HP Storage Operations Manager

Software Version: 10.00
Windows® and Linux® operating systems

User Guide

Document Release Date: March 2015
Software Release Date: March 2015

Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notice

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

AMD is a trademark of Advanced Micro Devices, 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.

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.hp.com

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

https://hpp12.passport.hp.com/hppcf/createuser.do

Or click the **the Register** link at the top of the HP Software Support page.

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

Support

Visit the HP Software Support Online web site at: **https://softwaresupport.hp.com**

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

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

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

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

https://hpp12.passport.hp.com/hppcf/createuser.do

To find more information about access levels, go to:

https://softwaresupport.hp.com/web/softwaresupport/access-levels

HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is **http://h20230.www2.hp.com/sc/solutions/index.jsp**

Contents

Contents 4

Chapter 1: Getting Started with SOM 14

Configuring Web Browsers for SOM 14

Configure Mozilla Firefox for SOM 14

Configure Mozilla Firefox Timeout Interval 15

Configure Microsoft Internet Explorer for SOM 15

Configure the Microsoft Internet Explorer Title Bar 16

Configuring the SOM Interface 16

Using the SOM Console 17

Navigating the SOM Console 22

Display Views 22

Access More Information About an Object (Forms and Analysis Pane) 24

Invoke Actions 26

Use the Tools Menu 27

Search the Help Topics 28

Mark Your Favorite Help Topics 30

About Workspaces 30

About the Analysis Pane 33

Working with Objects 35

Access a Subset of the Available Information About a Related Object 36

Access All Information About a Related Object 37

Modify Object Attribute Values 38

Displaying Information About SOM 40

Displaying SOM Version and License Information 40

System Information: Product Tab 41

System Information: Health Tab 42

System Information: Server Tab 42

System Information: Extensions Tab 43

Chapter 2: Configuring SOM for your Storage Environment 44

Configuring Security 45

Summary of Security Tasks 45

Choose a Mode for User Authentication 47

SOM Configuration Settings 48
Lightweight Directory Access Protocol (LDAP) 49
LDAP External Mode 49
Different Ways to Configure Security 50
Configure a User Account 51
Create a User Account 52
Modify a User Account 53
Delete a User Account 55
Configure User Groups 56
Predefined User Groups 57
Determine the User Group 57
Create a User Group 58
Modify a User Group 60
Delete a User Group 60
User Account Mapping Tasks 60
Create a User Account Mapping 61
Delete a User Account Mapping 62
Configure Security Groups 63
Recommendations for Planning Security Groups 63
Create a Security Group 64
Modify a Security Group 65
Delete a Security Group 65
Methods for Assigning Nodes to Security Groups 66
Configure Security Group Mappings 66
Map User Groups to Security Groups 67
Default Object Access Privileges 70
Configure Security Using the Security Wizard 70
Control Menu Access 73
Check Security Configuration 76
Communicate Console Access Information to Your Team 76
Opening the SOM Console 77
Configuring Sign-In to the SOM Console 79
Signing Out from the SOM Console 79
Troubleshoot Access 79
Check Security Configuration 80
View Summary of Changes in the Security Wizard 80
View the Users who are Signed In to SOM 81
Restore the Administrator Role 81
Restore SOM Access to the System User 81

Configuring Node Groups 82
Node Groups Provided by SOM 83
Recommendations for Planning Node Groups 83
Create a Node Group 84
Using Additional Filters for Node Group Definitions 88
Guidelines for Creating Additional Filters for Node Groups 89
Add Boolean Operators in the Additional Filters Editor 92
Create an Additional Filters Expression 95
Modify a Node Group 103
Delete a Node Group 104
Discovering Devices 104
Recommendations for Planning Discovery 105
Prerequisites for Discovering a Device 106
Prerequisites for Agentless Discovery of Linux, Solaris, and AIX Hosts 106
Commands for a Linux Host as a Root User 108
Commands for a Solaris Host as a Root User 109
Commands for an AIX Host as a Root User 111
Commands for a Linux Host as a Non-Root User 112
Commands for a Solaris Host as a Non-Root User 114
Commands for an AIX Host as a Non-Root User 115
Agentless Discovery of Windows Hosts 117
Commands as an Administrator 118
Prerequisites to Discover Hosts with CIME Agent 119
Prerequisites to Discover Host Clusters 120
Prerequisites to Discover VMware ESX Servers and Virtual Machines 121
Prerequisites to Discover Brocade Switches 122
Prerequisites to Discover Cisco Switches 123
Prerequisites to Discover HP XP/P9500 Arrays 126
Prerequisites to Discover HP 3PAR Arrays 126
Prerequisites to Discover HP StorageWorks EVA Arrays 127
Prerequisites to Discover HDS and HUS Arrays 127
Prerequisites to Discover an EMC Isilon Cluster 128
Prerequisites to Discover EMC VNX Filer 129
Prerequisites to Discover EMC Symmetrix Arrays 129
Prerequisites to discover EMC CLARiiON and VNX Block Storage Systems 130
EMC VPLEX Clusters 130
NetApp 7-Mode Device 131
Discovery Tasks 131

Configure Addresses for Discovery ... 131
Delete an Address ... 134
Configure Address Ranges for Discovery ... 134
Considerations for Defining an Address Range ... 134
Configure a Range for Discovery ... 135
Scan an Address Range ... 137
Modify an Address Range ... 137
Delete an Address Range ... 137
Configure Credentials for Discovery ... 138
Modify a Discovery Credential ... 139
Delete a Discovery Credential ... 139
Configure Tenants ... 140
Tenant and Initial Discovery Security Group Assignments ... 141
Recommendations for Planning Tenants ... 141
Create a Tenant ... 142
Change Tenant Assignment for a Node ... 143
Start Discovery ... 144
Status of Discovery ... 145
Discovery Views ... 146
Inferring Hosts Based on Rules ... 146
Regular Expressions in Rules ... 147
Create a Rule ... 155
Modify a Rule ... 158
Delete a Rule ... 158
Run a Rule Manually ... 159
View Inferred Hosts ... 159
Delete an Inferred Host ... 159
Delete Hosts Inferred by a Rule ... 160
Reconciliation of Hosts ... 160
Configuring Data Collection Settings ... 160
Recommendations for Configuring Data Collection ... 161
Create a Data Collection Policy ... 162
Modify a Data Collection Policy ... 165
Delete a Data Collection Policy ... 165
Create a Blackout Period ... 166
Modify a Blackout Period ... 167
Delete a Blackout Period ... 167
Data Collection Control ... 168

Change the Data Collection Control for a Device Profile ... 169
Planning Licenses ... 171
License Types ... 171
Temporary Instant-On License ... 172
Obtain and Install New License ... 172
Install a Perpetual License ... 172
From the Command Line ... 172
Using Autopass to Install a Perpetual License ... 172
Extend a Licensed Capacity ... 173
View License Information ... 173
Viewing Consumed MAP Count for Each Element ... 173
MAP Count Calculation ... 174
Configure Performance Pack ... 177
Monitoring Performance ... 178
Recommendations for Monitoring Policies ... 179
Prerequisites for a Monitoring Policy ... 179
Create a Monitoring Group ... 180
Modify a Monitoring Group ... 181
Create a Monitoring Policy ... 181
View Collectors ... 182
Viewing Performance Data ... 183
Modify a Monitoring Policy ... 183
Delete a Monitoring Policy ... 184
Managing Storage Tiers ... 184
How Do Rule-Based Assignments Work? ... 185
Best Practices for Creating Storage Tiers ... 186
Create a Storage Tier ... 186
Modify a Storage Tier ... 188
Delete a Storage Tier ... 189
Chapter 3: Managing your Storage Environment with SOM ... 191
Dashboards ... 191
Environment Capacity Dashboard ... 192
Asset Dashboard ... 194
Collection Status Dashboard ... 194
Inventory Views ... 196
Using the Analysis Pane ... 197
Hosts Views ... 197

Switches View ... 198
Storage Systems Views ... 199
Fabrics View ... 200
Nodes View ... 201
Node Groups View ... 202
FC HBA View ... 204
HBA Ports View ... 204
Switch Ports View ... 204
Storage System Ports View ... 205
Viewing Device Capacity ... 205
Host Capacity ... 205
Switch Capacity ... 207
Storage System Capacity ... 207
Device-Specific Exceptions ... 207
Hosts ... 208
Switches ... 211
Storage Systems ... 212
Viewing Device Performance ... 220
Performance Collectors for Hosts ... 221
Physical Disk Collectors ... 221
ESX Server Performance Collectors ... 222
Performance Collectors for Switches ... 223
Best Practices ... 224
Switch Performance Issues ... 225
FC Errors ... 225
CRC Errors ... 225
Link Failure ... 225
I/O Traffic ... 225
Performance Collectors for HP 3PAR Arrays ... 226
3PAR SMI-S Storage System Collector ... 226
3PAR SMI-S Controller Collector ... 228
3PAR SMI-S Volume Collector ... 230
3PAR SMI-S Physical Disk Collector ... 232
3PAR SMI-S Fiber Channel Port Collector ... 233
Performance Collectors for HP StorageWorks EVA Arrays ... 234
EVA SMI-S Storage System Collector ... 234
EVA SMI-S Controller Collector ... 237
EVA SMI-S Volume Collector ... 239

EVA SMI-S Physical Disk Collector 241
EVA SMI-S Fiber Channel Port Collector 243
Performance Collectors for EMC Symmetrix DMX/VMAX Arrays 246
EMC Symmetrix DMX SMI-S Storage System Collector 246
EMC Symmetrix DMX SMI-S Controller Collector 249
EMC Symmetrix DMX SMI-S Volume Collector 251
EMC Symmetrix DMX SMI-S Fibre Channel Port Collector 256
Performance Collectors for CLARiiON and VNX Arrays 257
EMC CLARiiON and VNX SMI-S Storage System Collector 258
EMC CLARiiON and VNX SMI-S FrontEnd Controller Collector 259
EMC CLARiiON and VNX SMI-S Volume Collector 261
CLARiiON and VNX SMI-S Physical Disk Collector 263
EMC CLARiiON and VNX SMI-S FrontEnd Port Collector 265
Topology Maps 265
Port Connector Form 268
Storage System Topology 269
Host Topology 269
Switch Topology 270
Fabric Topology 270
Chapter 4: Common Tasks 272
Start or Stop SOM Services 272
Delete Elements 273
Quarantine/Un-quarantine Elements 274
Launch Topology 275
Create an Asset Record 276
Appendix A: Inventory Views Tabs and Forms 278
Block Storage Systems View 278
Capacity Information of Block Storage Systems 279
File Storage Systems View 284
Capacity of File Storage Systems 285
Cluster Storage Systems View 286
Forms 287
Host Forms 287
Filesystem Form 287
HBA Card Form 287

HBA Port Form ... 288
Host Disk Drive Form ... 288
Multipath Disk Form ... 288
Volume Manager Volume Form ... 289
Disk Partition Form ... 289
Switch Forms ... 289
Switch Form ... 289
Fibre Channel Port Types ... 290
Switch Ports View ... 291
Storage System Forms ... 291
Storage System Processor Form ... 292
Storage Pool Form ... 292
Pool Capabilities Form ... 292
Storage Volume Form ... 293
Storage Extent Form ... 294
SCSI Card Form ... 294
Storage Disk Drive Form ... 295
File Systems Form ... 296
NAS Extent Form ... 296
Fabric Forms ... 297
Zone Alias Form ... 297
Zone Set Form ... 298
Zone Form ... 298
Node Forms ... 299
Node Device Filter Form ... 299
Node Group Forms ... 300
Device Category Form ... 300
Device Vendor Form ... 301
Device Family Form ... 302
Device Profile Form ... 303
Author Form ... 304
Additional Node Form ... 305
Node Group Hierarchy Form ... 306
Tabs ... 306
Fibre Channel Port Types ... 307
Asset Record Tab ... 308
Host Tabs ... 309
Hosts View: Virtual Machines Tab ... 309

Hosts View: File Systems Tab ... 310
Hosts View: Cards Tab ... 311
Hosts View: Ports Tab ... 311
Hosts View: Target Mappings Tab ... 312
Hosts View: Multipathing Tab ... 312
Hosts View: Volume Management Tab ... 313
Hosts View: Disk Partitions Tab ... 313
Hosts View: Disk Drives Tab ... 313
Volume Management Tab ... 314
Disk Drives Tab ... 314
Storage System Tabs ... 314
Storage Systems View: Storage System Processors Tab ... 314
Storage Systems View: Volumes Tab ... 315
Storage Systems View: Pools Tab ... 316
Storage Systems View: Host Security Groups Tab ... 316
Host Security Groups on EMC CLARiiON Storage Systems ... 317
Host Security Groups on EMC Symmetrix Storage Systems ... 318
Host Security Groups on HDS Storage Systems ... 318
Host Security Groups on HP P6000 EVA Storage Systems ... 319
Storage Systems View: Storage Extents Tab ... 319
Storage Systems View: Replication Pairs Tab ... 320
Storage Systems View: Backend Storage Tab ... 321
Storage Systems View: SCSI Controller Tab ... 323
Storage Systems View: Disk Drives Tab ... 323
Storage Systems View: Masked Hosts Tab ... 324
Storage Systems View: Pools Logical Usage Tab ... 324
Storage Systems View: Thin Provisioning Data Tab ... 325
Storage Systems View: Volumes Tab ... 327
Storage Systems View: System Nodes Tab ... 327
Storage Systems View: File Systems Tab ... 328
Storage Systems View: Snapshots Tab ... 328
Quotas Tab ... 329
Qtrees Tab ... 330
Shares Tab ... 330
NAS Extents Tab ... 331
Storage Systems View: Initiator Groups Tab ... 331
NAS Replication Pairs Tab ... 331
Storage Systems View: NAS Network Interface Tab ... 333
Storage Systems View: Ports Tab ... 333

CheckPoints Tab 334
Component Storage Systems Tab 335
Fabric Tabs 335
Fabrics View: Switches Tab 335
Fabrics View: Device Aliases Tab 336
Fabrics View: Zone Aliases Tab 336
Fabrics View: Zone Sets Tab 336
Fabrics View: Zones Tab 337
Node Tabs 337
Nodes View: Capabilities Tab 337
Nodes View: Node Groups Tab 338
Nodes View: Registration Tab 338
Node Group Tabs 339
Node Groups View: Device Filters Tab 339
Node Groups View: Additional Filters Tab 340
Node Groups View: Additional Nodes Tab 340
Node Groups View: Child Node Groups Tab 341
Node Groups View: Custom Properties Tab 341
We appreciate your feedback! 342

Chapter 1: Getting Started with SOM

The following topics introduces you to the main features of the SOMconsole and how to navigate the console and configure your browser for SOM.

- "Configuring Web Browsers for SOM" below
- "Using the SOM Console " on page 17

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 16

Configure Mozilla Firefox for SOM

By default, the SOM help opens in a new browser window.

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.

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.

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.

To configure how Microsoft Internet Explorer responds to SOM requests:

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:

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 SOM Interface

You can configure the following user interface features:

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

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" below.)
3. Click **Save and Close** to apply your changes.

To apply your Console Timeout or Initial View configuration changes, sign out of the SOMconsole. 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.Use the value **None (blank)** to specify that you do not want a default view automatically displayed by default. Select a view from the drop-down menu list.

Using the SOM Console

The SOM console is the graphical user interface of the SOM application. The main features of the SOM console are shown in the following diagram and explained in the table below.

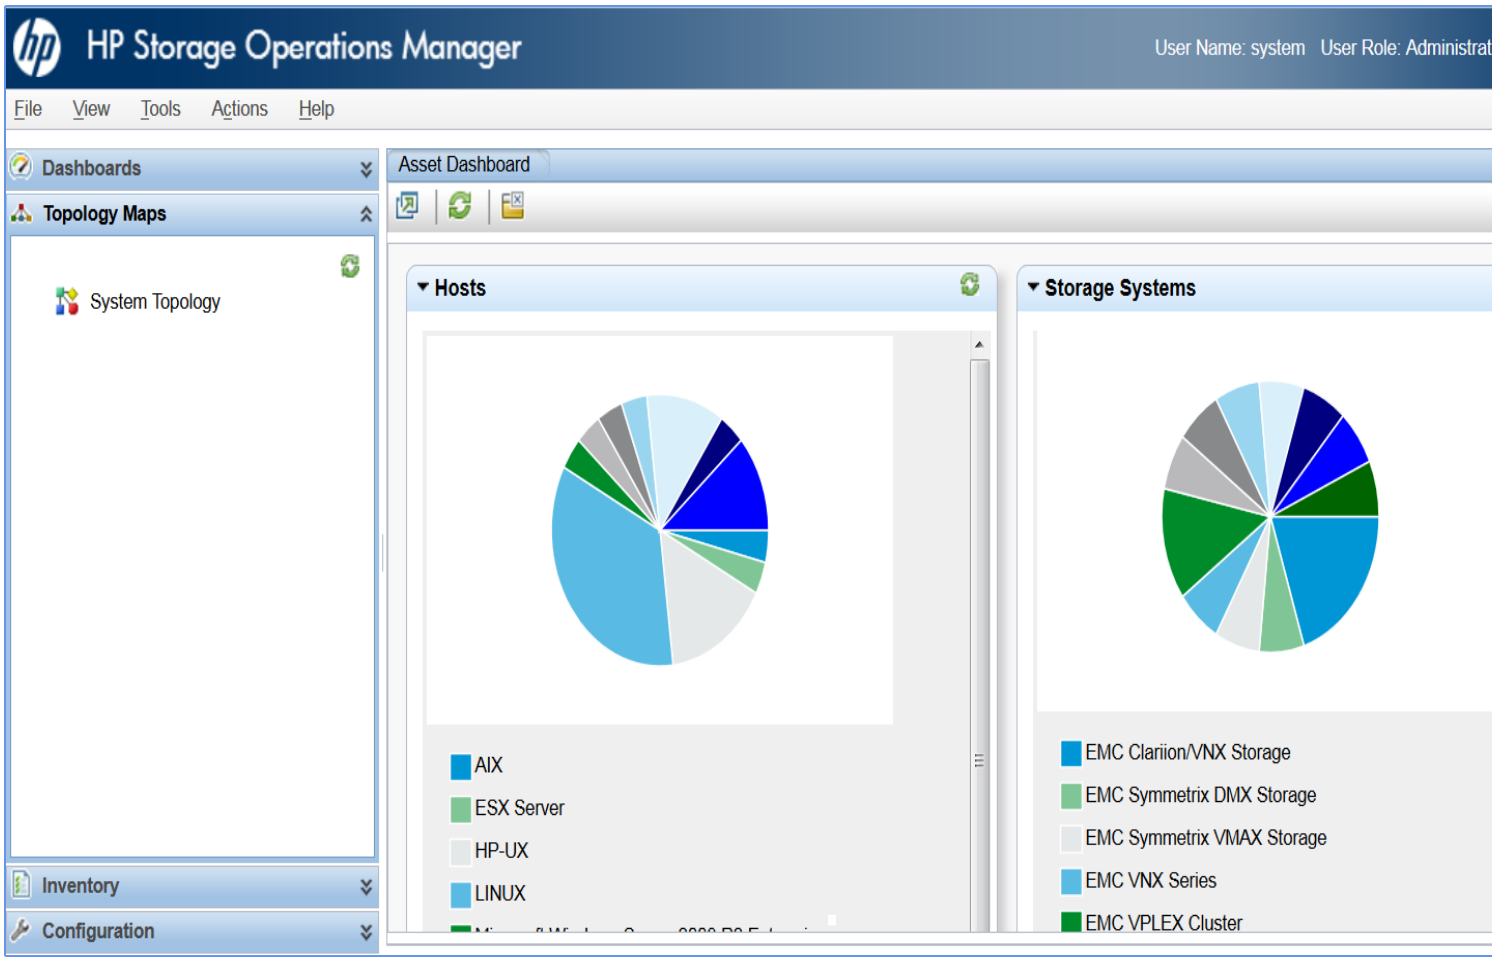

When using the SOM console, note the following:

- If you are using Microsoft Internet Explorer as your browser, you can sign into multiple SOM sessions. Use a different user name for each browser session.
- If you are using Mozilla Firefox as your browser, you can only sign into a single SOM session on each client system.
- You can bookmark the URL for the SOM console.
- Browser context menu might be displayed on right-click from the SOM Console. However, these options do not work and you can ignore them.

SOM Console Features

Feature	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 (see Refresh and Restore All Default View Settings) • Tools (see "Use the Tools Menu" on page 27) • Actions • Help (see "Search the Help Topics" on page 28) **Tip:** To expand SOM menus, you can click with the mouse or use Ctrl-Shift and the underlined character (if any). SOM uses Ctrl-Shift (instead of Alt) to avoid the browser's main menu behavior. For example, SOM provides Ctrl-Shift+H, then u for **Help → Using the SOM Console**. If the SOM menu does not expand as expected, your browser configuration already over-rides the SOM configuration for that keyboard combination of Ctrl-Shift+<*ASCII character*>.
Workspace navigation panel	Helps you navigate between workspaces and views. See "Display Views" on page 22 and "About Workspaces" on page 30.
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.
Console message bar	Alerts about any problems with the SOM 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 SOM console.

SOM Console Features , continued

Feature	Description
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 view title bar, 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.
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.

SOM Console Features , continued

Feature	Description
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.

Navigating the SOM Console

The main window of the SOM console is the starting point for navigation.

A *view* is a collection of related objects that are depicted as a table or map. A *form* provides the known details about a selected object.

From the main window, you can perform the following tasks:

- "Display Views" below
- "Access More Information About an Object (Forms and Analysis Pane)" on page 24
- "Invoke Actions" on page 26
- "Use the Tools Menu" on page 27
- "Search the Help Topics" on page 28
- "Mark Your Favorite Help Topics" on page 30

Display Views

Views contain information about the objects in your network. A view can be a table (a list of objects) or a map with graphical representation of connectivity information.

To display a view:

1. Click a workspace name in the workspaces navigation panel to display a group of views. The workspaces provided by SOM are shown below:

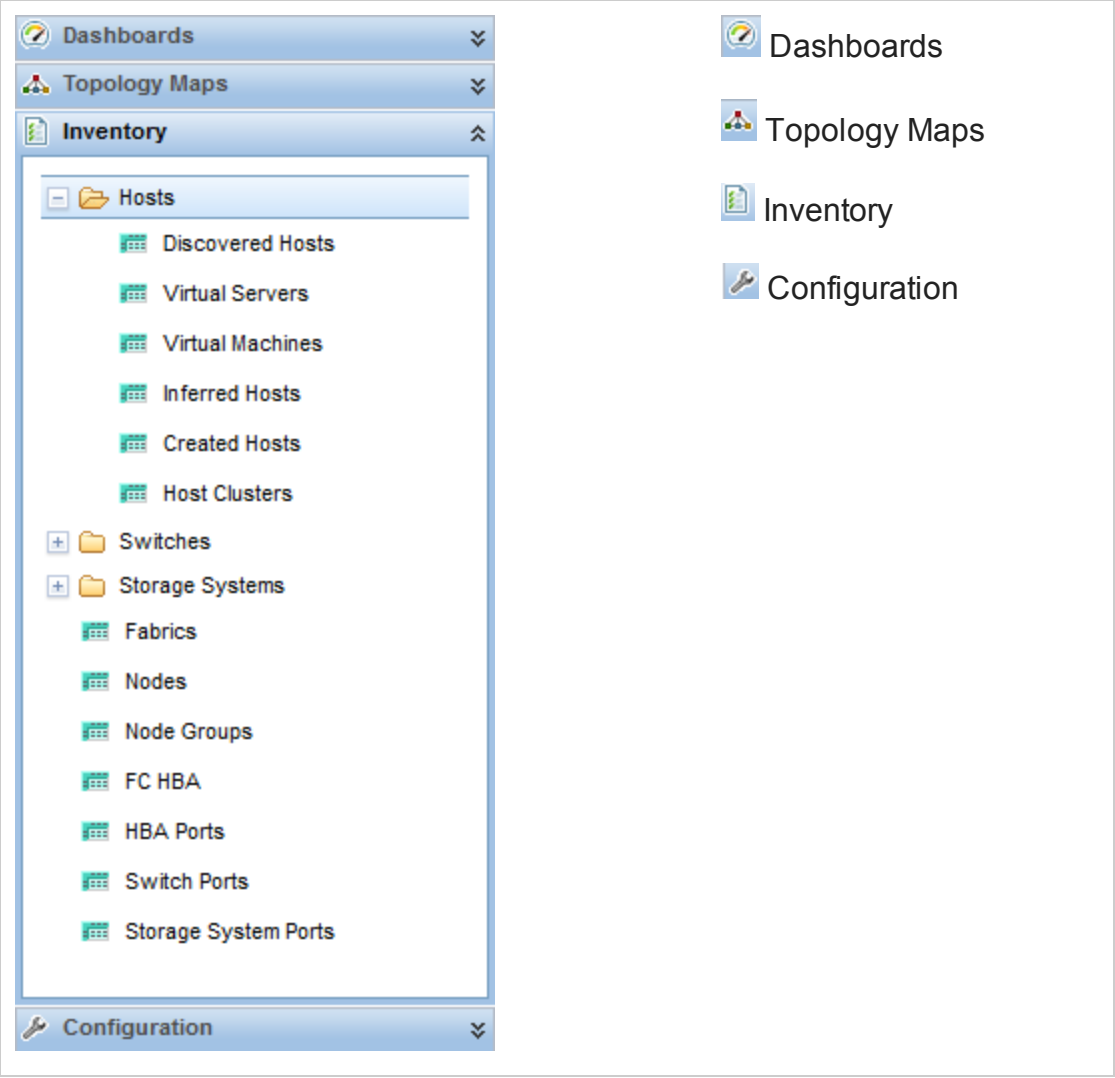

2. Select a view.

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.

Access More Information About an Object (Forms and Analysis Pane)

You can access more information about any object. For example, in your current view, you can obtain more information about physical or virtual switches. From within the physical or virtual switches form, you can access the information about that particular switch.

Access all object attributes and related objects by displaying 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 Node**:

  See the "Use the Tools Menu" on page 27 for more information.

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

  Double-click the row representing an object.

  The form appears, containing the details about the object. For more information, see "Working with Objects" on page 35.

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

  Do one of the following:

  - Select the map object and then click  **Open** on the toolbar.
  - Double-click the map object.

  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. For more information, see "Working with Objects" on page 35.

Access more details about an object using the analysis pane:

To access the analysis pane from a table view:

1. Select the workspace of interest (for example, **Inventory**).
2. Select the view that contains the object of interest (for example, the **Nodes** view).
3. Select the row that contains the object of interest.
4. SOM displays detailed information at the bottom of the view in the analysis pane.

To access the analysis pane in a map view:

1. Select the workspace of interest (for example, **Topology Maps**).
2. Select a map view (for example, select **System Topology**).
3. Click the map object of interest.
4. SOM displays detailed information at the bottom of the view in the analysis pane.

To access the analysis pane in a form:

- Click the form's toolbar **Show Analysis** icon to display information about the current form's top-level object in the analysis pane.

  Note: Show Analysis always displays the top-level object's information.

- Click a row in a table on one of the form's tabs to display detailed information about the selected object in the analysis pane.

SOM displays detailed information at the bottom of the form in the analysis pane.

Note the following:

- Look for one of the following at the bottom of the display area:

  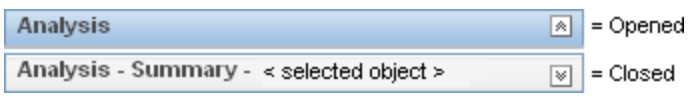

  Open the analysis pane if necessary by clicking the expand button.

- Place your mouse cursor over the title bar to display the ↕ symbol, then resize as necessary.
- The analysis pane remains empty until an object is selected.

- If you select multiple objects or clear a selection, SOM retains the analysis pane's contents.
- If you change views, SOM clears the analysis pane.
- Click any Refresh icon in the analysis pane to update a subset of displayed information.
- SOM automatically refreshes the entire analysis pane's contents when you save a form.

Invoke Actions

The actions available to you depend on your user role and 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.

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

- **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, single-click a row.
     - In a map view, single-click the object of interest.

     Tip: For multiple selections, use Ctrl-click.
  c. Select the **Actions** menu in the menu toolbar.

> **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 workspace navigation panel, select a table view.
  b. From the table view, double-click the row representing an object instance (for example, **node groups**).
  c. On the **Actions** menu, click 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 do not have to 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.)

Use the Tools Menu

There are certain tools provided beneath the **Tools** menu. The list of tools changes depending on the role to which you are assigned. The tools listed in the following table are available to Operator Level 2.

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

SOM Tools Menu Options

Tool	Description
Find Node	Searches the SOM database for the *case-sensitive* string of characters you provide. SOM finds the associated node. If multiple nodes match, SOM displays the Node form of the first match. SOM checks the following node attributes for a match: • Name • Hostname (*fully-qualified*) • System Name • IP Address
Signed In Users	View a list of the SOM users who are currently signed in to SOM.

Search the Help Topics

To search for specific information across all help topics

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 Variables , continued

Description	Variable	Example
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 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.
- Any time you delete your web browser cookies, your help topic favorites list is deleted.

About Workspaces

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

When you click the name of a workspace, the views associated with that workspace display below the workspace in the workspace navigation panel. After you select a view, the view display panel shows the requested data. See "Using the SOM Console " on page 17 and "Display Views" on page 22 for more information about the workspace navigation and view display panels.

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.

Note: Some views appear under folders. To access the list of views available for a folder, click the plus sign (+) that precedes the folder name.

SOM includes the following workspaces:

-  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.

  This workspace includes the following views:

  - Hosts
    - Discovered Hosts
    - Virtual Servers
    - Virtual Machines
    - Inferred Hosts
    - Created Hosts
    - Host Clusters
  - Switches
    - Physical Switches
    - Virtual Switches
  - Storage Systems

- Top Level Storage Systems
- All Storage Systems

- Fabrics
- Nodes
- Node Groups
- FC HBA
- HBA Ports
- Switch Ports
- Storage System Ports

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

About the Analysis Pane

The analysis pane displays related details about the selected object. SOM performs the appropriate analysis on the selected object to determine the most important information to display. Any hyperlink within the analysis pane displays more information about the selected detail.

To access the analysis pane from a table view:

1. Select the workspace of interest (for example, **Inventory**).
2. Select the view that contains the object of interest (for example, the **Nodes** view).
3. Select the row that contains the object of interest.
4. SOM displays detailed information at the bottom of the view in the analysis pane.

To access the analysis pane in a map view:

1. Select the workspace of interest (for example, **Topology Maps**).
2. Select a map view (for example, select **System Topology**).
3. Click the map object of interest.
4. SOM displays detailed information at the bottom of the view in the analysis pane.

To access the analysis pane in a form:

- Click the form's toolbar **Show Analysis** icon to display information about the current form's top-level object in the analysis pane.

  Note: Show Analysis always displays the top-level object's information.

- Click a row in a table on one of the form's tabs to display detailed information about the selected object in the analysis pane.

SOM displays detailed information at the bottom of the form in the analysis pane.

Note the following:

- Look for one of the following at the bottom of the display area:

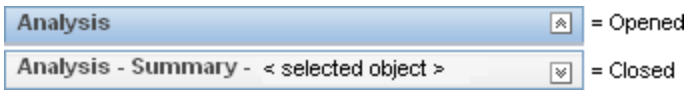

  Open the analysis pane if necessary by clicking the expand button.

- Place your mouse cursor over the title bar to display the ↕ symbol, then resize as necessary.
- The analysis pane remains empty until an object is selected.
- If you select multiple objects or clear a selection, SOM retains the analysis pane's contents.
- If you change views, SOM clears the analysis pane.
- Click any Refresh icon in the analysis pane to update a subset of displayed information.
- SOM automatically refreshes the entire analysis pane's contents when you save a form.

Tip: Some views are also accessible from 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.

Working 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 things, 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:

Do one of the following:

- Select the node of interest, and then click the **Open** icon.
- In most cases, double-click the object of interest.

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:

"Modify Object Attribute Values" on page 38

"Access All Information About a Related Object" on the next page

"Access a Subset of the Available Information About a Related Object" below

Access a Subset of the Available Information About a Related Object

While investigating the available information for an object (within that object's form), some information represents attributes of the object itself and some information is about related objects. 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**.

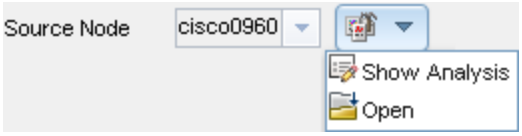

3. The analysis pane appears showing information about the related object. See "About the Analysis Pane" on page 33 for more information.

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

4. Mouse-over any Refresh icon to see the last time the details were updated.

   Click any Refresh icon to gather the most recent data.

Access All Information About a Related Object

While investigating the details for one object using a form, you can access information about another 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.

You can open another form from within a form for any object that is contained in the form you are viewing. Such objects are indicated using a Lookup icon.

This is an example Lookup Field:

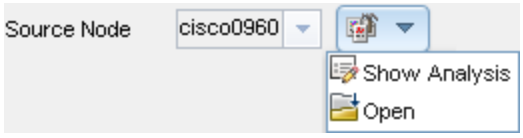

To open another form from within a form:

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

Modify Object Attribute Values

When viewing details for an object, such as a node, you can change one or more of its attribute values. This can be done only during configuration. For example, you can add notes to a node to explain steps that were taken to date to resolve the problem. Until the problem is resolved you can enter information related to a workaround. Finally, after a solution is determined, you can add information describing how the problem was resolved.

Note: If you have Guest user role, you cannot modify any attributes.

Two kinds of fields indicate that you can modify an attribute.

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

To modify information in a text box:

1. Move your cursor to the modifiable field of interest.
2. Type the new value. For example, the Notes attribute is a modifiable field:

   Tip: If the attribute appears to be a modifiable field, but it does not permit text entry, it is a memo field.

   

3. When you are finished with your edits:
   - Click **Save** to save your changes.
   - Click **Save and Close** to save your changes and close the form.

To modify information in a Lookup field:

1. Look for this icon to the right of a text box

. For example:

2. Do one of the following:

   - Start to type into the text box. SOM displays a list of all valid choices. You can select from the list to complete your choice.

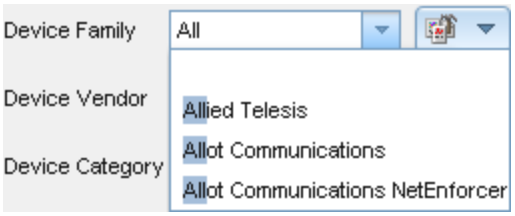

   - Click the Lookup icon, and choose Quick Find to display a list of valid choices:

3. When you are finished with your edits:

   - Click **Save** to save your changes.
   - Click **Save and Close** to save your changes and close the form.

Displaying Information About SOM

Two menu items provide current information about your installed SOM:

- **Help → System Information**

  The **System Information** window provides a wealth of 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.

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

  See "Displaying SOM Version and License Information" below.

Displaying SOM Version and License Information

Select **Help → About HP Storage Operations ManagerSoftware** to display the following information:

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

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

> **Tip:** See also "System Information: Product Tab" on the next page and "System Information: Extensions Tab" on page 43.

System Information: Product Tab

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

The **System Information** window provides a wealth of current information about SOM.

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

The **Product** tab displays information about SOM.

- Product name, version number, and date/time installed.
- Locale Information (language) for the current SOM session:
  - Client locale
  - Server locale
  - SNMP string encodings
  - Web browser
- SOM System Health shows the current status of SOM health:
  - Status
  - Last updated
- 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**.

  Type will be one of the following:

  - Instant-On
  - Premium
  - Ultimate

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

System Information: Health Tab

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

The **System Information** window provides a wealth of current information about SOM.

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

The **Health** tab displays information about the current health of the SOM management server.

The following table describes the possible SOM health status values.

SOM Overall Health Status

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.

System Information: Server Tab

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

The **System Information** window provides a wealth of current information about SOM.

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

The **Server** tab displays 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 **System Information** window provides a wealth of current information about SOM.

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

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

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 45	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 84	Define node groups based on device category, vendor, family and profile and assign nodes to node groups.	Administrators Only	All
" Discovery Tasks" on page 131	Configure IP address, IP address range, credentials for discovery, and tenant associations.	Administrators Only	All
"Inferring Hosts Based on Rules" on page 146	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 160	Configure policies and blackout periods for collecting data from managed elements.	Administrators only	All
"Monitoring Performance" on page 178	Collect performance metrics from managed elements.	Administrators only	Ultimate Perf-Pack Only
"Managing Storage Tiers" on page 184	Configure automated rules-based assignments for categorizing storage systems, volumes and pools into storage tiers.	Administrators only	All

Configuring Security

The SOMsecurity 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.

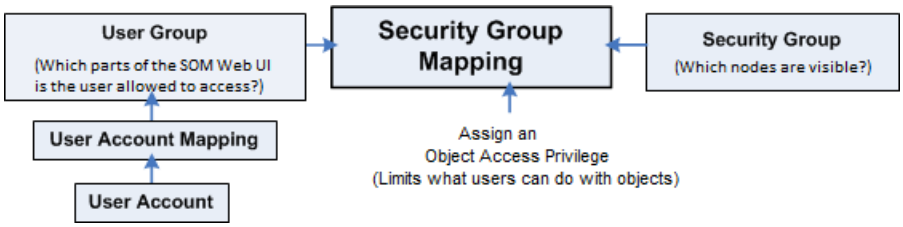

Summary of Security Tasks

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

Task	Description
"Choose a Mode for User Authentication " on the next page	Choose the type of user authentication – SOM Console Access or LDAP.
"Create a User Account " on page 52	You must create a user account for each SOM user.
"Create a User Group" on page 58	The administrator can create any number of user groups to meet the needs of your network environment. Examples of when additional user groups are needed 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 " Predefined User Groups " on page 57.
"Create a User Account Mapping " on page 61	If you created additional user groups, map the appropriate user accounts to each user group you created.

Task	Description
"Create a Security Group " on page 64	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 when you 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 66	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 Assigning Nodes to Security Groups" on page 66	By default, all SOM user groups have access to nodes assigned to the default security 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 .

SOM Configuration Settings

Configure the user names, passwords, and user group membership assignments in the SOM database.

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 52
Task 2	"Create a User Group" on page 58
Task 3	"Create a User Account Mapping " on page 61
Task 4	"Create a Security Group " on page 64
Task 5	"Map User Groups to Security Groups" on page 67

Lightweight Directory Access Protocol (LDAP)

If you have configured LDAP as the directory service for SOM, you must choose the external mode for user authentication.

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

LDAP External Mode

If you are using this 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.

External Mode - Security Configuration Tasks	
Task 1	Modify the `ldap.properties` file and create user accounts.

External Mode - Security Configuration Tasks	
Task 2	"Create a User Group" on page 58 User groups are stored in the SOM database. **Note**: Use the **Directory Service Name** attribute in the User Group form where you can record the *distinguished name*.
Task 3	Configure which objects are visible to each User Group: • "Create a Security Group " on page 64 • "Map User Groups to Security Groups" on page 67

Different Ways to Configure Security

You can configure security using the following methods:

- 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 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 is for security configuration only. It does not include tenant information.

Configure a User Account

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

- "Create a User Account " below
- "Modify a User Account" on the next page
- "Delete a User Account" on page 55

Create a User Account

User Account configurations includes 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 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 options to save the user account.
   - – To save the form.
   - – To save and open a new form.
   - – 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. See "SOM Configuration Settings " on page 48. ☑ indicates that SOM uses Lightweight Directory Access Protocol (LDAP) . Additional steps are required. See "Lightweight Directory Access Protocol (LDAP) " on page 49.
Password	Enter the **Password** value. Type any combination of alpha-numeric characters, punctuation, spaces, and underline characters. **Note:** If you 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 following Security Group 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)
	Re-type the **Password** value.

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 "Lightweight Directory Access Protocol (LDAP) " on page 49 .

To change the user name:

You must "Delete a User Account" on the next page, and then recreate the account mapping (see "SOM Configuration Settings " on page 48).

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. Retype 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 58 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. See "Create a User Account " on page 52 for more information.) ■ 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 in that way, to disable a user's access to SOM, you must use the appropriate process required by your environment's directory service software (see "Lightweight Directory Access Protocol (LDAP) " on page 49).

Caution: If you delete the last SOM user assigned to the SOM Administrators User Group, no one can access the Configuration workspace. See "Restore the Administrator Role" on page 81 for more information about how to recover from this mistake.

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 **OK** to delete the user account.

   - Click **Open**. The user account is displayed in the User Account form view. Click Delete User Account. The message "Are you sure you want to delete this item? This will also delete all contained objects and references." is displayed. Click **OK** to delete the user account.

   The user account configuration is automatically 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. " Predefined User Groups " on the next page

SOM provides " Predefined User Groups " on the next page. 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.

You can perform the following tasks for a user group:

- "Create a User Group" on page 58
- "Modify a User Group" on page 60
- "Delete a User Group" on page 60

Predefined User Groups

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

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

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 73 (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. See "About Workspaces" on page 30 for more information about workspaces. See Views Provided By SOM for more information about the views provided in each workspace.

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. You cannot modify User Group settings for 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

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.
   - – To save the form.
   - – To save and open a new form.
   - – 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 "Lightweight Directory Access Protocol (LDAP) " on page 49.
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 Delete User Group. 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.

User Account Mapping Tasks

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.
   - – To save the form.
   - – To save and open a new form.
   - – 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. (See "Create a User Group" on page 58 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 the **Lookup** icon. • To create new user account, click the **New** icon and provide the required information. See "Create a User Account " on page 52for more information.) • 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, you are only deleting the mapping between the two. You are not deleting the user account or user group 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 Delete User Account Mapping. 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 " on the next page
- "Modify a Security Group" on page 65
- "Delete a Security Group" on page 65

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.
   - – To save the form.
   - – To save and open a new form.
   - – 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 page 66.

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.

Security Group Attributes, continued

Attribute	Description
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 Delete Security Group. 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 70
- 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.

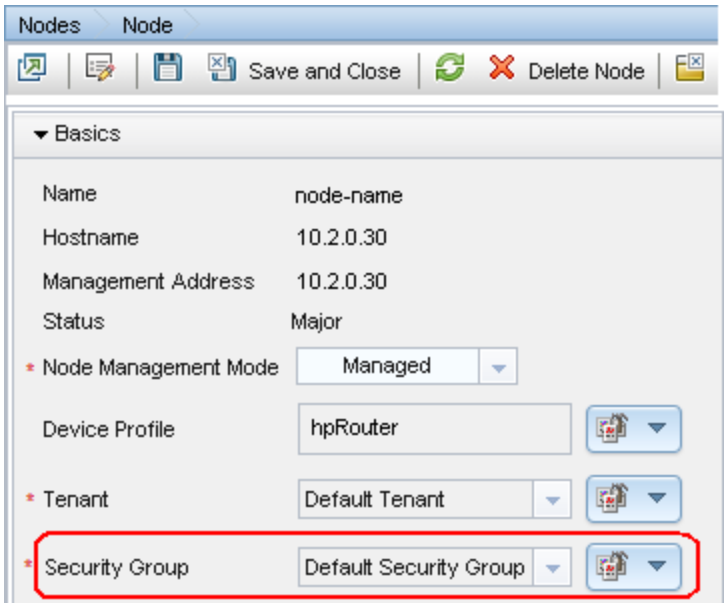

Tip: Administrators can use security groups in node group definitions that become filters in SOM views. If a SOM user cannot access any nodes in a particular node group, that filter dynamically disappears from the filter selection list in the SOM views.

Configure Security Group Mappings

Required only for Operator or Guest users:

Security Group Mappings control which nodes are visible to operators and guests, and what the operators and guests can do 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, so administrators can then map those 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 operator or guest is assigned to multiple security group mappings:

- Multiple predefined SOM User Groups, the SOM consoledisplays all the parts of SOM that are available to the highest User Group.
- 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 70
- "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.
   - – To save the form.
   - – To save and open a new form.
   - – 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 58 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. (See "Create a Security Group " on page 64for more information.) • 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 • 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. See "Control Menu Access" on page 73 for more information.

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. See "Create a User Account " on page 52for more information.

   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 User Account 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.

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

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.

7a. Assign Nodes to Security Groups. 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. See " Predefined User Groups " on page 57.

- 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 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 an SOM user group with *fewer privileges* than the Object Access Privilege, the user will not 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 will see only those actions 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. Click here for information about Tools Menu Access Limitations.

SOM Tools Menu Access Limitation

Tools Menu Item	SOM User Group	Object Access Privilege
Find Node	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
Node Group Details	Show Members (Include Child Groups)	SOM Level 1 Operators	Object Operator Level 1

URL Action Access Limitations, continued

Action Menu Item	Submenu Item	SOM User Group	Object Access Privilege
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.

Communicate Console Access Information to Your Team

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 79
- "Signing Out from the SOM Console" on page 79

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 your 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 SOM is using 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. See the CLI Reference Pages for more information.

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`

Communicate the following browser requirements for your team to use the SOM console:

- 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 all 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:

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" on the next page
- "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 into 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 into 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 such others. 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. See "Node Groups Provided by SOM" on the next page for information.

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.

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 the following default node groups. These are configured with specific information about your management domain. You can change them to meet your needs.

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.
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 ""Using Additional Filters for Node Group Definitions " on page 88" 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. 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 87 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.

Child Node Group Attributes	Description
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.

Using Additional Filters for Node Group Definitions

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 92
- " Create an Additional Filters Expression" on page 95

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	customAttrValue
node A	com.hp.som.capability.card.fru	location	datacenter1
node B	com.hp.som.capability.card.fru	<undefined>	<undefined>
node C	<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	customAttrValue
nodeA	com.hp.som.capability.card.fru	location	datacenter1
nodeB	com.hp.som.capability.card.fru	<undefined>	<undefined>
node C	<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 92 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:

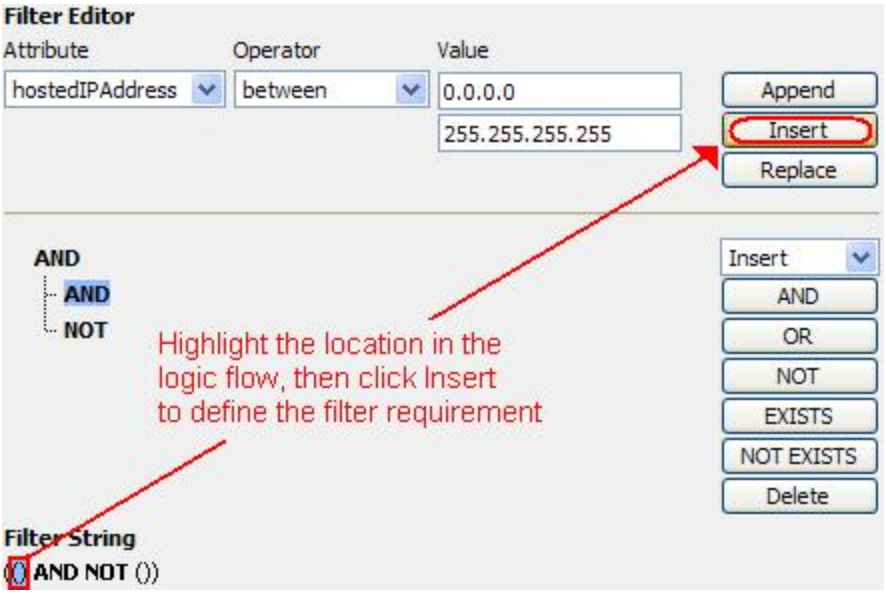

2. Click **Save and Close**.

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)

Additional Filters Editor Choices for Node Groups, continued

Attribute	Description
Operator	The standard query language (SQL) operations to be used for the search. **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**.

Additional Filters Editor Choices for Node Groups, continued

Attribute	Description
	• **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` 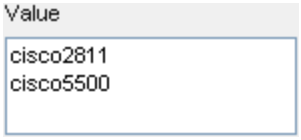 finds all systems with names that are **cisco2811** or **cisco5500**. **Note:** Each value must be entered on a separate line as shown in the example. SOM displays the list of attributes using comma-separated values enclosed in parentheses, for example (**cisco2811, cisco550**). However, the comma-separated 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:

Additional Filters Editor Choices for Node Groups, continued

Attribute	Description
	■ 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. Example: `sysName not in` 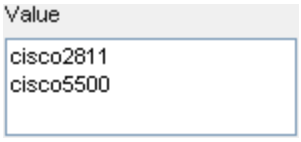

Additional Filters Editor Choices for Node Groups, continued

Attribute	Description
	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).

Additional Filters Editor Choices for Node Groups, continued

Attribute	Description
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. • 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.

Additional Filters Editor Buttons, continued

Button	Description
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 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)))`

Additional Filters Editor Buttons, continued

Button	Description
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 Health tab on the System Information page.

- 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, the discovery credential for the IP address or the FQDN.

Prerequisites for Agentless Discovery of Linux, Solaris, and AIX Hosts

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

- Linux
- Sun Solaris
- 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:

- 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.

- Use the root or non-root user account to access the host.

- **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.

  For 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, port and target mappings is not available.

  For an AIX host, the `MaxMediaSize` of disk drives and information related to the HBA, port and target mappings is not available.

- 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.

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:

- As a Root User
  - "Commands for a Linux Host as a Root User" below
  - "Commands for a Solaris Host as a Root User" on the next page
  - "Commands for an AIX Host as a Root User" on page 111
- As a Non-Root User
  - "Commands for a Linux Host as a Non-Root User" on page 112
  - "Commands for a Solaris Host as a Non-Root User" on page 114
  - "Commands for an AIX Host as a Non-Root User" on page 115

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.
`udevadm info -a`	Collects SCSI information about SUSE Linux.
`/usr/sbin/vxprint`	Provides information on Veritas Volume Manager's disk groups and their associations.

Command	Description
`/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.
`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.

Command	Description
`/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 "mpxio-disable"`
	Determines if Solaris native Mpxio is enabled.

`/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 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 a 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

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 discovered hosts' node name, kernel release and model details.
`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.
`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.

Commands	Description
`cat /proc/partitions`	Determines information about the disks and disk partitions of the host. The output of this command is used by`fdisk -l` command.
`udevinfo -a -p`	Collects SCSI information about Redhat Linux.
`ls -l`	Determines permission and ownership details. Also, determines permission details for 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 HBA related information.
`/usr/bin/systool -c scsi_host -v`	Collects information related to HBA ports.
`/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.
`/usr/sbin/metaset`	Provides information on Solaris native volume manager disk sets.

Command	Description
`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.

Command	Description

`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.
`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.

Command	Description
`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 server can be either Linux or Microsoft Windows.

Prerequisites for Windows agentless discovery with Windows SOM:

- Provide the IP address or DNS name of the host.
- Provide a user account with administrator privileges.
- Stop or uninstall the CIM extension on the host if applicable.
- 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.
  - `Admin$` and `IPC$` shares must be enabled.
  - The environment variable `PSEXEC_PATH` must be set to the installation path.
- Enable the WMI service on the remote Windows host.
- The HP MPIO binaries must be specified 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`

Prerequisites for Windows agentless discovery with Linux SOM:

- Provide the IP address or DNS name of the host.
- Provide a user account with administrator privileges.
- Stop or uninstall the CIM extension on the host if applicable.
- Enable the WMI service on the remote Windows host.
- Ensure that the TCP ports 135, 139, and 445 are open.
- Ensure that the `Admin$` and `IPC$` are enabled.
- The HP MPIO binaries must be specified 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`

Note: If the CIME agent is running, use the Agentless Discovery Hint (Configuration > Discovery Addresses) to discover a Windows host in the agentless mode.

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.

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.

Commands	Description
`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 = <NUMBER>`	Provides detailed information about HP MPIO paths to a 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

To discover a host with a CIM agent:

- Install the CIM extension on the host.
- Provide the IP address or DNS name of the host.
- 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
Automatic	CIM Extension	Pass
Automatic	Windows Agentless **Note**: If the registry key is modified to enable the agentless mode, discovery will pass with the Windows Agentless hint.	Fail
Automatic	No Hint	Discovered with CIM agent
Disabled	No Hint	Discovered as agentless
Disabled	CIM Extension	Fail

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 VirtualCenter)

Prerequisites to Discover VMware ESX Servers and Virtual Machines

ESX Servers and Virtual machines can be discovered through a Virtual Center (VC) 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 VMware Virtual Centre:

- 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.

- Provide the user name and password of the Virtual Center account that can view or access the ESX Servers or virtual machines to be discovered.

  The VirtualCenter account must have "Datastore Browse" privileges.

- Discover the Virtual Center.

  Notes:
  - All ESX Servers and virtual machines that the Virtual Center account can access are discovered automatically.
    Use the custom property `discovery.exclude.vmware.vm=true`, in the [*drive*:]`\ProgramData\HP\HP BTO Software\Conf\som\custom.properties` file to disable the automatic discovery.
  - You can discover a DRS cluster and its details only via a Virtual Center.
  - 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.

To discover via an ESX Server:

- 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.

- Provide the user name and password of an ESX server.
- 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 [*drive*:]`\ProgramData\HP\HP BTO Software\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 the Brocade Network Advisor (BNA).

To discover Brocade Switches, provide access details of the BNA server. Specify the IP address (Configuration > Discovery > Discovery Addresses) and authentication details of the BNA server.

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 inventory view.

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 (as it is not set by default).

  a. Type the following commands to set the community string on a Cisco Switch:
     i. To display the Cisco SNMP configuration settings

     ```
     cisco_switch# show snmp
     ```

     ii. To enter the configuration mode

     ```
     cisco_switch# config t
     ```

     iii. To enable the read only community string

     ```
     cisco_switch# snmp-server community public ro
     ```

     iv. To exit the configuration mode

     ```
     cisco_switch(config)# exit
     ```

     v. To save

     ```
     cisco_switch(config)# copy run start
     ```

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

  Or

- **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:

a. **Modify the Custom Properties File**
Use the `custom.properties.sample` file to create and edit the `custom.properties` file in the following locations:

- Windows - [*drive*:]`\ProgramData\HP\HP BTO Software\Conf\som`
- Linux - `/var/opt/OV/conf/som`

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.

Restart the `ovjboss` service after modifying these properties.

b. **Create an Account on a Switch**
To access a switch, you must create an account on a switch using either of the following:

- **Cisco Fabric Manager**
To create an account on all the switches in a fabric with the same credentials and security settings.
- **Cisco Device Manager**
To create an account on one switch.

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

Use the following CLI commands to create an account on a Cisco switch:

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

ii. To create a user account
`Cisco1-switch1(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.

Set Time-out Properties for Discovery

Use these properties in the `custom.properties` file to set the time-out value for the discovery of a Cisco switch.

- `cisco.snmp.timeout=15000`
- `cisco.snmp.retries=2`

To optimize discovery of Cisco switches, modify the following SNMP ping properties in the `custom.properties` file to ensure discovery of all the switches in the environment.

- `discovery.snmp.timeout=10000`
- `discovery.snmp.retries=3`

Set Cache Refresh Time

The default cache refresh time for Cisco switches is two hours. Therefore any modifications in the properties of a switch are not reflected if data collection is initiated before the cache is refreshed.

Modify the custom property `wbemcollector.collectioncachetimeout` in the `custom.properties` file to reduce the cache refresh time.

For example, `wbemcollector.collectioncachetimeout=1`, where the value indicates the number of hours. An hour is the minimum value for this property.

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 inventory view.

By default, disabled VSANs are not discovered. To discover disabled VSANs, set the following property in the custom.properties.sample file:

```
cisco.showDisabledVsans=true
```

Prerequisites to Discover HP XP/P9500 Arrays

SOM discovers HP XP arrays through the service processor (SVP) on the array using the RMI-API.

To discover HP XP arrays:

- **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.
- **Provide access to the SVP**
  Specify the IP address (Configuration > Discovery > Discovery Addresses) of the SVP.

By default, the following ports are used: 1099, and 51099.

The discovery hint (Configuration > Discovery > Discovery Addresses), **HDS/XP –Native API**, applies to both HP 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 HP 3PAR Arrays

SOM uses the 3PAR SMI-S server to discover the 3PAR array.

To discover a 3PAR array:

- **Start the 3PAR SMI-S server**
  By default, the SMI-S server is not started on the array. To start the SMI-S server, start the InForm CLI interface on the array and run the following command: `startcim`. This command starts the SMI-S server.

- **Provide access to the 3PAR array**
  Provide the IP Address or DNS of the 3PAR array along with the user name and password. The default authentication credentials are: 3paradm/3pardata.

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.
  The EVA is not discovered 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.

Prerequisites to Discover HDS and HUS Arrays

The management server uses the Hitachi HiCommand Device Manager (HDvM) and the built-in HDS provider to discover and collect data from an HDS and HUS 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 an HDS array:

- **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.

- **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 these ports are also open: 1099, 51099, and 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 management server communicates with the HiCommand Device Manager through a non-secure connection.

For more information about discovering HDS arrays, see " Discovery Tasks" on page 131.

Prerequisites to Discover an EMC Isilon Cluster

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.

- 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.

Or

- 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 non-root user account. The information obtained using either of the accounts is the same.

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:

- **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.

- **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.

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:

- Install EMC Solutions Enabler with the SMI-S package on a proxy server.
- 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.

- Ensure that these ports are open: 5988, and 5989.

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:

- Provide the IP address and authentication details of the VPLEX management console.
- 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.

NetApp 7-Mode Device

SOM uses the ONTAP API to discover NetApp 7-Mode devices.

To discover a 7-Mode device:

- Provide the IP address and authentication details of the NetApp device.
- Port 443 must be open.

Discovery Tasks

To discover a new device, follow these steps:

1. Configure an address for discovery.

   or

   "Configure a Range for Discovery" on page 135
2. (*optional*) "Configure Credentials for Discovery" on page 138
3. "Create a Tenant" on page 142
4. "Start Discovery" on page 144

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.
   - – To save the form.
   - – To save and open a new form.
   - – To save and close the form.

   The address is displayed in the Discovery Addresses view.

Attributes	Description
IP Address/DNS Name	Type the IP address or the FQDN of the device to be discovered. 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.

Attributes	Description
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.

Attributes	Description
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 Delete Discovery Addresses. 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.

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.
   - – To save the form.
   - – To save and open a new form.
   - – 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.

Attribute	Description
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.
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 Delete Discovery Addresses. 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" on the next page table.)
4. Click one of the options to save the credentials.
   - – To save the form.
   - – To save and open a new form.
   - – 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.

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.

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 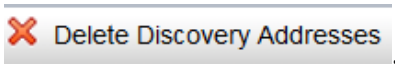

. The delete confirmation message is displayed. Click **OK** to delete the discovery credential.

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 95 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 143 for a Node for more information.

Related Topics

"Configure Security Groups " on page 63

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 45.

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.
   - – To save the form.
   - – To save and open a new form.
   - – 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.

Tenant Attributes, continued

Attribute	Description
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 page 141. **Caution:** 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. 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.

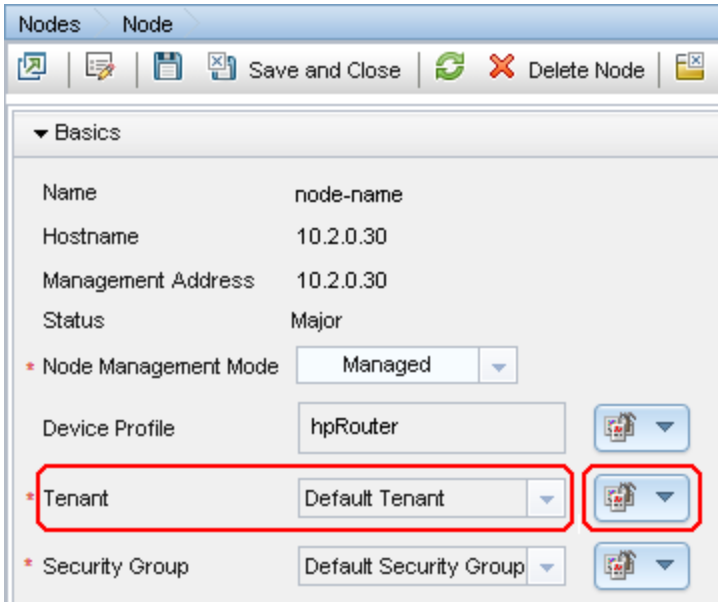

3. Click Save to change the tenant.

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 to see the complete list of discovery statuses.

Note: The Start Discovery option does not restart the process for devices that are being discovered.

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.

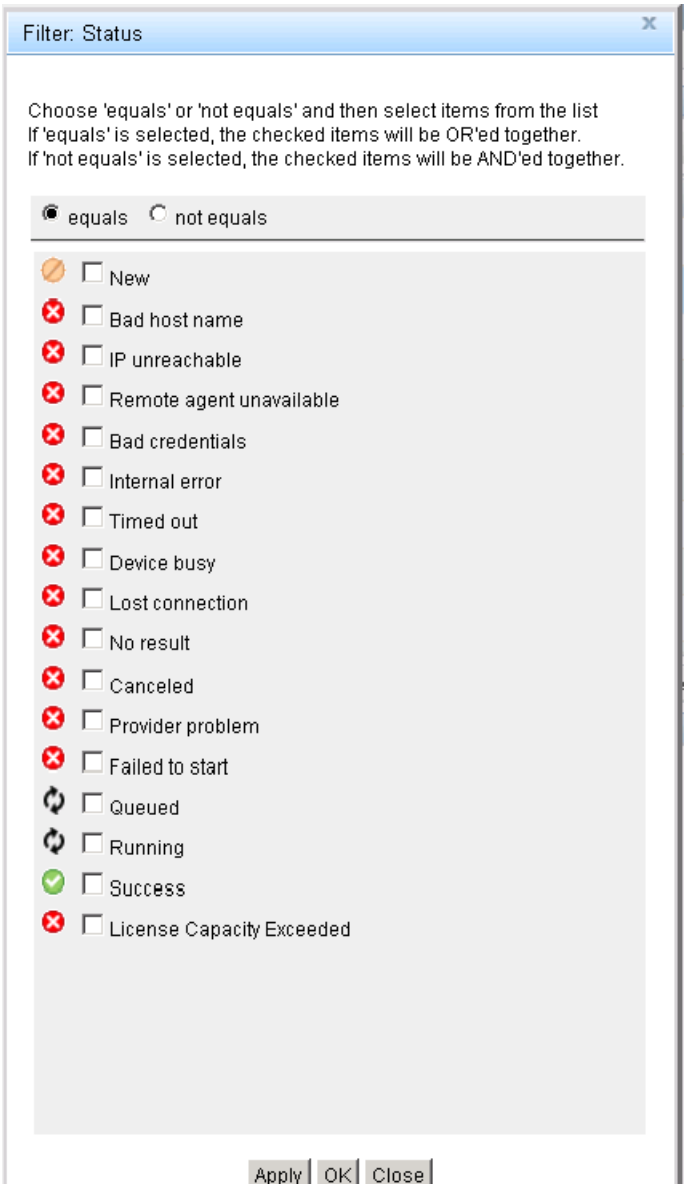

Discovery Views

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 144
- "Status of Discovery" on the previous page

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 or zone aliases

configured on storage systems and fabrics in the SAN. Rules probe your switch and storage configurations based on specific search parameters using regular expressions to infer hosts. For zone and zone aliases scope, the elements will be inferred only after the fabric connectivity information is available for the ports being inferred.

The following functionality is not available for hosts inferred through rule-based host inference:

- Automatic cluster membership detection
- Only presented storage information is available.

Virtual machines cannot be inferred using host inference rules.

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 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 the only the host names without the boston_ prefix. In such a case, you can use the following expression: boston_(.*)

Any host with a prefix of boston_ would be inferred, but only the text after boston_ would be 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.

Examples of Regular Expressions, continued

Environment	Regular Expression	Result
`Boston-HostName_com`	`.*?-(.*?)_.*`	Strings that match the pattern of text-text_text will be scanned. The text 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.

Examples of Regular Expressions, continued

Environment	Regular Expression	Result
`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.

Examples of Regular Expressions, continued

Environment	Regular Expression	Result

`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.
		For uppercase letters use [A-Q].
		You can change the range to match your environment; for example, a-s or N-Z.

Examples of Regular Expressions, continued

Environment	Regular Expression	Result

`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`	`.*_(.*)_.*|(.*-(.*)_.*)`	

Examples of Regular Expressions, continued

Environment	Regular Expression	Result
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.
`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.

Definition of Common Notation Used in Expressions, continued

Expression	Definition
?	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 HP 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.
	Note: Regular expressions are Java regular expressions and you must take care about using the greedy and reluctant quantifiers, as appropriate.

Definition of Common Notation Used in Expressions, continued

Expression	Definition

.	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 is displayed.

2. Click **New** on the view toolbar. The Host Inference Rule form is displayed.
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.

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.

Host Inference Rule Attributes	Description
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 the 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 the scope to infer hosts from the following options: **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. **Zone** – The rule searches the zone name for hosts on the fabrics. The discovery of the fabrics must be complete. **Zone Alias** – 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 Inference Rule Attributes	Description
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 147 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 is displayed.
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 is displayed.
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 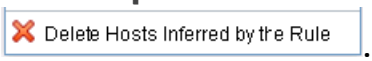

. The delete confirmation message is displayed. Click **OK** to delete the rule.

When you delete a rule, the hosts inferred from the rule are not deleted and will continue to appear in the Inferred Hosts view in the Inventory workspace. 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 is displayed.
2. Click a rule to select it, right-click and click **Run a Rule**.

   The inferred hosts are displayed in the Inferred Hosts view in the Inventory workspace.

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 to the host topology is 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 is displayed.

2. Click an inferred host that you want to delete. Right-click and select  . The delete confirmation message is displayed. 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 is displayed.

2. Select a rule from the table view, right-click and select 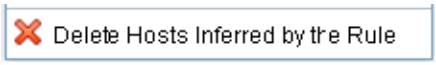

. The delete confirmation message is displayed. Click **OK** to delete the hosts inferred by the selected rule.

Reconciliation of Hosts

SOM performs reconciliation of hosts when you discover the inferred hosts after providing the credentials.

Reconciliation of hosts results in the following:

- The ports and cards are associated with the managed element after reconciliation.
- Inferred hosts wherein previously associated ports and cards have been reconciled and associated with managed elements will be 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 84 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 166 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.
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 Delete Data Collection Policy. 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 Delete Blackout Period. 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		
EMC Symmetrix VMAX Storage		
HP 3PAR Storage		
HP EVA 6000 Storage		
Hitachi Storage Series (also applies to HP XP/P9500 arrays)	Disk Drives, Storage Extents and Host Security Groups	

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 HP 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 SOMUltimate-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 HP 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 perpetual license, gather the following information:

- The Entitlement Certificate, which contains the HP product number and order number.
- The IP address of one of the SOM management servers.
- Your company or organization information.

Install a Perpetual License

You can install the perpetual license using the Autopass user interface or the command line interface.

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.

Using Autopass to Install a Perpetual License

To install a perpetual license, follow these steps:

1. At a 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 HP Sales Representative or your Authorized Hewlett-Packard Reseller for information about the SOM licensing structure. To obtain additional license keys, go to the HP License Key Delivery Service:

https://h30580.www3.hp.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.

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 a FC port	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.	

Element	Description	Number of MAPs	Comments
	AIX agentless discovery through SSH	1 MAP at a minimum or 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 with no FC ports.	1 MAP	The software assumes one MAP.

Element	Description	Number of MAPs	Comments
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. • 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	

Element	Description	Number of MAPs	Comments
HP 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/ Celerra		5 MAPs	Only single node supported.
EMC VNX Filer		5 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 performance of your storage environment using a monitoring policy. A monitoring policy enables you to configure collection of performance metrics for hosts, storage systems, and switches. You can configure specific set of metrics to be collected on a group of elements.

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 will be collected. For example, you can schedule to collect metrics at an interval of every 15 minutes from a device.
- **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 171 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

To create a monitoring group, follow these steps:

1. From the Configuration workspace, select **Object Groups** > **Monitoring Groups**. The 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. On the Collector Settings tab, click **New** to associate collectors to the monitoring group. The Collector Settings form is displayed.
4. Select the **Collector** from the drop-down list.
5. Click one of the save options.

> **Note**: You can associate multiple collectors to a monitoring group. You must have at least one collector associated with a monitoring group.

6. Click one of the save options to save the monitoring group.
   - – To save the form.
   - – To save and open a new form.
   - – To save and close the form.

Modify a Monitoring Group

You can modify a monitoring group to configure additional collectors or change collectors associated with it. However, a monitoring group must have at least one collector associated with it.

Create a Monitoring Policy

To create a Monitoring Policy, follow these steps:

1. From the workspace navigation panel, select the **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
      - – To save the form.
      - – To save and open a new form.
      - – 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 the previous page" 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.
   - – To save the form.
   - – To save and open a new form.
   - – To save and close the form.

6. Click one of the options to save the monitoring policy.
   - – To save the form.
   - – To save and open a new form.
   - – 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 SHR 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.

- – To save the form.
- – To save and open a new form.
- – 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 Delete Monitoring Policy. 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.

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 

. 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.

Chapter 3: Managing your Storage Environment with SOM

SOM provides the following features that enable you to manage your storage environment.

Feature	Description
"Inventory Views" on page 196	Provides a collection of views to access details of elements managed by SOM.
"Dashboards" below	Contains information panels pertaining to the entire storage environment, an element category, or an individual element.
"Topology Maps" on page 265	Displays the connectivity maps of the top level elements in the storage infrastructure.

Dashboards

SOM includes dashboards that provide the latest information about the number of discovered devices, data collection statuses, and the storage utilization in the environment. Dashboard views help compare and isolate the details required to analyze data.

A dashboard can contain 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 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.

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
  - 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.

- **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 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.

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/Un-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.

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**.

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 198
- "Storage Systems Views" on page 199
- "Fabrics View" on page 200
- "Nodes View" on page 201
- "Node Groups View" on page 202
- "FC HBA View" on page 204
- "HBA Ports View" on page 204

- "Switch Ports View" on page 291
- "Storage System Ports View" on page 205

Using the Analysis Pane

Use the Analysis pane to view the following information about a selected device:

- Summary
  Key information about a selected element.

  Note: The **Access Point** property displays the IP address that was used to discover and collect data from a device.

- Capacity
  Overall capacity utilization of a selected element. For more information, see "Viewing Device Capacity" on page 205.

- Performance
  Performance information about a selected element. For more information, see "Viewing Device Performance" on page 220.

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 291.
- **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.
- **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 308.

The **Properties** pane displays the properties of a selected switch.

The **Analysis** pane displays the "Switch Capacity" on page 207 and "Performance Collectors for Switches" on page 223 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.

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 of 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 335
- "Fabrics View: Device Aliases Tab" on page 336
- "Fabrics View: Zone Aliases Tab" on page 336
- "Fabrics View: Zone Sets Tab" on page 336
- "Fabrics View: Zones Tab" on page 337

The following shows an overview of a Fabric zoning structure.

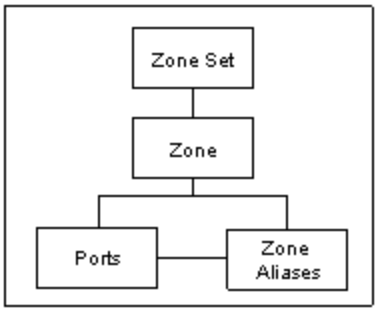

The **Properties** pane displays the properties of a fabric.

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).

Attribute	Description
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.

Use this view to see the switch ports and host initiator ports that a storage system port is connected to.

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.

Double-click a port in the **Connected Switch Ports** tab view to see its form view.

The **Properties** pane displays the properties of a storage system port.

Viewing Device Capacity

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

Host Capacity

The **Analysis** pane includes the following tabs with the overall capacity information of a selected host.

The aggregate capacity at the host level excludes network filesystems such as nfs, nfs4, cifs, smbfs, and ncpfs.

The tabs contain charts that are customizable. 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.
Presented Storage	The LUN size that each storage system presents to a host. The table displays the **Storage Systems** and **Size** of each LUN 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 in order to reclaim space. The table displays the following: • **Storage Volume** • **Size (GiB)** • **Storage System**
Unused Volume Group Capacity	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: • **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.
Volumes by % Used	Topmost volumes used by the selected host element. The table displays **Host Volume**, percentage of the volume capacity used (**Used %**), and a bar chart that reflects the percentage of capacity used.

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 Refresh to see the last time the details were 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 below:

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.

Device-Specific Exceptions

The following sections provide information that is specific to vendors, and about how SOM maps acquired device information.

Hosts

Switches

Storage Systems

Hosts

This topic captures information that is specific to certain hosts, how they are handled by SOM, or some vendor specific information.

Hosts Discovered using the Agentless Method

The following are limitations for hosts discovered using the agentless method, 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 target mapping 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:
       A. Vendor name
       B. Serial number
       C. Hardware version
       D. Port Type in the Ports tab

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.

- **Solaris Hosts**
  The **Hardware Version** of the HBA card (Properties pane of the Form view) is blank for Solaris hosts.

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) is blank.
- 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, [*drive*:] `\ProgramData\HP\HP BTO Software\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 VirtualCenter. If you discover individual ESX servers directly, the templates are not found.
- **Details of Virtual Hosts After Discovery**
  After discovery, the following details are available in the hosts inventory views:

Discovered Virtual Host	Inventory Details
VirtualCenter	The VirtualCenter's access point in the Summary tab with the associated virtual servers. ■ **Access Point** (of the VirtualCenter) ■ **Hosts Virtual Servers view** – Details and sub-components of the virtual servers managed by the VirtualCenter.
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 VirtualCenter's access point in the Summary tab. ■ **Access Point** (of the virtual server or VirtualCenter) ■ **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

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.

VMWare ESX Servers

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 Server.

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.

Switches

This section provides information that is specific to the vendors or to how SOM handles managed switches.

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

- Displaying the Slot and Port Number for Switches
  Set the `brocade.getSlotDetails` property to true in the properties file, `.properties` in the folder `<DATA_DIR>/conf/se` and restart jboss.
- Switches discovered through BNA appear differently in SOM.
  For information about changing the required settings, refer to 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 information that is specific to a certain device, how it is handled by SOM, or some vendor specific information.

HP 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, edit the

file `drive:\Program Files (x86)\HP\HP BTO Software\nmsas\common\deploy\transaction-jboss-beans.xml` and set the property "defaultTimeout" of "CoordinatorEnvironmentBean" to 600. Restart the ovjboss 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 HP XP terminology as follows:

Property	Continuous Access	HP Continuous Access Journal	HP Business Copy	HP 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 Post-RAID 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

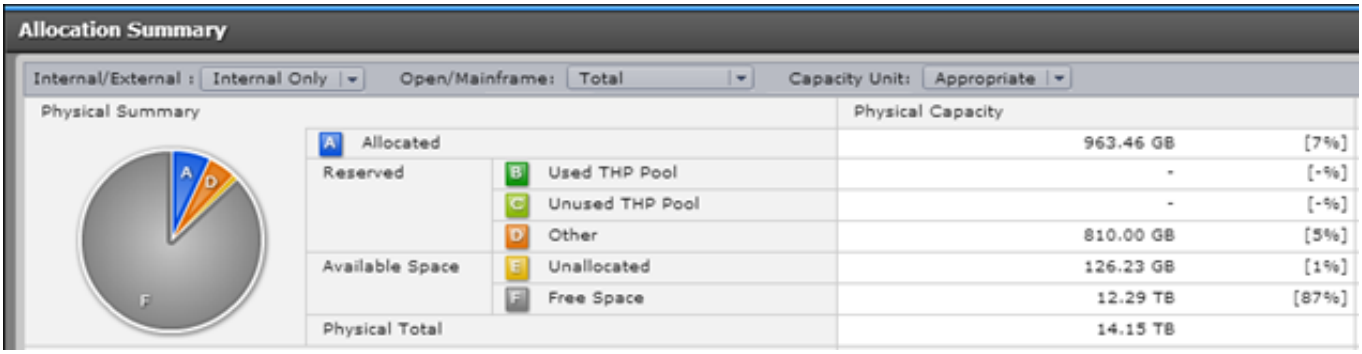

  The internal allocation metrics in RWC can be mapped to the following metrics in the Thin Provisioning Data tab:

  - **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

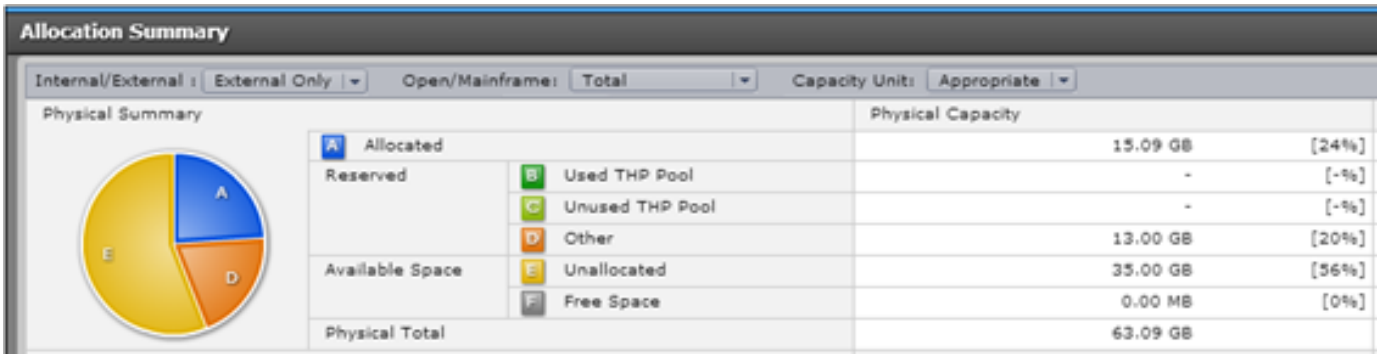

The external allocation metrics in RWC can be mapped to the following metrics in the Thin Provisioning Data tab:

- **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 **Post Raid 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.

HP 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 is due to a bug in the 3APR 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 restart data collection for the 3PAR array.

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

- **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.

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 Extents Tab**
  Displays storage composition information as volumes. For example, meta volumes, slice volumes, and so on. The Description property in the Properties pane identifies the volume type.

EMC Symmetrix/VMAX/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 Symmetrix 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.

- **Disk Drive Size of VMAX**
  The disk drive size is derived from the SMI-S property, `MaxMediaSize`. The value of this property does not match with the value obtained from the device vendor tool.

- **VMAX Disk Drive Form view**
  The Storage Disk Drive form view of VMAX disk drives shows the Thin and Thick volumes of a drive as obtained from the SMI-S. However, the EMC device vendor tool, shows only the Thin volumes.

- **Total Raw Capacity of VMAX**
  The Total Raw capacity in the Raw Capacity tab, is derived from the SMI-S property, `TotalManagedSpace` of the primordial pool and does not match with the value obtained from the device vendor tool.

- **Raw Available Capacity of VMAX**
  The SMI-S property `EMCRemainingRawCapacity` returns a zero value for the Raw Available capacity of the primordial pool. SOM derives the raw available capacity by the sum total of the remaining set of extents via the SystemDevice association.

- **Pools Tab of VMAX**
  The Pools tab displays Thin pools, Disk Groups, and SMI Disk Sparing profile pools.

  - **Disk Groups**
    Since the SMI-S provider returns the raw capacity of disk groups, the Post-RAID capacity of a disk group is calculated by the sum of the volumes in the disk group. Therefore, the Post-RAID capacity displayed in the Analysis pane of the Pools tab does not match with the value obtained from the device vendor tool. Consequently, the total capacity displayed in the Post-RAID Allocation tab does not match with the device vendor tool.

  - **Disk Sparing Profile Pools**
    The EMC provider implements the SMI Disk Sparing profile, which results in three pools:

    - "AVAILABLE_FOR_FAILOVER" - the extents associated with this pool represent the Spare disk drives.

- "FAILED-REPLACED_BY_SPARE" - the extents associated with this pool represent the failed disk drives
- "FAILED-NOT_REPLACED_BY_SPARE" - the extents associated with this pool represent the failed disk drives that could not be replaced.

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.

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.

Viewing Device Performance

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 show hourly data that is a rollover of data taken at intervals of 15 minutes. The charts are refreshed daily.

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 223
- "Performance Collectors for HP 3PAR Arrays" on page 226
- "Performance Collectors for HP StorageWorks EVA Arrays" on page 234
- "Performance Collectors for EMC Symmetrix DMX/VMAX Arrays" on page 246
- "Performance Collectors for CLARiiON and VNX Arrays" on page 257

Performance Collectors for Hosts

The following performance collectors (Configuration > Monitoring Settings > Collectors) are available for managed hosts:

- "Physical Disk Collectors" below
  - Windows Host Physical Disk Collector
  - Linux Agent Physical Disk Collector
  - HPUX Physical Disk Collector
- "ESX Server Performance Collectors" on the next page
  - ESX Host CPU Collector
  - ESX Host Memory Collector
  - ESX Host Physical Disk Collector

Physical Disk Collectors

The Physical Disk Collectors (Configuration > Monitoring Settings > Collectors) for hosts, comprise metrics to measure the performance of the disk drives on a host.

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Common Use
Disk IO Rate		
Disk Read (KBytes/Sec)	Average time in seconds to read data from disk	Compare read times for a given application (for example, read compared to writes).
Disk Total (KBytes/Sec)	Total read and write requests in seconds	Test maximum throughput.
Disk Write (KBytes/Sec)	Average time in seconds to write data to disk	Compare write times for a given application (for example, writes compared to reads).
Disk Utilization Percent		

Metric	Description	Common Use
Disk Utilization (%)	Based on the IRP (I/O request packets) round trip times the Average Disk Sec/Transfer. Indicates how busy a physical disk is over time.	Determine the average disk utilization for a given application or known number of processes. Utilization indicates how busy a disk is.

Disk performance metrics might not be available for all supported operating systems.

ESX Server Performance Collectors

The ESX host performance collectors (Configuration > Monitoring Settings > Collectors), comprise the following collectors to measure the performance of the disk drives, memory and CPU utilization of an ESX server.

- ESX Host CPU Collector
- ESX Host Memory Collector
- ESX Host Physical Disk Collector

The performance metrics are grouped into the following tabs of the Analysis pane:

Metric	Description	Common Use
CPU		
CPU Utilization (%)	The percentage of total CPU utilization for all processes running on a host.	Identify CPU bottlenecks.
Memory		
Free Physical Memory (KBytes)	Amount of physical memory available.	Measure available main memory for additional processes and threads.
Used Physical Memory (%)	Percentage of physical memory being consumed by all processes.	Indicates physical memory optimization and availability over a period of time.
Disk IO Rate		

Metric	Description	Common Use
Disk Read (KBytes/Sec)	Average time in seconds to read data from disk.	Compare read times for disks on an ESX Server.
Disk Total (KBytes/Sec)	Total read/writes in seconds.	Test maximum throughput.
Disk Write (KBytes/Sec)	Average time in seconds to write data to the disk.	Compare write times for disks on an ESX Server.

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 performance information at a switch level as well as port 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 following tabs of the Analysis pane:

Metric	Description	Common Use
Switch level metrics - Data Rate Tab		
Aggregated Port Bytes Received (MBytes/Sec)	Sum of bytes received for all ports in a switch over an interval	Measure inbound traffic for all ports on the switch.
Aggregated Port Bytes Transmitted (MBytes/Sec)	Sum of bytes transmitted for all ports in a switch over an interval	Measure outbound traffic for all ports on the switch.
Port Metrics		
Communication Tab		

Metric	Description	Common Use
CRC Errors	Number of Cyclic Redundancy Check errors over a period of time	Isolate CRC errors on a specific initiator or between devices
Link Failures	Number of link Failures over a period of time	Isolate connection failures and the effect on performance
Data Rate Tab		
Bytes Received (MBytes/Sec)	Number of bytes received over a given interval	Measure inbound traffic for specific ports on the switch.
Bytes Transmitted (MBytes/Sec)	Number of bytes transmitted over a given interval	Measure outbound traffic for specific ports on the switch.

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
  - Disconnects from FC-attached storage.
  - 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 with "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 HP 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 228
- "3PAR SMI-S Volume Collector" on page 230
- "3PAR SMI-S Physical Disk Collector" on page 232
- "3PAR SMI-S Fiber Channel Port Collector" on page 233

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		
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		

Metric	Description	Formula
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		
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		

Metric	Description	Formula
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
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 HP 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

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		
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 HP 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
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

Metric	Description	Formula
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
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

Metric	Description	Formula
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

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 HP 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 (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		

Metric	Description	Formula
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		
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 HP 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 HP 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		
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

Metric	Description	Formula
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		
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

Metric	Description	Formula
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
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

Metric	Description	Formula
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 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 (Req/Sec)	Number of write requests per second completed to a virtual disk that were received from all hosts	Δ WriteIOs / Δ Time
Read Rate (Req/Sec)	Rate at which data is read from each host port	Δ ReadIOs / Δ Time

Metric	Description	Formula
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

The following metrics track performance 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
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

Metric	Description	Formula
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)
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

Metric	Description	Formula
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		

Metric	Description	Formula
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)
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		

Metric	Description	Formula
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 Fiber Channel Port Collector

SOM monitors the following 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.	–

Metric	Description	Formula
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. Indicator 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 Characters (count)	Number of bad receive characters in the bit stream. Use to determine the number of bad frames associated with the Bad CRC metric above.	–
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.	–

Metric	Description	Formula
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 in 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
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 249
- "EMC Symmetrix DMX SMI-S Volume Collector" on page 251
- "EMC Symmetrix DMX SMI-S Fibre Channel Port Collector" on page 256

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

Metric	Description	Formula
Write Flush Data Rate (Bytes/Sec)	Number of tracks written per sec 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 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 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)

Metric	Description	Formula
Percent Writes (%)	Ratio of write I/Os to total I/Os.	PctWriteIOs = 100 x (deltaTotalWriteIOsRandomAndSeq / deltaTotalIOsRandomAndSeq)
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 (Req/Sec)	I/O rate which includes random and sequential reads and writes.	TotalIORate = deltaTotalIOsRandomAndSeq / 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

Metric	Description	Formula
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
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 each second that 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 each second that 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		
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

Metric	Description	Formula
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 Size (Bytes)	Average size of an I/O operation performed by the Symmetrix device.	AvgIOSize = (deltaKBytesTransferred x 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)

Metric	Description	Formula
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)
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

Metric	Description	Formula
Read Hit Rate Random (Req/Sec)	Number of random read hit I/O operations performed each second by the Symmetrix 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.	ReadHitRateRandom = deltaReadHitIOs / 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

Metric	Description	Formula
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
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 259
- "EMC CLARiiON and VNX SMI-S Volume Collector" on page 261
- "CLARiiON and VNX SMI-S Physical Disk Collector" on page 263
- "EMC CLARiiON and VNX SMI-S FrontEnd Port Collector" on page 265

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 / deltaTotalIOsRandomAndSeq)

Metric	Description	Formula
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

Metric	Description	Formula
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 Reads (%)	Percentage of CPU time dedicated to reads.	PctReadIOs = 100 x (deltaReadIOs / 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

Metric	Description	Formula
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

Metric	Description	Formula
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)
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

Metric	Description	Formula
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		
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

Metric	Description	Formula
Total Data Rate (Bytes/Sec)	Number of Bytes transferred through the CLARiiON/VNX Director each 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 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
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		
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

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. A direct line indicates that the path is known and discovered between the devices. A dotted line indicates that the relation between the devices is not directly discovered but is computed.

To view the topology map of your entire storage infrastructure, from the workspace navigation panel, click **Topology** > **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.

Note: If you are a new user, you might not be able view the System Topology.

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** on the toolbar 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) device topology, the connectivity is established between the NAS client hosts discovered in SOM and the NAS device.

If the NAS device uses devices from the storage network, then the NAS device will be shown as connected to the Fabric as explained above.

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. Select a particular fabric to see the connectivity among the elements within the fabric or select **Show All** to view the connectivity among all the discovered elements. By default, System Topology displays the topology of the topmost Fabric.
- **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, only the connectivity between the hosts and switches is displayed.

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.
	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
Port WWN	The unique 64-bit World Wide Name identifier of the port.
Port Speed	The port speed in Gbps.
Port State	Indicates the state of the port.
Port Type	Indicates the type of switch port. For example, F, E, FL, and so on.

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 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.

Host Topology

The Host Topology map displays an overview of the storage systems and switches connected to a selected host.

To navigate 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.

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.

Chapter 4: Common Tasks

This section describes procedures that are common to many HP Storage Operations Managerconfiguration and maintenance tasks. It includes the following topics:

- "Start or Stop SOM Services" below
- "Delete Elements" on the next page
- "Quarantine/Un-quarantine Elements" on page 274
- "Launch Topology" on page 275
- "Create an Asset Record" on page 276

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

To start SOM services

- **Windows**: Open the **Services** control panel. In the list of services, right-click each of the services, and then click **Start**.
- **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: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Stop**.

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)

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 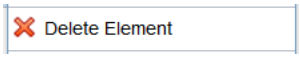

. The delete confirmation message is displayed. Click **OK** to delete the element.

Quarantine/Un-quarantine Elements

Quarantine is the state of an element where the data collection for the element is stopped. Elements for which data collection fails 3 times consecutively goes to a 'Quarantined' state.

Use the **Quarantine** option from the Inventory views to prevent data collection for the element. You can consider placing an element under quarantine in the following situations:

- Repeated data collection failures on an element.
- Element is under maintenance for firmware/hardware/software upgrade.

To quarantine an element:

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 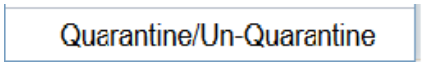

. The quarantine confirmation message is displayed. Click **OK** to quarantine the element.

Un-quarantine an Element

You must un-quarantine an element to resume data collection for the element. When you un-quarantine an element data collection for the element is triggered immediately.

To un-quarantine an element:

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 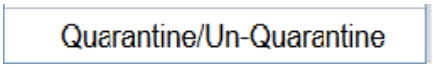

.

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.
Offering	Type of offering.
Vendor	The company that supplied the element.

Attributes	Description
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.

Appendix A: Inventory Views Tabs and Forms

This appendix contains the tabs and forms referenced in this guide.

Block Storage Systems View

Block storage systems display the following tabs:

- Storage System Processors
- Volumes
- Disk Drives
- Pools
- Ports
- Host Security Groups
- Storage Extents
- Replication Pairs
- Backend Storage
- SCSI Controllers
- Pools Logical Usage
- Masked Hosts
- Thin Provisioning Data
- Asset Record

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" on page 220.

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.

The list of pools and details of the used and unused raw space are available in the "Storage Systems View: Pools Tab" on page 316.

Double-click a pool to see its details and associated volumes and storage extents in the "Storage Pool Form" on page 292.

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.

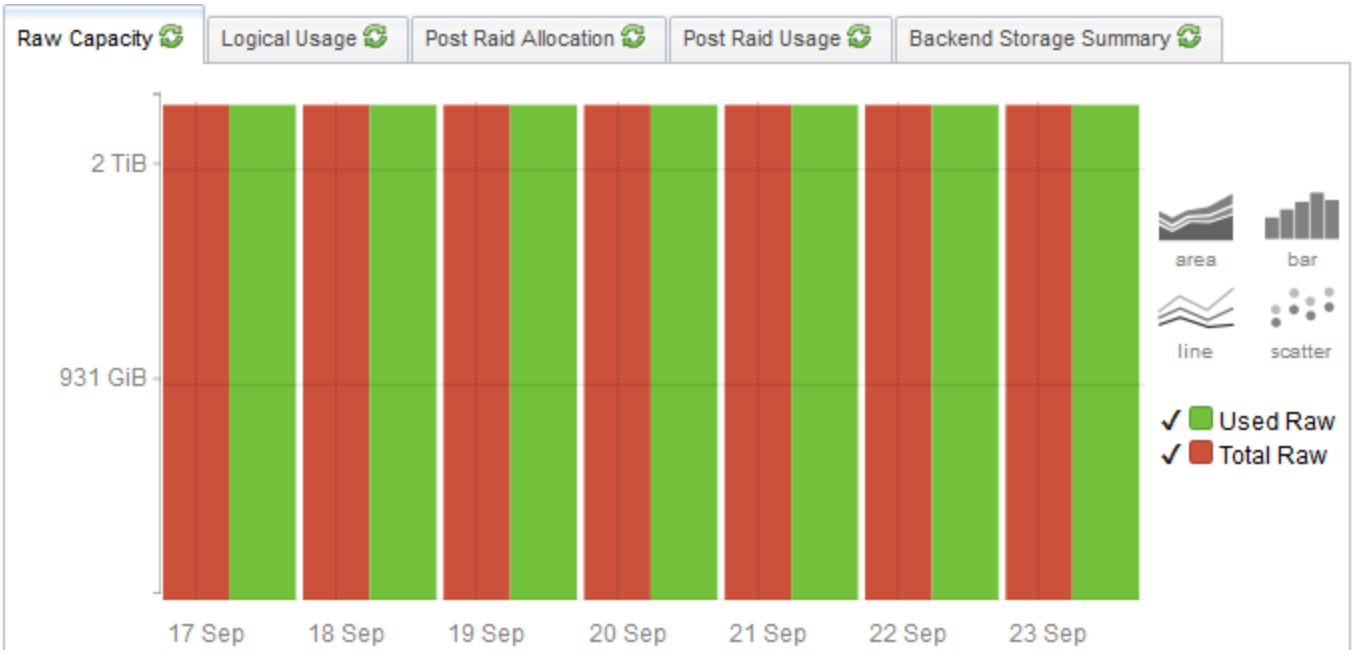

Logical Usage

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 Usage Tab" on page 324.

Double-click a storage pool to see its details and associated volumes and storage extents in the "Storage Pool Form" on page 292.

Example

The total (mapped) space visible to hosts is 294.45 GiB from the 1021.08 GiB space allocated (mapped + unmapped) post raid.

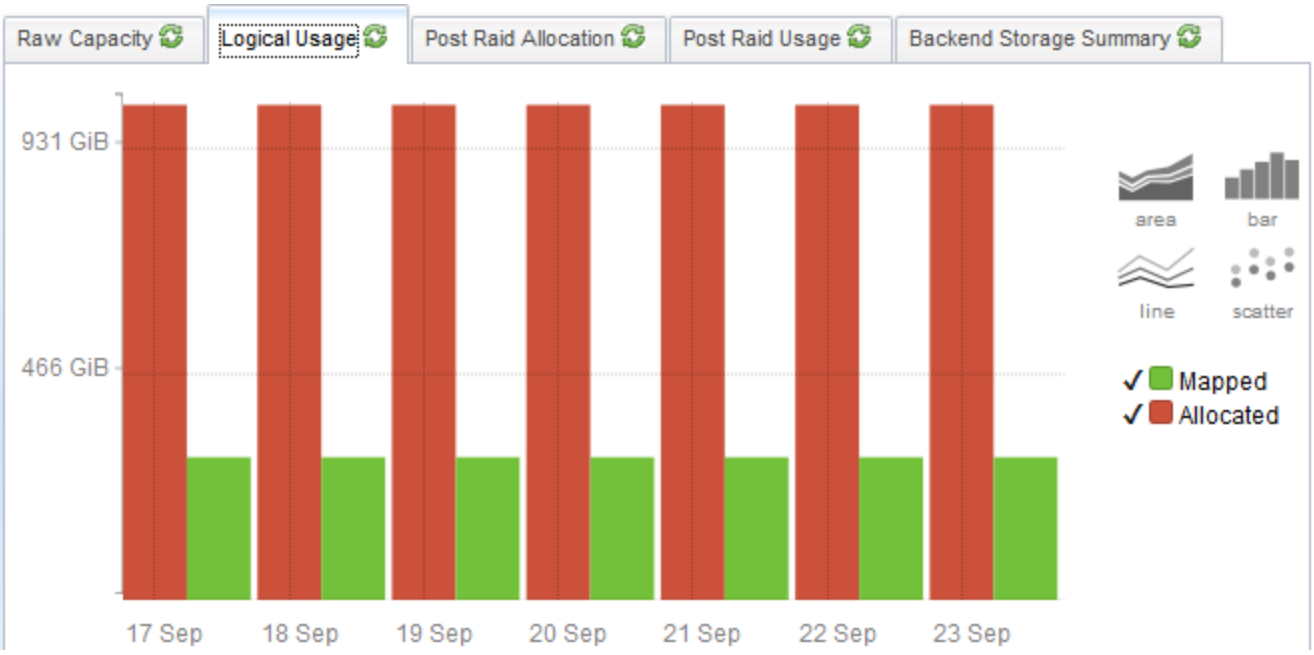

Post RAID Allocation

The Post Raid 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 325.

Double-click a storage pool to see its details and associated volumes and storage extents in the "Storage Pool Form" on page 292.

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.

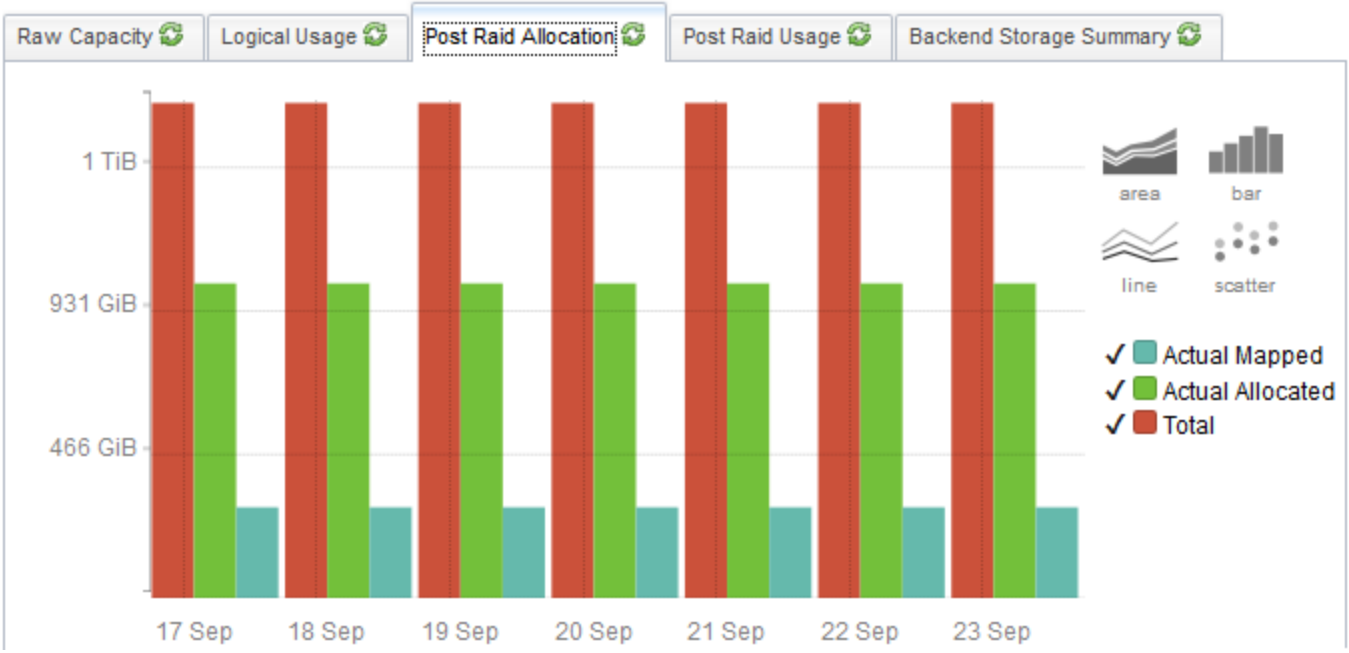

Post RAID Usage

The Post Raid 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.

Details of the capacity utilization of the individual physical pools are available in the "Storage Systems View: Thin Provisioning Data Tab" on page 325.

Double-click a storage pool to see its details and associated volumes and storage extents in the "Storage Pool Form" on page 292.

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.

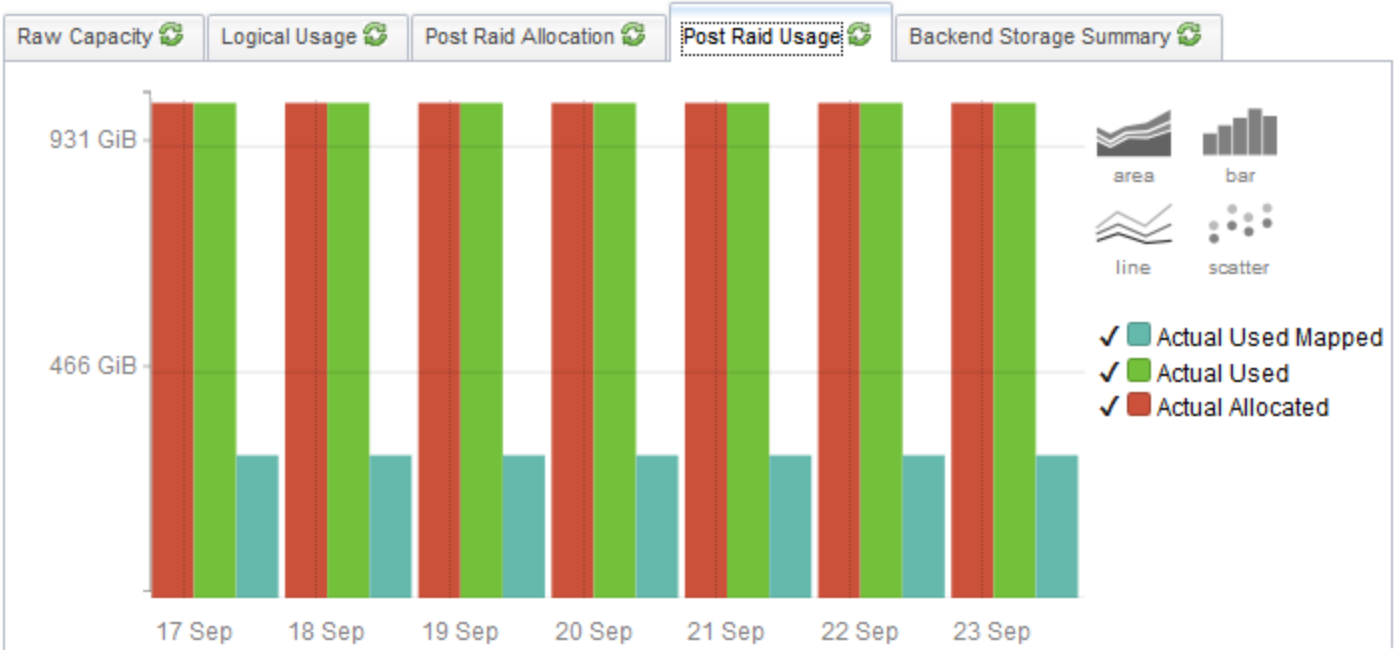

Backend Storage Summary

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 292.

- 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 293.

- 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

- Snapshots
- Shares
- Qtrees
- Quotas
- NAS Extents
- NAS Replication Pairs
- Volumes
- Disk Drives
- Initiator Groups
- NAS Network Interface
- Asset Record
- Ports
- CheckPoints

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" on page 220.

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**

  Note: This tab displays the raw capacity of NetApp cluster nodes when selected in the form view (Component Storage Systems tab) of the cluster.

- **Raw Capacity**

  The physical capacity of a NAS device using the following metrics:
  - **Used Capacity**
  - **Total Capacity**

  Note: 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 displayed.

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.

Forms

The SOM console includes forms for the following categories:

- "Host Forms" below
- "Switch Forms" on page 289
- "Storage System Forms" on page 291
- "Fabric Forms" on page 297
- "Node Forms" on page 299
- "Node Group Forms" on page 300

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.

General Ports	Description
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
- 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.

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

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 323
- "Storage Systems View: File Systems Tab" on page 328

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

Attribute	Description
Record Created	The time when the zone alias was first contacted.
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 336
- **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.
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.
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, HP, 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.

Attribute	Description
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.

Attribute	Description
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. 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.

Attribute	Description
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 **HP SOM Manager** implies that SOM created the object.

Caution: Each time a SOM upgrade is installed, objects with an Author attribute value of HP 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 HP.

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 309
- "Storage System Tabs" on page 314
- "Fabric Tabs" on page 335
- "Node Tabs" on page 337
- "Node Group Tabs" on page 339

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.

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 VirtualCenter or through the individual ESX Servers. Discovering the VirtualCenter results in one access point for all the ESX Servers managed by that VirtualCenter.

If you discover the VirtualCenter, and you also discover an individual ESX Server that is managed by the VirtualCenter, the ESX Server will have a separate access point and is not included in the list of ESX Servers associated with the VirtualCenter.

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.

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.
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.

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.

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.

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).

The **Analysis** pane displays the summary 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.

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: 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.

Attribute	Description
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, 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 Switch Port	The fibre-channel port of the switch connected to the front-end storage system. 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.

Attribute	Description
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: 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 internal SCSI controllers of a selected storage system.

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: 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).

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 Usage 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.

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.

Attribute	Description
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. 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.

Attribute	Description
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: 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.
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.

Attribute	Description
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.

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 198.

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 298.

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 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.

We appreciate your 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, March 2015 (Storage Operations Manager 10.00)

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@hp.com.