add multiple nodes to a node group, create a Custom Attribute for the nodes. Use the Additional Filters tab with the Custom Attribute value to group the nodes together.

Node Group Hierarchy Form

The Node Group Hierarchy form relates a parent node with a selected child node group.

The **Basics** pane displays the following properties:

Attribute	Description
Child Node Group	If you are an administrator, click **Lookup** for additional options. • **Show Analysis** - To view analysis information of a child node group. • **Open** - To open the Node Group Form of a child node group.
Expand Child in Parent Node Group Map	Indicates whether the nodes of a child node group are expanded and displayed in the parent node group map (Administrators only). If enabled, each node in the child node group appears on the parent node group map. If disabled, a hexagon represents a child node group on the parent node group map. Multiple child node groups if any are also displayed in the same manner. If a child node group is also a parent, its member nodes and child groups are displayed in the parent node group map if the Expand Child in Parent Node Group Map option is selected for each child node group. **Note**: This attribute appears in the Child Node Groups tab of the Node Group Form.

Tabs

The SOM console includes tabs for the following categories:

- "Host Tabs" on page 479
- "Storage System Tabs" on page 486
- "Fabric Tabs" on page 517
- "Node Tabs" on page 519
- "Node Group Tabs" on page 520

Fibre Channel Port Types

Understanding FC port types can help to identify ports along a storage path. The following table describes the different types of Fibre Channel ports:

Node Ports	Description
N_	Port on the node (such as, host or storage device) used with both FC-P2P or FC-SW topologies;

Node Ports	Description
port	also known as Node port.
NL_ port	Port on the node used with an FC-AL topology; also known as Node Loop port.
F_ port	Port on the switch that connects to a node point-to-point (for example, connects to an N_port); also known as Fabric port. An F_port is not loop capable
FL_ port	Port on the switch that connects to an FC-AL loop (such as, to NL_ports); also known as Fabric Loop port.
E_ port	Connection between two fibre channel switches. Also known as an Expansion port. When E_ports between two switches form a link, that link is referred to as an inter-switch link (ISL).
EX_ port	Connection between a fibre channel router and a fibre channel switch. On the side of the switch it looks like a normal E_port, but on the side of the router it is an EX_port.
TE_ port	Cisco addition to Fibre Channel, now adopted as a standard. It is an extended ISL or EISL. The TE_port provides not only standard E_port functions but allows for routing of multiple VSANs (Virtual SANs). This is accomplished by modifying the standard Fibre Channel frame (vsan tagging) upon ingress/egress of the VSAN environment. The TE_port is also known as Trunking E_ port.

General Ports	Description
Auto	Auto or auto-sensing port found in Cisco switches, can automatically become an E_, TE_, F_, or FL_port as needed.
Fx_port	Generic port that can become an F_port (when connected to a N_port) or a FL_port (when connected to an NL_port). Found only on Cisco devices where over-subscription is a factor.
G_port	G_port or generic port on a switch that can operate as an E_port or F_port. The G_port is found on Brocade and McData switches.
L_port	Loose term used for any arbitrated loop port, NL_port or FL_port. L_port is also known as Loop port.
U_port	Loose term used for any arbitrated port. U_port is also known as Universal port and is found only on Brocade switches.

Asset Record Tab

The Asset Record tab displays general asset information about a device, such as, departmental ownership, geographic location, contact information, and so on.

The information in this tab appears only if an asset record is created for a device so that the device can be tracked. The asset record can be created from the context menu in the inventory view. This is helpful to locate a device during troubleshooting.

The tab displays the following properties:

- Name
- Description
- Created Date
- Modified Date
- Status
- Storage System Type
- Offering - Dedicated or Leveraged. The value entered for this property while associating the device with a tier.
- Vendor
- Model
- Serial Number
- Bar Code
- Asset Code
- Asset Type
- Asset Tag
- Asset Category
- Geographic Location

Host Tabs

Hosts View: Virtual Machines Tab

The Virtual Machines tab displays the list of virtual machines hosted on a selected virtual server. Virtual machines can be discovered through the vCenter or through the individual ESX Servers. Discovering the vCenter results in one access point for all the ESX Servers managed by that vCenter.

If you discover the vCenter, and you also discover an individual ESX Server that is managed by the vCenter, the ESX Server will have a separate access point and is not included in the list of ESX Servers associated with the vCenter.

To view the properties and related components (filesystems, disk drives, and collector schedules) of a virtual machine, double-click or **Open** a selected virtual machine to see its Host form.

The tab displays the following properties:

- Name
- DNS Name
- Virtual Machine Name
- Description
- Vendor
- Model
- IP Address
- Operating System
- OS Version
- Size on Server (GiB)
- Virtual Machine State
- VM Tools
- Node

Hosts View: File Systems Tab

The File Systems tab displays the list of file systems mounted on a host.

A file system (also written as filesystem) is the allocation and management of files on a storage drive to facilitate efficient storage and retrieval.

The tab displays the following properties of a file system:

- Name
- Description
- Drive Type
- File System Type
- Total Size (GiB)

For additional properties of a file system and its related components (disk drives, VM Volumes, Disk Partitions), double-click or **Open** a selected file system to see the Filesystem Form.

The **Analysis** pane displays the filesystem summary details and the topology (the path for a host volume) of a selected host volume. For example, the path could be, host volume > HBA card > HBA port > switch port. And for a switch port, switch port > storage system port > storage volume.

Hosts View: Cards Tab

The **Cards** tab displays the list of Host Bus Adapter (HBA) cards for a selected host.

The tab displays the following properties:

Attribute	Description
Name	The name of the HBA card as collected from the host.
Node WWN	The unique 64-bit node worldwide name (WWN) identifier of the HBA card which is shared by all ports on the card.
Vendor	The vendor of the HBA card.
Model	The model name of the HBA card.
Serial Number	The serial number of the HBA card.

For additional properties and ports connected to an HBA card, double-click or **Open** a selected card to see the HBA Card Form.

Hosts View: Ports Tab

The Ports tab displays the list of Host Bus Adapter (HBA) ports for a selected host.

The tab displays the following properties:

Attribute	Description
Name	The name of the HBA port as collected from the host.
Port WWN	The unique 64-bit worldwide name identifier of the HBA port.
Connected Port WWN	The WWN of the switch port to which the HBA port is connected. This information is available only when the connected switch is discovered.
HBA Card	The HBA card that contains the port.
Port Speed in Gpbs	The speed of the HBA port.

For the properties and components of a selected HBA port, double-click or **Open** a selected port to see the HBA Port Form.

Hosts View: Target Mappings Tab

The **Target Mappings** tab displays the list of target mappings for a selected host.

Each target mapping represents a visible storage path to the host in terms of the initiator port on the host, the target port on the storage system and the LUN on the storage system.

The tab displays the following properties:

- HBA Port
- OS Lun Id
- Target Lun Id
- Target Port WWN
- Persistent
- SCSI Bus
- SCSI Target ID

For additional properties of a target mapping, double-click or **Open** a selected target mapping to see the HBA Port Target Form.

Hosts View: Multipathing Tab

The **Multipathing** tab displays information about the multipathing software configured on a host. This is based on the capability of a host and is visible only if a host supports multipathing.

The tab displays the following properties:

- Name
- Multipathing Type
- Multipathing Software
- Version of Software

For additional properties and related components (volume management, and disk drives) of a host path, double-click or **Open** a selected path to see the Multipath Disk Form.

Volume Management Tab

The Volume Management tab displays the logical volume manager(s) configured on a selected multipath host.

For additional properties and related components (disk partitions, file systems, multipath disks, and disk drives) of a volume manager, double-click or **Open** a selected volume manager to see the Volume Manager Volume Form.

Hosts View: Volume Management Tab

The **Volume Management** tab displays details of the logical volume manager(s) configured on a selected host. This is based on the capability of a host and is visible only if a host supports volume management.

The tab displays the following properties:

- Name
- Volume Management Software
- Version of Software

For additional properties and related components (disk partitions, file systems, multipath disks, and disk drives) of a volume manager, double-click or **Open** a selected volume manager to see the Volume Manager Volume Form.

Hosts View: Disk Partitions Tab

The **Disk Partitions** tab displays information about the disk partitions on a host. This is based on the capability of a host and is visible only if a host supports partitions.

The tab displays the following properties:

- Name
- Total Space (GiB)
- Description

For additional properties and components (disk drives and file systems) of a disk partition, double-click or **Open** a selected partition to see the Disk Partition Form.

Hosts View: Disk Drives Tab

The **Disk Drives** tab displays the list of disk drives on a host.

The tab displays the following properties:

Attribute	Description
Name	The name of the disk drive as discovered from the host.
Description	The type of disk drive. For example, Local Fixed Disk, Virtual Disk, Logical Volume SCSI disk drive, etc.

Attribute	Description
SCSI Bus	The number of the SCSI interconnect used by the disk drive.
Size (GiB)	The size of the disk drive.
OS Lun	The OS identifier of the logical volume on the host.

For additional properties of a disk drive, and its related components (file systems and disk partitions), double-click or **Open** a selected disk drive to see the Host Disk Drive Form.

Hosts View: NAS System Dependencies Tab

This NAS System Dependencies tab displays the list of file shares that are presented from NAS systems connected to a selected host.

In addition, you can see the mount point to which a file share is mapped on a selected host.

Hosts View: Storage System Dependencies Tab

The Storage System Dependencies tab displays the list of storage systems that a selected host is dependent on. This implies that the host uses volumes on the storage systems that are listed.

For each presented storage volume, information about its size, LUN ID (Logical unit number of the volume), the corresponding mount point, and the Volume Manager volume on the host is available.

Details of the connected ports—on the HBA and the storage system—and the visibility of the storage path are useful if a LUN is not available to the host at any given point in time.

Note: For a host, if there are fewer entries in the Switch Dependencies tab than in the Storage System Dependencies tab, the following conditions might exist:

- Switch not discovered
- Connected switch port is in different fabric
- Zoning misconfigurations

For more information about the possible cause, see the Path Visibility column in the Storage System Dependencies tab, and take appropriate action.

Open Incidents Tab

The Incidents tab displays a table view of the incidents associated with the selected object. These incidents are sorted by creation time so that you can view the incidents in chronological order. Use this view to determine which incidents are still open for the selected entity.

The Incident tab displays the following details about the selected incident:

- Severity
- Life Cycle State
- Last Occurrence Time
- Message

Double-click the row representing an incident to view the "Incident Form" on page 527. For more information, see "Incident Browsing Views" on page 27.

This tab does not display the incidents in the Closed state.

Hosts View: Switch Dependencies Tab

The Switch Dependencies tab displays details of the initiator and target switches that a selected host is dependent on to access presented storage volumes.

Storage paths are obtained from the host security groups (HSG) configured for a storage system after successful data collection of the storage system. Therefore, when a new host is discovered, data collection must be rerun for the connected storage systems to see information about connected switches in this tab view.

For each presented storage volume, information about the storage system, storage system port, LUN ID, and the target switch port details are available.

In addition, information about the VM volume, mount point, and the HBA port on the selected host and details of the initiator switch (port name, number, and slot number) are available.

Details of the connected ports—on the HBA and the storage system—and the visibility of the storage path are useful if a LUN is not available to the host at any given point in time.

Note: For a host, if there are fewer entries in the Switch Dependencies tab than in the Storage System Dependencies tab, the following conditions might exist:

- Switch not discovered
- Connected switch port is in different fabric
- Zoning misconfigurations

For more information about the possible cause, see the Path Visibility column in the Storage System Dependencies tab, and take appropriate action.

Disk Drives Tab

The Disk Drives tab displays the names of the disk drives on a host.

For additional properties and components (disk partitions and file systems) of a multipath host disk drive, double-click or **Open** a selected disk drive to see the Host Disk Drive Form.

Hosts View: Presented Storage Details Tab

The **Presented Storage Details** tab displays the list of storage volumes (LUNs) that are presented to a selected host.

For each LUN, details about its individual size, storage system, and storage pool are available.

The properties available in this tab view indicate the following about a LUN:

- **Is LUN Masked** - is True if a LUN is seen by at least one initiator port
- **Is LUN Zoned** - is True if a LUN has at least one visible path
- **is LUN mapped** - is True if a LUN is seen by at least one storage system port

Hosts View: Zone Membership Tab

The Zone Membership tab displays the list of all the member ports and related information, which belong to the same zones as the selected Host. The ports may be host ports or storage system ports.

The tab displays the following properties:

Attribute	Description
Port WWN	The unique 64-bit worldwide name identifier of the port.
Zone Alias	Zone Alias is a group of ports and it is a member of the zone. The Zone Alias might help to identify the devices.
Zone Name	Zone name of the selected host.
Zone Active	Indicates whether the zone is enabled.
Fabric	The fabric to which the zone of the selected host belongs.

Storage System Tabs

Storage Systems View: Storage System Processors Tab

The **Processors** tab displays information about the list of front-end controllers/adapters on the storage system.

Double-click a storage system processor to see its properties and connected ports in the Storage System Processor Form.

The tab displays the following properties of a storage system processor:

Attribute	Description
Name	The name of the front-end controller/adapter as discovered from the storage system.
Description	A description of the front-end controller/adapter.

Storage Systems View: Volumes Tab

The **Volumes** tab displays the list of volumes and associated pools on a selected storage system.

A volume is a virtual disk. Volumes are created in sizes that are desirable to be shown as LUNs to a host. A volume can be associated with more than one fibre channel port, resulting in multiple LUNs for the same volume. The defining characteristics of a LUN are the volume, port, and LUN number.

A storage pool is a group of disks associated together through a RAID configuration. The pool's capabilities define the level of protection for the associated volumes and LUNs.

For additional properties of a storage volume and its related components (storage system ports, storage extents, and disk drives), double-click or **Open** a selected volume to see the Storage Volume Form.

The Volumes tab displays the following properties:

Attribute	Description
Name	The name of the storage volume as discovered from the storage system.
Storage Pool	The name of the storage pool that the storage volume belongs to. For more details about the volumes and storage extents in a storage pool, see the Storage Pool Form.
File System Name	Applicable to file storage systems (NAS).
Dedup	The ratio of the physical storage space consumed without deduplication to the physical storage space consumed with thinly deduped virtual volumes. The dedup ratio does not include savings from inline zero detection.
Storage System Processor	The name of the storage processor associated with the storage volume. **Note:** This attribute is applicable only for storage devices where the storage system processor is the container for volumes.

The **Properties** pane displays properties of the volumes associated with the selected storage system.

The **Analysis** pane displays the summary (including properties such as Device Id, Total Size, Host Security Groups, and so on) and performance information of a selected volume.

Storage Systems View: Pools Tab

The **Pools** tab displays information about the storage pools associated with the selected storage system.

A storage pool is a group of disks associated together through a RAID configuration. The pool's capabilities define the level of protection for the associated volumes and LUNs. You should create at least one storage pool before provisioning a volume.

For additional properties of a storage pool and its related components (volumes, storage extents, and pool settings), double-click or **Open** a selected storage pool to see the Storage Pool Form.

The Pools tab displays the following properties of a storage pool:

Attribute	Description
Name	The name of the storage pool as discovered from the storage system.
Total Space (GiB)	The total space in gibibyte of the storage pool.
Available Space (GiB)	The space in gibibyte that is available in the storage pool.
Used Space (GiB)	The space in gibibyte that is utilized in the storage pool.
Dedup	The ratio of the physical storage space consumed without deduplication to the physical storage space consumed with thinly deduped virtual volumes. The dedup ratio does not include savings from inline zero detection.

The **Properties** pane displays properties of the storage pools associated with the selected storage system.

The **Analysis** pane displays the summary (Name, and Pool Type), capacity (Used and Available space), and performance information of a selected storage pool.

Storage Systems View: Host Security Groups Tab

The Host Security Group tab displays the list of defined host security groups and the host mode for each group.

A host security group is associated with a set of fibre-channel storage system ports and is created to secure access between HBA initiator ports and the storage volumes presented to a host from a selected storage system.

SOM uses the mapping definition to refer to the capacity that is accessible by one or more hosts external to a selected storage array (aggregated capacity of volumes that are accessible from hosts external to the subsystem).

For the properties of a host security group and its related components (storage system ports, volumes, initiator storage ports, and initiator HBA ports), double-click or **Open** a selected host security group to see the Host Security Group Form.

The host security group tab displays the following properties:

Attribute	Description
Name	The name of the host security group.
Host Mode	Displays the port settings for your operational environment. The settings for the host mode vary as per the storage system model. Host mode settings enable visibility of LUNs on the port to certain servers and HBAs.

Note: Incorrect provisioning operations can break the connection between an array and a host. If you rezone a device, make sure that no users or applications are using the device. For example, assume that ports of a storage system are members of zone set A, which is active. If you make zone set A inactive and the ports on the storage system are not members of the new active zone set, then the storage system becomes unavailable.

Expand for more information on how each storage system treats host security groups.

Host Security Groups on EMC CLARiiON Storage Systems

Host Security Groups on EMC Symmetrix Storage Systems

Host Security Groups on HDS Storage Systems

Host Security Groups on HP P6000 EVA Storage Systems

Host Security Groups on EMC CLARiiON Storage Systems

Keep in mind the following rules for host security groups on EMC CLARiiONstorage systems

- When a volume is created, it is assigned to one of the two controllers by default. Even though this volume is mapped to a controller, it is not visible to a host. The management server reports this volume as unmapped since it is not visible to a host initiator.

- Volumes can be only on SP_A or SP_B because CLARiiON is active/passive storage, which means it can have only one active path to a volume. Addition of initiators to any of the ports on a storage processor is listed for all ports of that storage processor.
- The host security group is created on all ports of the processor you select unless you select an initiator that uses a different processor and does not belong to a host security group. For example, assume you select processor SP_A, and then you select an initiator that belongs to SP_B buts does not belong to a host security group. The host security group is created for all ports on SP_B.
- Host security groups can consist of initiators (WWNs) only. You do not need to specify volumes. The initiator is shown in both host security groups SP_A and SP_B.
- Host security groups can consist of volumes (LUNs) only. You do not need to specify initiators.
- When you select an initiator for the host security group, the initiator has to be registered with the CLARiiON storage system.
- You can have more than one initiator in a security group if you have the proper multipathing software installed on the particular host where the initiator is located.

Host Security Groups on EMC Symmetrix Storage Systems

Keep in mind the following rules for host security groups on EMC Symmetrix Storage Systems.

- If LUN security is not turned on for an FA port, all volumes assigned to the FC port are visible to hosts that are on the SAN and have been zoned by the SAN. All volumes assigned to the FC port appear in the mapped category.
- When you create a host security group on a Symmetrix storage system, you are creating LUN mapping and masking in one step. In the native tools for Symmetrix storage systems, you will not see the host security group you created by using the management server. Instead you will see a volume bound to a port and a masked LUN bound to a host in the native tools.
- Host security group is associated with individual ports.
- Host security groups only allow one initiator for host security masking.
- To create a host security group, you must specify a port, initiator, and a volume.
- Every port has a LUN host security group, even if no LUNs are defined for that port. To bind a LUN to a port, edit the host security group and add the desired LUN to a port.
- You can also add LUNs to a Mask host security group. To add initiators, you must create the host security group.

Host Security Groups on HDS Storage Systems

Keep in mind the following rules for host security groups on HDS storage systems.

- FC port contains only volumes but no initiators (HBA WWN) assignment, the management server displays these volumes as unmapped since no external host can see these volumes yet.

- You can have zero to multiple initiators in a host security group.
- A host security group can be on only one port on the array. You can have host security groups with the same name, as long as they are on different ports.
- Host security groups appear in the native tool for HDS storage systems. In the logical view, the host security groups are listed by LDEV; in the physical view, they are listed by port.
- In the native tool for HDS storage systems, host security groups are referred to as a host security domain.

Host Security Groups on HP P6000 EVA Storage Systems

Keep in mind the following rules for host security groups on HP P6000 EVA storage systems.

- You can have multiple initiators per host security group.
- You can have zero to multiple volumes in a host security group.
- A host security group spans all ports on the array.

Storage Systems View: Storage Extents Tab

The **Storage Extents** tab displays information about the list of storage extents configured for a selected storage system.

For additional properties and related components (disk drives, source storage extents, target storage extents, volumes, and pools) of a storage extent, double-click or **Open** a selected storage extent to see the Storage Extent Form.

The tab displays the following properties of a storage extent:

Attribute	Description
Name	The name of the storage extent as discovered from the storage system.
CLPR	The number of Cache Logical Partitions (CLPR) on the storage extent.
Controller Name	The back-end controller that routes I/O from cache slots to the extent.

Storage Systems View: Dependent Storage Virtualizers Tab

The Dependent Storage Virtualizers tab view displays the list of front-end storage systems (virtualizers) that consume storage provided from a selected storage system.

To see the properties of a front-end storage system, double-click or **Open** a storage system.

Storage Systems View: Dependent Hosts Tab

The Dependent Hosts tab view displays the list of hosts that consume storage or shares from a selected storage system (block or NAS system).

To see the properties and related components of a host, double-click or **Open** a selected host to see its "Viewing Details of Discovered Hosts" on page 447.

Storage Systems View: Volume Presentation Details Tab

The Volume Presentation Details tab displays the list of storage volumes (LUNs) presented to the hosts that are connected to a selected storage system.

For each LUN, details about its individual size, storage pool, storage system, and the host to which it is presented are available.

The following properties in this tab view provide additional information about the storage volumes:

- **Is LUN Masked** - is True if a LUN is seen by at least one initiator port
- **Is LUN Zoned** - is True if a LUN has at least one visible path
- **is LUN mapped** - is True if a LUN is seen by at least one storage system port

Storage Systems View: Connected Switches Tab

The Connected Switches tab view displays the list of switches that are connected to a selected storage system.

This view is useful to see the connectivity details of the storage system ports and all the connected switch ports.

To see the properties and related components of a switch, double-click or **Open** a selected switch to see its "Switches View" on page 292.

Storage Systems View: Backend Switches Tab

The Backend Switches tab view displays the list of switches that are connected to the backend storage systems that provide storage to a selected front-end storage system (virtualizer).

This view is useful to see details of the switch ports on both the storage systems that is the backend and the virtualizer.

To see the properties and related components of a switch, double-click or **Open** a selected switch to see its "Switches View" on page 292.

Storage Systems View: Replication Pairs Tab

The **Replication Pairs** tab displays information about the list of volume replication pairs for a selected storage system.

The tab displays the following properties of a replication pair:

Attribute	Description
Source Storage Volume	The source storage volume for the replication pair. For details about the source storage volume and its components (Storage System Ports, Storage Extents, and Disk Drives), click to link to the Storage Volume Form.
Target Storage Volume	The target storage volume for the replication pair. For details about the target storage volume and its components (Storage System Ports, Storage Extents, and Disk Drives), click to link to the Storage Volume Form.
Copy Type	An SMI-S term used to describe the Replication Policy. Values are: • Async: Creates and maintains an asynchronous copy of the source. • Sync: Creates and maintains a synchronized copy of the source. • UnSyncAssoc: Creates an unsynchronized copy and maintains an association to the source.
Replica Type	An SMI-S term that provides information about how the Replica is being maintained. Values include: • Full Copy: Generates a full copy of the source object. • Before Delta: Maintains the source object from the Replica as delta data . • After Delta: Maintains the Replica from the source object as delta data. • Log: Maintains a log file of the changes from the Replica to the source object. • Not Specified: Indicates that the method of maintaining the copy is not specified.
When Synced	The date when the replication pair was last synchronized. Not all devices report this value.
Sync State	The synchronized state of the replication pair.
Sync Maintained	Specifies whether the synchronization of the replication pair is maintained.
Locality	Specifies whether the replication pair spans two devices and, if it does, whether the target or source is on this device.
Remote System Identifier	The IDs of remote devices if the replication pair spans several devices. This is useful if SOM has not yet discovered the other device.
Sync State Collection Time	The last time the sync state field was updated.

Storage Systems View: Backend Storage Tab

The **Backend Storage** tab displays details of the storage volumes that are consumed by the selected front-end storage system in a virtualized storage environment.

In a virtualized storage environment, a front-end storage array (acting as a storage virtualizer) serves as the access point for several storage arrays from which it can consume volumes (called the backend storage). Virtualized storage extents enable administrators to efficiently manage storage volumes and data access for improved performance and cost reductions.

If the backend storage system is not discovered, only the front-end storage extent and the storage system port is displayed in the following columns:

- Storage Extent
- Initiator Port
- Initiator Switch Port (if the switch is discovered)
- Initiator Switch (if the switch is discovered)

Subsequently, after the associated backend storage is discovered, data is populated in the following columns:

- Target Switch (if the switch is discovered)
- Target Switch Port (if the switch is discovered)
- Target Port
- Backend Volume
- Backend Storage System

For the connectivity information between the front-end and backend storage systems, double-click or **Open** a storage extent to see the Storage Extent Connection Form. The form view allows you to navigate to the analysis information and the form view of each component.

The Backend tab displays the following properties:

Attribute	Description
Storage Extent	The name of the storage extent that is created on the selected front-end storage system. For details about the storage extent and its components (disk drives, source storage extents, target storage extents, volumes, and pools), see the Storage Extent Form.
Initiator Port	The fibre-channel port of the front-end storage system virtualizer. For the properties of the front-end storage system port and its connected switch ports, see the Storage System Port Form.
Initiator	The fibre-channel port of the switch connected to the front-end storage system.

Attribute	Description
Switch Port	For the properties of the initiator switch port and its connected front-end storage system ports see the Switch Port Form.
Initiator Switch	The switch that is connected to the front-end storage system. For the properties of the initiator switch and its connected ports, see the Switch Form.
Target Switch	The switch that is connected to the backend storage system. For the properties of the target switch and its connected ports, see the Switch Form.
Target Switch Port	The fibre-channel port of the switch connected to the backend storage system. For the properties of the target switch port and its connected backend storage system ports see the Switch Port Form.
Target Port	The fibre-channel port of the backend array to which the volume is mapped. For the properties of the backend storage system port and its connected switch ports, see the Storage System Port Form.
Backend Volume	The unique volume name of the SCSI LUN on the fabric that is exposed by the storage controller (typically a RAID array) to the SAN Volume Controller. For details about the backend storage volume and its components (Storage System Ports, Storage Extents, and Disk Drives), see the Storage Volume Form.
Backend Storage System	The name that uniquely identifies the backend storage system. For details about the backend storage system and its components, see the Storage System Form.

The **Analysis** pane displays the Storage Extent Connection Summary tab with details of the switch ports, backend storage system and the target LUN ID of a selected storage extent.

Storage Systems View: Backend Storage Tab

The **Backend Storage** tab displays details of the storage volumes that are consumed by the NAS gateways.

The Backend tab displays the following properties:

Attribute	Description
Disk Drives	The LUNs presented from backend storage array.
Initiator Port	The fibre channel port of the NAS gateways. For the properties of the storage system port and its connected switch ports, see the

Attribute	Description
	Storage System Port Form.
Initiator Switch Port	The fibre channel port of the switch connected to the NAS gateways. For the properties of the initiator switch port and its connected system ports see the Switch Port Form.
Initiator Switch	The switch that is connected to the NAS gateways. For the properties of the initiator switch and its connected ports, see the Switch Form.
Target Switch	The switch that is connected to the backend storage system. For the properties of the target switch and its connected ports, see the Switch Form. **Note:** Data is populated in this column after the associated backend storage is discovered.
Target Switch Port	The fibre channel port of the switch connected to the backend storage system. For the properties of the target switch port and its connected backend storage system ports see the Switch Port Form. **Note:** Data is populated in this column after the associated backend storage is discovered.
Target Port	The fibre channel port of the backend array to which the volume is mapped. For the properties of the backend storage system port and its connected switch ports, see the Storage System Port Form. **Note:** Data is populated in this column after the associated backend storage is discovered.
Backend Volume	The unique volume name that is exposed by the backend storage system to the NAS Gateways. **Note:** Data is populated in this column after the associated backend storage is discovered.
Backend Storage System	The name that uniquely identifies the backend storage system. SOM supports HPE StoreEasy Storage with HPE 3PAR as backend storage.

Attribute	Description
	For details about the backend storage system and its components, see the Storage System Form. **Note:** Data is populated in this column after the associated backend storage is discovered.

Storage Systems View: SCSI Controller Tab

The **SCSI Controller** tab displays the SCSI information that is internal to the disks drives on a selected storage system. This view shows the names of the available internal back-end SCSI controllers of a selected storage system. Back-end controllers route I/O from cache slots to the disk.

For additional properties of a SCSI controller and its connected disk drives, double-click or **Open** a selected SCSI card to see the SCSI Card Form.

Storage Systems View: Sub-LUN Tier Policies Tab

The Sub-LUN Tier Policies tab displays the list of Sub-LUN Tier policies configured for a selected storage system.

Sub-LUN Tier policies are associated with storage groups or storage pools (that comprise storage volumes) and are applied to storage tiers. Storage tiers are a combination of drive technology and RAID protection. A storage tier comprises storage pools of the same disk drive type. Disk drive types such as, low-cost SATA drives, high-performance Fibre Channel drives, or Enterprise Flash drives (that are both high-performance and cost effective) are grouped into different tiers.

The Sub-LUN Tier Policies tab displays the following properties:

Attribute	Description
Name	The name of the Sub-LUN Tier policy.
Mode	The type of storage tier where the data is moved to, as per the Adaptive Optimization policy. The mode property displays the following values: • **Performance** mode moves more data to the high-performance tier. • **Cost** mode moves more data to the low-performance, less-expensive tier. • **Balanced** mode moves data such that performance and cost are balanced. **Note:** This property is applicable only for storage systems that support Adaptive Optimization.

For more information about the storage groups and tiers associated with a selected Sub-LUN Tier policy, double-click or **Open** a policy to see the "Sub-LUN Tier Policy Form" on page 466.

Storage Systems View: Disk Drives Tab

The **Disk Drives** tab displays the list of disk drives on a selected storage system.

The tab displays the following properties:

Attribute	Description
Name	The name of the storage disk drive as discovered from the storage system.
Size (GiB)	The size of the disk drive.
Status	Indicates the status of the storage system disk drive.
System Node	The name of a NAS system node. **Note**: This property is not relevant for block storage systems.
SCSI Card	The name of the storage controller card. **Note**: This property is not relevant for file storage systems (NAS).
Disk Type	The type of disk, internal or external.

For additional properties of a disk drive and its related components (storage extents and volumes), double-click or **Open** a selected disk drive to see the Storage Disk Drive Form.

Storage Systems View: Masked Hosts Tab

The **Masked Hosts** tab displays the list of client hosts that can see and access the volumes of a selected storage system.

For additional properties and related components of a host, double-click or **Open** a host to see the Host Form.

The Masked Hosts tab displays the following properties:

Attribute	Description
Name	The name of the masked host.
IP Address	The IP address of the masked host.

Storage Systems View: Pools Logical Data Tab

The Pools Logical Usage tab displays the storage capacity of the storage pools that is seen by hosts.

For additional properties and related components (volumes, storage extents, and pool settings) of a storage pool, double-click or **Open** a selected pool to see the Storage Pool Form.

The tab displays the following properties of a storage pool:

Attribute	Description
Name	The name of the storage pool.
Mapped Space	The sum of volumes visible to hosts. For a volume to be mapped, it must have a logical mapping to at least one host initiator.
Unmapped Space	The sum of volumes not visible to hosts. An unmapped volume is storage committed as a single volume but not visible or potentially visible to any initiator.
Allocated Space	The sum of Mapped and Unmapped that is the sum of logical volumes allocated from a storage pool.

The **Analysis** pane displays the summary (name and pool type), and capacity information of a selected storage pool.

Open Incidents Tab

The Incidents tab displays a table view of the incidents associated with the selected object. These incidents are sorted by creation time so that you can view the incidents in chronological order. Use this view to determine which incidents are still open for the selected entity.

The Incident tab displays the following details about the selected incident:

- Severity
- Life Cycle State
- Last Occurrence Time
- Message

Double-click the row representing an incident to view the "Incident Form" on page 527. For more information, see "Incident Browsing Views" on page 27.

This tab does not display the incidents in the Closed state.

Storage Systems View: Thin Provisioning Data Tab

The **Thin Provisioning Data** tab displays the allocation and usage of the physical capacity of the storage pools that are configured for a selected storage system.

Storage systems that support Thin Provisioning are capable of extending volumes to a host until a volume reaches the configured maximum size. For Storage Systems that do not have this capability, the Thin Provisioning Data tab is not shown.

For additional properties and related components (volumes, storage extents, and pool settings) of a storage pool, double-click or **Open** a selected storage pool to see the Storage Pool Form.

The tab displays the following properties:

Attribute	Description
Name	The name of the storage pool.
Total Capacity (GiB)	The sum of the configured storage pools in a storage system. This excludes raw disk space and external storage that is not configured in the storage pools.
Unallocated (GiB)	Available storage capacity that can be allocated. The data shown varies depending upon the RAID type that is used for allocation.
Actual Mapped (GiB)	The sum of physical storage that is allocated in the storage pools and visible to the hosts.
Actual Unmapped (GiB)	The sum of physical storage that is allocated in the storage pools but not visible to the hosts.
Actual Allocated (GiB)	The sum of physical storage that is allocated in the storage pools. Storage that is allocated cannot be used for creating volumes.
Actual Allocated (%)	The percentage of physical storage that is allocated in the storage pools to the Total Capacity of the storage pools.
Virtual Allocated (GiB)	The sum of storage that is virtually allocated for a storage pool.
Over Allocation (GiB)	The difference between virtual and physical allocation. A non-zero value indicates the amount storage that is allocated above the physical storage of a storage pool. Physical allocation is the sum of Actual Allocated and Unallocated storage.
Over Allocation	The percentage of storage that is over allocated in a storage pool.

Attribute	Description
(%)	If the percentage is non-zero, the storage is over allocated. Otherwise it is under allocated.
Actual Used Mapped (GiB)	The sum of allocated physical storage in the storage pools that is actually used and visible to the hosts.
Actual Used Unmapped (GiB)	The sum of allocated physical storage in the storage pools that is actually used but not visible to the hosts.
Actual Used (GiB)	The sum of the capacity that is actually used by the volumes in a storage pool.
Actual Unused (GiB)	The actual capacity that is not used.
Used (%)	The percentage of raw disk capacity that is used.

The **Analysis** pane displays the summary (name, description, and pool type), and capacity information of a selected storage pool.

Storage Systems View: Volumes Tab

The Volumes tab displays the LUNs configured on a selected file storage system (NAS) device.

The tab displays the following properties of a LUN:

- Name
- Storage Pool - Applicable only to block storage systems.
- File System Name

For more information about a selected LUN, double-click or **Open** a LUN to see the Storage Volume Form.

Storage Systems View: System Nodes Tab

SOM displays the following components of NAS devices in the NAS System Nodes tab:

- NetApp 7 mode: vFilers
- Celerra: Data Movers
- Ibrix (X9000): File Server nodes
- Isilon: Nodes
- Store Easy (X380): Nodes

For more information about a selected NAS system node, double-click or **Open** a node to view its properties and related components in the System Node Form. The tabs displayed for an individual NAS System Node are similar to those available in the File Storage Systems View.

Storage Systems View: File Systems Tab

The File Systems tab displays the list of file systems on a selected storage system.

A file system (also written as filesystem) is the allocation and management of files on a storage drive to facilitate efficient storage and retrieval.

The tab displays the following properties of a file system:

- Name
- Filesystem Type
- Description
- Total Size (GiB)
- Used Size (GiB)
- Available Size (GiB)

For additional properties of a file system and its related components (disk drives, NAS extents, and Snapshots/Checkpoints), double-click or **Open** a file system to see the File Systems Form.

Storage Systems View: Deduplication and Compression Tab (NetApp Devices Only)

For NetApp devices, SOM displays the **Deduplication and Compression** tab in the inventory view, which provides information about the total disk space saved on the volume because of the deduplication and compression features.

The **Deduplication and Compression** tab displays the following properties:

Attribute	Description
Deduplication State	Indicates whether the deduplication feature is enabled or disabled.
Compression State	Indicates whether the compression feature is enabled or disabled.
Compression Savings (GiB)	The total disk space saved by compression.
Percent Compression Savings (%)	The percentage of disk space saved by compression.
Deduplication Savings	The total disk space saved by deduplication.

Attribute	Description
(GiB)	
Percent Deduplication Savings (%)	The percentage of disk space saved by deduplication.
Total Savings (GiB)	The total disk space saved by deduplication and compression.
Percent Total Savings (%)	The percentage of disk space saved by deduplication and compression.
File System Name	The name of the file system.

Note: If deduplication or compression feature is not configured on a file system, the applicable attributes display blank values for that file system.

The **Deduplication and Compression Summary** tab in the **Analysis** pane displays summary information about the disk space saved by deduplication and compression features on volumes.

Note: For NetApp 7-mode devices, at the storage system level, the deduplication and compression information is displayed in both the inventory view and the analysis pane.

For NetApp C-mode devices, at the cluster level, the deduplication and compression information is displayed only in the analysis pane. SOM calculates the metrics by aggregating the file systems information from all the Vservers. At the Vserver level, the information is displayed in both the inventory view and the analysis pane.

Storage Systems View: Snapshots Tab

The Snapshots tab displays the list of snapshots that are created of the file systems of a selected NAS device.

A snapshot is an image (backup copy) of a file system and can be used to restore a file system if data gets corrupted. It is a set of reference markers, or pointers, to the data stored on a disk drive.

Snapshots differ from checkpoints in the following ways:

- Can reside locally as well as remotely
- Are read-only
- Are transient
- Cease to exist after being unmounted
- Track changed blocks at the file system level

To view the following properties of a snapshot, double-click or **Open** a snapshot to see the Snapshot/Checkpoint Form.

- Name
- File System Name
- Description
- Total Size (GiB)
- Status
- Snapshot ID
- Record Created
- Storage System

Quotas Tab

The Quotas tab displays the list of quotas configured for a selected file storage system.

A quota (user and group quotas) limits the amount of disk space and the number of files that a particular user or group can write to a file system. Directory tree quotas determine how much space is available for a specific directory and/or how many files can be written to it.

To view the following details of a selected quota, double-click or **Open** a quota to see the Quota form:

- Space Soft Limit (GiB)
- Space Hard Limit (GiB)
- File Soft Limit
- File Hard Limit
- Quota Type
- Quota Target
- Threshold
- Space Usage (GiB)
- File Usage
- File System
- Storage System
- Record Created

Qtrees Tab

The Qtrees tab displays the list of qtrees configured on a selected NAS device. A qtree is a subdirectory under the root volume directory.

To view the following details of a selected qtree and the quotas configured on a qtree, double-click or **Open** a qtree to see the Qtree form:

- Name
- FileSystem
- Status
- Storage System
- Record Created

Shares Tab

The Shares tab displays the list of static file systems shares (of type SMB/CIFS) configured on a selected NAS file system.

To view the following details of a selected file system share, double-click or **Open** a share to see the Share form:

- Name
- Mount Point
- Share Type
- Description
- System Node
- FileSystem
- Storage System
- Record Created

NAS Extents Tab

SOM displays the following components of NAS devices in the NAS Extents tab:

- NetApp Aggregates
- Celerra meta/pool volumes
- X9000 logical volumes

To view the disk drives from which a NAS extent is created, and the file systems created on a NAS extent, double-click or **Open** a NAS extent to see the NAS Extent Form.

Storage Systems View: Initiator Groups Tab

The Initiator Groups tab displays the list of initiator groups configured on a selected storage system. Each initiator group consists of host initiators and LUNs that the hosts can access.

The tab displays the following properties:

- Name
- Type - Indicates the protocol used within the group
- Operating System

For more information about the Initiators (host WWNs) and volumes that belong to an initiator group, double-click or **Open** a group to see the Initiator Group form.

NAS Replication Pairs Tab

The **NAS Replication Pairs** tab displays information about the list of file system replication pairs for a selected NAS system.

The tab displays the following properties of a NAS replication pair:

Attribute	Description
Source File System	The source file system of the replication pair. For details about the source file system, navigate to the File Systems tab in the inventory view of the source NAS device.
Target File System	The target file system for the replication pair. For details about the target file system, navigate to the File Systems tab in the inventory view of the target NAS device.
Copy Type	An SMI-S term used to describe the Replication Policy. Values are: • Async: Creates and maintains an asynchronous copy of the source. • Sync: Creates and maintains a synchronized copy of the source. • UnSyncAssoc: Creates an unsynchronized copy and maintains an association to the source.
Replica Type	An SMI-S term that provides information about how the Replica is being maintained. Values include: • Full Copy: Generates a full copy of the source object. • Before Delta: Maintains the source object from the Replica as delta data . • After Delta: Maintains the Replica from the source object as delta data. • Log: Maintains a log file of the changes from the Replica to the source object. • Not Specified: Indicates that the method of maintaining the copy is not specified.
When Synced	The date when the replication pair was last synchronized. Not all devices report this value.

Attribute	Description
Sync State	The synchronized state of the replication pair.
Sync Maintained	Specifies whether the synchronization of the replication pair is maintained.
Locality	Specifies whether the replication pair spans two devices and, if it does, whether the target or source is on this device.
Remote System Identifier	The IDs of remote devices if the replication pair spans several devices. This is useful if SOM has not yet discovered the other device.
Sync State Collection Time	The last time the sync state field was updated.

Storage Systems View: NAS Network Interface Tab

The NAS Network Interface tab displays the list of Ethernet ports and network cards on a NAS System Node.

To view the following properties of a selected Ethernet port, double-click or **Open** a port to see the NAS Network Interface Form:

- Name
- Description
- Status
- Port Type
- IP Address
- Mac Address
- NIC Name
- Port
- Storage System
- Record Created
- Role
- Data Protocol Access

Storage Systems View: Ports Tab

The Ports tab displays the list of FC ports of a selected storage system.

For additional properties of a port and its connected ports, double-click or **Open** a selected port to see Storage System Ports.

The tab displays the following properties:

Attribute	Description
Name	The name of the FC port as discovered from the storage system.
WWN	The unique 64-bit worldwide name identifier of the FC port.
Port Type	Indicates the type of FC port. For example, N, F, E, NL, FL, and so on.
Port State	Indicates the state of the FC port.
Storage System Processor	The front-end controller that contains the port. **Note**: The Storage System Processor property is not relevant for NAS devices.
Port Speed in Gbps	The port speed.

CheckPoints Tab

The CheckPoints tab displays the list of checkpoints that are created of the file systems of a selected NAS device.

A checkpoint is an image (backup copy) of a file system that can be used to restore a file system if data gets corrupted. It is a set of reference markers, or pointers, to the data stored on a disk drive.

Checkpoints differ from snapshots in the following ways:

- Reside on the same device as the original file system
- Can be read-only or read-write
- Are persistent
- Can exist and be mounted on their own
- Track changed blocks on each file in the file system

To view the following properties of a checkpoint, double-click or  **Open** a checkpoint to see the Snapshot/Checkpoint Form.

- Name
- File System Name
- Description
- Total Size (GiB)
- Status
- Snapshot ID
- Record Created
- Storage System

Component Storage Systems Tab

The Component Storage Systems tab displays the storage systems that comprise a storage cluster. The storage systems in a storage cluster could be any of the following:

- vservers
- nodes
- block storage systems
- file storage systems

For additional properties and related components of a cluster member, double-click or **Open** a selected component storage system to see its form view.

Backup Storage System

Backup Systems View: File Systems Tab

The File Systems tab displays the list of file systems on a selected backup storage system.

Double-click or **Open** a selected file system to view the following properties in its form view:

- Name
- Filesystem Type
- Description
- Total Size (GiB)
- Used Size (GiB)
- Available Size (GiB)
- Record Created
- Storage System

The **Capacity** tab in the **Analysis** pane displays the summary information about the amount of space utilized on the selected file system.

Backup Systems View: MTree Tab

MTrees are logical partitions of the EMC Data Domain file system that enable more granular management of the EMC Data Domain file system. For more information about MTrees, see the device documentation.

For EMC Data Domain storage system, SOM displays the **MTree** tab in the inventory view, which provides information about the MTree, its status, and the amount of raw data that is written into the MTree.

Double-click or **Open** a selected MTree to view the following properties in its form view:

Attribute	Description
Name	The pathname of the MTree.
Status	Indicates the status of the MTree. For example, Deleted, Read-only, Read/Write, Replication destination, and so on.
Pre-Comp Used(GiB)	The amount of raw data from the backup application that has been written to the MTree before compression.

The MTree form view displays the following tabs for the selected MTree:

- Quotas tab
- Shares tab
- Snapshots tab

Note: The MTree form view tabs are dynamic tabs, they are displayed depending on the availability of data.

The **Compression (Daily Written)** and **Compression Factor** tabs in the **Analysis** pane displays summary information about the data compression and compression factors for the selected MTree.

Quotas Tab

In EMC Data Domain storage system, MTree quotas set the storage space restriction for an MTree to prevent it from consuming excess space. There are two kinds of quota limits: hard limits and soft limits. For more information about Quotas, see the device documentation.

Backup Systems View: Replication Pairs Tab

The **Replication Pairs** tab displays information about the list of MTree replication pairs for a selected backup system.

Double-click or **Open** a selected replication pair to view the following properties in its form view:

Attribute	Description
Source	Source MTree.
Destination	Destination MTree.
Pre-Comp Remaining (GiB)	Amount of pre-compressed data remaining to be replicated.
Copy Type	An SMI-S term used to describe the Replication Policy. Values are:

Attribute	Description
	• Async: Creates and maintains an asynchronous copy of the source. • Sync: Creates and maintains a synchronized copy of the source. • UnSyncAssoc: Creates an unsynchronized copy and maintains an association to the source.
When Synced	The date when the replication pair was last synchronized. Not all devices report this value.
Sync State	The synchronized state of the replication pair.
Locality	Specifies whether the replication pair spans two devices and, if it does, whether the target or source is on this device.
Remote System Identifier	Remote backup system name.
Sync State Collection Time	The last time the sync state field was updated.

Backup Systems View: DD Boost

For EMC Data Domain storage system, SOM displays the **DD Boost** tab in the inventory view, which provides information about the DD Boost storage units.

Double-click or **Open** a selected DD Boost storage unit to view the following properties in its form view:

- Name
- Pre-Comp Used(GiB)
- Description

For more information about DD Boost properties, see the device documentation.

The **Compression (Daily Written)** and **Compression Factor** tabs in the **Analysis** pane displays summary information about the data compression and compression factors for the selected DD Boost storage unit.

Backup Systems View: Shares Tab

The Shares tab displays the list of file systems shares (of type NFS/CIFS) configured on a selected backup system.

Double-click or **Open** a selected Share to view the following properties in its form view:

- Name
- Mount Point

- Share Type
- Description
- System Node
- FileSystem
- Storage System
- Record Created

Backup Systems View: Quotas Tab

The Quotas tab displays the list of MTree quotas configured for the selected backup system.

The MTree quotas are the storage space limit for an MTree, Storage Unit, or VTL pool. Quotas help to prevent them from consuming extra space. There are two kinds of quota limits: hard limits and soft limits. You can set either a soft or hard limit or both.

For more information about Quotas tab properties, see the device documentation.

Backup Systems View: Snapshots Tab

The Snapshots tab displays the list of snapshots that are created of the MTrees in EMC Data Domain system.

A snapshot is an image (backup copy) of an MTree and can be used to restore an MTree if data gets corrupted. It is a set of reference markers, or pointers, to the data stored on a disk drive.

Double-click or **Open** a selected Snapshot to view the following properties in its form view:

- Name
- MTree
- Pre-Comp Size(GiB)
- Created Time
- Expires On
- Status

For more information about the Snapshots tab properties, see the device documentation.

Backup Systems View: Virtual Tape Library Tab

EMC Data Domain Virtual Tape Library (VTL) is a disk-based backup system that helps backup applications to connect and manage EMC Data Domain system storage using a feature that imitates the use of physical tape library. For more information about Virtual Tape Libraries, see the device documentation.

Double-click or **Open** a selected VTL to view the following properties in its form view and its associated Tapes, Tape Drives, Changers, and Access Groups in the "Tapes Tab" below, "Tape Drives Tab" below, "Changers Tab" on the next page, and "Access Groups Tab" on the next page:

- Name
- Status
- Description

Tapes Tab

The Tape tab provides information about the tapes supported by the selected EMC Data Domain Virtual Tape library (VTL). The tab displays information such as number of tapes supported by the VTL, their capacity, state, location, and so on.

The **Tape Drives** tab displays the following properties:

Attribute	Description
Name	The name of the Tape supported by the VTL
Capacity(GiB)	The total capacity of the tape.
Used Space(GiB)	The amount of space used on the tape.
State	The state of the tape.
Location	The location of the tape
Compression	The amount of compression performed on the data on a tape.

For more information about the properties for the tape, see the device documentation.

Tape Drives Tab

The Tape Drives tab provides information about the drives supported by the selected EMC Data Domain Virtual Tape library (VTL). The tab displays information such as number of drives supported by the VTL, their model, status, and so on.

The **Tape Drives** tab displays the following properties:

Attribute	Description
Name	The name of the VTL and the drive supported by it.
Model	The product name of the drive.

Attribute	Description
Vendor	The manufacturer or vendor of the drive.
Status	Whether the drive is Empty, Open, Locked, or Loaded.
Description	Information about the drive.

For more information about the properties for the tape drives, see the device documentation.

Changers Tab

The Changers tab provides information about the Changers supported by the selected EMC Data Domain Virtual Tape library (VTL). For each VTL, only one changer is available. For more information about changers, see the device documentation.

Double-click or **Open** a selected Changer to view the following properties in its form view

- Vendor
- Serial Number
- Status
- Description

For more information about the properties for the Changers, see the device documentation.

Access Groups Tab

EMC Data Domain Access groups contain a list of initiator WWPNs (worldwide port names) or aliases and the associated drives and changers they are allowed to access. For more information about EMC Data Domain access groups, see the device documentation.

Double-click or **Open** a selected Access Group to view the following properties in its form view:

Attribute	Description
Name	The name of the VTL and the associated access group.
LUNID	Device address – maximum number is 16383.
In Use Endpoints	Set of endpoints currently being used: primary or secondary.
Primary Endpoints	Initial (or default) endpoint used by backup application. In the event of a failure on this endpoint,

Attribute	Description
	the secondary endpoints may be used, if available.
Secondary Endpoints	Set of fail-over endpoints to use if primary endpoint fails.

For more information about the properties for the EMC Data Domain access groups, see the device documentation.

Backup Systems View: Network Interface Tab

The Network Interface tab displays the list of Ethernet ports and network cards on a System Node of a backup storage system.

Double-click or **Open** a selected Network Interface to view the following properties in its form view:

- Name
- Description
- Status
- Port Type
- IP Address
- Mac Address
- Storage System
- Record Created

Note: SOM does not display all IP aliases because it collects network interface details from the SNMP table `.iso.org.dod.internet.mgmt.mib-2.interfaces.ifTable`, which does not contain all the IP aliases created in the EMC Data Domain system.

Backup Systems View: Disk Drives Tab

The **Disk Drives** tab displays the list of disk drives on a selected backup storage system.

Double-click or **Open** a selected Disk Drive to view the following properties in its form view:

- Name
- Description
- Model
- Vendor
- Architecture

- Hardware Version
- Serial Number
- Enabled Status
- Status
- RPM
- Maximum Access Time
- Compression Methodology
- Maximum Media Size(GiB)
- Default Block Size
- Maximum Block Size
- Minimum Block Size
- Storage System
- Record Created
- Disk Type

Backup Systems View: Initiator Groups Tab

The Initiator Groups tab displays the list of initiator groups configured on a selected backup system.

For more information about the Initiators (host WWNs) that belong to an initiator group, double-click or **Open** a group to see the Initiator Group form.

Backup Storage Systems View: Ports Tab

The Ports tab displays the list of FC ports of a selected backup system.

Double-click or **Open** a selected Port to view the following properties in its form view:

Properties	Description
Name	The name of the FC port as discovered from the backup system.
WWN	The unique 64-bit worldwide name identifier of the FC port.
Port Type	Indicates the type of FC port (Physical FC or Endpoint).
Port Status	Indicates the operational status of the port.
Port Speed in Gbps	The port speed.
Description	Provides additional information about the port.

Backup Systems View: Open Incidents Tab

The Incidents tab displays a table view of the incidents associated with the selected object. These incidents are sorted by creation time so that you can view the incidents in chronological order. Use this view to determine which incidents are still open for the selected entity.

The Incident tab displays the following details about the selected incident:

- Severity
- Life Cycle State
- Last Occurrence Time
- Message

Double-click the row representing an incident to view the "Incident Form" on page 527. For more information, see "Incident Browsing Views" on page 27.

This tab does not display the incidents in the Closed state.

Fabric Tabs

Fabrics View: Switches Tab

The **Switches** tab displays the names of the FC switches that comprise the selected fabric.

For more information about the properties and ports of a fabric switch, double-click or **Open** a switch to display the "Switches View" on page 292.

The **Analysis** pane displays the summary details and performance information of a selected switch.

Fabrics View: Device Aliases Tab

An administrator uses a device alias to associate a Port WWN to a user friendly name. They are not VSAN specific, and can be used for other features besides zoning. Device Aliases can be configured manually for each switch, or can be propagated via Cisco Fabric Services. By default, device alias distribution is enabled. The device alias feature uses the coordinated distribution mechanism to distribute the modifications to all the switches in a fabric.

The **Device Aliases** tab displays the list of aliases configured for Cisco switch ports in a selected fabric.

For more information about the properties of a device alias, double-click or **Open** an alias to display the Device Alias form.

Fabrics View: Zone Aliases Tab

The **Zone Aliases** tab displays the names of the zone aliases in a selected Fabric.

A zone alias is a collection of zone members. A zone is a logical group of ports (N_Ports and NL_Ports or both) that are permitted to communicate with each other via the fabric. Ports and devices in a zone are called zone members. Ports that are members of a zone can communicate with each other, but they are isolated from ports in other zones. Devices, however, can belong to more than one zone. A zone alias can be added to one or more zones.

For more information, about the ports that are associated with a zone alias, double-click or **Open** a zone alias to display the Zone Alias Form.

Fabrics View: Zone Sets Tab

The **Zone Sets** tab displays the list of zone sets for a selected fabric element.

A zone set is a set of zone definitions for a fabric. A zone set can contain one or more zones, and a zone can be a member of more than one zone set. A zone set, can be activated or deactivated as a single entity across all switches in the fabric. A switch fabric can have multiple zone sets, but only one zone set can be active.

The tab displays the following properties:

Attribute	Description
Name	The name of the fabric zone set.
Active	True or False. Indicates whether the zone set is active.

To see the properties of a zone set and the list of zones within a zone set, double-click or **Open** a zone set to display the Zone Set Form.

Fabrics View: Zones Tab

The **Zones** tab displays information about the list of zones in the selected fabric.

A zone is a logical group of ports (N_Ports and NL_Ports or both) that are permitted to communicate with each other via the fabric. Using zoning, you can automatically or dynamically arrange fabric-connected devices into logical groups across a physical fabric. Zoning applies only to the switched fabric topology (FC-SW).

The tab displays the following properties:

Attribute	Description
Name	The name of the Fabric zone.
Active	True or False. Indicates whether a zone is active.

To see the properties of a zone and details of the aliases and ports in a zone, double-click or **Open** a selected zone to view the "Zone Form" on page 470.

Node Tabs

Nodes View: Capabilities Tab

The Capabilities Tab displays the list of capabilities that are predefined for a node based on the device that a node is associated with after the discovery of the device.

Capabilities help distinguish nodes from one another. Capabilities enable SOM and application programmers to provide more information about a node than is initially stored in the SOM database.

Note: Capability values cannot be modified as they are generated by SOM.

The Capability tab displays the following properties:

Attribute	Description
Label	A system defined label.

Nodes View: Node Groups Tab

The Node Groups tab displays the node groups to which a selected node belongs.

For additional information about a node group, double-click or **Open** a selected node group to display the Node Group Form.

To view the entire list of node groups provided by SOM and those that are created by the administrator, see the Node Groups View of the Inventory workspace.

Nodes View: Registration Tab

The Registration tab displays the registration properties and identifiers for a selected node.

Registration Attributes

Attribute	Description
Created	Date and time the selected node instance was created. SOM uses the locale of the client and the date and time from the SOM management server. **Note**: This value does not change when a node is rediscovered. This is because the Node

Registration Attributes, continued

Attribute	Description
	instance is modified, but not created.
Last Modified	Date the selected node instance was last modified. SOM uses the locale of the client and the date and time from the SOM management server. **Note the following:** • When a node is rediscovered, the Last Modified time is the same as the Discovery Completed time. This is because the node's Discovery State changes from Started to Completed. • When a Node is initially discovered, the Last Modified time is slightly later than the Created time. This is because node discovery does not complete until after the Node is created.

Object Identifiers Attributes

Attribute	Description
ID	The Unique Object Identifier, which is unique within the SOM database.
UUID	The Universally Unique Object Identifier, which is unique across all databases.
Node Object Access Role	Indicates the access permission for the selected node.

Node Group Tabs

Node Groups View: Device Filters Tab

The Device Filters tab displays a list of device filters that are specified for a selected node group. Device filters such as, Device Category, Device Vendor, Device Family, or Device Profile can be used to determine node group membership.

Note: Only administrators can set device filters for node groups.

For more information about a device filter, double-click or **Open** a device filter to see the Node Device Filter Form.

SOM ascertains the following for a node to belong to a node group:

- Evaluates Device Filters. If any exist, nodes must match at least one specification to belong to the node group.
- Evaluates Additional Filters. Nodes must also pass all specifications for Additional Filters if any to belong to the node group.
- Additional Nodes. If specified are always included in the node group, regardless of any filters.
- Child node groups. If added, are treated the same as Additional Nodes.

Node Groups View: Additional Filters Tab

The Additional Filters tab enables an administrator to use Boolean expressions to refine the requirements for membership to a node group based on device attributes.

Note: If a SOM administrator creates additional filters for a selected node group, SOM displays the Additional Filters expression.

Use the Filter Editor to create expressions that refine the requirements for membership to a node group. Make sure to design complex additional filters offline as a Boolean expression first. This method can help to minimize errors when entering your expressions using the Filter Editor.

Nodes must also match the expression specified in Additional Filters to belong to a node group.

SOM combines the results of all node group configuration settings in the following manner:

- Evaluates Device Filters. If any exist, nodes must match at least one specification to belong to the node group.
- Evaluates Additional Filters. Nodes must also meet all additional filter specifications to belong to the node group.
- Evaluates Additional Nodes that are specified and includes them in the node group, regardless of any filters.
- Evaluates Child Node Group results and treats them as Additional Nodes.

Note: The **Filter Editor** requires that your user name be assigned an administrator role.

Node Groups View: Additional Nodes Tab

The Additional Nodes tab lists case-sensitive Hostnames of the additional nodes that are added (SOM administrators only) as members of a selected node group. Node hostnames that are added are always included in the node group regardless of any filters.

For more information about a node hostname, double-click or **Open** a selected node hostname to see the Additional Node Form.

> **Note**: You can also add member nodes to a node group by specifying its address if the hostname is not available.

Node Groups View: Child Node Groups Tab

The Child Node Groups tab displays a list of node groups that belong to a selected parent node group. SOM provides the All Elements parent node group that comprises four child node groups: FC Fabrics, FC Switches, Hosts, and Storage Systems. Child node groups if added, are always included in a node group, regardless of any filters.

A set of node groups can be hierarchically configured, for example, based on geographical location. The parent node group might be named North America to represent the nodes in that continent. Additional node groups might exist for each country in which your business offices reside (for example, Canada, Mexico, and the United States). Each of these individual node groups is configured as a child node group of the North America node group.

The tab view displays the following:

- **Name** of a child node group
- **Expand Child in Parent Node Group Map** - displays the nodes of a selected child node group in its parent node group map if selected in the Node Group Hierarchy form of a child node group.

To view analysis information or edit a child node group, double-click or **Open** a selected child node group to see the Node Group Hierarchy Form (SOM Administrators only).

Node Groups View: Custom Properties Tab

The Custom Properties tab displays a list of custom properties/fields that are created by an administrator. These properties can be used for storage tiers or as filter criteria to generate custom reports.

Sudo User for Linux

Sudo allows users (permitted by the administrator) the ability to run commands as a root user. Using a non-root user account, information about disk drives, disk partitions, volume management, multipath path information, and the serial number of a host is not available to SOM. Discovery and data collection of a Linux host is based upon the privileges of the sudo user account as configured in the `/etc/sudoers` configuration file. For more information about creating and configuring a sudo user, see "Creating and Configuring a Sudo User" on the next page.

You can also log on to a host, with the sudo user account, and run the commands at the command line interface to get the required information. See "Commands for a Linux Host as a Sudo User" on page 126.

Creating and Configuring a Sudo User

As an administrator, you can either create a new user or configure an existing user as a sudo user to access a Linux host.

Create a New User

1. To create a new user, run the following command:
   `useradd <username>`

   For example, `useradd John`
2. Add the new user to the list of users in the `/etc/sudoers` file.

Configure a Sudo User

Appropriate privileges must be granted to the sudo user based on which the user can discover a Linux host without a CIM extension installed on the host.

To configure a sudo user:

1. Open the `sudoers` file in the edit mode using the following command:
   `visudo`
2. Add the required privileges to the user based on the set of commands that the sudo user might run.
   Examples:
   - `John ALL = (ALL) NOPASSWD:ALL`
     The user `John` is provided with privileges equivalent to a root user to run all commands.
   - `Michael ALL = (root) /sbin/lvdisplay`
     The user `Michael` is provided with the privilege to run only the `lvdisplay` command.
   - `Eric ALL = (root) /usr/sbin/dmidecode,/sbin/fdisk, /sbin/vgdisplay, /sbin/dmsetup, /sbin/pvdisplay, /sbin/lvdisplay, /sbin/multipath, /bin/ls`
     The user `Eric` is provided with the privilege to run a particular set of commands.

Reset the `requiretty` flag

Ensure that the `requiretty` flag in the `sudoers` file is disabled. If this flag is set, sudo will only run when the user is logged in to a real tty. This flag is off by default.

Incidents

SOM generates incidents based on third-party SNMP traps received from the storage infrastructure and management events generated by SOM.

This section includes the forms that comprise the configurations required to enable SOM to process and forward the incoming SNMP traps.

About Pairwise Configurations

Often two incidents have a logical relationship to each other, for example, `CiscoLinkDown` followed by `CiscoLinkUp`. There is no need for both incidents to take up room in your Incident view. Nesting the two together helps you do your job quickly and efficiently.

Use the Pairwise Configuration to pair the occurrence of one incident with another subsequent incident. When the second incident in the pair occurs, the first incident becomes a correlated child incident within the parent incident.

SOM provides default pairwise configurations for devices that SOM supports. You can view the default pairwise configurations in a table when you navigate to the Pairwise Configurations form in SOM.

When using Pairwise Configurations, note the following:

- You can use Payload Filters (for example, with trap varbinds) to identify the first and second incidents in a Pairwise Configuration.
- You can specify the same incident (for example, the same trap OID) as both the first and second incident configuration for a Pairwise Configuration.
- Using the Payload Filter to distinguish the first and second incidents (the first could represent a non-normal state and the second a normal state), different instances of the same incident configuration can cancel one another.
- You can also set up the Payload Filters such that the same incident instance cancels itself.
- You can use the same incident configuration in multiple Pairwise Configurations. For example:
  - Incident configuration A cancels both incident configuration B and incident configuration C
  - Incident configuration A cancels incident configuration B and incident configuration B cancels incident configuration C.
- Single incident instance can cancel multiple incident instances (for example, one Link Up trap cancels multiple instances of a Link Down trap).

  Note: If multiple Link Up/Link Down trap pairs are received within a 30 seconds, SOM investigates only once.

- Use the Duration time to specify the time in which the second incident configuration cancels the first incident configuration. This Duration is calculated from the `originOccurrenceTime` of the second incident backwards in time, canceling any number of first incidents within the Duration specified.

- You can also specify whether to delete any incidents that were canceled according to the Pairwise Configuration and that occurred within the time period specified by the Duration attribute.
- When matching incidents, SOM automatically takes into account the following values:
  - **SNMP Trap incidents**. SOM takes into account from which device the trap originated using the `cia.address` value of the source address of the trap.
  - **Management Event incidents**. SOM takes into account the name of the incident's Source Object and Source Node.

    Tip: SOM displays the Name value used to identify the Source Node and Source Object in the **Source Node** and **Source Object** attribute for each incident in the Incident form.

Tip: When configuring the Matching Criteria, you do not need to specify any of the cia names that SOM automatically takes into account. See "Matching Criteria Configuration Form (Identify Incident Pairs)" on page 537 for more information.

Related Topics:

"Pairwise Incidents Prerequisites" on page 539

"Correlate Pairwise Incidents" on page 207

Assign Incidents

If you are an SOM user with a Level 1 Operator (with more limited access privileges than Level 2 Operators), Level 2 Operator, or Administrator role, you can assign an incident to yourself or to another operator. If the incident is already assigned to another operator, you can change the assignment or "Unassign Incidents" on page 546.

To assign or change assignment for one incident:

1. Navigate to the Incident form of interest.
   a. From the workspace navigation panel, select the **Incident Browsing** workspace.
   b. Select any Incident view.
   c. Select the row representing the incident you want to assign.
2. Select **Actions** → **Assign** → **Assign Incident**.

   Tip: You can also right-click any object in a table or map view to access the items available within the **Actions** menu.

3. In the incident form's Basics tab, locate the **Assigned To** field.
4. From the dropdown list, select the required operator (assignee).
5. Click **Save** to save your changes or **Save and Close** to save your changes and exit the form..

The user name you entered or selected appears in the **Assigned To** column in any Incident views that include that incident.

To assign or change assignment for multiple incidents:

1. Navigate to the Incident view of interest.
   a. From the workspace navigation panel, select the **Incident Browsing** workspace.
   b. Select any Incident view.
2. Press Ctrl-Click to select each row that represents an incident you want to assign.
3. Select **Actions** → **Assign** → **Assign Incident**.
4. Select the user name.

The user name you selected appears in the **Assigned To** column in any Incident views that include those incidents.

Custom Incident Attribute Form

The Custom Incident Attributes (CIAs) form provides extended information that SOM gathered about the incident. For example, if the incident is reporting an SNMP trap, the varbind values are stored as CIAs. Each CIA includes a name, type, and value group that can be populated differently for different types of incidents.

To view custom incident attribute information, follow these steps:

1. Navigate to the **Incident** form.
   a. From the workspace navigation panel, select the **Incident Browsing** workspace.
   b. Select the incident view that contains the incident of interest; for example, **SNMP Traps**.
   c. To open the Incident form, double-click the row representing an incident. The Incident Form displays all details about the selected incident.
2. In the **Incident** form, select the **Custom Attributes** tab.
3. Double-click the row representing the Custom Incident Attribute (CIA) of interest.

See the table below for an explanation of the Name, Type, and Value attributes displayed.

> **Note:** All varbind values are stored as CIAs in SOM.

Custom Incident Attributes

Attribute	Description
Name	Name used to identify the CIA. The Custom Incident Attribute (CIA) name limit is 80 characters. If this limit is exceeded, SOM truncates the value from the left. **Note:** If different varbinds have the same oid, SOM appends a number to the original oid; for example: .1.2.3.4.5.6.2.7.1_1 and .1.2.3.4.5.6.2.7.1_2
Type	Describes the type of data that is stored for the CIA. Examples of types include: **Double** - Used to describe real numbers; for example 12.3 **Integer** - Used for integer numeric values; for example 1, 2, or 3 **String** - Used for character values **Boolean** - Used to store true or false values **Note:** All SNMP Trap types begin with **asn**. If the CIA represents a varbind value, SOM might provide additional types, such as **Counter**.
Value	For management events that are generated from SOM, this value is the CIA value in the incident that was provided by SOM. The Custom Incident Attribute value limit is 2000 characters. If this limit is exceeded, SOM truncates the value from the right.

Incident Form

The Incident form provides details for troubleshooting purposes. From this form you can access more details about the node involved, and the Source Object attribute provides more information about the interface, IP Address, connection, or SNMP Agent that is contributing to the problem.

Basic Attributes

Attribute	Description
Message	A description of the problem that you want SOM to display.
Severity	Seriousness that SOM calculates for the incident. Possible values are: • No Status • Normal

Basic Attributes, continued

Attribute	Description
	• Warning • Minor • Major • Critical • Disabled • Unknown
Priority	Used to communicate the urgency of resolving the selected incident. SOM sets this value to null by default. The lower the number the higher the priority. Possible values are: • None • Low • Medium • High • Top
Lifecycle State	Identifies where the incident is in the incident lifecycle. You control this value. • **Registered** – Indicates that an incident arrived in the queue stored in the SOM database. • **In Progress** – State selected by someone on your team to indicate that they are taking responsibility for investigating the problem. • **Completed** – State selected by someone on your team to indicate completion of the incident investigation and implementation of a solution. • **Closed** – Indicates that SOM determined the problem reported by this Incident is no longer a problem. • **Dampened** – Indicates that, within the configured *acceptable time period*, SOM does not submit the incident to the queue until after the time period (configured by the SOM administrator).
Source Node	Displays the Name attribute value of the node associated with the incident. **Note:** If the SOM database does not contain any Node object for this device, the source node value is **<none>**.
Source Object	Name used to indicate the configuration item that is malfunctioning on the source node.

Basic Attributes, continued

Attribute	Description
Assigned To	Name of the user to which this incident is assigned. This value must be a valid user name (determined by the SOM administrator).
Notes	Provided for communication among your team (for example, explanations or workarounds). Information might include reasons why the status was changed, what has been done to troubleshoot the problem, or who worked on resolving the incident. Type a maximum of 1024 characters. Alpha-numeric, spaces, and special characters are permitted. **Note:** You can sort your incident table views based on this value. Therefore, you might want to include keywords for this attribute value.

Incident Form: General Tab

The Incident Form provides details for troubleshooting purposes.

General Attributes

Attribute	Description
Name	Name of the rule used to configure the incident. This name is initially created by SOM.
Category	Generated by SOM to indicate the problem category. Possible values include: **Accounting** - Used to indicate problems related to usage statistics and allocation of costs associated with the billing of time and services provided by devices. This category is not used by SOM with default configurations, but it is available for incidents you define. **Application Status** - Indicates there is a problem with the health of the SOM software. Examples of these kinds of events include license expiration or that a certain SOM process lost connection to the Process Status Manager. **Configuration** - Indicates there is a problem with the configuration of a managed device. For example, there is a physical address mismatch. **Fault** – Indicates a problem with the network; for example, Node Down. **Performance** – Indicates an exceeded threshold. For example, a utility exceeds 90 percent.

General Attributes, continued

Attribute	Description
	Security – Indicates there is a problem related to authentication; for example, an SNMP authentication failure. **Status** - Often indicates some status change occurred on a device. For example, when a Cisco device powers up or powers down. **Note:** The icons are only in table views.
Family	Used to further categorize the types of incidents that might be generated. Possible values are: **Node** – Indicates the incident is related to a node problem. **Correlation** – Indicates the incident has additional incidents correlated beneath it. These incidents are associated with a duplicate count so that you can determine the number of correlated incidents associated with it. **Note:** The icons are only in table views. Use **Node** or **Correlation** to create a filter on Family column in an incident browsing view. Other options are not supported.
Origin	Identifies how the incident was generated. Possible values are: **SOM** – Indicates the incident was generated by SOM processes. **Manually Created** – SOM does not use this Origin with default configurations. It is available for incidents you define. **SNMP Trap** – Indicates the incident was forwarded from an SNMP Agent. **Other** – Indicates the incident was generated by a source other than the Origin categories provided. **Note:** The icons are only in table views.
Correlation Nature	This incident's contribution to a root-cause calculation, if any. Possible values are: **Root Cause** – Indicates the Incident is determined by SOM. **User Root Cause** – Indicates the Incident is configured by your SOM administrator to make SOM always treat this Incident as Correlation Nature: Root Cause. **Secondary Root Cause** – Indicates the Incident is related to Root Cause but is not the

General Attributes, continued

Attribute	Description
	primary problem. Secondary Root Cause Incidents are the Child Incidents of Parent Incidents and often begin as primary Root Cause incidents.Whenever a primary Root Cause Incident is correlated under another Incident, its Correlation Nature becomes Secondary Root Cause. For example, if an **Interface Down** incident is followed by a **Node Down** Incident on a neighboring device, the **Interface Down** Incident becomes a Child Incident to the Parent **Node Down** incident. Its Correlation Nature becomes Secondary Root Cause. Use the **All Incidents** view to examine both Secondary Root Cause and primary Root Cause Incidents. Use the **Root Cause** view to see only the primary Root Cause Incidents. In the Root Cause Incidents view, any Secondary Root Cause Incident is correlated under its associated primary Root Cause Incident. **Symptom** – Indicates any incidents that were generated from a trap notification related to the root cause incident. For example, a Link Down incident generated from a Link Down trap notification might appear as a **Symptom** to an Interface Down incident in the root cause incidents view. **Service Impact** - Indicates a relationship between incidents in which a network service is affected by other incidents. For example, an Interface Down incident can affect a Router Redundancy Group that is part of an HSRP service. **Stream Correlation** – Stream correlations are created as SOM analyzes events and traps to determine the root cause incident for a problem. Examples of stream correlations include Dedup (duplication of events) and Rate (occurrence of events by time). **None** – Indicates there is no incident correlation for the incident. **Info** – Indicates the incident is informational only. **Dedup Stream Correlation** – Stream correlations are created as SOM analyzes events and traps to determine the root cause incident for a problem. Dedup Stream Correlation indicates the Incident is a Deduplication Incident. Deduplication Incident configurations determine what values SOM should match to detect when an Incident is a duplicate. Duplicate Incidents are listed under a Duplicate Correlation Incident. SOM tracks the number of duplicates generated. This value is

General Attributes, continued

Attribute	Description
	captured as the Duplicate Count attribute and is incremented on the Duplicate Correlation Incident. **Rate Stream Correlation** – Stream correlations are created as SOM analyzes events and traps to determine the root cause incident for a problem. Rate Stream Correlation indicates the Incident is a Rate Incident. Rate Incidents track incident patterns *based on the number of incident reoccurrences within a specified time period*. After the count within the specified time period is reached, SOM emits a Rate Correlation Incident and continues to update the Correlation Notes with the number of occurrences within that rate. **Note:** The icons are only in table views.
Duplicate Count	Lists the number of duplicate incidents that SOM encountered for the selected incident. This number increments in the associated deduplication incident that SOM generates to inform the operator of incidents needing attention. The incidents are reoccurring according to the deduplication criteria specified in the incident's deduplication configuration. For example, by default, incidents generated from SNMP traps will not have their deduplication count incremented. If the SOM administrator defines a deduplication criteria for the SNMP trap, SOM generates an incident specifying that the SNMP trap is reoccurring according to the criteria specified in the incident's associated deduplication configuration. This incident is the one that increments and displays the **Duplicate Count** value. Note the following: • By default, SOM updates the **Duplicate Count** every 30 seconds. This interval cannot be changed. • SOM continues to update the duplicate count regardless of an incident's Lifecycle State. For example, if an incident's **Lifecycle State** is set to **Closed**, the duplicate count continues to be incremented. This behavior helps you identify situations in which the incident is not yet fixed. Take note if the Duplicate Count is incremented after a lengthy time period has elapsed; this might indicate there is a new problem with the node, interface, or address. • Duplicates are configured by the SOM administrator using the **SNMP Trap Configuration** or **Management Event Configuration** form available from the **Configuration** workspace.
RCA	Used by SOM to identify whether SOM considers the incident to be active or inactive. If set

General Attributes, continued

Attribute	Description
Active	to **True**, the incident is considered to be active. If set to **False**, the incident is considered to be inactive. SOM considers an incident to be active when the root cause analysis (RCA) engine is actively evaluating the problem reported by this incident. SOM considers an incident to be inactive when SOM confirmed that the problem reported by this incident is no longer a problem. For example, the device is now functioning properly. SOM initially sets an incident's **RCA Active** attribute to True and the incident's Lifecycle State to **Registered**. When SOM sets the **RCA Active** attribute to **False**, it also sets the incident's Lifecycle State to **Closed**. Examples of when an incident's **RCA Active** attribute is set to **False** include: • When an interface goes up, SOM closes the InterfaceDown incident. • When a node goes up, SOM closes the NodeDown incident.
Correlation Notes	Stores notes about the correlation status of the incident. SOM provides the following information in the **Correlation Notes** field when it sets an incident's **Lifecycle State** to **Closed**: • The Conclusion information identifying the reason SOM changed the incident's Lifecycle State to Closed. For example, SOM might include an Interface Up Conclusion as the reason an Interface Down incident was closed. • The time measured between when SOM detected a problem with one or more storage devices to the time the problem was resolved. • The time when SOM first detected the problem associated with the incident. • The time when SOM determines the problem associated with the incident is resolved. SOM inserts the information in front of any existing information provided. **Note:** SOM provides Correlation Notes information only when it has analyzed and Closed the incident. Software that is integrated with SOM might also provide information identifying the reason an incident was closed. Any time an incident is closed manually (for example, by the network operator), SOM does not provide Correlation Notes information.
First Occurrence Time	Used when suppressing duplicate incidents or when specifying an incident rate. Indicates the time when the duplicate or rate criteria were first met for a set of duplicate incidents or for a set of incidents that has a rate criteria that was met.

General Attributes, continued

Attribute	Description
Last Occurrence Time	Used when suppressing duplicate incidents or specifying an incident rate. Indicates the time when the duplicate or rate criteria were last met for a set of duplicate incidents or for a set of incidents that has a rate criteria that was met. If there are no duplicate incidents or incidents that have a rate criteria that were met, this date is the same as the First Occurrence Time.
Origin Occurrence Time	The time at which an event occurred that caused the incident to be created; for example, the time held in the trap.

Incident Form: Correlated Parents Tab

Correlated Parents Table

Attribute	Description
Correlated Parents	If the current incident is a child incident, any correlated parent incidents of the child appears in this table view. For example, parent incidents are created when a root cause problem is detected. A Node Down root cause incident is a parent of an Interface Down incident. Therefore, on an Interface Down Incident form, a Node Down incident might appear under the Correlated Parents tab. Double-click the row representing an incident. The Incident Form displays all details about the selected incident.

Incident Form: Correlated Children Tab

Correlated Children Table

Attribute	Description
Correlated Children	If the current incident is a parent incident, any correlated child incident of the parent appears in this table view. For example, an Interface Down incident would be correlated as a child under a Node Down root cause incident. Therefore, on a Node Down incident form, an Interface Down incident would appear on the Correlated Children tab. Double-click the row representing an incident. The Incident Form displays all details about the selected incident.

Incident Form: Custom Attributes Tab

Note: SOM lists the Custom Attributes for incidents in the order in which they are received from the SNMP trap. If you sort or filter the Custom Attribute table, click the Restore Default Settings icon to restore the Custom Attribute order for the selected incident.

Custom Attributes Table

Attribute	Description
Custom Incident Attributes	Used by SOM to add additional information to the incident that SOM makes available for viewing. Each custom incident attribute (CIA) includes a name, type, and value group that can be populated differently for different types of incidents. Varbind values that accompany SNMP traps are a common use for this attribute. Double-click the row representing the Custom Incident Attribute that has the "Custom Incident Attribute Form" on page 526 you want to see.

Incident Form: Registration Tab

Registration Attributes

Attribute	Description
Created	Date and time the selected object instance was created. SOM uses the locale of the client and the date and time from the SOM management server. **Note:** This value does not change when a node is rediscovered. This is because the Node object is modified, but not created.
Last Modified	Date the selected object instance was last modified. SOM uses the locale of the client and the date and time from the SOM management server. Note the following: • When a node is rediscovered, the Last Modified time is the same as the Discovery Completed time. This is because the node’s Discovery State changes from Started to Completed. • When a Node is initially discovered, the Last Modified time is slightly later than the Created time. This is because node discovery does not complete until after the Node is created.

Object Identifiers Attributes

Attribute	Description
ID	The Unique Object Identifier, which is unique within the SOM database.
UUID	The Universally Unique Object Identifier, which is unique across all databases.

Include Custom Incident Attributes in Your Message Format (SNMP Trap Incident)

You can use CIAs in your message format to extend the amount of information presented.

To determine which CIAs are available for any particular incident type, follow these steps:

1. Open an Incident view.
2. Double-click a row to open the Incident form.
3. Navigate to the **Custom Attributes** tab. A complete list of available CIAs (for that incident type) appears in the table.

To include a CIA in your message format, type the dollar-sign character (`$`) plus any of the following:

- Varbind position number or asterisk (*) to include all varbind values
- Name of the CIA
- Object identifier (oid) of the CIA (useful when the varbind position number is not consistent among vendors)

Note: A single incident cannot include two CIAs with the same name. However, two incidents can contain CIAs having the same names and values.

The following table presents some example formats with the subsequent output.

Example Incident Message Formats

Example Message Format	Output in Incident View
Possible trouble with $3	`Possible trouble with` <varbind 3>
Possible trouble with $11	`Possible trouble with` <varbind 11>
Possible trouble with $77 (where the varbind position 77 does not exist)	`Possible trouble with <Invalid or unknown cia> 77`
Possible trouble with $*	`Possible trouble with` <cia1_name: cia_value>, <cia2_ name; cia_value>,< cia*n*_name: cia_value>
Possible trouble with $3x	`Possible trouble with` <varbind 3>x

Example Incident Message Formats, continued

Example Message Format	Output in Incident View
Possible trouble with $1.2.3.4.5	`Possible trouble with` <value of the CIA with oid of 1.2.3.4.5>
Possible trouble with $cia.sourceObject.UcmdbId	`Possible trouble with` <value of the CIA with name of cia.sourceObject.UcmdbId>

Tip: SOM provides an error message when a CIA cannot be found. For example, if you enter an unavailable varbind position, name, or object identifier (oid), SOM returns an "Invalid or unknown cia" error message.

Matching Criteria Configuration Form (Identify Incident Pairs)

To configure which attributes SOM uses to verify incident identity:

1. Complete the steps in "Pairwise Incidents Prerequisites" on page 539 so your choices for this Item Pair configuration are displayed in the SOM console. (Two Incident forms should be open before you proceed to step 2.)
2. Navigate to the **Matching Criteria Configuration** form.
   a. From the workspace navigation panel, select the **Configuration** workspace.
   b. Expand the **Incidents** folder.
   c. Select **Pairwise Configurations**.
   d. Do one of the following:
      - To create a new pairwise configuration, click the New icon.
      - To edit a pairwise configuration, double-click the row representing the configuration you want to edit.
   e. Navigate to the **Matching Criteria** tab.
   f. Do one of the following:
      i. To create a new matching criteria configuration, click the New icon.
      ii. To edit a matching criteria configuration, double-click the row representing the configuration you want to edit.
3. Specify the Object Identifier (OID) or trap varbind position number you want SOM to use to confirm the identity of the pair of incidents (see table).
4. Click **Save and Close** to save your changes and return to the previous form.

5. Repeat steps 1-3 any number of times. The incidents must pass all Matching Criteria, plus have identical Source Node and Source Object attribute values.

Matching Criteria Configuration

Attribute	Description
First Incident Criterion	Type the specification required to confirm the identify of the first incident in this logical pair of incidents. Provide one of the following: • The SNMP trap varbind Abstract Syntax Notation value - ASN.1 (OID) • The SNMP trap varbind position number • The Custom Attribute **Name** value or the Name column in the table on the "Incident Form: Custom Attributes Tab" on page 535 of the Incident you are configuring as a member of this logical pair).
Second Incident Criterion	Type the specification required to confirm the identify of the second incident in this logical pair of incidents. Provide one of the following: • The SNMP trap varbind Abstract Syntax Notation value - ASN.1 (OID) • The SNMP trap varbind position number • The Custom Attribute **Name** value or the Name column in the table on the "Incident Form: Custom Attributes Tab" on page 535 of the Incident you are configuring as a member of this logical pair).

Own One or More Incidents

SOM lets you own incidents. When you specify that you want to own an incident, the incident is assigned to you.

To own one or more incidents:

1. Navigate to the incident view of interest.
   a. From the workspace navigation panel, select the **Incident Browsing** workspace.
   b. Select the incident view of interest; for example **Open Incidents**.
2. Press Ctrl-Click to select each row that represents an incident you want to own.
3. Select **Actions** → **Assign** → **Own Incident**.

   > **Tip:** You can also right-click any object in a table or map view to access the items available within the **Actions** menu.

Your user name appears in the **Assigned To** column in any incident views that include the incident.

Pairwise Incidents Prerequisites

When matching incidents, SOM automatically takes into account the following values:

- **SNMP Trap incidents**. SOM takes into account from which device the trap originated using the `cia.address` value of the source address of the trap.
- **Management Event incidents**. SOM takes into account the name of the incident's Source Object and Source Node.

> **Tip:** SOM displays the Name value used to identify the Source Node and Source Object in the **Source Node** and **Source Object** attribute for each incident in the Incident form.

If you must provide more details to accurately identify the logical pair of incidents (from among all possible incidents related to that source node), complete the Optional step 6 below.

Complete the following steps before attempting to set up a Pairwise Configuration:

1. Identify the incidents or SNMP traps that consist of the logical relationship that makes the pair.

   > **Note:** The incident configurations you select can be the same or different for each pair.

2. Configure those two incidents or traps within SOM, if they are not already configured.
3. Generate one of each of the incidents or SNMP traps so you can see an example of each in one of SOM Incident views.
4. To display the Incident form, double-click the row representing the first sample incident for the pair.
5. Repeat the previous step with the second sample incident for the pair.

Payload Filter Details

Payload Filter Editor enables you to create expressions that refine the filters. Make sure to design any complex Payload Filters offline as a Boolean expression first. This method can help to minimize errors when entering your expressions using the Payload Filter editor.

When creating a Payload Filter, note the following:

- Payload Filter expressions for the `like` and `not like` operators use the syntax defined for java regular expressions (java.util.regex Pattern class).
- You must use a `ciaName` that already exists in the trap or event you are configuring.
- Each set of expressions associated with a Boolean Operator is treated as if it were enclosed in parentheses and evaluated together.
- View the expression displayed under **Filter String** to see the logic of the expression as it is created.
- The `AND` and `OR` Boolean Operators must contain at least two expressions as shown in the example below.

The following example filters incidents on voltage state:

```
AND
    ciaName = .1.3.6.1.4.1.9.9.13.1.2.1.7
    ciaValue = 5
```

SOM evaluates the expression above as follows:

`(ciaName = .1.3.6.1.4.1.9.9.13.1.2.1.7 AND ciaValue = 5)`

SOM finds all incidents with a varbind `.1.3.6.1.4.1.9.9.13.1.2.1.7` value of **5**.

> **Note:** When you use `ciaName` and `ciaValue` in a Payload Filter, you must enter the `ciaName` and `ciaValue` as a pair as shown in the preceding example.

- The placement of your cursor and the subsequent text that is selected is important when performing operations using the Payload Filter Editor. For example, you append to, replace, or change the indentation of the expression that is selected.
- The placement of your cursor and the subsequent text that is selected is especially important when adding your Boolean operators.
- You can include more than one varbind in the same Payload Filter expression as shown in the following example:

  ```
  ((ciaName like \Q.1.3.6.1.4.1.9.9\E.* AND ciaValue = 25) AND (ciaName like
  \Q.1.3.6.1.2.1.2.2.1.1.3\E.* AND ciaValue = 3))
  ```

  In this example, a given trap must meet each of the following criteria:

  - Contain a varbind whose Object Identifier (OID) matches the regular expression `\Q.1.3.6.1.4.1.9.9\E.*` and has a value of `25`.
  - Contain a varbind whose OID matches the regular expression `\Q.1.3.6.1.2.1.2.2.1.1.3\E.*` and has a value of `3`.

Payload Filter Editor Settings

Attribute	Description
Attribute	The attribute name on which SOM searches. Filterable attributes include the following: • ciaName • ciaValue **Note:** When you use `ciaName` and `ciaValue` in a Payload Filter, you must enter the `ciaName` and `ciaValue` as a pair. For example: `(ciaName =.1.3.6.1.4.1.9.9.13.1.2.1.7 ) AND ( (ciaValue = 4) OR ( ciaValue = 5))` is not supported.

Payload Filter Editor Settings, continued

Attribute	Description
Operator	Valid operators are described below. • **=** Finds all values equal to the value specified. Click here for an example. Example: `ciaName=.1.3.6.1.4.1.9.9.13.1.2.1.7` matches any incident that contains a varbind with the name value of **.1.3.6.1.4.1.9.9.13.1.2.1.7**. • **!=** Finds all values not equal to the value specified. Click here for an example. Example: `ciaName! = .1.3.6.1.4.1.9.9.13.1.2.1.7` matches any incident that contains a varbind with a name value other than **1.3.6.1.4.1.9.9.13.1.2.1.7**. • **<** Finds all values less than the value specified. Click here for an example. Example: `ciaValue < 6` matches any incident that contains a varbind with a value less than **6**. • **<=** Finds all values less than or equal to the value specified. Click here for an example. Example: `ciaValue <= 6` matches any incident that contains a varbind with a value less than or equal to **6**. • **>** Finds all values greater than the value specified. Click here for an example. Example: `ciaValue > 4` matches any incident that contains a varbind with a value greater than **4**. • **>=** Finds all values greater than or equal to the value specified. Click here for an example. Example: `ciaValue >= 4` matches any incident that contains a varbind with values greater than or equal to **4**. • **between** Finds all values equal to and between the two values specified. Click here for an example. Example: `ciaValue between` 1 and 4 matches any incident that contains a varbind value equal to or greater than **1** and equal to or less than **4**. **Note:** Each value must be entered on a separate line. • **in** Finds any match to at least one value in a list of values. Click here for an example.

Payload Filter Editor Settings, continued

Attribute	Description
	Example: `ciaValue in` 4 and 5 matches any incident with a varbind value of either **4** or **5**. **Note:** Each value must be entered on a separate line. SOM displays the list of attributes using comma-separated values enclosed in parentheses, for example (**4, 5**). However, the comma-separated list is used only for display purposes. The actual delimiter is the new line. • **is not null** Finds all non-blank values. Click here for an example. Example: `ciaValue is not null` matches any incident with a varbind that contains a value. • **is null** Finds all blank values. Click here for an example. Example: `ciaValue is null` matches any incident with a varbind that does not contain a value. • **like** Finds matches using the syntax defined for java regular expressions. Click here for more information. The period asterisk (.*) characters mean *any number of characters of any type at this location*. The period (.) character means *any single character of any type at this location*. **Note:** To include literal string values in the Value attribute, enclose the string value in \Q<*literal_value*>\E as shown in the Examples listed below. Example: `ciaName like \Q.1.3.6.1.4.1.9.9\E.*` matches any incident that contains a varbind name value that begins with 1.3.6.1.4.1.9.9 and (optionally) ends with any number of characters. `ciaValue like .*Chicago.*` finds all traps or events that contain a varbind value that includes the string **Chicago**. • **not between** Finds all values except those between the two values specified. Click here for an example.

Payload Filter Editor Settings, continued

Attribute	Description
	Example: `ciaValue not between 5 8` matches an incident that contains a varbind with the values less than **5** or greater than **8** . • **not in** Finds all values except those included in the list of values. Click here for an example. Example: `ciaValue not in` 1 or 2 matches any incident that contains a varbind with values other than **1** and **2**. **Note:** Each value must be entered on a separate line. SOM displays the list of attributes using comma-separated values enclosed in parentheses, for example, (**1, 2**). However, the comma-separated list is used only for display purposes. The actual delimiter is the new line. • **not like** Finds all that do not have the values specified using the syntax defined for Java regular expressions. See the Pattern (Java Platform SE6) API documentation at: Click here for an example. The period asterisk (.*) characters mean *any number of characters of any type at this location*. The period (.) character means *any single character of any type at this location*. **Note:** To include literal string values in the Value attribute, enclose the string value in \Q<*literal_value*>\E as shown in the Examples listed below. Example: `ciaName not like \Q.1.3.6.1.4.1.9.9\E.*` matches any incident that contains a varbind name value that does not begin with **.1.3.6.1.4.1.9.9** and (optionally) ends with any number of characters. `ciaValue not like .*Chicago.*` finds all traps or events that do not contain a varbind value that includes the string **Chicago**.
Value	The value for which you want SOM to search. Note the following: • The values you enter are case sensitive.

Payload Filter Editor Settings, continued

Attribute	Description
	• SOM displays a variable number of value fields depending on the Operator selected. For example, the `between` Operator causes two value fields to be displayed. • The `between, in` and `not in` operators require that each value be entered on a separate line.

Payload Filter Editor Buttons

Button	Description
Append	Appends the current expression (Attribute, Operator,and Value) to the selected expression already included in the filter string.
Insert	Inserts the current expression (Attribute, Operator,and Value) in front of the cursor location within the Filter String.
Replace	Replaces the selected expression with the expression displayed in the Attribute, Operator, and Value fields.
AND	Inserts the AND Boolean Operator in the selected cursor location. **Note:** View the expression displayed under **Filter String** to see the logic of the expression as it is created.
OR	Inserts the OR Boolean Operator in the current cursor location. **Note:** View the expression displayed under **Filter String** to see the logic of the expression as it is created.
NOT	Can be used in any part of the Filter String to specify that SOM should exclude interfaces with values that pass the expression that immediately follows the `NOT`. For example, when evaluating the following expression, SOM includes interfaces with (interface description) `ifDesc` containing **VLAN**, and excludes any Interfaces that have **VLAN10** for the (interface name) `ifName` value: `(ifDesc like VLAN AND NOT (ifName=VLAN10))` **Note:** View the expression displayed under Filter String to see the logic of the

Payload Filter Editor Buttons, continued

Button	Description
	expression as it is created.
EXISTS	Used for filters that include Capabilities or Custom Attribute names and values in the Filer String. Indicates that you want SOM to consider interfaces that have Capabilities or Custom Attributes when evaluating the Filter String. **Tip:** When creating complex Filter Strings that include `customAttrName` and `customAttrValue` pairs as one component of an "***or***" statement, to prevent SOM from excluding Nodes that have zero Custom Attributes, use **EXISTS** or **NOT EXISTS** criteria for the `customAttrName` and `customAttrValue` pair definitions. Otherwise Nodes that do not have any Custom Attributes are automatically excluded even if they have values that pass other aspects of your filter. For example, when evaluating the following Filter String, SOM includes interfaces with (interface description) `ifDesc` containing **VLAN**, as well as any Interfaces Custom Attribute Role value is **LAN Connection to Oracle Server**: `(ifDesc like VLAN OR EXISTS((customAttrName=Role AND customAttrValue=LAN Connection to Oracle Server)))` **Note:** View the expression displayed under Filter String to see the logic of the expression as it is created.
NOT EXISTS	Used for filters that include Capabilities or Custom Attribute names and values in the Filer String. Indicates that you want SOM to consider interfaces that do not have any Capabilities or Custom Attributes when evaluating the Filter String, but exclude the interfaces that match the expression that follows the **NOT EXISTS**. **Tip:** When creating complex Filter Strings that include `customAttrName` and `customAttrValue` pairs as one component of an "***or***" statement, to prevent SOM from excluding Nodes that have zero Custom Attributes, use **EXISTS** or **NOT EXISTS** criteria for the `customAttrName` and `customAttrValue` pair definitions. Otherwise Nodes that do not have any Custom Attributes are automatically excluded even if they have values that pass other aspects of your filter.

Payload Filter Editor Buttons, continued

Button	Description
	For example, when evaluating the following expression, SOM includes interfaces with (interface description) `ifDesc` containing **VLAN**, and excludes any Interfaces that have the Custom Attribute Role and that Role value is **LAN Connection to Oracle Server**: `(ifDesc like VLAN OR NOT EXISTS((customAttrName=Role AND customAttrValue=LAN Connection to Oracle Server)))` **Note:** View the expression displayed under Filter String to see the logic of the expression as it is created.
Delete	Deletes the selected expression. **Note:** If the Boolean Operator is selected, the Payload Filter Editor deletes all expressions associated with the Boolean Operator.

Unassign Incidents

If you are an SOM user with a user role of Level 1 Operator (with more limited access privileges than Level 2 Operators), Level2 Operator, or Administrator, you can unassign an incident for yourself or for another user.

To unassign one Incident:

1. Navigate to the incident form of interest.
   a. From the workspace navigation panel, select the **Incident Browsing** workspace.
   b. Select any incident view.
   c. Select the row representing the incident you want to unassign.
2. Select **Actions** → **Assign** → **Unassign Incident**.

   > **Tip:** You can also right-click any object in a table or map view to access the items available within the **Actions** menu.

3. Click **Save** to save your changes or **Save and Close** to save your changes and exit the form..

The **Assigned To** column is empty in any incident views that include that incident.

To unassign multiple Incidents:

1. Navigate to the incident view of interest.

   a. From the workspace navigation panel, select the **Incident Browsing** workspace.
   b. Select any incident view.
2. Press Ctrl-Click to select each row that represents an incident you want to unassign.
3. Select **Actions** → **Assign** → **Unassign Incident**.

The **Assigned To** column is empty in any incident views that include that incident.

Send Documentation Feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on User Guide (Storage Operations Manager 10.20)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to storage-management-doc-feedback@hpe.com.

We appreciate your feedback!