HP Storage Operations Manager

Software Version: 10.00
Windows® and Linux® operating systems

Deployment 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

Chapter 1: About this Guide

This guide contains a collection of information and best practices for administering SOM. This guide is for an expert system administrator or HP support engineer with experience in deploying and managing SOM installations. Read this guide before you start installing SOM.

Note: This document is updated as new information becomes available. 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/group/softwaresupport

For more information, see "Documentation Updates" on page 3.

Chapter 2: Planning a SOM Deployment

Planning the deployment is a critical activity to ensure that the SOM server can manage your storage environment effectively. Use the following guidelines for planning a successful deployment of SOM in your environment:

- **Sizing SOM and SHR servers**
  The size of the environment you want to manage decides how the servers should be sized and configured. To decide the SOMserver configuration that is suitable for your environment, see the "Performance and Sizing for the SOM Management Server" in the *SOM Support Matrix*.

- **Gather the system pre-requisites**
  Ensure that all the system pre-requisites are met before attempting to install SOM. Not meeting system requirements could result in an installation failure. For information about installation pre-requisites, see "Planning for Installation" in the *SOM Interactive Installation Guide*.

- **Check the firewall port configuration**
  SOM server uses various ports for communicating with the managed environment, the browser clients and the SHR Report server. The port configuration is largely decided by the proxy configuration of the managed devices. Ensure that the required configurations are enabled in the firewall configuration before starting the product installation. This will remove delays in discovering the managed environment once the product is deployed. For port configuration details, see "Ports and Firewall " on page 46.

- **Plan Tenancy**
  If you plan to have multiple tenants in your environment, it is a good idea to configure tenants before you start discovering your environment. You can associate the tenants to a discovery address. Elements discovered via this discovery address will automatically get associated with the configured tenant. It is easier to configure tenants before discovery as compared to moving elements to tenants after discovery. For information about planning tenants, see "Recommendations for Planning Tenants" on page 78.

- **Plan your Node Groups**
  In SOM, many of the management primitives like data collection and monitoring policies are applied to elements that are grouped instead of applying to individual elements. This means that group definitions can be created in advance, and then the discovery process will distribute the elements across different groups accordingly. For information about node group creation, see "Recommendations for Planning Node Groups" on page 56.

- **Determine the Data Collection Policies**
  Data collection policies can be pre-defined and applied to node groups before discovering your environment. After the elements are discovered, they will be categorized into the groups as described in "Planning your Node Groups" and then appropriate policies are applied to them. For example, if you follow a convention of naming all windows hosts as 'win*', then you can create a group definition based on that and apply a data collection policy with a custom freshness as required. This is a onetime definition and thereafter, as your windows hosts get discovered, they are already covered by data collection policies without additional administration overheads.

  Also,you may want to set the level of data to collect for elements in your environment. By default, the system is configured NOT to collect all data on all devices in the environment. If you want to collect deeper data for a set of devices, you can configure this using the Data Collection Control feature. By planning this in advance you can avoid additional data collection cycles to get more data.

  For information about best practices on data collection policies, see "Recommendations for Data Collection Policies" on page 64

- **Configure Performance Monitoring**
  As a rule of thumb, wait for one data collection to complete in your environment for all (or a majority) of devices before configuring monitoring policies. You can refer to the Collection Status Dashboard to monitor the number of running collections. In large environments, configuring monitoring policies can overlap with the data collection process and might result in missing statistics. For information about best practices on creating monitoring policies, "Recommendations for Monitoring Performance " on page 65.

- **Decide your host discovery strategy – Rule-based inference/Agentless/Agent**
  Host discovery needs to be planned carefully, since they add the maximum bulk in terms of the number of elements discovered. SOM uses the following mechanisms to discover hosts in your environment:

  a. Rule-based inference – Use the configurations in your environment like Zones, Zone Aliases, and Host Security Groups to understand the distribution of storage to the hosts.

  b. Agentless discovery of hosts – Discover hosts without deploying agents on them by using mechanisms like WMI, SSH, or native API (for example, VMWare).

  c. Discovery by deploying agents on hosts – Deploy agents on hosts to discover hosts.

  Typically, agent deployment incurs administrative overheads. At the same time, agents provide the maximum depth of information on a host.

To reduce this administrative overhead, you can use a combination of the methods listed above. Rule-based inference can be configured in your environment as soon as you have discovered your switches (fabric) and storage systems. Use the presented storage views and reports to understand the storage distribution in your environment. Once you understand the distribution of storage in your environment, you will be able to identify the top hosts that are consuming storage in your environment and then decide to discover them using agentless or agent based mechanism for further analysis.

The choice between agentless discovery or agent-based discovery is driven by the depth of information you require on the host.

For information on details about host discovery, see"Host Discovery" on page 59.

- **Configure the SOM reporting server**
  Service Health Reporter (SHR) is the reporting engine for SOM and needs to be installed on a separate server. Ensure that the ports for communication between SOM and SHR are available. For information about the required port, see "Ports and Firewall " on page 46.

  For information about sizing the SHR server to suit your environment, see "Performance and Sizing for the SOM Reporting Server" in the *SOM Support Matrix.*

  After installing the SOMreporting server and deploying the SOM content packs you must configure certificates for enabling file transfer between the SOM server and the reporting server. SOM server ., see "Configuring Connections Between SOM Management Server and SOM Reporting Server" in the *Storage Resource Management Reports Guide*.

Chapter 3: 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.

Obtaining and Installing 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

Viewing 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	

Chapter 4: CIM Extensions

The Common Information Model (CIM) standard specifies a structure of information about managed elements. CIM provides for consistent data structure and access regardless of device vendor. CIM is maintained by the Distributed Management Task Force (DMTF).

The Storage Management Initiative Specification (SMI-S) enables consistent management of heterogeneous storage elements. SMI-S is based on the Common Information Model (CIM) and Web-Based Enterprise Management (WBEM) standards for accessing management information over HTTP. SMI-S is maintained by the Storage Networking Industry Assocation (SNIA).

A SOM CIM extension is a collection agent that runs on a storage host to gather information about that host. The SOM management server communicates with the CIM extension while discovering and managing the host.

For the SOM management server to obtain information from the host, the CIM extension must be running. The CIM extension starts automatically after installation and whenever the host boots. On an HP-UX host, the CIM extension uses `/sbin/rc2.d` scripts.

SOM can manage some storage hosts using an agentless process. The agentless approach, however, limits the information that is available to SOM. For more information, see the *SOM Device Support Matrix*.

The default location of the CIM extension is:

- *Windows*: `<Drive:>\Program Files (x86)\APPQcime\CimExtensions`
- *UNIX or Linux*: `/opt/APPQcime/`

Installing CIM Extensions

A CIM extension communicates with a host bus adapter (HBA) using the fibre channel host bus adapter application programming interface (FC-HBA API) created by the Storage Network Industry Association (SNIA). The SOM management server only supports communication with HBAs that are compliant with the HBA API. For more information about the HBA API, see the SNIA web page: **http://www.snia.org**.

The `hbatest` program on the SOM installation media outputs the name and number of all HBAs on the host that support the FC-HBA API. In some instances, `hbatest` might report that it cannot find

an HBA driver even though an HBA driver is installed. In this case, try installing a different, SNIA-compliant version of the HBA driver.

The SOM installation media includes operating system-specific CIM extensions in the `CIMExtensionsCD1` directory.

To install a CIM extension

1. Verify that at least one host bus adapter (HBA) on the host supports the FC-HBA API. Follow the procedures that apply to your environment:
   - "Verify FC-HBA API Support on a Windows Host" below
   - "Verify FC-HBA API Support on an HP-UX Host" on the next page
   - "Verify FC-HBA API Support on a Linux Host" on page 21
2. Verify that port 4673 is available on the host and reachable by the SOM management server.

   Alternatively, identify a different port for the CIM extension. After installation, configure the CIM extension to use that port as described in "Change the CIM Extension Port Number" on page 29.
3. Install the CIM extension software on the host. Follow the procedures that apply to your environment:
   - "Install the CIM Extension Software on a Windows Host" on page 22
   - "Install the CIM Extension Software on an HP-UX Host" on page 23
   - "Install the CIM Extension Software on a Linux Host" on page 24

   Tip: If your security environment requires that you customize the CIM extensions or the CIM extension installation process, you might need to use a third-party tool to deploy CIM extensions. Third-party tools are commonly used in large environments that require the use of a request for change (RFC) process.

Verify FC-HBA API Support on a Windows Host

To verify that at least one host bus adapter on a Windows host supports the FC-HBA API

1. In a command window, change to the `CimExtensionsCD1/Windows/tools` directory on the SOM installation media.
2. Enter the following command:

```
hbatest.exe -v
```

The beginning of the command output should be similar to the following example:

```
hbaapi.dll, version XXXXXXXXXXXXXX will be used to get HBA
information.
HBA API Library version is 2
hbatest build date: Jun 26 2014:20:14:26
Number of HBA's is 2
*****************************************
```

After the header, the command output lists each HBA present on the host.

Return to the installation procedure.

Verify FC-HBA API Support on an HP-UX Host

To verify that at least one host bus adapter on an HP-UX host supports the FC-HBA API

1. Go to the `CimExtensionsCD1/HPUX/tools` directory on the SOM installation media.
2. Run the following command:

   ```
   ./hbatest
   ```

   The program runs its diagnostics.

HP SNIA adapters AXXXXA come from fileset FC-FCD, FC-TACHYON-TL. Unless separated purposely during the installation of the operating system, filesets are there by default. To view the location of the library, enter the following at the command prompt:

```
more /etc/hba.conf
```

The `hba.conf` file includes the following lines:

```
com.hp.fcms32 /usr/lib/libhbaapihp.sl #32 bit lib names end in 32
com.hp.fcms64 /usr/lib/pa20_64/libhbaapihp.sl #64 bit lib names
end in 64
com.hp.fcd32 /usr/lib/libhbaapifcd.sl
com.hp.fcd64 /usr/lib/pa20_64/libhbaapifcd.sl
```

Return to the installation procedure.

Verify FC-HBA API Support on a Linux Host

To verify that at least one host bus adapter on a Linux host supports the FC-HBA API

1. Go to the `CimExtensionsCD1/linux/tools` directory on the SOM installation media.
2. Run the following command:

   ```
   ./hbatest
   ```

   The program runs its diagnostics.

Driver Information for Verifying Emulex SNIA Adapters (Red Hat Linux Only)

The Emulex driver does not contain the library that is required by the SOM management server. You must install Emulex HBAnywhere software so that the management server can discover hosts configured with HBAnywhere and the HBATool can detect the Emulex host bus adapter.

After you install the HBAnywhere software, you can find the location of the libraries in the `/etc/hba.conf` file.

To view the `hba.conf` file on a Linux host, run the following command:

```
cat /etc/hba.conf
```

The output lists the library name and then the path, as shown in the following examples:

- Linux 64-bit host Emulex driver example output

  ```
  com.emulex.emulexapilibrary /usr/lib64/libemulexhbaapi.so
  com.emulex.emulexapilibrary /usr/lib/libemulexhbaapi.so
  ```

  Note: The HBAnywhere CLI must be used for IA64 Linux.

- Linux 32-bit host Emulex driver example output

  ```
  com.emulex.emulexapilibrary /usr/lib/libemulexhbaapi.so
  ```

Return to the installation procedure.

Install the CIM Extension Software on a Windows Host

You must have administrator privileges to install the the CIM extension on a Windows host.

If a firewall is enabled on the Windows host, open the CIM extension port before installing the CIM extension. The default CIM extension port is 4673. For information about configuring the Windows firewall, see the documentation for the Microsoft Windows operating system.

The Windows CIM extension can be installed interactively or in silent mode. Use silent mode to install the Windows CIM extension with the default settings and no user intervention.

Interactive Mode

To install the CIM extension using interactive mode

1. Log on to the Windows host as a user with administrator privileges.
2. Insert the SOM installation media into the DVD drive.
3. In Windows Explorer, change to the `CimExtensionsCD1\Windows` directory, and then double-click `InstallCIMExtensions.exe`.
4. Follow the instructions on the screen.

Silent Mode

To install the CIM extension using silent mode

1. Verify that no other programs are running.
2. Remove the previous version of the CIM extension as described in "Remove the CIM Extension from a Windows Host" on page 41.
3. Log on to the Windows host as a user with administrator privileges.
4. Insert the SOM installation media into the DVD drive.
5. In a command window, change to the following directory:

   ```
   CimExtensionsCD1\Windows
   ```

6. Enter the following command:

   ```
   InstallCIMExtensions.exe -i silent
   ```

Return to the installation procedure.

Install the CIM Extension Software on an HP-UX Host

The following instructions apply to a local installation of the CIM extension.

You must install the CIM extension for HP-UX to the default directory. If there are space issues, such as large CIM extension binary files, create a symbolic link to a folder with more space.

To install the CIM extension

1. Log on to the HP-UX host as the `root` user.
2. Insert the SOM installation media into the DVD drive.
3. Create the `/DVD` directory by running the following command:

   ```
   mkdir /DVD
   ```

4. Mount the SOM installation media by enter the following at the command prompt:

   ```
   mount /dev/dsk/c#t#d#   /DVD
   ```

   In this instance, the c, t, and d numbers correspond to DVD device numbers.

   To find out c#t#d# for your DVD drive, run the `ioscan -fnC disk` command on the HP-UX host.

5. Run the following command:

   ```
   swinstall -x mount_all_filesystems=false -s
   /cdrom/HPUX/APPQcime.depot APPQcime
   ```

   The installation is complete when a message similar to the following appears:

   ```
   analysis and execution succeeded
   ```

6. Unmount the DVD by running the following command:

```
umount /DVD
```

In this instance, `/DVD` is the name of the directory where you mounted the DVD.

Return to the installation procedure.

Install the CIM Extension Software on a Linux Host

The following instructions apply to a local installation of the CIM extension.

The installation is a two-step process where a "requires" rpm is run to check for dependencies, and then the full rpm is installed.

You must install the CIM extension for Linux to the default directory. If there are space issues, such as large CIM extension binary files, create a symbolic link to a folder with more space.

To install the CIM extension

1. Log on to the Linux host as the `root` user.
2. Insert the SOM installation media into the DVD drive.
3. Change to the `CIMExtensionCD1/linux/requires_rpm` directory on the SOM installation media.

   ```
   cd /DVD/linux/requires_rpm
   ```

   In this instance, `/DVD` is the name of the DVD drive.
4. When running the "required" rpm returns only the one expected dependency error, run the following command:

   ```
   rpm -idvh <rpm_package_name>
   ```

   In this instance `<rpm_package_name>` is the name of the rpm package listed in the following table.

Operating System	RPM
64-bit Red Hat versions 6 and later	APPQcime-<Version>-<Release>-x86_64.rpm

Operating System	RPM
■ Red Hat 32-bit installations on x86 ■ 64-bit installations earlier than Red Hat version 6 ■ SUSE installations on x86 or x64	APPQcime-<Version>-<Release>-i386.rpm
(Red Hat and SUSE Linux) IA64-based installations	APPQcime-<Version>-<Release>-ia64.rpm

The following output is displayed:

```
Preparing... ########################################### [100%]

1:APPQcime ########################################### [100%]
```

The installation is done when you are returned to the command prompt.

5. *Optional*. Verify that the packages were installed:

```
rpm -qa | grep APPQcime-Requires
rpm -qa | grep APPQcime
```

Return to the installation procedure.

Configuring a CIM Extension

The `cim.extension.parameters` file determines the CIM extension behavior. The CIM extension reads this file at startup.

The `cim.extension.parameters-sample` file provides a template configuration.

These files are located in the following directory:

- *Windows*: `[Installation_Directory]\CimExtensions\conf`
- *UNIX/Linux*: `/opt/APPQcime/conf`

The default behavior of a CIM extension is as follows:

- The SOM management server must use the administrator or root account on the host for communications with the CIM extension.
- The CIM extension sends and receives communications on port 4673.
- The CIM extension listens on the loopback address of the host.

To change this behavior, create the `cim.extension.parameters` file by copying and customizing the provided template file (`cim.extension.parameters-sample`).

To configure the CIM extension

1. Log on to the host as a user with administrator or root privileges.
2. Change to the CIM extension configuration directory:
   - *Windows*: `[Installation_Directory]\CimExtensions\conf`
   - *UNIX/Linux*: `/opt/APPQcime/conf`
3. Save a copy of the `cim.extension.parameters-sample` file as `cim.extension.parameters` in the same directory.
4. In a text editor, edit the file `cim.extension.parameters` file as required.

   For information about commonly changed parameters, see the table of CIM extension parameters.

   For information about configuring log files, see "Log File Properties" on page 37.
5. Save and close the file.
6. Restart the CIM extension.
   - *Windows*:

     Restart the AppStorWin32Agent service from the **Services** window or reboot the host.
   - *UNIX/Linux*:

     ```
     /opt/APPQcime/tools/stop
     /opt/APPQcime/tools/start
     ```

Commonly Configured CIM Extension Parameters

Parameter	Description

`-users`	Restricts the discovery of the host to a list of valid host users. Each user defined in this parameter must be a valid existing user on the host, and the user name must match one of the user names used on the discovery page to discover the host for authentication to occur. The user does not need to have root authority. Use a colon (`:`) to separate multiple users.
	
	The format of the user name depends on the operating system:
	
	• *Windows*: Specify the domain name and the user name, for example:
	
	`-users domain_name\user_name`
	
	• *UNIX*: Specify the user name without the domain name, for example:
	
	`-users user_name`
	
	For more information, see "Restrict the Users Who Can Discover the Host" on the next page.

`-credentials <username>:<password>`	Specifies a user name and password on the host to facilitate communication between the SOM management server and the managed host. This configuration eliminates the need to use the local operating system user/password database for credential verification. This user name / password pair is known only to the CIM extension and does not identify a real user on the host. The specified account name might not exist on the host. The `-users` parameter always takes precedence over the `-credentials` parameter. To use the `-credentials` parameter when the `-users` parameter has been added to the `cim.extension.parameters` file, comment out the `-users` parameter by inserting the number sign character (`#`) at the beginning of the `-users` line.

Parameter	Description
`-mgmtServerIP <ip address>`	Restricts the CIM extension to listen only to the specified SOM management servers. Use commas to separate multiple address values. For example: `-mgmtServerIP 127.0.0.1,192.168.0.1`
`-port <new port>`	Specifies the port that the CIM extension accesses. For example: `-port 1234` See "Change the CIM Extension Port Number" on the next page.

`-on <ip address1>`
`-on <ip address2:port>`
	For multi-homed systems, restricts the CIM Extension to listen only on designated IP address.
	Use multiple entries for multiple addresses. For example:
	`-on <15.218.125.12>`
	`-on <15.218.125.123:5432>`
	See "Configure the CIM Extension to Listen on a Specific IP Address" on page 30.

UNIX/Linux only. For command line help about CIM extension configuration, run the following command:

```
/opt/APPQcime/tools/start -help
```

Restrict the Users Who Can Discover the Host

The `-users` parameter increases security by restricting access to the CIM extension. When you use the SOM management server to discover the host, provide one of the user names that was specified in the `-users` parameter.

To use the management server to discover a host without using the root account, provide the password to another valid user account with fewer privileges on the host.

First, add the user to the parameters file. Next, log on to the management server, access the Discovery page, and provide the user name and password for jsmythe. Only the user name and password for jsmythe can be used to discover the host.

To add a user to the parameters file

1. Back up the CIM extension configuration directory to a location outside the CIM extension installation directory:
   - *Windows*: `[Installation_Directory]\CimExtensions\conf`
   - *UNIX/Linux*: `/opt/APPQcime/conf`
2. In a text editor, open the `cim.extension.parameters` file.
3. Add the following line:

   ```
   -users myname
   ```

   In this instance, myname is a valid user name on the host.

   To enter multiple users, separate them with a colon; for example, `-users myname:jsymthe`.
4. Save the file.
5. Restart the CIM extension.
   - *Windows*:

     Restart the AppStorWin32Agent service from the **Services** window or reboot the host.
   - *UNIX/Linux*:

     ```
     /opt/APPQcime/tools/stop
     /opt/APPQcime/tools/start
     ```

Change the CIM Extension Port Number

By default, the CIM extension uses port 4673. If this port is already in use, change the CIM extension port as follows:

1. Back up the CIM extension configuration directory to a location outside the CIM extension installation directory:
   - *Windows*: `[Installation_Directory]\CimExtensions\conf`
   - *UNIX/Linux*: `/opt/APPQcime/conf`
2. In a text editor, open the `cim.extension.parameters` file.

3. Add the following line:

   ```
   -port <port_number>
   ```

   Replace *<port_number>* with the port number to use.

4. Save the file.

5. Restart the CIM extension.

   - *Windows*:

     Restart the AppStorWin32Agent service from the **Services** window or reboot the host.

   - *UNIX/Linux*:

     ```
     /opt/APPQcime/tools/stop
     /opt/APPQcime/tools/start
     ```

6. Update the SOM management server with the new port number for this host.

   a. Open the **Discovery Addresses** form (**Configuration > Discovery > Discovery Addresses**) for this host.

   b. In the **IP Address** box, enter the IP address with a colon followed by the new port number. For example:

      ```
      192.168.1.2:1234
      ```

      In this instance, 192.168.1.2 is the IP address of the host, and 1234 is the new port number.

      If you already added the host to the discovery list (**Configuration > Discovery > Discovery Address**) on the management server, you must remove it and then re-add it. You cannot have more than one listing of the host with different ports.

Configure the CIM Extension to Listen on a Specific IP Address

To configure the CIM extension to listen on a specific IP address

1. Back up the CIM extension configuration directory to a location outside the CIM extension installation directory:

- *Windows*: `[Installation_Directory]\CimExtensions\conf`
- *UNIX/Linux*: `/opt/APPQcime/conf`

2. In a text editor, open the `cim.extension.parameters` file.
3. For each IP address to listen on, add the following line:

   ```
   -on <IP_address>
   ```

   Replace *<IP_address>* with one IP address. Optionally, add a port. For example, to listen on port 3456 of IP address 192.168.2.2, use the following text:

   ```
   -on 192.168.2.2:3456
   ```

4. Save the file.
5. Restart the CIM extension.
   - *Windows*:

     Restart the AppStorWin32Agent service from the **Services** window or reboot the host.

   - *UNIX/Linux*:

     ```
     /opt/APPQcime/tools/stop
     /opt/APPQcime/tools/start
     ```

6. Update the SOM management server with the new port number for this host.
   a. Open the **Discovery Addresses** form (**Configuration > Discovery > Discovery Addresses**) for this host.
   b. In the **IP Address** box, enter the IP address with a colon followed by the new port number. For example:

      ```
      192.168.1.2:1234
      ```

      In this instance, 192.168.1.2 is the IP address of the host, and 1234 is the new port number.

      If you already added the host to the discovery list (**Configuration > Discovery > Discovery Address**) on the management server, you must remove it and then re-add it. You cannot have more than one listing of the host with different ports.

Configuring CIM Extensions to Run Behind Firewalls (UNIX Only)

To discover a host behind a firewall, use the following table as a guideline. Assume the management server wants to discover HostA, which has three network interface cards on three separate networks with three separate IP addresses: 10.250.250.10, 172.31.250.10, and 192.168.250.10. The following table presents configuration options.

- The “Manual Start Parameters for CIM Extensions” column provides the values you would enter to start the CIM extension manually on the host. For more information about how to start a CIM extension manually, see "Starting a CIM Extension Manually" on page 39.
- The “If Mentioned in cim.extension.parameters” column provides information about modifying the `cim.extension.parameters` file (see "Change the CIM Extension Port Number" on page 29).
- The “Step 1 Discovery and RMI Registry Port” column provides information about the IP addresses that are required for the discovery list. The CIM extension uses the RMI registry port. When a port other than 4673 is used for the CIM extension, the port must be included in the discovery IP address; for example, 192.168.1.1:1234. In this instance, 192.168.1.1 is the IP address of the host, and 1234 is the port the CIM extension uses.

Troubleshooting Firewalls

Configuration	Manual Start Parameters for CIM Extension	If mentioned in cim.extension.parameters	Step 1 Discovery and RMI Registry Port
Firewall port 4673 opened between host and management server.	start		10.250.250.10 *OR* 172.31.250.10 *OR* 192.168.250.10 Communication Port: 4673

Troubleshooting Firewalls, continued

Configuration	Manual Start Parameters for CIM Extension	If mentioned in cim.extension.parameters	Step 1 Discovery and RMI Registry Port
Firewall port 1234 opened between host and management server.	start -port 1234	-port 1234	10.250.250.10:1234 *OR* 172.31.250.10:1234 *OR* 192.168.250.10:1234 Communication Port: 1234
Firewall port 4673 opened between host and management server on the 172.31.250.x subnet.	start -on 172.31.250.10	-on 172.31.250.10	172.31.250.10 Communication Port: 4673
Firewall port 1234 opened between host and management server on the 192.168.250.x subnet.	start -on 192.168.250.10:1234	-on 172.31.250.10:1234	172.31.250.10:1234 Communication Port: 1234

Troubleshooting Firewalls, continued

Configuration	Manual Start Parameters for CIM Extension	If mentioned in cim.extension.parameters	Step 1 Discovery and RMI Registry Port
With 3 firewall ports opened on different ports respectively 1234, 5678, 9012.	start -on 10.250.250.10:1234 -on 172.31.250.10: 5678 -on 192.168.250.10: 9012	-on 10.250.250.10:1234 -on 172.31.250.10: 5678 -on 192.168.250.10: 9012	10.250.250.10:1234 *OR* 172.31.250.10:5678 *OR* 192.168.250.10:9012 Communication Port: 1234, 5678, 9012
With firewall port 4673 opened between host and management server. NAT environment, where 10.250.250.10 subnet is translated to 172.16.10.10 when it reaches the other side of the firewall.	start		172.16.10.10 Communication Port: 17001

Troubleshooting Firewalls, continued

Configuration	Manual Start Parameters for CIM Extension	If mentioned in cim.extension.parameters	Step 1 Discovery and RMI Registry Port
With firewall port 1234 opened between a host and management server. NAT environment, where 10.250.250.10 subnet is translated to 172.16.10.10 when it reaches the other side of the firewall.	start -port 1234	-port 1234	172.16.10.10 Communication Port: 17001
With 3 firewall ports opened on different ports respectively 1234, 5678, 9012. NAT environment, where all 3 NICs are translated to different 172.16.x.x subnets.	start -on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10:9012	-on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10:9012	172.16.10.10:1234 *OR* 172.16.20.20:5678 *OR* 172.16.30.30:9012 Communication Port: 1234, 5678, 9012

Troubleshooting Firewalls, continued

Configuration	Manual Start Parameters for CIM Extension	If mentioned in cim.extension.parameters	Step 1 Discovery and RMI Registry Port
False DNS or IP is slow to resolve.		jboss.properties, cimom.Dcxws.agency.firstwait=200000 cimom.Dcxws.agency.timeout=200000	Any IP that is reachable Communication Port: 4673
No DNS, never resolve.		jboss.properties cimom.Dcxws.agency.firstwait=200000 cimom.Dcxws.agency.timeout=200000	Any IP that is reachable Communication Port: 4673
No firewall. Discover with a non-existent user for security reasons.	start -credentials string1:string2 In this instance, string1 is supplied in discovery as the “username” and string2 is supplied as the “password”.	-credentials username:password	Specify username and password in the discovery list. Communication Port: 4673

Troubleshooting Firewalls, continued

Configuration	Manual Start Parameters for CIM Extension	If mentioned in cim.extension.parameters	Step 1 Discovery and RMI Registry Port
With 3 firewall ports opened on different ports, respectively 1234, 5678, 9012. Discover with a nonexistent user for security reasons.	start -on 10.250.250.10:1234 -on 172.31.250.10:5678 -on 192.168.250.10:9012 -credentials string1:string2 In this instance, string1 is supplied in discovery as the “username” and string2 is supplied as the “password”.	-on 10.250.250.10:1234 -on 172.31.250.10: 5678 -on 192.168.250.10: 9012 -credentials username:password	10.250.250.10:1234 *OR* 172.31.250.10:5678 OR 192.168.250.10:9012 Specify username and password in the discovery list. Communication Port: 1234, 5678, 9012

Log File Properties

The `cim.extension.parameters` file contains the following properties for each log file:

- `<log name>.log.File` – Sets the name and location of the log file.
- `<log name>.log.MaxFileSize` – Sets the maximum file size in MB.
- `<log name>.log.MaxBackupIndex` – Sets the maximum number of files created before the files are overwritten.

The default location of the CIM extension log files is:

- *Windows*: `[Installation_Directory]\CimExtensions\tools`
- *UNIX/Linux*: `/opt/APPQcime/tools`

Log files roll over upon reaching the configured size. Each log consists of a configured number of files.

For example, the `cxws.log` file collects most of the CIM extension logging information. The CIM extension appends start time, stop time, and unexpected error conditions to the existing `cxws.log` file. The default `cxws.log` file configuration in the `cim.extension.parameters` file is as follows:

```
-D cxws.log.File=cxws.log
-D cxws.log.MaxFileSize=30MB
-D cxws.log.MaxBackupIndex=3
```

By default, the `cxws.log` file rolls over each time it becomes larger than 30 MB. The `cxws.log` file is renamed `cxws.log.1`, and a new `cxws.log` file is created. When the `cxws.log` file rolls over again,`cxws.log.1` is renamed `cxws.log.2`, the `cxws.log` file is renamed `cxws.log.1`, and a new `cxws.log` file is created and so on for a maximum of three backup log files:

- `cxws.log`
- `cxws.log.1`
- `cxws.log.2`
- `cxws.log.3`

Finding the Version of a CIM Extension

To find the version number of a CIM extension

- *Windows*: In the **Programs and Features** control panel, examine the value in the **Status** column for the AppStorWin32Agent service.
- *UNIX or Linux*: Run the following command:

  `/opt/APPQcime/tools/status`

To find the version number of a CIM extension, run the following command:

`/opt/APPQcime/tools/start -version`

The output displays the CIM extension version number and build date. For example:

```
Starting CIM Extension for HP-UX
CXWS for mof/cxws/cxws-HPUX.mof
CXWS version x.x.x.x, built on Fri 12-March-xxxx 12:29:49 by
dmaltz
```

Checking the Status of a CIM Extension

To determine the status of a CIM extension

- *Windows*: In the **Services** window, examine the value in the **Status** column for the AppStorWin32Agent service.
- *UNIX or Linux*: Run the following command:

```
/opt/APPQcime/tools/status
```

Starting a CIM Extension Manually

The SOM management server can only gather information about a host when the installed CIM extension is running.

You must have administrator or root privileges to start a CIM extension. The CIM extension only provides the information within the privileges of the user account that started the CIM extension. Only administrator or root has enough privileges to provide the information the management server needs. If you do not start the CIM extension with administrator or root privileges, the management server display messages similar to the following:

```
Data is late or an error occurred.
```

To start a CIM extension

- *Windows*: Start the AppStorWin32Agent service from the **Services** window.
- *UNIX or Linux*: Run the following command:

```
/opt/APPQcime/tools/start
```

Tip: You can use any of the options in the table of CIM extension parameters when starting a CIM extension from the command line.

Stopping a CIM Extension

The management server can only gather information about a host when the installed CIM extension is running.

You must have administrator or root privileges to stop a CIM extension.

To stop a CIM extension

- *Windows*: Stop the AppStorWin32Agent service from the **Services** window.
- *UNIX or Linux*: Run the following command:

```
/opt/APPQcime/tools/stop
```

Customize JVM Settings for a CIM Extension

To customize the Java virtual machine (JVM) configuration for a CIM extension, create the `wrapper.user` file by copying and customizing the provided template file (`wrapper.user-sample`). Place the configuration file in the following directory:

- *Windows*: `[Installation_Directory]\CimExtensions\conf`
- *UNIX/Linux*: `/opt/APPQcime/conf`

The CIM extension retains and uses the customized `wrapper.user` file after each future upgrade of the CIM extension.

To configure a CIM extension JVM

1. Log on to the host as a user with administrator or root privileges.
2. Change to the CIM extension configuration directory:
   - *Windows*: `[Installation_Directory]\CimExtensions\conf`
   - *UNIX/Linux*: `/opt/APPQcime/conf`
3. Save a copy of the `wrapper.user-sample` file as `wrapper.user` in the same directory.

4. In a text editor, edit the file `wrapper.user` file according to the comments in the file.
5. Save and close the file.
6. Restart the CIM extension.
   - *Windows*:

     Restart the AppStorWin32Agent service from the **Services** window or reboot the host.
   - *UNIX/Linux*:

     ```
     /opt/APPQcime/tools/stop

     /opt/APPQcime/tools/start
     ```

Removing CIM Extensions

To remove a CIM extension from a host, follow the applicable procedure:

- "Remove the CIM Extension from a Windows Host" below
- "Remove the CIM Extension from an HP-UX Host" on the next page
- "Remove the CIM Extension from a Linux Host" on the next page

Remove the CIM Extension from a Windows Host

If you remove a CIM extension from a Windows host where there is a service that is using WMI (such as Microsoft Exchange), you will see a message that the WMI service could not be stopped. Continue with the removal of the CIM extension, and then reboot the host after the removal process completes.

To remove the CIM extension from a Windows host

1. Log on to the Windows host as a user with administrator privileges.
2. Open the **Programs and Features** or the **Add or Remove Programs** control panel.
3. In the list of installed programs, right-click **Windows CIM Extension**, and then click **Uninstall**.
4. Follow the instructions on the screen.

5. After the uninstaller completes, delete the CIM extension installation directory.

   The default location is:

   ```
   <Drive:>\Program Files (x86)\APPQcime\CimExtensions
   ```

6. It is recommended to reboot the host.

Remove the CIM Extension from an HP-UX Host

To remove the CIM extension from an HP-UX host

1. Log on to the HP-UX host as the `root` user.
2. Stop the CIM extension by running the following command:

   ```
   /opt/APPQcime/tools/stop
   ```

3. To ensure that you are not in the `/opt/APPQcime` directory, change to the root directory.
4. Run the following command:

   ```
   swremove APPQcime
   ```

   Expected output is similar to the following example:

   ```
   * Beginning Execution
   * The execution phase succeeded for hpuxqaX.dnsxxx.com:/".
   * Execution succeeded.
   ```

5. To remove the `APPQcime` directory, run the following command:

   ```
   rm -r APPQcime
   ```

Remove the CIM Extension from a Linux Host

To remove the CIM extension from a Linux host

1. Log on to the Linux host as the `root` user.
2. Stop the CIM extension by running the following command:

   ```
   /opt/APPQcime/tools/stop
   ```

3. Uninstall the "requires" rpm. For example:

   ```
   rpm –e APPQcime-Requires-XX-224
   ```

4. Uninstall the CIM extension:

   ```
   rpm -e APPQcime
   ```

5. To remove the `APPQcime` directory, run the following command:

   ```
   rm -r APPQcime
   ```

Troubleshooting CIM Extensions

The following topics describe some common approaches to troubleshooting a CIM extension:

- "Agent Service Does Not Start (Windows Only)" below
- "CIM Extension Hangs Because of Low Entropy (Linux Only)" on the next page

Agent Service Does Not Start (Windows Only)

The CIM agent service, `AppStorWin32Agent`, might not start after you install the agent on a Windows Server 2003/2008 R2 IA64 platform.

This issue appears if the JVM exits because of a memory allocation issue during the start of the agent on Intel® Itanium®-based computers.

To resolve this issue:

1. Open the following file in a text editor:

   `[Installation_Directory]\CimExtensions\conf\win32agent.conf`

2. Decrease the value of the property `wrapper.java.maxmemory`. For example, if the current value is 1024, reduce the value to 512.

3. Restart the AppStorWin32Agent service from the **Services** window or reboot the host.

CIM Extension Hangs Because of Low Entropy (Linux Only)

At times, the Linux CIM extension might hang on startup due to low entropy.

The Linux kernel uses keyboard timings, mouse movements, and IDE timings to generate entropy for `/dev/random`. Entropy gathered from these sources is stored in an "entropy pool," and random values returned by `/dev/random` use this pool as a source. This means that `/dev/random` does not return any values if the entropy counter is too low, and programs reading from `/dev/random` are blocked until there is enough collected entropy. This behavior can happen on servers with no keyboards, no mice, and no IDE disks.

1. To determine whether the Linux agent is hung due to this problem, run the following command:

   ```
   kill -3 java_process_id
   ```

   In this instance, java_process_id is the process ID of the Java process for the Linux agent. It is not the process ID returned by the `status` command.

   The preceding command generates the stack trace, which should be similar to the following example:

   ```
   INFO | jvm 1 | 2006/11/22 10:56:58 | at
   java.security.SecureRandom.next(Unknown Source)
   INFO | jvm 1 | 2006/11/22 10:56:58 | at
   java.util.Random.nextInt(Unknown Source)
   INFO | jvm 1 | 2006/11/22 10:56:58 | at
   com.sun.net.ssl.internal.ssl.SSLContextImpl.engineInit(Unknown
   Source)
   INFO | jvm 1 | 2006/11/22 10:56:58 | at
   javax.net.ssl.SSLContext.init(Unknown Source)
   INFO | jvm 1 | 2006/11/22 10:56:58 | at
   com.appiq.cxws.agency.agent.AgentMessageDispatcher.
   createServerSocket(AgentMessageDispatcher.java:1
   INFO | jvm 1 | 2006/11/22 10:56:58 | at
   com.appiq.cxws.agency.agent.AgentMessageDispatcher.
   startAccepting(AgentMessageDispatcher.java:74)
   ```

2. To fix the problem, in the `/opt/APPQcime/conf/wrapper.conf` file, in the Java Additional Properties section, search for the property, `wrapper.java.additional.N=-`

`Djava.security.egd=file:/dev/random` and change `random` to `urandom`.

After the change, the property should be similar to:

```
wrapper.java.additional.N=-
Djava.security.egd=file:/dev/urandom
```

3. Restart the CIM extension:

```
/opt/APPQcime/tools/stop
/opt/APPQcime/tools/start
```

Chapter 5: Configuration

This chapter contains an introduction to concepts, initial configurations required, defaults provided by SOM, some best practices and planning information that will help you implement SOM in your environment.

Ports and Firewall

The following table shows the ports SOM used on the management server.

Legend	
I/O	The port must be opened on both SOM server and the target device.
O	The port must be opened on the target device.
I	The port must be opened on the source server; for example, the SOM management server.

Ports Used on the SOM Management Server

Port	Type	Name	Purpose	Change Configuration	In/Out
80	TCP	nmsas.server.port.web.http	Default HTTP port used for Web UI and Web Services; after this port is open, it becomes bi- directional.		I/O
443	TCP	nmsas.server.port.web.https	Default secure HTTPS port (SSL); used for Web UI and Web Services.	Modify the nms-local.properties file	

Ports Used on the SOM Management Server, continued

Port	Type	Name	Purpose	Change Configuration	In/Out
1098	TCP	nmsas.server.port.naming.rmi	• Used by SOM command line tools to communicate with a variety of services used by SOM • HP recommends configuring the system firewall to restrict access to these ports to localhost only	Modify the nms-local.properties file	
1099	TCP	nmsas.server.port.naming.port	• Used by SOM command line tools to communicate with a variety of services used by SOM. • HP recommends configuring the system firewall to restrict access to these ports to localhost only	Modify the nms-local.properties file	

Ports Used on the SOM Management Server, continued

Port	Type	Name	Purpose	Change Configuration	In/Out
3873	TCP	nmsas.server.port.remoting.ejb3	• Used by SOM command line tools to communicate with a variety of services used by SOM. • HP recommends configuring the system firewall to restrict access to these ports to localhost only	Modify the nms-local.properties file	
4444	TCP	nmsas.server.port.jmx.jrmp	• Used by SOM command line tools to communicate with a variety of services used by SOM. • HP recommends configuring the system firewall to restrict access to these ports to localhost only.	Modify the nms-local.properties file	

Ports Used on the SOM Management Server, continued

Port	Type	Name	Purpose	Change Configuration	In/Out
4445	TCP	nmsas.server. port.jmx.rmi	• Used by SOM command line tools to communicate with a variety of services used by SOM. • HP recommends configuring the system firewall to restrict access to these ports to localhost only	Modify the nms-local.properties file	
4446	TCP	nmsas.server. port.invoker.unified	• Used by SOM command line tools to communicate with a variety of services used by SOM. • HP recommends configuring the system firewall to restrict access to these ports to localhost only.	Modify the nms-local.properties file	
4712	TCP	nmsas.server. port.ts.recovery	Internal transaction service port .	Modify the nms-local.properties file	
4713	TCP	nmsas.server. port.ts.status	Internal transaction service port.	Modify the nms-local.properties file	
4714	TCP	nmsas.server. port.ts.id	Internal transaction service port.	Modify the nms-local.properties file	

Ports Used on the SOM Management Server, continued

Port	Type	Name	Purpose	Change Configuration	In/Out
5432	TCP	com.hp.ov.nms.postgres.port	This PostgreSQL port is the port the embedded database listens on for this SOM management server.	Modify the nms-local.properties file	
8886	TCP	OVsPMD_MGMT	SOM ovspmd (process manager) management port.	Modify the /etc/services file	
8887	TCP	OVsPMD_REQ	SOM ovsmpd (process manager) request port.	Modify the /etc/services file	
8989	TCP	com.hp.ov.nms.events.action.server.port	Enables the action server port to be configurable.	Modify the nnmaction.properties file	

Ports Used for Communication Between the SOM Management Server and Other Systems

Port	Type	Purpose	Client, Server	In/Out
80	TCP	Default HTTP port for SOM; used for Web UI and Web Services.	Server	
80	TCP	Default HTTP port for SOM connecting to other applications. The actual port depends on SOM configuration.	Client	
389	TCP	Default LDAP port.	Client	
443	TCP	Default secure HTTPS port for SOM connecting to other applications; the actual port depends on SOM configuration. Default HTTPS port for HP OM on Windows.	Client	

Ports Used for Communication Between the SOM Management Server and Other Systems, continued

Port	Type	Purpose	Client, Server	In/Out
443	TCP	Default secure HTTPS port; used for Web UI and Web Services .	Server	
636	TCP	Default secure LDAP port (SSL).	Client	
135	TCP	psexec port, Windows Agentless on the management server.	Server	
445	TCP	psexec port, Windows Agentless on the management server.	Server	
139	TCP	winexe port, Windows Agentless on the management server.	Server	
383	TCP	LCore communication port on CMS used for communication with the SOM reporting server.	Server	

Security Recommendations for the SOM Management Server

This section provides information for increasing the security of the SOM management server.

It is recommended to limit traffic to the SOM web server to only those users who should have access. Possible ways to limit this traffic include:

- Configure a firewall in front of the SOM management server.
- Isolate user access to the SOM management server on specific network interfaces only.

SOM installs with the JMX console enabled. It is recommended to disable the JMX console by moving the JMX console definition file to another location, for example up one directory level. The default location of the JMX console definition file is:

- *Windows*: `<SomInstallDir>\nmsas\common\deploy\jmx-console.sar`
- *Linux*: `$SomInstallDir/nmsas/common/deploy/jmx-console.sar`

Node Groups

A Node Group is a collection of nodes (elements) or child node groups that have the same device filter criteria. After discovery elements are automatically assigned to node groups based on predefined attributes.

Node groups can be used for any or all of the purposes:

- For categorization that enables you to identify basic categories in the system, such as hosts, storage systems, switches, and fabrics.
- Categorization enables you with easier monitoring and administration. It helps you apply settings to a group and avoid dealing with elements on an individual basis. For example, you can implement a data collection policy on a node group rather than on individual elements.
- As a primary filtering technique for customizing different views.
- User access control to limit access to a set of nodes through security mappings.

Default Node Groups

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.

- All Elements
- FC Fabrics
- FC Switches
- Hosts
- Storage Systems

These are based on device categories derived from the system object ID during the discovery process.

Node Group Membership

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" below
- "Additional Filters" below
- "Additional Nodes" on the next page
- "Child Node Groups" on the next page

Device Filters

Device filters provide categories such as device category, vendor, family, or device profile. Nodes must match at least one specification to belong to the node group.

During discovery, SOM collects direct information through SNMP queries and derives other information from that through device profiles. By gathering the system object ID, SOM can index through the correct device profile to derive the following information:

- Vendor
- Device category
- Device family within the category

These derived values, in addition to the device profile itself, are available for use as filters. For example, you can group all objects from a specific vendor, regardless of device type and family. Or you can group all devices of a type such as router, across vendors.

Additional Filters

With this option you can specify additional filters using Boolean expressions based on a list of object attributes.

Use the additional filters editor to create custom logic to match fields including:

- tenantName (Name)
- securityGroupName (Security Group)
- sysName (System Name)
- sysLocation (System Location)
- sysContact (System Contact)

- hostname (Hostname, case-sensitive)
- hostedIPAddress (Address)
- mgmtIPAddress (Management Address)
- nodeName

Filters can include the AND, OR, NOT, EXISTS, NOT EXISTS, and grouping (parentheses) operations. See "Specify Node Group Additional Filters" in the SOM Online help for more information.

Additional Nodes

This option enables you to add additional nodes to the node group regardless of any filters.

It is better to use Additional Filters to qualify nodes for node groups. If the environment contains critical devices that are too difficult to qualify using filters, add them to a group by individual host name. Add nodes to a node group by individual host names only as a last option.

Child Node Groups

Enables you to add node groups to the node group to establish hierarchical containers. Child node groups are treated similarly as additional nodes.

Node Groups Evaluation

SOM evaluates each discovered node to determine its node group membership using the following criteria:

- Any node that matches one or more entries (if any exist) on the Device Filters tab and the filter specified on the Additional Filters tab is a member of the node group.
- All nodes specified on the Additional Nodes tab are members of the node group.
- All nodes that are members of at least one node group specified on the Child Node Groups tab are members of the node group.

Group Overlap

Regardless of the intended uses for group definitions, the first step is to define which nodes are members of a group. Because you can create groups for different purposes, each object can be included in multiple groups. Consider the following example:

- You might want to group all HP 3PAR arrays into single group using the Device Profile filter.
- Top elements are automatically assigned to the default node group of Storage Systems.
- You might want to collect data from all storage arrays regardless of device vendor or device family.

The 3PAR array with an IP address 10.10.10.3 would qualify for all three groups. You want to find the balance between having a usably rich set of groups available for configuration and viewing, and overloading the list with superfluous entries that will never be used.

Hierarchies/Containment

You can create simple, reusable, atomic groups and combine them hierarchically for monitoring or visualization. Using hierarchical containers for nodes greatly enhances map views by providing cues about the location or type of object at fault. SOM gives you complete control of the definition of the groups and their drill-down order.

You can create simple, reusable atomic groups first, and then specify them as child groups as you build up. Alternatively, you can specify your largest parent group first and create child groups as you go.

For example, your environment might contain EMC ClARiiON storage systems and VNX Filer. You can create parent groups for EMC devices and for all file storage. Because the hierarchy is specified when you create the parent and designate its children, each child group, such as EMC devices, can have multiple parents.

Hierarchies work well for the following situations:

- Types of nodes with similar monitoring needs
- Types of nodes to be quarantined together
- Groups of nodes by operator job responsibility
- When you use groups in map views and table views

Note: Keep in mind that as you use group definitions to specify monitoring configuration, hierarchy does not imply ordering for settings. The settings with the lowest ordering number apply to a node. By carefully incrementing ordering numbers, you can emulate inheritance concepts for settings.

Planning Node Groups

SOM provides a default collection of node groups to simplify your configuration tasks. You can use existing groups, modify them, or create your own. Over time HP might add more default groups to simplify your configuration tasks.

Interaction with Device Profiles

When each device is discovered, SOM uses its system object ID to index into the list of available

Device Profiles. The Device Profile is used to derive additional attributes of the device, such as vendor, product family, and device category.

As you configure node groups, you can use these derived attributes to categorize devices to apply data collection settings. For example, you might want to collect data from all devices regardless of vendor throughout your environment at a certain interval. You can use the derived device category, Storage System, as the defining characteristic of your node group. All discovered devices whose system object ID maps to the category, Storage Systems, will receive the configured settings for the node group.

Considerations for Planning

Determine the criteria by which you want to group nodes. Following are some factors you can consider while planning node groups:

- Which are the critical devices that you want to collect data?
- Do you want to differentiate data collection intervals or data gathered by device type?
- Can you use the default node groups provided by SOM?

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.

Discovery

The devices that comprise your Storage Area Network (SAN) must be discovered so that they can be monitored and managed by SOM. To discover devices in your network, you must configure the addresses for discovery and provide credentials, if required.

Notes on discovery before you begin planning:

- SOM does not perform any default discovery. You must configure discovery before any elements appear in the Inventory views.
- Discovery is handled on an individual address basis. The status of each address configured for discovery indicates whether the discovery is successful or not.
- The process of initial discovery takes some time depending on the number of addresses you have configured for discovery.
- You can create credentials and then associate them to multiple addresses.
- The Discovery Hint option enables you to select a value, based on which SOMinvokes only the selected provider for discovering the device instead of invoking all the providers thereby reducing discovery time.

Methods of Discovery

SOM provides the following methods of discovery.

Method	Notes
Automatic Discovery (*only initial discovery*)	Default method for initial discovery. Multiple elements can be discovered at once.
Manual Discovery	Only one element can be discovered.
Importing Addresses from a file	Discovery settings from a previous installation.

Automatic Discovery

This is the default and recommended method for initial discovery. This is best suited when you have a large bulk of addresses to be discovered. The discovery addresses that you add or import get into the queue for discovery after a pre-configured time.

Notes on automatic discovery

- Runs only run once during initial discovery
- Allows for multiply devices to be discovered at once, as well as scanning on a range of IP addresses.

Manual Discovery

This method is best suited when you have to add a single element or you have a small number of elements to be discovered. You must associate the device-specific credentials before you begin discovery. Though this method provides a tighter control over discovery, this is time-consuming if you have a large number of addresses to be discovered.

Importing Discovery Settings from a File

If you have discovery settings from a previous installation, you can import it into the management server rather than re-enter the information. The import discovery settings feature enables you to import the following information:

- IP addresses to be discovered
- Default user names and passwords, which are encrypted
- Agentless host inference rules

Notes on import:

- To prevent re-entering the information for each management server instance, you can import the same file for multiple management server instances.
- When you import a file, your previous settings are overwritten.

- If you receive an error message when you try to import the discovery settings, verify that you are using the right password. If you are using the correct password, there is a possibility that the file is corrupt.

When you import discovery settings file, it triggers automatic discovery of addresses. If you do not want to use automatic discovery, you can disable the option.

Host Discovery

SOM provides the following methods to discover and manage hosts and their associations to storage devices.

Discovery Method	Description
Discovery with a CIM extension	Manage hosts by installing a CIM Extension on the host.
Agentless discovery	Manage hosts without installing a CIM extension.
Inferred agentless discovery	Gather information from hosts based on host security groups, zones, and zone aliases without installing a CIM extension to be installed.

Discovery with a CIM extension

A SOM CIM extension is a collection agent that runs on a storage host to gather information about that host. The SOM management server communicates with the CIM extension while discovering and managing the host. Install the CIM extension on each host that you want to manage. The CIM extension must be running for the management server to obtain information from the host.

If you change the password of a host after you discover it, you must change the password for the host in the discovery list, and then you must stop and restart the CIM Extension running on that host before you run a discovery.

Agentless Discovery

Agentless discovery provides management server the capability to discover hosts without installing the CIM extension on the host. The management server supports the agentless discovery for hosts running on Microsoft Windows, Linux operating systems and Solaris systems.

The management server uses the following to discover a host:

- The Windows Management instrumentation (WMI) for discovering Windows host.
- Secure Shell (SSH) for discovering Linux hosts.

Agentless discovery works only if a CIM extension is not running on the host to be discovered. If the management server finds a CIM extension running on the host, by default it prefers discovery using a CIM extension over the agentless discovery.

You can also rediscover hosts, which are already discovered using the CIM extension in the management server, using the agentless discovery. However, all the history information associated with the host and applications on the host is deleted from the management server.

Data collected from the host depends on the discovery method used to gather information from the host. The following table summarizes the data collected for hosts based on the discovery method. Use the table as a guide to plan your approach to host discovery.

Inferred Agentless Discovery

SOM can display and gather information from hosts without CIM extensions. You can infer agentless hosts by creating rules based on host security groups, zones or zone aliases configured on storage systems and fabrics in the SAN. After inferring hosts, you can discover the hosts by providing the credentials. If the discovery is successful, the hosts are reconciled and the inferred hosts become managed hosts.

Data collected from hosts varies based on the discovery method. You can plan host discovery based on the type of data you want to collect from the hosts.

Capabilities of Agentless Discovery

The management server gathers following information from a host discovered using the agentless discovery:

- Host associations to the applications, storage, and network devices.
- IP/DNS related information
- Gathers detailed configuration information for every host.
- Logical storage volume information, including mount points, physical devices, drive types, and file system details.
- Disk partition information, including disk partition names, mapped logical volumes, mapped physical drives, and total capacity.
- Disk drive information, including drive names, SCSI bus information, and mapped disk partitions.

- Multipathing and Volume Manager Configuration details.
- Information related to HBAs.

Limitations of Agentless Discovery

Although, agentless discovery enables the management server to discover and find extensive information related to the hosts, it has some limitations.

Note: All the mentioned limitations can be overcome by installing a CIM extension on the host.

SOM will discover Windows host using agentless configuration on a host running CIM Extension if Discovery hint for agentless is provided during discovery.

Following are the limitations for an agentless host, based on the operating system it is running on:

Limitations for Windows hosts

- A user account with non-administrator privileges cannot discover a Windows host.
- 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. It is because the management server does not support the native volume manager software that is the Microsoft Virtual Disk Service Dynamic Provider.

Limitations for Linux hosts

- Following information is not available for a non-root user account:
  - Information related to Veritas DMP devices is not available.
  - Information related to serial number and manufacturer of the system is not available.
  - Information related to disk drives and disk partitions is not available.
- The following performance metrics are not available for a Linux host:
  - Disk Read
  - Disk Total
  - Disk Utilization

- Disk Write
- Processor utilization

- The number of target mappings obtained by the agentless discovery 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.
- Following issues are observed for the Linux hosts containing HBAs discovered in the management server:
  - Following information is not available for HBAs:
    - Vendor name
    - Serial number
    - Hardware version
    - Information for Port Type on HBA Port Properties page.
  - When you try to rediscover the agentless hosts using the CIM extension, the management server does not reconcile the HBA information obtained during the agentless discovery against the information obtained using CIM extension. The old HBA data obtained using the agentless discovery is deleted and new information is collected for HBAs using the CIM extension discovery. Thus, all the custom information related to HBAs is deleted when the host is rediscovered using the CIM extension.
  - For a Linux host containing HBAs with dual port adapter, each port is displayed as an individual adapter on the HBA adapter page with each adapter mapped with its port on the HBA port page.
  - Bindings page is not updated when the following is performed:
    - Paths to the LUNs are disabled.
    - HBA port is disabled.
    - Subsequent data collection is run.

    This limitation can be overcome by rebooting the host. The Bindings page is automatically updated on rebooting the host.

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:

When providing an address for discovery, specify a tenant for each discovery address. A node is automatically created for an IP address that is discovered successfully. When you define a tenant, then you must 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. Administrators 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 the 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 SOM 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 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.

Host Clusters

The management server provides full support for managing clusters. Cluster support includes the following features:

- Clusters are recognized as managed elements.
- Cluster capacity utilization is accurately reported.
- The management server supports automatic discovery of several popular cluster servers.

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.

Recommendations for Data Collection Policies

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.

Recommendations for Monitoring Performance

Following are some recommendations to consider while configuring monitoring policies:

- 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.

LDAP-Based Authentication

This chapter contains information about integrating SOMwith a directory service for consolidating the storage of user names, passwords, and, optionally, SOM user group assignments. It contains the following topics:

SOM User Access Information and Configuration Options

Together, the following items define an SOM user:

- The **user name** uniquely identifies the SOM user. The user name provides access to SOM and receives incident assignments.
- The **password** is associated with the user name to control access to the SOM console or SOM command.
- **SOM user group** membership controls the information available and the type of actions that a user can take in the SOM console. User group membership also controls the availability of SOM commands to the user.

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

User Accounts	Password	User Group	User Group Membership
Directory Service	Directory Service	SOM	Directory Service

External Mode: All SOM User Information in the Directory Service

With this option, SOM accesses a directory service for all user access information, which is defined externally to SOM and is available to other applications. Membership in one or more directory service groups determines the SOM user groups for the user.

The configuration and maintenance of SOM user access information is a joint effort as described here:

- The directory service administrator maintains the user names, passwords, and group membership in the directory service.
- The SOM administrator maps the directory service groups to SOM user groups in the SOM console
- The SOM administrator configures the SOM `ldap.properties` file to describe the directory service database schema for user names and groups to SOM.

HP recommends the following configuration process:

1. Configure and verify SOM user name and password retrieval from the directory service.
2. Configure SOM user group retrieval from the directory service.

For information about integrating with a directory service for all user information, see the rest of this chapter and the *SOM Help*.

Configuring SOM to Access a Directory Service

Directory service access is configured in the `ldap.properties` file. To configure user access from the directory service, follow the appropriate procedure for your directory service.

- Steps for Microsoft Active Directory
- Steps for Other Directory Services

- "Map the Directory Service Groups to SOM User Groups" on page 70

Steps for Microsoft Active Directory

1. Back up the `ldap.properties` file that was shipped with SOM, and then open the file in any text editor.
2. Overwrite the file contents with the following text:

   ```
   java.naming.provider.url=ldap://<myldapserver>:389/

   bindDN=<mydomain>\\<myusername>

   bindCredential=<mypassword>

   baseCtxDN=CN=Users,DC=
   <myhostname>,DC=<mycompanyname>,DC=<mysuffix>

   baseFilter=CN={0}

   defaultRole=guest

   #rolesCtxDN=CN=Users,DC=
   <myhostname>,DC=<mycompanyname>,DC=<mysuffix>

   roleFilter=member={1}

   uidAttributeID=member

   userRoleFilterList=admin;level2;level1
   ```

3. Specify the URL for accessing the directory service. In the following line:

   ```
   java.naming.provider.url=ldap://<myldapserver>:389/
   ```

   Replace *<myldapserver>* with the fully-qualified hostname of the Active Directory server (for example: `myserver.example.com`).

   > **Tip:** To specify multiple directory service URLs, separate each URL with a single space character ( ).

4. Specify credentials for a valid directory service user. In the following lines:

   ```
   bindDN=<mydomain>\\<myusername>
   bindCredential=<mypassword>
   ```

   Make the following substitutions:

- Replace *<mydomain>* with the name of the Active Directory domain.
- Replace *<myusername>* and *<mypassword>* with a user name and password for accessing the Active Directory server.

5. Specify the portion of the directory service domain that stores user records. In the following line:

   ```
   baseCtxDN=CN=Users,DC=<myhostname>,DC=<mycompanyname>,
     DC=<mysuffix>
   ```

   Replace *<myhostname>*, *<mycompanyname>*, and *<mysuffix>* with the components of the fully-qualified hostname of the Active Directory server (for example, for the hostname `myserver.example.com`, specify: `DC=myserver,DC=example,DC=com`).

Steps for Other Directory Services

1. Back up the `ldap.properties` file that was shipped with SOM, and then open the file in any text editor.
2. Specify the URL for accessing the directory service. In the following line:

   ```
   #java.naming.provider.url=ldap://<myldapserver>:389/
   ```

   Do the following:

   - Uncomment the line (by deleting the # character).
   - Replace *<myldapserver>* with the fully-qualified hostname of the directory server (for example: `myserver.example.com`).

     Tip: To specify multiple directory service URLs, separate each URL with a single space character ( ).

3. Specify the portion of the directory service domain that stores user records. In the following line:

   ```
   baseCtxDN=ou=People,o=myco.com
   ```

   Replace `ou=People,o=myco.com` with the portion of the directory service domain that stores user records.

4. Specify the format of user names for signing in to SOM. In the following line:

```
baseFilter=uid={0}
```

Replace `uid` with the user name attribute from the directory service domain.

Map the Directory Service Groups to SOM User Groups

Replicate the DN of the LDAP groups in SOM. Map the admin or level1 or level2 roles in SOM to the LDAP groups through the directory service name.

1. In the SOM console, map the predefined SOM user groups to their counterparts in the Directory service:
   a. From the workspaces navigation panel, select the **Configuration**> **Security** > **User Groups**. The User Groups view is displayed.

   b. Double-click the admin row.

   c. In the Directory Service Name field, enter the full distinguished name (DN) of the Directory Service group for SOM administrators.

   d. Click **Save and Close**.

   e. Repeat step b through step d for each of the guest, level1, and level2 rows.

   Tip: These mappings provide SOM console access. Every user who will access the SOM console must be in a directory service group that is mapped to one of the predefined SOM user groups named in this step.

2. For other groups containing one or more SOM users in the directory service, create a new user group in the SOM console:
   a. From the workspaces navigation panel, select the **Configuration**> **Security** > **User Groups**. The User Groups view is displayed.

   b. Click **New**and then enter the information for the group:
      - Set Unique Name to any unique value. Short names are recommended.
      - Set Display Name to the value users should see.
      - Set Directory Service Name to the full distinguished name of the directory service group.
      - Set Description to text that describes the purpose of this SOM user group.

   c. Click **Save and Close**.

   d. Repeat step b and step c for each additional directory service group of SOM users.

Security

In SOM, security and multi-tenancy provide for restricting user access to information about the objects in the SOMdatabase. This restriction is useful for customizing the views of operators to their areas of responsibility. It also supports service providers with per-organization configuration of SOM.

By default, all console users can see information for all objects in the SOM database. If this default configuration is acceptable for your environment, you do not need to read this section.

This section focuses on the SOM security and tenant models and provides suggestions and examples of configuration. The following topics are covered:

- "The SOM Security Model " below
- "The SOM Tenant Model" on page 77
- "Some Examples of Security Configuration" on page 82

The SOM Security Model

The SOM security model provides user access control to the objects in the SOM database. This model is appropriate for use by any network management organization that wants to limit SOM user access to specific objects. The SOM security model has the following benefits:

- Provides a way to limit a SOM console operator's view of the network. Operators can focus on specific device types or network areas.
- Provides for customizing operator access to the SOM topology. The level of operator access can be configured per node.
- Simplifies the configuration and maintenance of node groups that align with the security configuration.
- Can be used independently of the SOM tenant model.

Security Groups

In the Storage Operations Manager security model, user access to nodes is controlled indirectly though user groups and security groups. Each node in the topology is associated with only one

security group. A security group can be associated with multiple user groups.

Each user account is mapped to the following user groups:

- One or more of the following default user groups:
  - Administrator
  - Global operator
  - Level 1 operator
  - Level 2 operator
  - Guest user

  This mapping is required for SOMconsole access and determines which actions are available within the SOMconsole. If a user account is mapped to more than one of these SOMuser groups, the user receives the superset of the permitted actions.

  Note: The Global Operators user group grants access to topology objects only. A user must be assigned to one of the other user groups (Level 1, Level 2 or Guest) to access the console.

  The administrator should not map the Global Operators user group to any security group because this user group is, by default, mapped to all security groups.

- (*Optional*) Custom user groups that are mapped to security groups.
  These mappings provide access to objects in the Storage Operations Manager database. Each mapping includes an object access privilege level that applies to the nodes for a security group.

Default Security Group

In a new installation, the Default Security Group is the initial security group assignment for all nodes. By default, all users can see all objects in the Default Security Group. You can control the configuration of nodes to the Default Security Group and users' access to the objects in the Default Security Group.

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.

A Sample Approach to Plan Security Groups

Following are an outline of high-level steps for planning the configuration of security groups:

1. Analyze the managed network topology to determine the groups of nodes to which the users need access.
2. Remove the default associations between the default user groups and the default security group and the Unresolved Incidents security group.
   Doing this step ensures that users do not inadvertently obtain access to nodes they should not be managing. At this point, only administrators can access objects in the topology.
3. Configure a security group for each subset of nodes. Remember that a given node can belong to only one security group.
   a. Create the security groups.
   b. Assign the appropriate nodes to each security group.
4. Configure custom user groups.
   a. For each security group, configure a user group for each level of user access.
      - If you are if storing user group membership in the Storage Operations Manager database, no users are mapped to these user groups yet.

- If you are storing user group membership in a directory service, set the Directory Service Name field for each user group to the distinguished name of that group in the directory service.

b. Map each custom user group to the correct security group. Set the appropriate object access privilege for each mapping.

5. Configure user accounts.
   - If you are storing user group membership in the Storage Operations Manager database, do the following:
     - Create a user account object for each user who can access the console. (The process of configuring user accounts depends on whether you are using a directory service for Storage Operations Manager console logon.)
     - Map each user account to one of the default user groups (for access to the console).
     - Map each user account to one or more custom user groups (for access to topology objects).
   - If you are storing user group membership in a directory service, verify that each user belongs to one of the default user groups and one or more custom user groups.
6. Verify the configuration.
7. Maintain the configuration.
   - Watch for nodes added to the default security group, and move these nodes to the correct security groups.
   - Add new console users to the correct user groups.

Example Security Group Structure

The three ovals in the following diagram indicate the primary groupings for which users need to view the nodes in this example Storage Operations Manager topology. For complete user access control, each of the four unique subgroups corresponds to a unique security group. Each unique security group can be mapped to one or more user groups to represent the available levels of user access to the objects in that security group.

Example Security Group Mappings lists the mappings between the security groups and the possible custom user groups for this topology. (An actual implementation of this security model might not require all of these custom user groups.) Example User Account Mappings lists the mappings for several user accounts and the user groups for this topology.

Example Topology for User Access Requirements

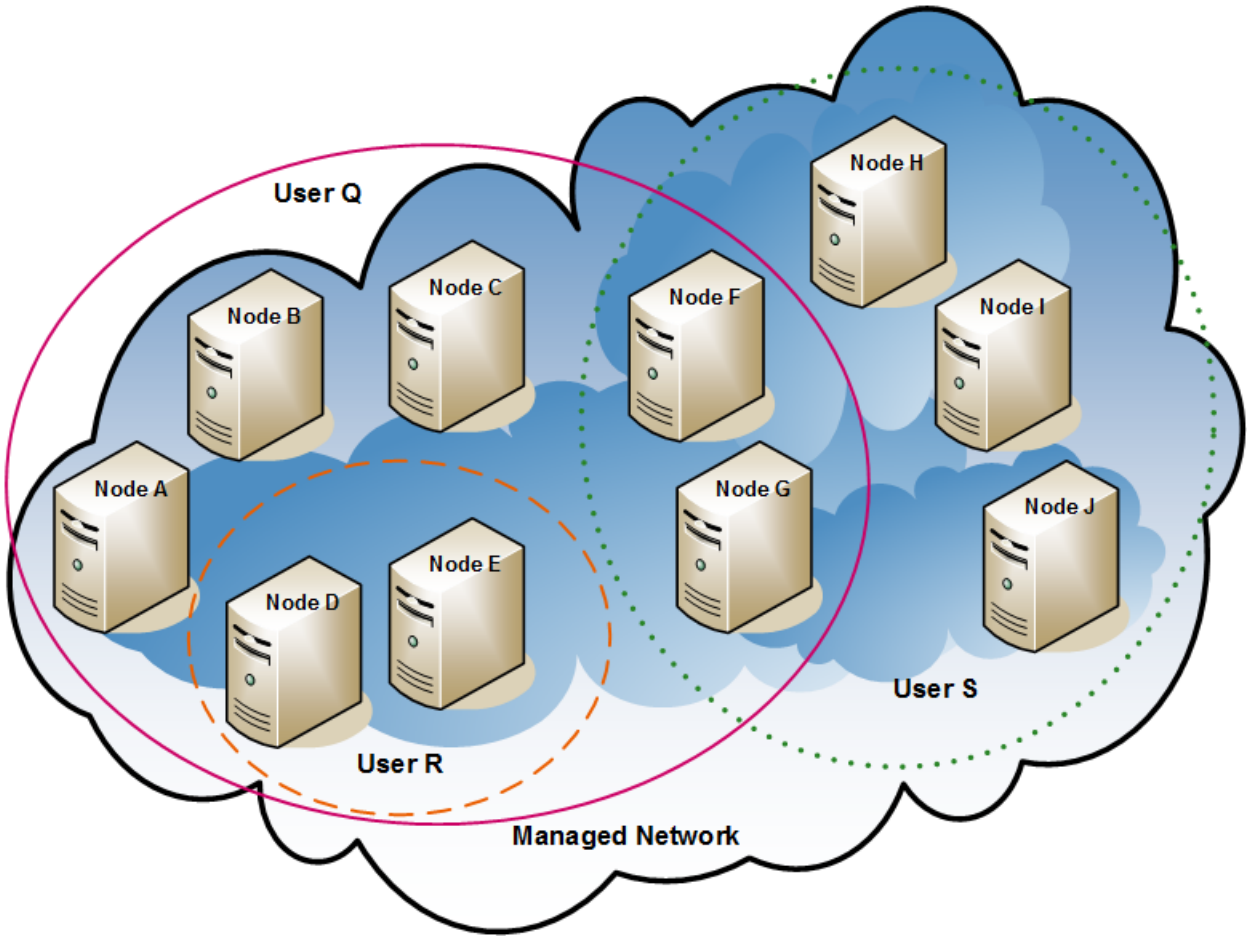

Example Security Group Mappings

Security Group	Nodes of Security Group	User Group	Object Access Privilege
SG1	A, B, C	UG1 Administrator	Object Administrator
		UG1 Level 2	Object Operator Level 2
		UG1 Level 1	Object Operator Level 1
		UG1 Guest	Object Guest

Example Security Group Mappings, continued

Security Group	Nodes of Security Group	User Group	Object Access Privilege
SG2	D, E	UG2 Administrator	Object Administrator
		UG2 Level 2	Object Operator Level 2
		UG2 Level 1	Object Operator Level 1
		UG2 Guest	Object Guest
SG3	F, G	UG3 Administrator	Object Administrator
		UG3 Level 2	Object Operator Level 2
		UG3 Level 1	Object Operator Level 1
		UG3 Guest	Object Guest
SG4	H, I, J	UG4 Administrator	Object Administrator
		UG4 Level 2	Object Operator Level 2
		UG4 Level 1	Object Operator Level 1
		UG4 Guest	Object Guest

Example User Account Mappings

User Account	User Groups	Node Access	Notes
User Q	SOM Level 2 Operators	none	This user has operator level 2 access to the nodes in the pink oval (solid line).
	UG1 Level 2	A, B, C	
	UG2 Level 2	D, E	
	UG3 Level 2	F, G	

Example User Account Mappings, continued

User Account	User Groups	Node Access	Notes
User R	SOM Level 1 Operators	none	This user has operator level 1 access to the nodes in the orange oval (dashed line).
	UG2 Level 1	D, E	
User S	SOM Level 2 Operators	none	This user has operator level 2 access to the nodes in the green oval (dotted line).
	UG3 Level 2	F, G	
	UG4 Level 2	H, I, J	
User T	SOM Level 2 Operators	none	This user has access (with varying privilege levels) to all nodes in the example topology. This user has administrative access to nodes D and E but cannot see the menu items for tools that require administrative access. If this user has access to the management server, this user can run command-line tools that require administrative access against nodes D and E only.
	UG1 Guest	A, B, C	
	UG2 Administrator	D, E	
	UG3 Level 2	F, G	
	UG4 Level 1	H, I, J	

The SOM Tenant Model

The Storage Operations Manager tenant model provides strict segregation of topology discovery and data into tenants, also called organizations or customers. This model is appropriate for use by service providers, especially managed service providers and large enterprises.

The Storage Operations Manager tenant model has the following benefits:

- Marks the organization to which each node belongs.
- Meets regulatory requirements for separating operator access to customer data.
- Simplifies the configuration and maintenance of node groups that align with the tenant configuration.
- Simplifies configuration of security.

Tenants

The SOM tenant model adds the idea of an organization to the security configuration. Each node in the topology belongs to only one tenant. The tenant provides logical separation in the Storage Operations Manager database. Object access is managed through security groups.

For each node, the initial discovery tenant assignment occurs when the node is first discovered and added to the Storage Operations Manager database. Storage Operations Manager assigns all the discovered nodes to the default tenant. Therefore, if you use the security model without configuring any tenants, all nodes are assigned to the default tenant. By default, all users have access (through the default security group) to all objects associated with this tenant. An administrator can change the tenant for a node at any time after discovery.

Each tenant definition includes an initial discovery security group, the default security group. Storage Operations Manager assigns the node to the default security group along with the default tenant. An administrator can change the security group for a node at any time after discovery.

Note: When you change the tenant for a node, it does not automatically change the security group of the node.

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.

A Sample Approach to Plan Tenants

The following steps outline the high-level approach to planning and configuring multi-tenancy:

1. Analyze your customer requirements to determine how many tenants are required in the Storage Operations Manager environment.
   It is recommended that tenants be used only when managing multiple separate networks with a single management server.
2. Analyze the managed topology to determine which nodes belong to each tenant.
3. Analyze the topology of each tenant to determine the groups of nodes to which Storage Operations Manager users need access.
4. Remove the default associations between the default user groups and the default security group and the Unresolved Incidents security group.

   Doing this step assures that users do not inadvertently obtain access to nodes they should not be managing. At this point, only administrators can access objects in the topology.
5. Create the identified security groups and tenants.
   For each tenant, set the Initial Discovery Security Group to either the default security group or a tenant-specific security group with restricted access. This approach ensures that new nodes for the tenant are not generally visible until the administrator configures access.
6. Prepare for discovery by assigning tenants to seeds.

   Tip: After discovering a group of nodes, you can change the value of the Initial Discovery Security Group. Using this approach limits the manual re-assignment of nodes to security groups.
7. After discovery completes, do the following:
   a. Verify the tenant for each node and make changes as necessary.
   b. Verify the security group for each node and make changes as necessary.

Example Tenant Structure

The following diagram shows an example Storage Operations Manager topology containing two tenants, represented by the rectangles. The three ovals indicate the primary groupings for which users need to view the nodes. The topology for Tenant 1 is managed as a single group, so it needs

only one security group. The topology for Tenant 2 is managed in overlapping sets, so it is separated into three security groups.

Example Security Group Mappings for Multiple Tenants lists the mappings between the security groups and the possible custom user groups for this topology. (An actual implementation of this security model might not require all of these custom user groups.) Example User Account Mappings for Multiple Tenants lists the mappings for several user accounts and the user groups for this topology.

Example Topology for Multiple Tenants

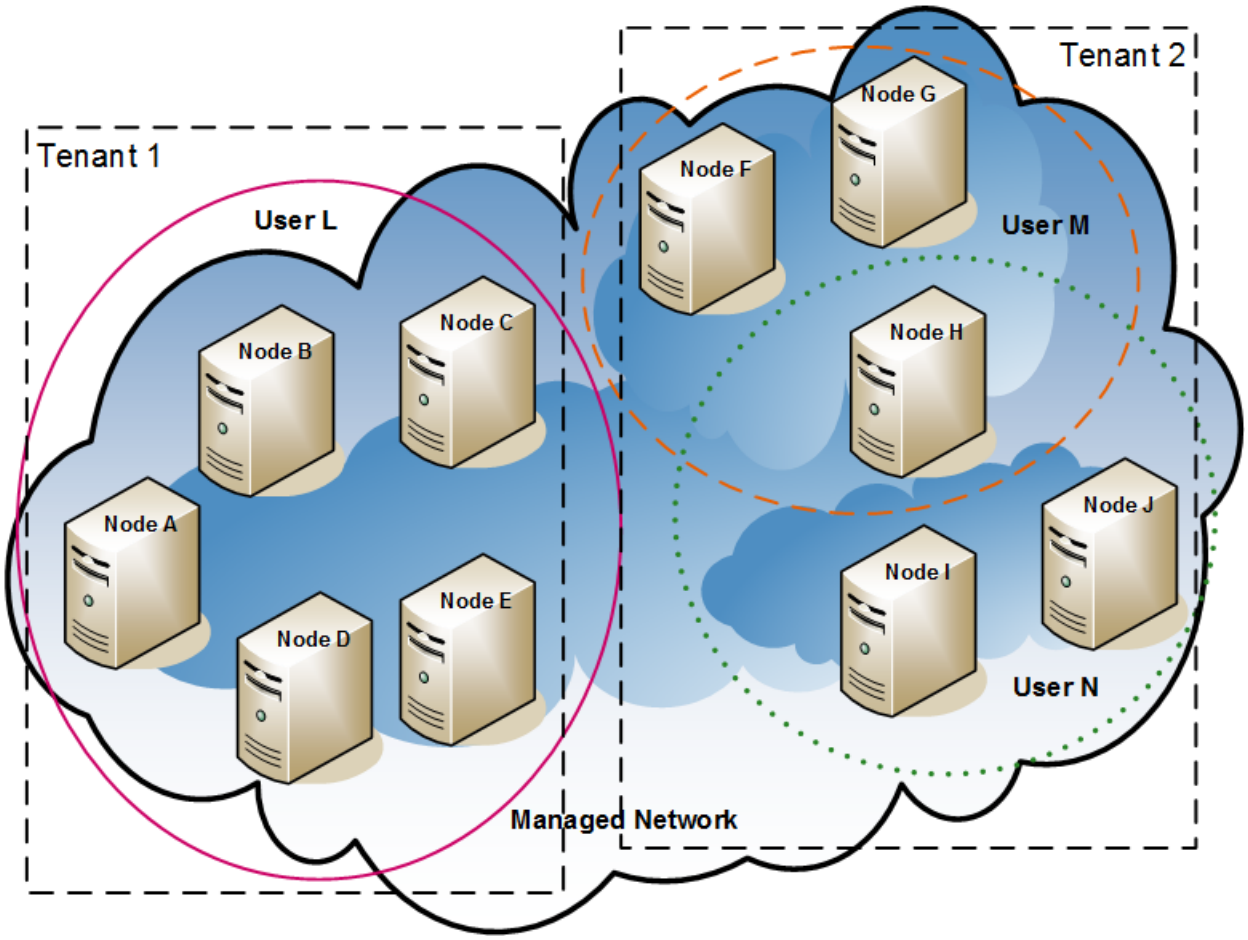

Example Security Group Mappings for Multiple Tenants

Security Group	Nodes of Security Group	User Group	Object Access Privilege
T1 SG	A, B, C, D, E	T1 Administrator	Object Administrator
		T1 Level 2	Object Operator Level 2
		T1 Level 1	Object Operator Level 1
		T1 Guest	Object Guest
T2 SGa	F, G	T2_a Administrator	Object Administrator
		T2_a Level 2	Object Operator Level 2
		T2_a Level 1	Object Operator Level 1
		T2_a Guest	Object Guest
T2 SGb	H	T2_b Administrator	Object Administrator
		T2_b Level 2	Object Operator Level 2
		T2_b Level 1	Object Operator Level 1
		T2_b Guest	Object Guest
T2 SGc	I, J	T2_c Administrator	Object Administrator
		T2_c Level 2	Object Operator Level 2
		T2_c Level 1	Object Operator Level 1
		T2_c Guest	Object Guest

Example User Account Mappings for Multiple Tenants

User Account	User Groups	Node Access	Notes
User L	SOM Level 2 Operators	none	This user has operator level 2 access to the nodes in the pink oval (solid line), which groups all nodes in Tenant 1.
	T1 Level 2	A, B, C, D, E	
User M	SOM Level 1 Operators	none	This user has operator level 1 access to the nodes in the orange oval (dashed line), which groups a subset of the nodes in Tenant 2.
	T2_a Level 1	F, G	
	T2_b Level 1	H	
User N	SOM Level 2 Operators	none	This user has operator level 2 access to the nodes in the green oval (dotted line), which groups a subset of the nodes in Tenant 2.
	T2_b Level 2	H	
	T2_c Level 2	I, J	

Some Examples of Security Configuration

The following examples present possible security strategies. Use them as a guideline for configuring security. Select the example that best matches your security configuration requirements:

- "Example: Divide Node Access Between Two or More User Groups" below
- "Example: Allow a Subset of Users to Access a Subset of Nodes" on page 85

Example: Divide Node Access Between Two or More User Groups

This example configures security to divide the responsibility for network monitoring based on the following locations:

- Chicago
- Detroit

Each location includes a Level 1 Operator (with more limited access privileges than Level 2 Operators) and a Level 2 Operator. Tina, the Administrator, handles both locations. Kevin is a backup for both Chicago and Detroit and must access the nodes in both Chicago and Detroit.

The following diagram illustrates the security requirements:

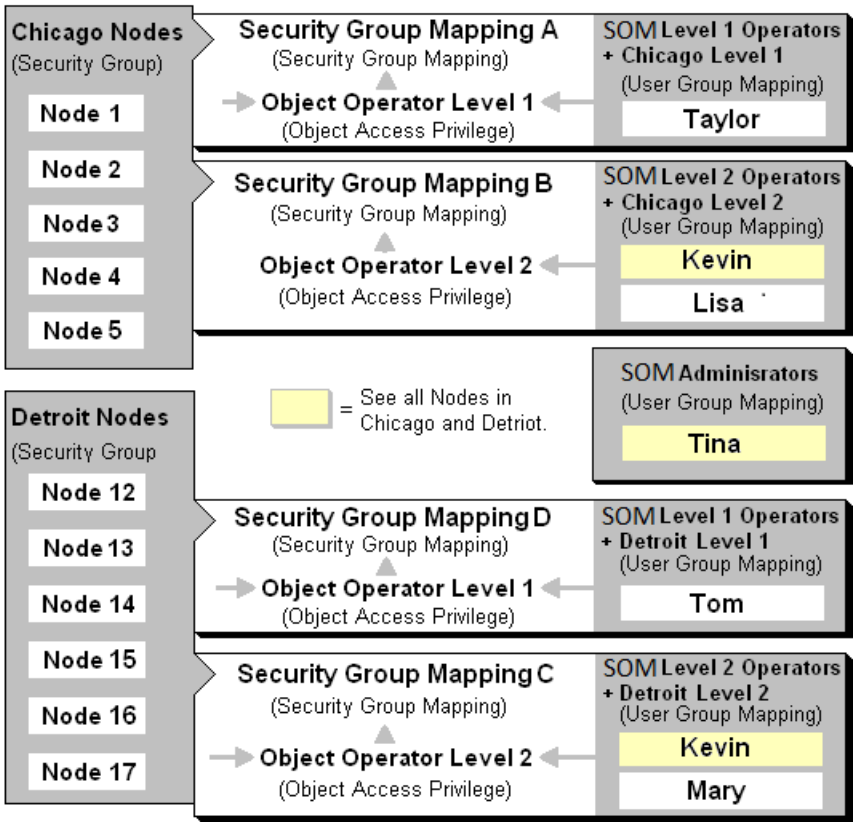

The following table lists the SOM console (SOM User Group) and node access requirements (User Group, Object Access Privilege and Security Group) for each location.

Note: You can place all operators into the SOM Level 2 Operators if you want all operators to see all menu options, but only have the ability to run them based on their Object Access Privilege.

Example Security Configuration

User Accounts	SOM User Groups	User Groups	Object Access Privileges	Security Groups
Tina	SOM Administrator	Not Applicable. The SOM Administrator can access all nodes.	Not Applicable. The SOM Administrator has Administrator privileges to all nodes.	Not Applicable. The SOM Administrator can access all nodes.

Example Security Configuration, continued

User Accounts	SOM User Groups	User Groups	Object Access Privileges	Security Groups
Kevin	SOM Level 2 Operators	Chicago Level 2 Detroit Level 2	Object Operator Level 2	Chicago Nodes, Detroit Nodes
Lisa	SOM Level 2 Operators	Chicago Level 2	Object Operator Level 2	Chicago Nodes
Taylor	SOM Level 1 Operators	Chicago Level 1	Object Operator Level 1	Chicago Nodes
Mary	SOM Level 2 Operators	Detroit Level 2	Object Operator Level 2	Detroit Nodes
Tom	SOM Level 1 Operators	Detroit Level 1	Object Operator Level 1	Detroit Nodes

To set up security for the Chicago and Detroit locations follow these procedures:

- Remove the Default Security Group Mapping to the default User Groups - Level 1 Operators, Level 2 Operators, and Guest user .

  Note: The default user groups are provided for those administrators who are not concerned with Security configuration. After you remove these Security Group Mappings, the user groups provide access only to the SOM console rather than to the SOM consoleand to all nodes.

- Create the User Accounts. See the Example Security Configuration table.
- Create the additional user groups required for the Chicago and Detroit Security Groups (Chicago Level 2, Chicago Level 1, Detroit Level 2, Deteroit Level 1). (See the Example Security Configuration table.)
- Map User Accounts to SOM User Groups. (See the Example Security Configuration table.)
- Create the Security Groups for Detroit and Chicago location.
- Map each security group to the new User Groups. (See the Example Security Configuration table.)
  - **ChicagoLevel1** User Group to the **Chicago Nodes**
  - **DetroitLevel1** User Group to the **Detroit Nodes**
  - **DetroitLevel2** User Group to the **Detroit Nodes**

- Assign the nodes to the appropriate Security Group.
- View a summary of your configuration changes.
- Save you configuration changes.

Example: Allow a Subset of Users to Access a Subset of Nodes

This example configures security to allow a subset of users to access only those nodes in Building 5. The remaining users can access all nodes discovered by SOM.

This location includes a Level 1 Operator (with more limited access privileges than Level 2 Operators) and a Level 2 Operator. Jeff is a Level 2 Operator who can access only the nodes in Building 5.

Note: Be sure to create a user account that is mapped to the SOM Administrator User Group so that one person has access to the Configuration workspace and all the nodes in the network. See the topic "Restore the Administrator Role" in the Online Help for more information.

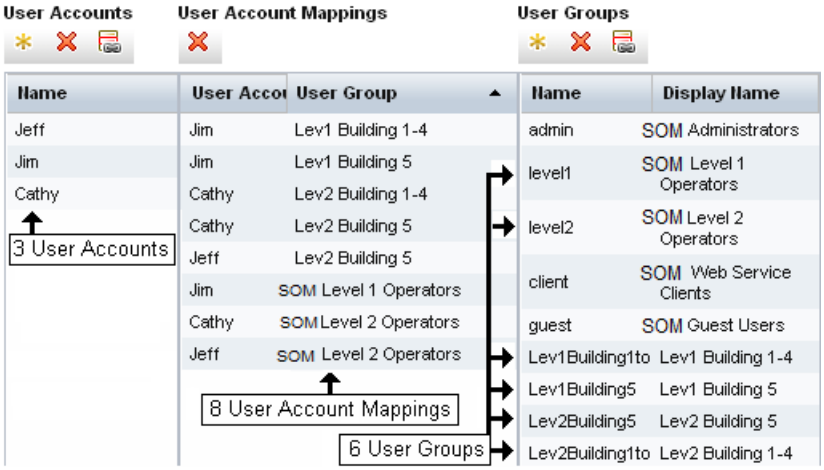

The following table lists the SOM console access requirements (SOM User Group) and node access requirements (User Group, Object Access Privilege and Security Group) for each User Account.

Note: You can place all operators into the SOM Level 2 Operators if you want all operators to see all menu options, but only have the ability to run them based on their Object Access Privilege.

Example Security Configuration

User Accounts	SOM User Groups	User Groups	Object Access Privileges	Security Groups
Jim	SOM Level 1 Operators	Lev1Buildings1-4 Lev1Building5	Object Operator Level 1	Default Security Group
Cathy	SOM Level 2 Operators	Lev2Buildings1-4 Lev2Building5	Object Operator Level 2	Default Security Group
Jeff	SOM Level 2 Operators	Lev2Building5	Object Operator Level 2	Building 5 Nodes

To set up security for this location follow these procedures:

- Remove the Default Security Group Mapping to the user groups - Level 1 Operators, Level 2 Operators, and Guest

  Note: The SOM User Groups are provided for those administrators who are not concerned with security configuration. After you remove these Security Group Mappings, the SOM User Groups provide access to the SOM console only rather than to the SOM console and to all nodes.

- Create the User Accounts. (See the Example Security Configuration table.)
- Create Additional User Groups. (See the Example Security Configuration table.)
- Map User Accounts to SOM User Groups. (See the Example Security Configuration table.)
  - Assign **Jim** to the **Lev1Building1-4** and **Lev1Building5** User Group
  - Assign **Cathy** to the **SOM Level 2 Operators**, **Lev2Building1-4**, and **Lev2Building5** User Groups
  - Assign **Jeff** to the **SOM Level 2 Operators** and **Lev2Building 5** User Groups.
- Create the Building 5 Security Group.
- Map each Security Group to the new User Groups. (See the Example Security Configuration table.)
  - **Lev1Building5** User Group to the **Building 5 Nodes**.
  - **Lev2Building1-4** User Group to the **Default Security Group**

    - **Lev2Building5** User Group to the **Building 5 Nodes**.
- Assign the nodes to the appropriate Security Group.
- View a summary of your configuration changes.

Chapter 6: Backup and Restore of the SOM Embedded Database

SOM provides the following commands to back up and restore the SOM embedded database. This functionality is useful for creating a snapshot of the data and restoring it.

Ensure that the `somdbmgr` service is running before you begin the backup and restore operations.

Commands and Description

Command

`sombackupembdb.ovpl [-?|-h|-help] [-force] [-noTimeStamp] - target <directory>`

Creates a complete backup of the SOM embedded database (excluding the file system data) while SOM is running.

Parameter	Description

`-?`
`-h`
`-help`
	Displays syntax and usage of the `sombackupembdb.ovpl` command.

`-force`	Starts SOM if it is not already running.
`-noTimeStamp`	Removes the time stamp from the back up name.
`-target <directory>`	(*Required*) Specifies the target directory where the data needs to be backed up.

`somrestoreembdb.ovpl [-?|-h|-help] [-force] -source <file>`

Restores a backup that was created by using the `sombackupembdb.ovpl` script.

Parameter	Description

`-?`
`-h`
`-help`
	Displays syntax and usage of the `serestoreembdb.ovpl` command.

`-force`	Stops or starts SOM as required.
`-source <file>`	(*Required*) Specifies the source file name where the data that needs to be restored is saved.

`somresetembdb.ovpl`

Drops the SOM embedded database tables. Runs the ovstart command to recreate the tables.

When you reset the database, 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

Parameter	Description

`-?`
`-h`
`-help`
	Displays syntax and usage of the `somrestoreembdb.ovpl` command.

`-force`	Stops or starts SOM as required.
`-source <file>`	(*Required*) Specifies the source file name where the data that needs to be restored is saved.

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 Deployment 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.