HP Storage Essentials

Software Version: 9.4.1

User Guide

Document Release Date: Friday, February 24, 2012

Software Release Date: July 2011

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 Notices

Trademark Notices

Java is a registered trademark of Oracle and/or its affiliates.

Microsoft® and Windows® are registered trademarks of Microsoft Corporation.

Oracle is a registered trademark of Oracle Corporation.

UNIX® is a registered trademark of the Open Group.

Acknowledgements

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

This product includes software developed by the JDOM Project (http://www.jdom.org/).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/).

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

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:

http://h20230.www2.hp.com/selfsolve/manuals

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

http://h20229.www2.hp.com/passport-registration.html

Or click the **New users – please register** link on the HP Passport login 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:

http://www.hp.com/go/hpsoftwaresupport

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:

http://h20229.www2.hp.com/passport-registration.html

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

[This page intentionally left blank]

1 Overview

About This Product

This product simplifies your complex environment and lowers your cost of management with CIM-based integrated storage management. The management software integrates the management of applications, servers, storage networks, and storage systems in a single, easy-to-implement and intuitive solution.

The management server integrates the various components in the storage area network infrastructure into a CIM/WBEM/SMI standards-based database, so you can eliminate any vendor dependencies and view and manage your infrastructure as a whole. A SAN is a network configuration that is dedicated to transporting storage data among network devices, such as storage systems, servers, tape libraries, and switches. Since the SAN is dedicated to transporting storage data, it frees up the data network for regular TCP/IP traffic.

This product gives your administrators a single, integrated console to manage tactical activities such as provisioning storage, managing real-time events, installing new applications, and migrating servers and storage, as well as strategic activities such as forecasting, planning, and cost analysis. The management server's integrated storage management lowers your cost of acquiring and managing a heterogeneous storage environment.

Key Benefits

- More efficient use of existing assets
- Increased application availability and performance
- Quicker deployment of storage infrastructure and business applications
- Protection of customer flexibility and investments with a standards-based interface

Storage Management Terms

- **CIM** – A common data model of an implementation-neutral schema for describing overall management information in a network/enterprise environment.
- **Web-Based Enterprise Management (WBEM)** – An initiative based on a set of management and Internet standard technologies developed to unify the management of enterprise computing environments.
- **Storage Management Initiative (SMI)** – A SNIA standard for implementing data storage management using CIM.

See the glossary in this user guide for additional definitions.

Suggested Topics for First-Time Users

As a first-time user, you should first become familiar with discovery and Get Details in the management server. You must configure the management server to discover the devices on the network, so that the management server becomes aware of them. Then you must run Get Details, so that the management server is aware of the various types of elements on the network.

To learn more about discovery and Get Details, see the following topics:

- Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries on page 71
- Discovering Applications, Backup Hosts, and Hosts on page 179

After you discover and obtain details about the devices in the network, begin adding users and adding them to roles and organizations. See Managing Security on page 287 for more information.

When you are done adding users and roles, use the following list as a guideline for the topics you should learn about:

- Provisioning Manager on page 675
- Event Management on page 529
- Performance Manager on page 595

Product Components

This product ships with the following software:

- **Management server** – The management server provides various tools to let you monitor and manage your SAN devices. See Management Server Components on the facing page for more information about these tools.
- **Report Optimizer** – Report Optimizer provides detailed reporting on the infrastructure, such as statistics and usage trends. If you want to use Report Optimizer to create reports, contact support for a license that grants you this additional permission. To create reports, login to Report Optimizer directly. Refer to the Quick Start Guide for Report Optimizer for more information about creating reports. The documentation for Report Optimizer can be found by clicking the Help links in Report Optimizer, which might be installed on a separate server depending on your configuration.
- **CIM extensions** – A CIM extension gathers information from the operating system and host bus adapters. It then makes the information available to the management server. See the installation guide for information on how to install the CIM extensions.
- (*Optional*) **Module for managing Microsoft Exchange Server** – Management software for Microsoft Exchange lets the administrator actively manage the data storage requirements for Microsoft Exchange.
- (*Optional*) **Module for managing Oracle Database** – Management software for Oracle reduces database storage cost and improves performance, availability, and reliability by

assisting the Oracle database administrator in administering Oracle instances, particularly storage.

- (*Optional*) **Module for managing Microsoft SQL Server** – Management software for Microsoft SQL Server lets the administrator manage and monitor Microsoft SQL Server.
- (*Optional*) **Module for managing Sybase Adaptive Server Enterprise** – Management software for Sybase lets the administrator actively manage the data requirements for Sybase.
- (*Optional*) **Module for managing Caché** – Management software for Caché lets the administrator actively manage the data storage requirements for Caché.
- (*Optional*) **Module for managing DB2** – Management software for DB2 lets the administrator actively manage the data storage requirements for DB2.
- (*Optional*) **Module for managing Informix** – Management software for Informix lets the administrator actively manage the data storage requirements for Informix.

Management Server Components

You might not have access to all of the features described in this section, depending on the following:

- The type of license you have. Depending on your license, all features might not be available. See the List of Features to determine if you have access to all features. The List of Features is accessible from the Documentation Center (**Help** > **Documentation Center**).
- The role assigned to your user account. For example, users assigned to the Help Desk role by default only have access to Application Viewer and Event Manager.

The following is the full feature set available:

- **Application Viewer** – Enables you to monitor and display data from applications. You can access Application Viewer by clicking **Application Viewer**.
- **Backup Manager** – Enables you to keep track of element backups. See Using Backup Manager on page 729 for more information.
- **Capacity Manager** – Provides a graphical representation of an element's storage capacity in the storage network. See Finding an Element's Storage Capacity on page 553 for more information.
- **Chargeback Manager** – Enables you to manage departmental ownership, track cost, and assemble business reports making inquiries, such as audits and inventory reviews, easier. See Chargeback Manager on page 803 for more information.
- **Command Line Interface (CLI)** – Provides an alternate way for you to manage elements that the management server monitors. You can use the CLI commands in scripts to manage your storage.
- **File System Viewer** – Does a recursive lookup on the file system and stores the information in an embedded database. File System Viewer can scan files very quickly, because of its structure in the database, and because it uses a multi threaded process. More than one

process can be used at a time to scan the files. Refer to the File Servers Guide for more information.

- **Element Manager** – Provides a fast and contextualized way to find information about elements, allowing you to quickly verify information and troubleshoot problems. Element Manager also enables you to use folders to create hierarchical groups of elements. See Using Element Manager on page 519.
- **Event Manager** – Enables you to view, clear, sort, and filter events from managed elements. An event can be anything that occurs on the element, for example, a device connected to a Brocade switch has gone off-line. See About Event Manager on page 529 for more information.
- **Path Provisioning** – Enables you to schedule a provisioning task, such as creating zones, to run at a later time. See Path Provisioning on page 757 for more information.
- **Performance Manager** – Provides a graphical representation of the performance history of an element, such as bytes transmitted per second for a switch. For more information, see Performance Manager on page 595.
- **Policy Manager** – Can automatically send an e-mail, generate an event, or run a remote script when an element is being overused or when one of the following occurs:
  - A new element is discovered
  - Successful provisioning occurs
  - An event occurs on one or more specified elements

  See Managing Policies on page 579 for more information.
- **Provisioning Manager** – Enables you to create zones, zone sets, and zone aliases, in addition to storage pools, volumes, and host security groups. For more information, see Provisioning Manager on page 675.
- **Reporter** – Provides detailed reporting on the infrastructure, such as statistics and usage trends. If you want to use Report Optimizer to create reports, contact support for a license that grants you this additional permission. You can only create reports if you login to Report Optimizer directly. Refer to the Quick Start Guide for Report Optimizer for more information about creating reports. The documentation for Report Optimizer can be found by clicking the Help links in Report Optimizer, which might be installed on a separate server depending on your configuration.
- **System Manager** – Is the gateway to many features that let you view details about the discovered elements. System Manager provides a topology that lets you view how the devices in your network are connected. For more information, see Viewing Element Topology and Properties on page 413.

User Interface

The Home page provides a gateway into the functionality for the product. The user interface for the Home page is split into several panes:

- **Top pane** – Provides access to discovery and configuration features, in addition to features that let you change your login profile. The management server can be configured to display the icons for the management server's utilities.
- **Left pane** – The status light and the buttons for the management server's utilities are displayed in the left pane.
- **Right pane** – The output of a feature, such as the topology in System Manager are displayed in the right pane.

Top Pane

The menus and button in the upper-right corner of the page provide the following functionality.

- **Element Manager Search Box** – The Element Manager search box enables you to access Element Manager from anywhere in the product.
  To use the Element Manager search box, enter text in the box, and click the  button. The Element Manager window opens and lists all of the elements whose name includes the string of text that you entered.
- **Configuration** – This menu provides the tools for you to manage the management server, such as saving the database. See Configuring the Management Server on page 333.
- **Security** – This menu enables you to manage users, organizations, roles, and licenses. See Managing Security on page 287.
- **Discovery**– This menu provides the tools for the management server to discover and obtain information from elements in your network. See Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries on page 71 and Discovering Applications, Backup Hosts, and Hosts on page 179.
- **Sign out** – Use this button to log out of the management server (see Signing Out of the Management Server on page 69).
- **Help** – This menu provides access to the online help and the copyright information.

The **Configuration**, **Security**, and **Discovery** buttons only appear if you belong to a role that has System Configuration selected on the Edit Role page (such as the Domain Administrator role). See Managing Security on page 287 for more information.

The links in the upper-left corner let you modify your profile and sign out of the management server.

The status light in the upper-left corner indicates the status of the management server. Normally, it is green.

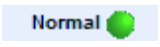

The status light settings are the following:

Scenario	Status Light
Normal	
Discovery	
Getting Topology	
Backup Topology Details	
Include infrastructure details	

When the status light is yellow or red, click the text to the left of the light to access discovery logs quickly (**Discovery** > **View Logs**). If you have other web pages open, refresh those pages for the latest changes.

Note: If the left pane is closed, the status light appears on top of the page.

Left Pane

The buttons used to access the management server's main utilities are displayed in the left pane.

You might not see all of the following utilities, depending on the role assigned to your user account. For example, users assigned to the Help Desk role by default have access only to Application Viewer and Event Manager.

- **Application Viewer** – Enables you to monitor and display information obtained from discovered databases.
- **Backup Manager** – Enables you to monitor the overall status of the backup process, and visualize the backup configuration and recoverability of a file, directory, volume, or server (see About Backup Manager on page 729).
- **Capacity Manager** – Provides a graphical representation of an element's storage capacity in the storage network (see Finding an Element's Storage Capacity on page 553).
- **Chargeback Manager** – Enables you to manage ownership by department, track costs, and assemble business reports. You can view data gathered by Chargeback Manager by element, department, or the entire enterprise. You can show upper management reports with the data that Chargeback Manager gathered. These reports can be e-mailed on a regular schedule (see About Chargeback Manager on page 803 for more information).
- **Event Manager** – Enables you to view, clear, sort, and filter API-generated events (see About Event Manager on page 529).
- **File System Viewer** – Does a recursive lookup on the file system and stores the information in an embedded database. Refer to the File System Viewer for more information.
- **Performance Manager** – Provides a graphical representation of the results obtained from monitoring your elements (see Creating Performance Charts on page 613).

- **Policy Manager** – Enables you to set up rules so that an automated response occurs when a particular event happens, or a value triggers the system (see About Policy Manager on page 579).
- **Provisioning Manager** – Enables you to create zones, zone sets, and zone aliases, in addition to storage pools, volumes, and host security groups (see About the Provisioning Manager on page 675).
- **Reporter** – Provides detailed reporting on the infrastructure, such as statistics and usage trends. To use Report Optimizer to create reports, contact support for a license that grants you this additional permission. To create reports, log on to Report Optimizer directly. For more information about creating reports, see the *Quick Start Guide for Report Optimizer*. The documentation can be found by clicking the Help links in Report Optimizer, which might be installed on a separate server depending on your configuration.

  If you click the Reporter icon and are running HP Storage Essentials from a secure website, you will be told you are leaving a secure Internet connection and be asked if you want to continue.

  If you do not want your users to see this message, change the `SSLOnly` property from false to true, as described in "You Are About to Leave a Secure Connection" Message when Accessing Reporter on page 851
- **System Manager** – Enables you to access systems, and view assets by fabric and logical path (see About System Manager on page 413.

Opening and Closing the Left Pane

The management server enables you open and close the left pane.

You can create additional screen space for the main pane by closing the left pane. Just click the section between the arrows at the upper-right border.

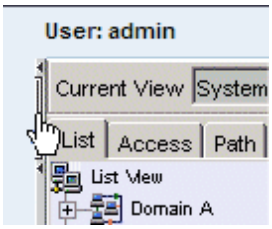

The buttons then move to the top of the page.

When the left pane is closed, you can see the following buttons:

Button	Enables Access to the Following
	Home Page

Button	Enables Access to the Following
	Element Manager
	Application Viewer
	File System Viewer
	Backup Manager
	System Manager
	Capacity Manager
	Performance Manager
	Event Manager
	Provisioning Manager
	Policy Manager
	Chargeback Manager
	Report Optimizer

To open the left pane, click the section between the arrows in the upper-left border.

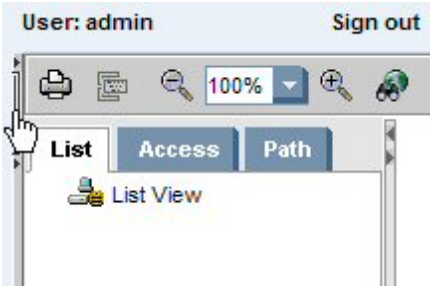

Home Page

After you log into the management server, the Home page appears. The Home page provides an overview of the main features for the management server. To access a feature, click its icon.

Feature	Description	Where to Find More Information
 Element Manager	Element Manager provides a fast and contextualized way to find information about elements, enabling you to quickly verify information and troubleshoot problems.	Management Server Components on page 47
 File System Viewer	File System Viewer provides a single integrated system for managing entire file storage infrastructure by combining file, logical, and physical storage resource monitoring and reporting with active SAN management.	*File System Viewer Guide*
 Policy Manager	Policy Manager can send you e-mail, generate an event, or run a custom script when policy conditions are met. You can define policies to monitor provisioning status, element creation, storage utilization, and incoming events.	Managing Policies on page 579
 icon" /> Application Viewer	Application Viewer lists the applications and their instances running in the storage area network (SAN) and any events associated with them.	"Accessing Information About Applications" in the *Application Guide*.

Feature	Description	Where to Find More Information
 Backup Manager	Backup Manager enables you to keep track of element backups.	Using Backup Manager on page 729
 System Manager	System Manager shows you the topology of your SAN and gives you the ability to explore details about each element.	Viewing Element Topology and Properties on page 413
 Performance Manager	Performance Manager provides detailed performance-management capabilities, enabling you to visualize what you have and how it is performing.	Performance Manager on page 595
 Capacity Manager	Capacity Manager provides a graphical representation of an element's storage capacity and utilization in the storage network.	Finding an Element's Storage Capacity on page 553
 Event Manager	Event Managerkeeps you informed of what is happening with your managed elements. Its filter and report format enables you to easily view, clear, and sort the events you are interested in.	About Event Manager on page 529

Feature	Description	Where to Find More Information
 Provisioning Manager	Provisioning Manager simplifies your SAN-zoning and storage management tasks. The format minimizes errors by giving you easy-to-follow instructions and step-by-step screens.	Provisioning Manager on page 675
 Chargeback Manager	Chargeback Manager enables you to manage departmental ownership, track costs, and assemble business reports, making inquires, such as audits and inventory reviews, easier.	About Chargeback Manager on page 803
 Reporter	Report Optimizer provides reports using data collected by the management server. Some of these reports give you enterprise views of your hosts, switches, storage systems, or applications, while others give you an at-a-glance analysis based on assets, ownerships, chargeback, or performance information. If you want to use Report Optimizer to create reports, contact support for a license that grants you this additional permission. If you want to use Report Optimizer to create reports, you must log on to Report Optimizer directly. For more information, see the *Quick Start Guide for Report Optimizer*.	*Quick Start Guide for Report Optimizer* The documentation for Report Optimizer can be found by clicking the **Help** links in Report Optimizer, which might be installed on a separate server depending on your configuration.

If you click the Reporter icon when you are running HP Storage Essentials from a secure website, you will be told you are leaving a secure Internet connection and will be asked if you want to continue.

If you do not want your users to see this message, change the `SSLOnly` property from false to true, as described in "You Are About to Leave a Secure Connection" Message when Accessing Reporter on page 851

You might not see all of these features, depending on the following:

- The type of license you have. All features might not be available. See the "List of Features" to determine if you have access to all features. The List of Features is accessible from the

Documentation Center (**Help** > **Documentation Center**).

- The role assigned to your user account. For example, users assigned to the Help Desk role have access to Application Viewer and Event Manager, but not to System Manager, Provisioning Manager, Backup Manager and Policy Manager. See Security for the Management Server on page 287 for more information.

Launching the Backup Host Configuration and Discovery Wizard

If you installed the Data Protector Reporter Edition, the Backup Host Configuration and Discovery Wizard is available to you. The Backup Host Configuration and Discovery Wizard assists you perform the initial discovery and configuration tasks using a single user interface. You can invoke the **Backup Host Configuration and Discovery Wizard** from the **Getting Started** page.

Caution: Before you can discover Data Protector, you must complete the requirements provided in Prerequisites for Agentless Discovery of Data Protector on page 194.

The Backup Host Configuration and Discovery Wizard page displays the following tabs:

- Discovery – Helps you discover the hosts running the Data Protector server. It also provides options to configure the discovery details and backup server schedule. See Step 1 – Discover Backup Host Address below.
- Backup – Enables you to set values to retain the backup sessions in the database. See Step 2 – Set Retention Value for Backup Session Data on page 58.
- System – Helps you configure email notifications on reports and policies. You can assign an SMT server from which the management server can send email notifications. Step 2 – Set Retention Value for Backup Session Data on page 58.
- Reports – Provides options to schedule the Report Cache Refresh and configure the Reporter Login. It also provides options to configure the Report Optimizer email and FTP server. See Step 4 – Configure Report Optimizer Settings on page 59.

Step 1 – Discover Backup Host Address

The **Discovery** tab of the configuration wizard helps you configure and discover single or multiple backup servers. Before you discover the backup hosts, you must add and configure the backup hosts.

Data Protector Reporter Edition, by default, does not come with MAPs. Therefore, you cannot discover devices that have MAPs, such as switches, arrays and CIM extension, even though this functionality is displayed in the product and mentioned in the documentation. If you are running Data Protector Reporter without MAPs, you can only discover the backup servers without a CIM extension installed, as described inPrerequisites for Agentless Discovery of Data Protector on page 194.

To configure a backup host:

1. Provide the backup host's IP address, user name, and password as follows:
   - **Single server**:

     In the IP Address/ DNS Name box, type the IP address of the device and provide the host's user credentials.
   - **Multiple servers**:
     - In the **From IP address** box, type the lowest IP address in the range of elements you want to discover.
     - In the **To IP address** box, type the highest IP address of the range of elements you want to discover.
     - Provide the host's user credentials (optional); otherwise, the default credentials will be used.
   - Select **Import** to import the IP addresses for discovery, and do one of the following:
     - Click **Browse** to find an XML file containing the list of IP addresses to be discovered.

       Or
     - In the **Filename** box, provide a complete path to the file.
     - In the **Password** box, type the password for the discovery list. If the discovery list does not have a password assigned to it, leave this field blank.
2. Configure the Discovery Details Schedule as follows:
   - Select **Add the Address to this schedule** option.
   - Select a name from the **Schedule Name** list, or select **New Schedule** to create your own schedule name. Provide a name for the schedule.
   - Type a description for the schedule.
   - Set **Next Schedule Run** date and time. Click the calendar icon to select a date and time.
   - Set **Repeat Interval** period. Type a value for interval and select an unit of time from the list.

   However, you can choose to skip the above step.
3. Configure the Backup server schedule. You can enable the schedules for the following:
   - Image collection
   - Sessions collection
   - Media collection
   - Session monitoring
   - Drive monitoring

4. Click **Add**. This validates the backup configuration details and saves it to the database. The validated IP addresses of the Data Protector backup servers are listed in the **Addresses to Discover** table.

After you configure the backup hosts, you must discover them. You can also edit or delete the backup hosts.

To discover the IP addresses from the **Address to Discover** table:

1. Select the IP addresses you want to discover.
2. Click **Discover**. The following message appears: "Are you sure you want to discover the selected IP addresses?"
3. Click **OK** to start the discovery process. This initiates Discovery Step 1 and Backup Data Collection. The discovery status is displayed as "Discovery is in progress.." You can click on the link to view the discovery logs.

To edit IP addresses from the Address to Discover table:

1. Select the IP addresses you want to edit.
2. Click **Edit**. The Edit window opens.
3. Edit the settings, and then click **Save**. The changes will apply to all the selected backup servers.

   You can also reset your changes by clicking the **Reset** button.

To delete the IP addresses from the Address to Discover table:

1. Select the IP addresses you want to delete.
2. Click **Delete**. The following message appears: "Are you sure you want to delete the addresses?"
3. Click **OK** to delete the selected discovery addresses from the table.

After the configuration and discovery of backup hosts are complete, click **Next** to go to the **Backup** tab.

Step 2 – Set Retention Value for Backup Session Data

The **Backup** tab of the configuration wizard provides options to set the retention value for the Sessions to be stored in the database.

To set the retention value:

1. Type the number of days (a value between 30 and 1098) in the box.
2. Click **Submit**.
3. Click **Next** to go to the **System** tab.

Step 3 – Set Up Email Notifications

The **System** tab of the configuration wizard helps you set notifications from the management server on reports and policies.

To configure email notification:

1. Select **Enable**.
2. In the **Server Name or IP Address** box, type the DNS name or IP address of the Simple Mail Transfer Protocol (SMTP) server, you want to use to send the email notification.
3. In the **Port** box, type the Port number.
4. In the User Name box, type the user name for the SMTP server.
5. In the Password box, type the password of the above user.
6. In the Verify Password box, re-type the password.
7. In the Sender box, type the email address of the sender. This address is displayed in the From box in the email.
8. If you want the replies to go to an email address other than the one specified In the Sender box, type an email address you want to receive the replies to in the Reply box.
9. Click **Save**.

Click **Next** to go to the **Reports** tab.

Step 4 – Configure Report Optimizer Settings

The Reports tab enables you to schedule a reports cache refresh and configure the reporter login. You can also specify the email server to be used for sending the reports and the FTP server to post the reports.

To schedule a reports cache refresh:

1. Select Enable.
2. Click the calendar icon to set the date and time for a scheduled task.
3. In the **Time** box, type the time in 24-hour format with the hour and minutes separate by a colon. For example, 22:15. Click the date on which you want the task to run.
4. Click **Set**.
5. In the **Repeat Interval** box, type an interval. Select a unit of time from the list.
6. Click **Save**.

To configure the reporter login settings:

1. In the **Host Name or IP** box, type the IP of the Reporter Optimizer system.
2. In the **Port Number** box, type the port number.

3. Click **Save**.

You can also reset or change the password. When you click **Reset the password**, the password is set to default.

To configure the Report Optimizer E-mail server:

1. Select a Job Server from the list.
2. In the Domain Name box, type the domain name.
3. In the Host box, type the IP address of the host.
4. In the Port box, type the port number.
5. In the User name box, type the user name.
6. Click **Save**.

To specify the Report Optimizer FTP server:

1. In the Host box, type the IP address of the host.
2. In the Port box, type the port number.
3. In the Account box, type the user name.
4. In the User name box, re-type the user name as above.
5. Type password for the user.
6. Click **Save**.

Click **Close** to complete the discovery and configuration tasks and exit the wizard.

- Select **Do not automatically display this page again** option if you do not want to invoke the Backup Host and Configuration wizard each time you log on to the management server.
- Click **Close** to exit the wizard without completing your configuration tasks. You can, at a later stage, access the wizard by using the **Discovery** menu (**Discovery** > **Wizard**) or **Configuration** menu (**Configuration** > **Wizard**).

Accessing the Management Server

Keep in mind the following:

- If you do not have a license installed, you are asked to install one. If you do not have a valid license, contact Customer Support (go to **Help** > **Documentation Center**).

  To install the license, click the **Import License File** button on the Licenses tab (**Security** > **Licenses**).
- Make sure you do not have pop-up blocking software enabled. The management server uses pop-ups for dialog boxes.
- Verify that the following are enabled on your Web browser:

- Cookies
- JavaScript
- Java

- If you are using a Web browser on Sun Solaris, you must manually install the Java Plug-in to access several components on the client. For more information, see Installing the Java Plug-in on page 63.

  It might take more time to log into the management server during a topology recalculation.

- You can modify several properties so that when users logon to the management server they are told they are logging into a secure system subject to auditing. For more information see Displaying a Banner at Logon on next page.

To access the management server:

1. Enter one of the following in a Web browser:
   - For secure connections:

     `https://machinename`

     In this instance, machinename is the name of the management server.

     To stop receiving a Security Alert message each time you use the HTTPS logon, install the security certificate as described in Installing the Software Security Certificate on page 66. Install the security certificate after you completed all the required steps.

     Or

   - For nonsecure connections:

     Enter the DNS name of the computer in the URL instead of localhost, even if you are running a Web browser directly on the management server. If you use https://localhost to access the management server, you will receive a “Hostname Mismatch” error when you attempt to use System Manager or Performance Manager in the management server.

     `http://machinename`

     In this instance, machinename is the name of the management server.
2. If you receive an error message when you attempt to connect to the management server, the AppStorManager service might be still starting. Wait for it to complete its start script.

   If you see a message like the following one, see Receiving HTTP ERROR: 503 When Accessing the Management Server on page 847:

   ```
   Receiving HTTP ERROR: 503 javax.ejb.EJBException: null;
   CausedByException is: Unexpected Error; nested exception is:
   java.lang.NoClassDefFoundError
   ```
3. In the management server login page, enter `admin` in the **Name** box and `password` in the **Password** box, and then click **Login**.
4. If you are shown the software license agreement and you agree with its terms, click **Accept**.

To prevent the license agreement from being displayed each time you log on to the management server, select **Do not show me this again**.

5. If the management server does not detect a license, you are asked to import the license. Click **Import License File** to install the license.

   The license file can be obtained from customer support.

Displaying a Banner at Logon

You can modify the following properties so that when users logon to the management server they are told they are logging onto a secure system subject to auditing.

- `showBannerInfo` - Controls whether the message will be displayed at logon time. The default value is ShowBannerInfo false, which means no banner is provided at logon time.
- `secureBannerInfo` - Controls the content of the message. The text is formatted through HTML.

Modify the properties as follows:

1. Select **Configuration** > **Product Health** and click **Advanced** in the Disk Space tree.
2. Paste the following in the Custom Properties box:

   ```
   showBannerInfo=true
   ```

   The banner information assigned to the secureBannerInfo property is now displayed when users log on to the management server. To view the message that will be displayed, click **Show Default Properties** at the bottom of the page.
3. To modify the message displayed, enter the following in the Custom Properties box:

   ```
   secureBannerInfo=<Your message here.>
   ```

   Replace `<Your message here.>` with the message you want to provide to users. You can also use HTML to format the message. For example, to span two rows, enter the following:

   ```
   secureBannerInfo=<p>This system is a secure system.</p><p>It is subject to auditing.</p>
   ```

Display Requirements for the Management Server

HPrecommends using a minimum screen resolution of 1024 pixels by 768 pixels. Refer to the documentation for your operating system for information on how to change the screen resolution.

Installing the Java Plug-in

Java 6 Runtime Environment is required to access several features in the management server, such as System Manager. If your Web browser is running on Sun Solaris or Linux, you must manually install the Java plug-in. If you are running another operating system, you can download the plug-in by accessing any page that uses the plug-in. The page will detect you do not have the plug-in installed and it will ask you if you want to install it.

HP recommends that Firefox 2 be upgraded to the latest version. Firefox 3 has been tested with JRE 1.6.

(Windows only clients) If you do not have the Java plug-in already installed and are running Firefox, you must use a web browser other than Firefox to install the plug-in. Firefox is unable to find the missing plug-in. After you install the plug-in, you can use Firefox to run the plug-in.

For information on how to install the Java plug-in on Solaris or Linux operating systems, see the following.

Solaris Clients

Some later versions of the Firefox browser might not work properly if Firefox is upgraded on a system modified previously or with the documented workaround provided in this section. If your web browser does not run properly, reverse the steps provided in this section by unlinking the `libnpjp2.so` file and relinking the `libjavaplugin_oji.so` file.

To install the Java plug-in:

1. Go to the following URL and download the installation file for the Sun JRE when asked:

   http://<management_server>/servlet.html?page=JavaPluginSolaris

   In this instance, <management_server> is the hostname of the management server.

2. Set the executable permission of the downloaded file:

   ```
   # chmod +x downloaded_file_name
   ```

3. In a terminal window, run the downloaded file in a directory where you want the JRE installed. The file installs the Sun JRE on your computer.

   The Java plug-in for your Web browser is available in the following file:.

   ```
   $JRE_HOME/plugin/i386/ns7/libjavaplugin_oji.so
   ```

   In this instance, $JRE_HOME is the directory containing the JRE installation.

4. In a terminal window, go to the `$HOME/.mozilla/plugins` directory. Create a plugins directory if it does not exist in this directory.

5. Remove any existing links in this directory to the Java plug-in.

6. Create a symbolic link to the Java plug-in with the following command:

   ```
   ln -s $JRE_HOME/plugin/sparc/ns7/libjavaplugin_oji.so .
   ```

You must include the dot at the end of the command.

7. If you are a root user on the server and want to make the plug-in available to all users, create a symbolic link in the plugins directory under the browser's installation directory, typically `/opt/SUNWns/plugins`.

   Any existing plug-ins in a user's home directory take precedence over this system-wide plug-in.

8. Restart your Web browser.

Linux Clients

Firefox 3.0.11 on Linux might not work properly if Firefox is upgraded on a system modified previously or with the documented workaround provided in this section. If your web browser does not run properly, reverse the following steps by unlinking the `libnpjp2.so` file and relinking the `libjavaplugin_oji.so` file.

To install the Java plug-in on a Linux 32-bit client:

1. Go to the following URL and download the installation file for the Sun JRE when asked:

   http://<management_server>/servlet.html?page=JavaPluginLinux

   In this instance, <management_server> is the hostname of the management server.

2. Set the executable permission of the downloaded file:

   ```
   # chmod +x downloaded_file_name
   ```

3. In a terminal window, run the downloaded file in a directory where you want the JRE to be installed. The executable installs the Sun JRE on your computer.

   The Java plug-in for your Web browser is available in the following file:

   ```
   $JRE_HOME/plugin/i386/ns7/libjavaplugin_oji.so
   ```

   In this instance, $JRE_HOME is the directory containing the JRE installation.

4. In a terminal window, go to the `$HOME/.mozilla/plugins` directory. Create a `plugins` directory if it does not exist.

5. Remove any existing links to the Java plug-in that are in this directory. You can use the `rm libjavaplugin_oji.so` command in a terminal window to remove an existing symbolic link to the Java plug-in.

6. Create a symbolic link to the Java plug-in by using the following command:

   ```
   # ln -s $JRE_HOME/plugin/i386/ns7/libjavaplugin_oji.so .
   ```

   You must include the dot at the end of the command.

   If you create this symbolic link in any directory other than `$HOME/.mozilla/plugins`, your browser will not be able to use this new Java plug-in.

7. If you are a root user on the server and you want to make the plug-in available to all users,

create a symbolic link to the Java plug-in that is in the `plugins` directory under the browser's installation directory.

Any existing plug-ins in a user's home directory take precedence over this system-wide plug-in.

8. Restart your Web browser.

   At times, the Linux agent might hang on startup on systems 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 source. This means that /dev/random will not return any values if the entropy counter is too low, and programs reading from /dev/random will be blocked until there is enough collected entropy. This can happen on servers with no keyboards, no mice, and no IDE disks.

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

   ```
   # kill -3 java_process_id
   ```

   In this instance, java_porcess_id is the process id of the Java process for the Linux agent. This is not the process id returned by the #./status command.

   The preceding command will generate the stack trace, which should look like the following:

   ```
   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)
   ```

10. To fix the problem, in the `/opt/APPQcime/conf/wrapper.conf` file, under 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 look like the following:

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

Installing the Software Security Certificate

To stop receiving a Security Alert message each time you use the HTTPS logon, install the software security certificate, as described in this section.

Keep in mind the following:

- Enter the DNS name of the computer in the URL instead of localhost. If you use https://localhost to access the management server, you are shown a "Hostname Mismatch" error.

  Management server installation on Linux requires a hardcoded IP of the server's IP address to start the appstormanager service. (Linux requires the Fully Qualified Domain Name and the IP address on separate lines on `/etc/hosts` for the management server to start. This is the OS default.) For more details, see the *Installation Guide*.

- To receive a trusted certificate, you need to purchase a certificate from a trusted entity. (Most browsers have trust relationships set up for Verisign, Entrust, and Thawte, among others.) Set the Common Name (CN) to the name of your management server. Note that the Common Name in the certificate must match the name in the URL.

Installing the Certificate Using Microsoft Internet Explorer

The following steps apply to Microsoft Internet Explorer 6.0. For other versions, see the product documentation.

To install the certificate:

1. Access the management server by entering the following:

   `https://<machinename>`

   In this instance, machinename is the name of the management server.

2. When the security alert message appears, click **OK**.

   If you do not want the Web browser to warn you about a secure connection at any Web site, select the **In the future, do not show this warning** option.

3. When you are told there is a problem with the site's security certificate, click **View Certificate**.

4. When you are shown the certificate information, click **Install Certificate** at the bottom of the screen.

5. When you are shown the Certificate Import Wizard, click **Next** to continue the installation process.

6. Select one of the following:

   - **Automatically select the certificate store based on the type of certificate** – Places the certificate automatically in the appropriate location.

- **Place all certificates in the following store** – Enables you to pick the store where the certificate will be stored.

7. Click **Finish**.
8. When asked if you want to install the certificate, click **Yes**.

Installing the Certificate Using Firefox 1.5

1. Access the management server by entering the following:

   `https://machinename`

   In this instance, machinename is the name of the management server.

2. When the security alert message appears, click the **Accept this certificate permanently** radio button.
3. Click **OK**.

Changing the Security Certificate to Match the Name of the Server

If your users are shown a Security Alert window with the following message, you might want to modify the security certificate so users feel more comfortable with installing the certificate:

```
The name of the security certificate is invalid or does not match the name of the site.
```

You can change the security certificate so that users receive the following message instead:

```
The security certificate has a valid name matching the name of the page you are trying to view.
```

When you change the certificate, you must use the `generateAppiqKeystore` program to delete the original certificate, and the `generateAppiqKeystore` program to create a new certificate and to copy the new certificate to the management server.

Linux

1. Go to the `<Install_Dir>/install` directory and run the following command:

   `eval './usersvars.sh'`

   The quotes must be entered as left single quotes.

2. Go to the `<Install_Dir>/Tools` directory.

   In this instance, Install_Dir is the directory where you installed the management server.

3. To delete the original certificate, enter the following at the command prompt:

   `perl generateAppIQKeyStore.pl del`

   The original certificate is deleted.

If you see an error message when you enter this command, a previous certificate might not have been created. You can ignore the error message.

4. To create a new certificate containing the DNS name of the management server, enter the following at the command prompt:

   ```
   perl generateAppIQKeyStore.pl
   ```

5. If the program is unable to detect a DNS name, enter the following at the command prompt:

   ```
   perl generateAppIQKeyStore.pl create <mycomputername>
   ```

   In this instance, mycomputername is the DNS name of the computer.

6. To copy the new certificate to the management server, enter the following at the command prompt:

   ```
   perl generateAppIQKeyStore.pl copy
   ```

   The new certificate is copied to the correct location.

Windows

1. Go to the `%MGR_DIST%\Tools` directory.
2. To delete the original certificate, enter the following at the command prompt:

   ```
   %MGR_DIST%\Tools> generateAppiqKeystore.bat del
   ```

   The original certificate is deleted.

3. To create a new certificate containing the DNS name of the management server, enter the following at the command prompt:

   ```
   %MGR_DIST%\Tools> generateAppiqKeystore.bat
   ```

4. If the program is unable to detect a DNS name, enter the following at the command prompt:

   ```
   %MGR_DIST%\Tools> generateAppiqKeystore.bat <mycomputername>
   ```

   In this instance, mycomputername is the DNS name of the computer.

5. To copy the new certificate to the management server, enter the following at the command prompt:

   ```
   %MGR_DIST%\Tools> generateAppiqKeystore.bat copy
   ```

   The new certificate is copied to the correct location.

Restarting the Service for the Management Server

By default, the software runs as a service from the time the management server has restarted. If you must restart the service, it is recommended that you restart the service from the Services window, rather than from the command prompt window.

The service must be running for users to access the management server.

To restart the AppStorManager service on a management server:

Linux

Linux management servers require a fixed IP address for starting the AppStorManager service.

1. Open a command prompt window.
2. To stop the management server, enter the following at the command prompt:

   ```
   /etc/init.d/appstormanager stop
   ```

3. To start the management server, enter the following at the command prompt:

   ```
   /etc/init.d/appstormanager start
   ```

4. To see the status of the management server, enter the following at the command prompt:

   ```
   /etc/init.d/appstormanager status
   ```

Windows from the Control Panel

1. Go to the **Administrative Tools** > **Services** window, usually accessible from the Control Panel.
2. Right-click **AppStorManager**
3. Select **Stop** from the menu.
4. To start the management server, right-click **AppStorManager**, and select **Start** from the menu.

Windows from the Command Line

1. To stop AppStorManager, enter the following command at the command prompt:

   ```
   net stop AppStorManager
   ```

2. To start AppStorManager, enter the following command at the command prompt:

   ```
   net start AppStorManager
   ```

Signing Out of the Management Server

Sign out of the management server to prevent unauthorized users from accessing the SAN.

To sign out, click the **Sign Out** link in the upper-left corner of the page.

2 Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries

Before you can use the management server, you must execute the discovery process to make the software aware of the elements on your network, such as switches, storage systems, NAS devices, and tape libraries. Discovery obtains a list of discovered elements and information about their management interface and dependencies.

The management server can discover only elements with a suitable management interface. For information about supported hardware, see the support matrix for your edition.

This section consists of the following information:

Overview of Discovery Steps

Discovery for switches, storage systems, tape libraries and NAS devices consists of several actions:

1. Discover your switches. See Discover Switches on page 84.
2. Discover your storage systems, tape libraries, and NAS devices. See Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries above.
3. To view the topology quickly in System Manager, obtain the topology as described in Building the Topology View on page 150 (*optional*). Keep in mind this step only gathers the information necessary for displaying the topology.

4. Perform Get Details. Get Details is required to obtain detailed information from the elements you discovered, including provisioning information. See Get Details on page 151.

   Running Get Details takes time. You might want to perform this process when the network and the managed elements are not busy. See Get Details on page 151.

Overall Discovery Tasks

Before you begin the discovery process, note the following:

- Get Details does not default to an automatic schedule. In most cases, HP recommends running Get Details once a day during off-peak hours. For more information, see Get Details on page 151.
- Make sure the credentials you enter are correct. When credentials are not supplied, the default user names and passwords are tried for the element.
- In a discovered multipathing configuration where a switch port is set to disabled, after a subsequent Detailed Discovery the host's dependent switches are not shown correctly, and multipathing meta data is incorrect.
- A Discovery or Test Discovery operation by the management server might show a successful contact for elements such as CXWS host operating systems, SMI-S providers and arrays for which you might not be licensed. Detailed information for these unsupported elements is not gathered; however, and there will be no detailed information about the unsupported elements in the management server user interface.
- For elements that support multiple discovery protocols (for example, SNMP and SMI-S), only one protocol at a time is supported for a given element. To change the protocol used to discover an element that has already been discovered, delete the element before attempting to run Get Details gain with a different protocol. For more information, see Deleting Elements from the Product on page 157.
- The management server does not support legacy Fibre Channel arbitrated loop devices connected to switches. This includes devices that are attached through Brocade Quick Loop implementations.
- Elements discovered through SMI-S and hosts discovered with CIM extensions from version 5.1 and later of HP Storage Essentials cannot be added to discovery groups. These elements are listed separately and can be placed independently into scheduled Get Details tasks without being part of a discovery group. This allows you greater flexibility when gathering discovery data. For more information, see Creating Custom Discovery Lists on page 154.
- If you have a problem discovering an element, try enabling Troubleshooting Mode. For more information, see Troubleshooting Mode on page 861.
- Performing a single element refresh of any SAN element can be slow, especially elements with a large number of ports or volumes. A single element refresh can take even longer if performance collectors are running; the task can take several hours. For faster results, add the element to a discovery group and then perform Get Details (discovery groups let you specify elements for discovery). Refer to the installation and user guides for more information about discovery groups.

- To obtain information about the storage area network (SAN), include in the discovery the IP addresses for the following:
  - Fibre channel switch. The Fibre Channel switch contains a list of all elements within the fabric. The management server obtains a detailed listing of all elements connected to the switch fabric.
  - A host containing a Host Bus Adapter (HBA). All Fibre Channel host adapters look for available elements attached to the HBA. This information is gathered by CIM extensions and sent to the management server.

    Until the CIM extensions are installed, the management server is not able to obtain this data when you perform discovery for elements.
  - A proxy connected to the SAN – Include a proxy that has a direct connection or a SAN connection to the management server. An example of a proxy is the EMC Solutions Enabler or Hitachi HiCommand Device Manager. LSI storage systems do not require a proxy, as they can be accessed directly. Make sure the proxy service has started. On a computer running Windows, this can be determined by looking in the Services window. EMC Solutions Enabler requires additional steps for discovery. See Discovering EMC Solutions Enabler on page 107 for more information.
- In this management server version release, you can preserve discovery through the “Win32Provider”. This typically speeds up discovery, and is helpful if you do not want to put the CIM Extension on every Windows host that you want to discover but instead require their internal (WMI) discovery. The user interface has not changed to support this, but there are minor changes to how some information displays:
  - In the View Logs screen, the list of address/provider combinations being “probed” appears in a different order than previously.
  - There is a new property in jboss.properties that you can override with custom property values. This new property, with its default value is: discoveryThreads=10. This determines the number of different threads running simultaneously doing step 1 discovery. You can modify this number to provide a larger or smaller pool of threads used for this purpose. Generally, increasing this number will make Step 1 discovery go faster, within the limitations of system resources,. Use the user interface to change the value.
- Step 1 discovery no longer tests by default for certain device types using certain methods. These are
  - UNIX hosts using older CIM Extension versions (automatic testing is still performed with version 6.0 and later)
  - Other switches using SNMP (automatic testing is still performed via SMI provider)
  - If you still want these discovery options, modify the customProperties.properties file to override certain properties by changing their defaults from “true” to “false.” Use the user interface to change the “true” default to "false" to include these tests.

- discovery.exclude.SnmpSwitchProvider=true
- discovery.exclude.CiscoSNMPProvider=true

  It is strongly recommended you use the user interface to make these changes, (rather than editing the properties file directly). The user interface to do this is described in the “Configuring the Management Server” chapter of the User Guide in the “Managing Product Health, Advanced Settings” section. Be aware that changing the discovery options vary the speed of the discovery process and might affect whether certain devices are discovered.

- If there are device types that you do not have, and do not expect to discover, then you can speed up discovery by excluding other providers by using the user interface to change the corresponding relevant entries to "true":
  - #discovery.exclude.Win32Provider=false
  - #discovery.exclude.SunDotHillProvider=false
  - #discovery.exclude.LSISSI_Provider=false
  - #discovery.exclude.HdsProvider=false
  - #discovery.exclude.ClariionProvider=false
  - #discovery.exclude.EmcProvider=false
  - #discovery.exclude.NetAppFilerProvider=false
  - #discovery.exclude.HPEVA_Provider=false
  - #discovery.exclude.VCProvider=false

  The biggest performance improvement will be realized by excluding the “WIn32Provider”. However, doing so means Windows hosts will only be discovered if a recent CIM Extension has been installed.

The process for making the management server aware of the elements on your network consists of four stages:

1. If you have several switches and storage systems that use the same password and user name, set that password and user name as the default (see Setting Default User Names and Passwords on the facing page).
2. Discover your switches. For information on how to discover the types of switches in your network, see Discover Switches on page 84.
3. Discover your storage systems, NAS devices and tape libraries (see Discover Storage Systems, NAS Devices, and Tape Libraries on page 105).
4. Perform Get Details (**Discovery** > **Details**), which is required to obtain information from your discovered elements.

   Running Get Details takes time. You might want to perform this process when the network and the managed elements are not busy (see Get Details on page 151).

Overview of Discovery Features

Discovery features enable you to:

- Provide up to three default user name and passwords for discovery.
- Import pre-existing discovery lists, so you do not need to re-enter discovery information.
- Save your existing discovery list.
- Modify a discovery entry.
- Remove elements from a discovery list.
- Import or save discovery settings to a file.

Setting Default User Names and Passwords

You can specify up to three default user names and passwords. If several of the elements in the same domain use the same user name and password, assign that user name and password as the default. The management server uses the default user names and passwords if a user name and password are not assigned to an element in the **Setup** screen.

For example, if you have several hosts using the same user name and password, you could enter the default user name and password. If one of the hosts is connected to a storage system with another user name and password, you would also enter this user name and password.

Do not specify the user name and password for the storage system in the individual range because that overrides the default user name and password.

To access a Windows-based device, prefix the user name with domain_name\, as shown in the following example. This is required by the Windows login mechanism.

```
domain_name\user_name
```

In this instance:

- domain_name is the domain name of the element
- user_name is the name of the account used to access that element

Instead of providing a user name and password for an element, you can enter credentials that were provided in the `cxws.default.login` file, as described in "Creating Default Logins for Hosts" in the *Installation Guide*.

To save time, before you begin, make sure the user names and passwords are correct. The software tries each of the default user names and passwords whenever it finds an element.

To add the default user name and passwords:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the**HP Storage Essentialswindow.**
2. Under Discovery Setup, select **Step 1** at the top of the screen.

3. Click **Set Default User Name and Password**. The Setting User Names and Passwords pane appears.

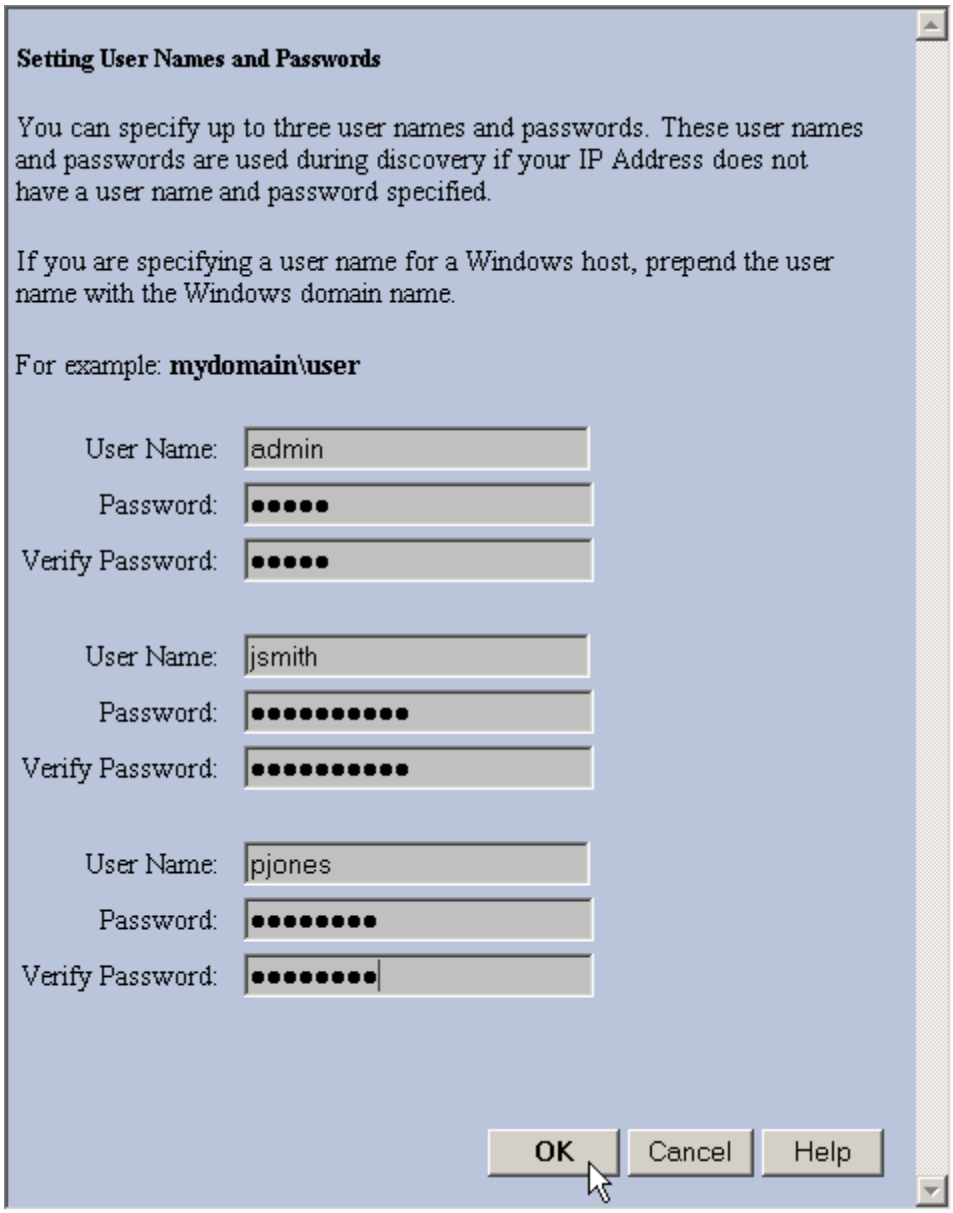

4. In the User Name box, enter the user name for one or more elements.
5. In the Password box, enter the corresponding password for the user name entered in the previous step.
6. In the Verify Password box, re-enter the password.
7. Repeat steps 4 through 6 for other default user names and passwords you want to add.
8. Click **Add System**.

Adding an IP Range for Scanning

The management server can be set up so that when scanning, instead of adding each IP address individually the server can detect a range of IP addresses, automatically populating the list of elements to be discovered.

Keep in mind the following:

- Include in the scanning a proxy server that has a direct connection or a SAN connection to the management server, such as the EMC Solutions Enabler. Make sure the proxy service has started. For Microsoft Windows systems, the check the proxy service status in the Services window.
- You cannot scan an IP range to discover an instance of HiCommand Device Manager that listens on a port other than port 2001. The management server does not allow port numbers in the scanning of IP ranges, so you are not able to specify the port. For more information, see Discovering HDS Storage Systems on page 116.
- Enter a range within the same subnet. The management server cannot scan IP ranges across subnets.
- If you enter an IP range that includes more than one subnet for Step1 Discovery, the discovery mechanism behaves as if the range is all in the same subnet. So if you discover, for example, the range 192.168.1.10-192.168.2.20, it will discover 192.168.1.10-192.168.1.20. Specify starting and ending IP addresses that are in the same subnet on the Discovery page.
- One way to detect multiple IP addresses at one time is to add an IP range for scanning. The management server scans the IP range for elements and populates the discovery list with the elements it could contact. You can then discover those elements.

To add an IP address range to scan:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the**HP Storage Essentialswindow.**
2. Click the **IP Ranges** tab.

   The IP ranges already added are listed.
3. Click **Add Range**.

   The Add Range for Scanning pane appears.

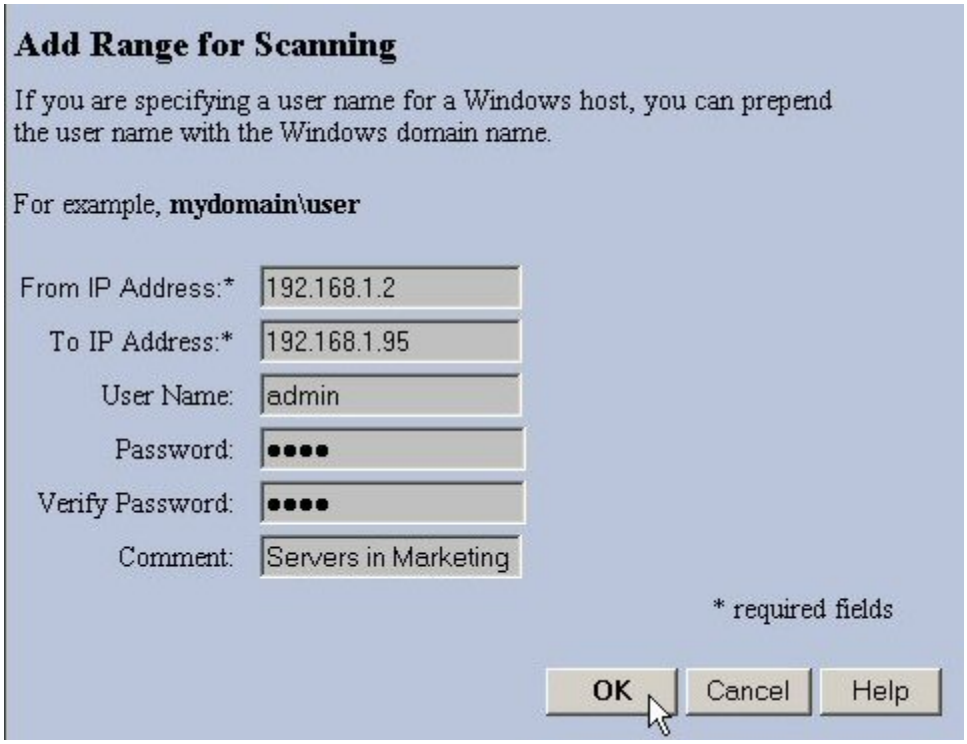

4. In the From IP Address box, enter a lowest IP address in the range to be scanned.
5. In the To IP Address box, enter the highest IP address in the range to be scanned.
6. In the User Name box (*optional*), enter a common user name for elements in the IP range.
7. In the Password box (*optional*), enter a common password for elements in the IP range.
8. In the Verify Password box, re-enter the password.
9. In the Comment box, enter a brief description of the servers; for example, “Servers in Marketing.”
10. Click **OK** to close the Add Range for Scanning pane.
11. Click the **Start Scanning** button on the IP Ranges tab.

The management server scans the IP range and populates the **Addresses to Discover** table on the IP Addresses tab.

Adding a Single IP Address or DNS Name for Discovery

The following steps provide general information on how to discover an element. For more information, see Discover Switches on page 84 and Discover Storage Systems, NAS Devices, and Tape Libraries on page 105.

To add a single IP address or DNS name to discover:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the **HP Storage Essentialswindow.**
2. Under Discovery Setup, select **Step 1** at the top of the screen.
3. On the IP Addresses tab, click **Add Address**.
4. In the IP Address/DNS Name box, enter the IP address or DNS name of the device you want to discover.
5. If you need to enter a port, type a colon (:) after the IP address or DNS name you entered in the **IP Address/DNS Name** box. Then enter a port number; for example:

   `DNSName.companyname.com:1234`

   In this instance, 1234 is the port number.
6. In the User Name box (*optional*), enter the user name. This box can be left blank if you are discovering an LSI storage system or if the element's user name and password are one of the default user names and passwords.

   You can also enter credentials that were provided in the `cxws.default.login` file, as described in "Creating Default Logins for Hosts" in the *Installation Guide*.
7. To set the password, take one of the following actions:
   - If you do not want to do provisioning on a storage system, leave the Password box blank. For LSI storage systems, you must also select the **Do Not Authenticate** option.

     Or
   - To do provisioning on a storage system, enter the corresponding password for controller or proxy and make sure the **Do Not Authenticate** option is not selected.

     Or
   - For all elements other than storage systems, provide the password if it is required for authentication. If the element does not require a password, leave the Password box blank.
8. If you entered a password in the previous step, re-enter the password in the **Verify Password** box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Click **OK**.
11. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Modifying a Single IP Address Entry for Discovery

You can change the user name and password the software uses to access an element. Whenever a user name or password changes on an element that the management server monitors, the management server must be made aware of the change. For example, if the password for a host changes, you must update the management server database with the new password.

The following steps only change the user name and password stored in the database. They do not change the device's user name and password.

To modify a user name or password for discovery:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the HP Storage Essentials> window.
2. Click the **Edit** () button for the element whose user name or password you want to modify.
3. To change the user name, enter the new user name in the User Name box.

   Any special characters can be entered in the User Name box.
4. To add or change a comment, enter a comment in the Comment box.
5. To change the password:
   a. Click **Change password**.
   b. Enter the new password in the New Password box.
   c. Enter the password again in the Verify Password box.
   d. Click **OK** in the Change Password page.
6. Click **OK** in the Edit Address for Discovery page.
7. Select the option **Step 2 – Topology: Select the discovered elements and build the topology view**.
8. Select the element for which you changed the user name and/or password.
9. Click **Get Topology**. The software updates its database with the new user name and/or password.

Removing Elements from the Addresses to Discover List

When you remove IP addresses and/or ranges from the Addresses to Discover list, the elements associated with those IP addresses are not removed from the management server. Only the information that was used to discover them is removed.

To remove items from the Discovery list:

1. Click the **Discovery** icon in the upper-right pane of the HP Storage Essentialshome page.
2. Click **Setup**.
3. Select **Step 1** at the top of the page.

4. Do one of the following:
   - Select the IP addresses and/or IP ranges you want to remove from the list, and then click **Delete**.

     Or

   - Click the **Delete** ( ) button corresponding to the elements you want to remove from the Addresses to Discover list.

The elements associated with these addresses are not removed from the management server. For information about how to remove an element from the management server, see Deleting Elements from the Product on page 157.

Importing Discovery Settings from a File

If you have a previous discovery list, you can import it rather than re-enter the information.

The import discovery settings feature enables you to import the following information to the Discovery list:

- IP addresses to be discovered
- Default user names and passwords, which are encrypted
- Discovery information for applications
- Agentless rules

Note the following:

- 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.
- The Run on Discovery column on the Rule tab (**Discovery** > **Agentless**) is cleared when a discovery list is imported. Run Discovery Step 3 to repopulate the column.
- When you save the discovery settings to a file, the management server is not included in the list and you must perform Discovery Step 1 and Step 3 (Get Details) against the management server. For instructions, see Importing a File below and Rediscovering the Management Server on next page.

Importing a File

To import a file:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the HP Storage Essentialswindow.
2. Click the **Import Settings from File** link.

3. In the Import Settings from File window, do one of the following:
   - Click **Browse** to find the file.

     Or

   - In the Filename box, enter a complete path to the file.
4. In the Password box, enter the password for the discovery list. If the discovery list did not have a password assign to it, leave this field blank.
5. Click **OK**.

The information on the following tabs is updated:

- IP Addresses
- IP Ranges
- Applications

See Rediscovering the Management Server below for adding the management server to the discovery list.

Rediscovering the Management Server

Run discovery Step 1 and Step 3 to rediscover the management server as follows:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the **HP Storage Essentials** window.
2. Click the **Monitoring Product Health** link. The Monitoring Product Health window opens.

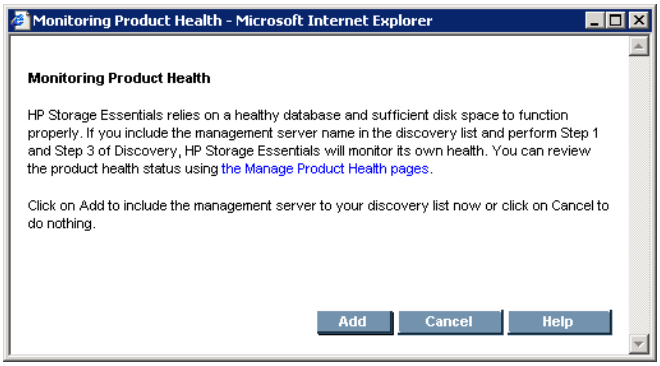

3. Click **Add**. The Discovery Setup, Step 1 – Setup page shows the HP Storage Essentialsmanagement server as localhost.

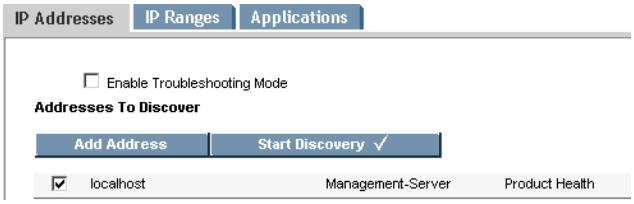

4. Select the check box next to localhost and click **Start Discovery**. When Step 1 discovery is finished, the management server is put into the cxws://localhost discovery group.
5. Select **Discovery** > **Details**.
6. Run **Get Details** for the cxws://localhost discovery group.

Saving Discovery Settings to a File

After you discover your elements, save the discovery settings of the elements in your discovery list.

The **Save Settings to File** link on the Discovery Targets tab enables you save the following information:

- IP addresses to discover
- Default user names and passwords, which are encrypted
- Oracle TNS Listener ports
- Microsoft Exchange configuration
- Agentless rules

To prevent re-entering the information for each instance of the management server, you can import the file for multiple instances.

To save the discovery settings to a file:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the HP Storage Essentials window.
2. Click **Setup** in the upper-right corner.
3. Click the **Save Settings to File** link.
4. In the Password box, enter the password for the management server.
5. In the Verify Password box, enter the password from the previous step, and then click **OK**.
6. When you are asked if you want to open or save the file, choose **Save**.

   The Downloading window appears.
7. Enter a name for the `*.xml` file and select the directory to which you want to save the file. The default name of the file is `DiscoverySettings.xml`.
8. In the Password box, provide a password for the discovery list.

   This password is required later when you import the file. Choose a password you will remember.
9. Click the **Save** button in the Save As window. The file is saved.

Discover Switches

If you have a set of switches managed by more than one SMI-S provider, discover the switches in the same fabric using only one of the SMI-S providers. If you discover the same set of switches through more than one SMI-S provider, the access point used to discover the switches will be deleted from the management server's Discovery pages. To recover from this situation, use the system topology screen and delete each individual switch that was managed by the access point that was deleted (if asked whether you want to delete the access point or the element, select the element). After all the affected switch elements are deleted from within the system topology, change the Discovery list so all switches in that fabric are discovered through the same SMI-S provider. You can now Discover the switches.

The following is an overview of the discovery requirements for switches.

Discovery Requirements for Switches

Element	Discovery Requirements	Additional Information
Brocade switches (SMI-S)	IP address or DNS name, and the user name and password from the Brocade SMI Agent security setup.	See Discovering Brocade Switches on page 88.
Cisco switches	• For **Cisco** switches with SNMPv1 or SNMPv2 connections: Enter the public or private community SNMP string for the switch in the User Name box. All switches in the fabric must have the same community string defined. Leave the Password box empty. *Or* • For **Cisco** switches with SMI-S or SNMPv3 connections: Provide the user name and password for the switch.	See Discovering Cisco Switches on page 90.
QLogic and HP M-Series switches (SNMP)	IP address/DNS name of the QLogic and HP M-Series switch. Enter the SNMP read-only community string as the user name. You do not need to enter a password.	See Discovering QLogic and HP StorageWorks M-Series Switches on page 98.
McDATA switches	Additional steps are required for discovering these switches, and the steps vary according to your network configuration.	See Discovering McDATA Switches on page 99.

Discovering Brocade and McDATA switches through BNA

You can discover Brocade and McDATA switches through Brocade Network Advisor (BNA). When HP Storage Essentials discovers McDATA switches through BNA, HP Storage Essentials treats them as Brocade switches. For example, you can use HP Storage Essentials to view zone aliases on McDATA switches when they are discovered through BNA.

McData and Brocade switches are see in the same access point as Brocade switches. HP Storage Essentials does not display the switch ID.

Migrating Brocade Switches

If you had previously discovered Brocade switches, you can now discover them through BNA, as described in Migrating Brocade Switches from the SMI Agent to BNA Discovery on next page.

Migrating McDATA Switches

If you had previously discovered a McDATA switch directly and you now want to discover it through BNA, migrate the McDATA switch as described in Migrating McDATA Switches from SMI-S to BNA Discovery on next page before using HP Storage Essentials to discover it through BNA.

To discover Brocade and McDATA switches through BNA:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the HP Storage Essentialswindow.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address of the BNA server.
6. In the User Name box, enter the user name for BNA. This field can be left blank if the element's user name and password are one of the default user names and passwords.
7. In the Password box, enter the password for BNA. This box can be left blank if the proxy server's user name and password are one of the default user names and passwords.
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Migrating Brocade Switches from the SMI Agent to BNA Discovery

You can migrate Brocade switches that had been discovered through the SMI agent to Brocade Network Advisor (BNA). When switches that were previously discovered through the SMI agent are discovered through BNA, the following occurs:

- The switches are now associated with the BNA access point.
- The SMI access point is deleted.
- No historical information is lost. See Restoring Statistics from Deleted Elements on page 159 for more information.

You do not have to migrate all of your Brocade switches to discovery through BNA. You can have a mixed environment where some of the switches are discovered through BNA and others are discovered through the SMI agent.

A switch though must be managed through BNA or SMI. Once you add a switch to BNA, you cannot continue to manage it through SMI.

Be sure you want to migrate the switch to BNA discovery before you migrate the switch. If you decide at a later point to return the switch to discovery through the SMI agent discovery, you will lose historical data. To migrate a switch from BNA discovery to discovery through the SMI agent, delete both the BNA access point and the switches discovered through BNA. Then, rediscover the switch through the SMI agent, as described in Discovering Brocade Switches on page 88.

To migrate the Brocade switches to discovery through BNA:

1. Use BNA to discover the Brocade switch that had previously been discovered through SMI-S. The old access points are deleted.
2. Discover the Brocade switch as described in Discovering Brocade and McDATA switches through BNA on previous page.

Migrating McDATA Switches from SMI-S to BNA Discovery

You can migrate your McDATA switches to BNA. When switches that were previously discovered through SMI-S are discovered through BNA, the following occurs:

- The switches are now associated with the BNA access point.
- The SMI-S access point is deleted.
- No historical information is lost.

McDATA switches can only be managed from one management appliance at a time: either McDATA Enterprise Fabric Connectivity Management (EFCM) or Brocade Network Advisor (BNA); however, you do not have to migrate all of your McDATA switches to discovery through BNA. You can have a mixed environment where some of the switches are discovered through BNA and others are discovered through the SMI-S agent, as long as they are management by the associated management appliance, as described in the following table.

Management Appliance	Discovery Method
McDATA Enterprise Fabric Connectivity Management (EFCM)	SMI-S
Brocade Network Advisor (BNA)	BNA

Be sure you want to migrate the switch to BNA discovery before you migrate the switch. If you decide at a later point to return the switch to discovery through SMI-S, you will lose historical data. To migrate a switch from BNA discovery to discovery through SMI-S, delete both the BNA access point and the switches discovered through BNA. Then, rediscover the switch through the SMI agent, as described in Discovering McDATA Switches on page 99.

To migrate McDATA switches to discovery through BNA:

1. Remove the previously discovered McDATA switch from EFCM by using the native tool.
2. Use BNA to discover the McDATA switch.
3. Discover the McDATA switch as described in Discovering Brocade and McDATA switches through BNA on page 85.

How Switches Discovered Through BNA Appear in the Product

Switches discovered through BNA appear differently in the product. You might need to resolve the following issues:

- If the physical name of the switch has never been set, it might display a default name, such as the switch model. See Setting the Physical Name of the Switch below.
- The logical switch might display the same name as the physical switch. See Setting the Virtual Name of a Switch below.
- The fabric name might display a World Wide Name. Setting the Virtual Name of a Switch below.

Setting the Physical Name of the Switch

If the physical name of the switch has never been set, it might display a default name that has a switch model such, such as Brocade 4100 or Silkworm. It is strongly recommended you set the physical name of the switch; otherwise, if you have five switches of the same model, they will all display the same name in System Manager.

You can set the physical name of the switch, by providing a value for the chassisname property on the switch. Refer to the documentation for your switch for more information.

Setting the Virtual Name of a Switch

For example, you might see three switches in Discovery Step 3: switch1, switch2, and switch3. The logical switch might also have the same name as the physical switch. In some cases though, it might not be clear which virtual switches are associated with a certain physical switch.

It is strongly recommended that you set name of the virtual switch so that you can identify the corresponding physical switch. You might also want to use location and job code in the name. You can set the name of the virtual switch by using the switchname command on the switch.

Setting the Name of the Fabric

You might need to set the fabric name for a switch discovered through BNA. When a switch is discovered through BNA, you are shown the name of the fabric in the List tab of System Manager. The fabric name that is displayed in the List tab of System Manager is the name that was set in the BNA discovery tool. If the fabric name was not set, the fabric name is automatically generated with the World Wide Name. The World Wide Name of the primary switch is usually used for the name of the fabric.

Discovering Brocade Switches

The management server uses the Brocade SMI-S Provider (also known as the Brocade SMI Agent) to discover Brocade switches. Before you can discover Brocade swit ches with SMI-S, you must download and install the Brocade SMI Agent software on the proxy server. Do not install the SMI-S provider on the management server. You can download the Brocade SMI Agent and documentation from the following site:

http://www.brocade.com/services-support/drivers-downloads/smi-agent/application_matrix.page

For more information on Brocade SMI Agent versions, see the support matrix for your edition. For information on how to exclude Brocade switches, see Excluding Brocade Switches from SMI-S Discovery on the facing page

To discover Brocade SMI-S switches:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the HP Storage Essentialswindow.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address of the proxy server that is running the SMI-S agent. (Some proxy servers require the following format http://IPADDRESS.)
6. In the User Name box, enter the user name for the SMI-S proxy server. This box can be left blank if one or more of the following conditions are fulfilled:
   - The element's user name and password are one of the default user names and passwords.
   - The element does not require authentication.
7. In the Password box, enter the password for the SMI-S proxy server. This box can be left blank if one or more of the following conditions exists:

- The proxy server's user name and password are one of the default user names and passwords.
- The proxy server does not require authentication.

8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Excluding Brocade Switches from SMI-S Discovery

When HP Storage Essentials discovers Brocade switches through SM-S, it discovers the switches in the fabric and adds the ports to the MAP count. To reduce MAP counts, restrict the number of Brocade switches discovered through SMI-S.

To exclude one or more Brocade switches from SMI-S discovery:

1. Find the serial numbers of the switches you want to exclude:
   - Discover the switches through Discovery Step 1 (**Discovery** > **Setup**). Do not do Discovery Step 2 or Discovery Step 3 (Get Details).
   - Go to the Discovery Step 3 (**Discovery** > **Details**) page, but do not click the **Get Details** button. You are only going to this page to obtain the serial numbers of the switches you want to exclude from discovery.
   - Click one of the switches you want to exclude. You are shown the Navigation page for the switch. The serial number is displayed in the table.
2. Select **Configuration** > **Product Health**, and then click **Advanced** in the **Disk Space** tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Paste the following text into the Custom Properties box:

   ```
   Brocade.smia.excludelist=
   ```

5. Add the serial numbers corresponding to the Brocade switch you want to exclude from discovery. Separate additional serial numbers with a comma, as follows:

   ```
   Brocade.smia.excludelist=ALJ0645D1BK,LX060003058
   ```

   In this instance, ALJ0645D1BK and LX060003058are serial numbers for Brocade switches. You can obtain the serial numbers from the Brocade webtool.

6. When you are done, click **Save**. The product notifies you if a restart of the AppStorManager service is required.
7. Remove the access point for the switches you want to exclude from discovery:
   - Go to the Discovery Step 3 (**Discovery** > **Details**) page, but do not click the **Get Details** button
   - Click the Delete ( ) button for the switches you want to exclude.
8. Restart the AppStorManager service.

Discovering Brocade Switches with Inter-Switch Links

Brocade Inter-Switch Link (ISL) trunking is a feature that enables traffic to be optimally shared across available inter-switch links (ISL). A trunk group logically joins two to four ISLs into one logical ISL. This minimizes congestion in the SAN by optimizing ISL usage, and by managing multiple ISLs as a group instead of individually.

When HP Storage Essentials discovers Brocade switches that utilize the same Inter-Switch Link (ISL), it recognizes the shared ISL trunk. Therefore, if you move one switch into a Discovery group, the other switch (linked by the ISL trunking) is also moved into the same Discovery group.

The topology screens, as well as other related displays, show the ISLs between switches and indicate the total number of ISLs and how many of them are trunked (for supported switches). For example, 6(3 trunked) means 6 is the total number and 3 is how many of them are trunked. ISL trunking information for supported switches is also provided by switch port Properties, switch port Detail table on the Navigation page, and by Reporter in various predefined reports. Note that to obtain ISL switch changes, you must perform a Discovery Get Details.

Discovering Cisco Switches

The management server discovers Cisco switches through SNMP and SMI-S connections depending on the switch model. See the support matrix for your edition for details on supported switch models and firmware revisions.

You must discover all of your Cisco switches using one of the discovery methods:

- SMI-S

  Or

- SNMPv1/SNMPv2

  Or

- SNMPv3

If you previously discovered Cisco switches through SMI-S, you can change the discovery method to SNMP, as described in Migrating Cisco Switches from SMI-S to SNMP Discovery on page 96. Likewise, you can change the discovery method from SNMP to SMI-S, as described in Migrating Cisco Switches using SNMP Discovery to SMI-S Discovery on page 97.
Cisco switches discovered through SMI-S do not show ports with non-Cisco SFP hardware by default. If the SFP or GBIC is not Cisco hardware, the port is not shown in the port table for the

switch. If you want the management server to manage third-party transceivers installed in Cisco switches, paste the following property and its value in the Custom Properties box, which can be found in **Configuration** > **Product Health**> **Advanced**:`cisco.smis.allow.incompatible.port=true`

Pre-Discovery Steps for Cisco SMI-S Discovery

To prepare Cisco switches for SMI-S discovery:

1. Download and install the Cisco cimserver software. For instructions, see the *HP StorageWorks C-Series* document at **http://www.hp.com/go/hpsim/providers**.
2. Enable the CIM Server for Cisco switches discovered through the SMI-S provider, as follows:
   a. On the Cisco switch, enter the following command to display the Common Information Models (CIM) configurations and settings:

      ```
      cisco_switch# show cimserver
      ```

   b. To enter the configuration mode, enter the following:

      ```
      cisco_switch# config t
      ```

   c. To enable access to the server, enter the following:

      ```
      cisco_switch# cimserver enableHttps
      ```

      Or

      ```
      cisco_switch# cimserver enableHttp
      ```

   d. To enable the CIM Server, enter the following:

      ```
      cisco_switch(config)# cimserver enable
      ```

   e. To exit configuration mode, enter the following:

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

For more information go to: http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/sw/san-os/smi-s/developer/guide/proced.html

For steps on how to discover Cisco switches, see Discovering Cisco Switches on previous page.

Pre-Discovery Steps for Cisco Switches Using SNMPv1 or SNMPv2

To prepare the Cisco switch using SNMPv1 or SNMPv2 for discovery:

1. Change the value of discovery.exclude.CiscoSNMPProvider from `true` to `false` as follows:
   a. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
   b. Click **Show Default Properties** at the bottom of the page.

c. Copy discovery.exclude.CiscoSNMPProvider=true.

d. Return to the Advanced page by going to **Configuration** > **Product Health**, and then clicking **Advanced** in the Disk Space tree.

e. Paste the copied text into the Custom Properties box.

f. Replace `true` with `false` so the property and its value are displayed as follows:

```
discovery.exclude.CiscoSNMPProvider=false/
```

g. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

2. Set the same community string for each of the Cisco SNMP switches in the fabric. The community string is not set by default on Cisco SNMP switches. To set the community string:

a. On the Cisco switch, enter the following command to display the Cisco SNMP configurations and settings:

```
cisco_switch# show snmp
```

b. To enter the configuration mode, enter the following:

```
cisco_switch# config t
```

c. To enable the read only community string:

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

d. To exit configuration mode, enter the following:

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

e. To save your changes:

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

For more information about Cisco SNMP, see the documentation at: http://cisco.com/en/US/docs/switches/datacenter/mds9000/sw/nx-os/configuration/guides/sysmgnt/sysmgnt_cli_4_2_published/snmp.html

For steps on how to discover Cisco switches, see Discovering Cisco Switches on page 90.

Pre-Discovery Steps for Cisco Switches Using SNMPv3

The pre-discovery steps for Cisco switches using SNMPv3 require you to create an account and to modify properties within HP Storage Essentials.

Creating Accounts

For account creation use Cisco Fabric Manager, which lets you create an account on all the switches in a fabric with the same credentials and security settings, or use Cisco Device Manager, which will let you create an account on just one switch.

If you create an account with the same credentials on all the switches, you only need to enter the credentials once for Step 1 discovery. If you create accounts with different credentials on each of the switches, you must enter the username and password for each of the different accounts on the Step 1 discovery page.

To use CLI commands for creating an account with Cisco switches:

1. Enter the following at the command prompt for the configuration setting:

   ```
   Cisco-switch1# config
   ```

2. Enter the following at the command prompt for the switch:

   ```
   Cisco1-switch1(config)# username <user> password <password>
   ```

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

Modifying Properties to Enable Discovery of SNMPv3 Switches

You must modify several properties to enable the discovery of Cisco switches using SNMPv3.

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
2. Set the discovery.exclude.CiscoSNMPProvider property to false by pasting the following example into the Custom Properties box:

   ```
   discovery.exclude.CiscoSNMPProvider=false
   ```

3. Set the cimom.cisco.useSNMPv3 property to true by pasting the following example into the Custom Properties box:

   ```
   cimom.cisco.useSNMPv3=true
   ```

4. Set the cimom.cisco.snmp.authenticationProtocol property to MD5 or SHA by pasting the following example into the Custom Properties box:

   ```
   cimom.cisco.snmp.authenticationProtocol=MD5
   ```

   Replace MD5 with SHA if the switches are using SHA.

Value of cimom.cisco.snmp.authenticationProtocol	Definition of Authentication Protocol
MD5	Message Digest 5
SHA	Secure Hash Algorithm -1

5. Set the cimom.cisco.snmp.privacyProtocol property to DES, AES or None by pasting the following example into the Custom Properties box:

   ```
   cimom.cisco.snmp.privacyProtocol=DES
   ```

   If the switches are using a privacy protocol other than DES, replace DES in the example with AES or None.

Value of cimom.cisco.snmp.privacyProtocol	Definition of the Privacy Protocol
DES	Data Encryption Standard
AES	Advanced Encryption Standard
None	No privacy protocol is used.

5. If the product requested that you restart the AppStorManager service after modifying any of the properties, restart the AppStorManager service.
6. See Steps for Discovering Cisco Switches below for information on discovering the switch.

Steps for Discovering Cisco Switches

Make sure to complete the pre-discovery steps according to the discovery type:

Discovery Type	Where to Find Pre-Discovery Steps
SMII-S	Pre-Discovery Steps for Cisco SMI-S Discovery on page 91
SNMPv1 or SNMPv2	Pre-Discovery Steps for Cisco Switches Using SNMPv1 or SNMPv2 on page 91
SNMPv3	Pre-Discovery Steps for Cisco Switches Using SNMPv3 on page 92

Keep in mind the following when discovering Cisco switches with SNMP:

- You can view zones, zone sets, and zone aliases on a Cisco switch, but you cannot use the management server to create, modify, or remove them from a Cisco switch.
- No ports are reported for uninstalled GBICs.
- If you have Cisco switches in multiple fabrics, you can avoid entering the community SNMP string (SNMPv1 and SNMPv2) or the username and password for the switch (SNMPv3 and SMI-S) each time you want to discover a switch in a fabric. Select **Discovery** > **Setup** > **Set Default User Name and Password** and enter the information as provided in the following list:
  - **SNMPv1 and SNMPv2**. Enter the SNMP string as the default user. All switches in the fabric must have the same community string defined. You do not need to provide a password.

- **SNMPv3 and SMI-S**. Enter the user name for the switch as a default user name and enter the password for the switch as the default password. All switches in the fabric must have the same user name and password.

Keep in mind the following when discovering Cisco switches with SMI-S:

- When you discover a Cisco SMI-S switch, you must provide a user name and password.
- Cisco switches discovered through SMI-S do not show ports with non-Cisco SFP hardware by default. If the SFP or GBIC is not Cisco hardware, the port is not shown in the port table for the switch. If you want the management server to manage third-party transceivers installed in Cisco switches, paste the following property and its value in the Custom Properties box, which can be found in **Configuration** > **Product Health**> **Advanced**:`cisco.smis.allow.incompatible.port=true`
- If you are using the SMI-S provider, you must discover all Cisco switches in a fabric. If you discover only one switch, the inactive zones and zone sets that reside on other switches are not displayed on the management server.

To discover Cisco switches:

1. Click **Discovery**, and then click **Setup** in the upper-right pane of the HP Storage Essentialswindow.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the DNS name or primary IP address of the Cisco switch you want to discover.
6. Take one of the following actions:
   - For **Cisco** switches with SNMPv1 or SNMPv2 connections:

     In the User Name box, enter the public or private community SNMP string for the switch. All switches in the fabric must have the same community string defined.

     Or

   - For **Cisco** switches with SMI-S or SNMPv3 connections:

     In the User Name box, enter the switch user name.
7. In the Password and Verify Password fields, take one of the following actions:
   - For **Cisco** switches with SNMPv1 or SNMPv2 connections:

     Leave the Password box blank.

     Or

   - For **Cisco** switches with SMI-S or SNMPv3 connections:

     In the Password box, enter the switch password.

8. Run Discovery Step 1.
9. Do one of the following depending on the discovery type you selected.

Discovery Type	Action
SMI-S	Repeat the previous steps to discover each switch in the fabric.
SNMPv1 or SNMPv2	All Cisco switches are discovered in the fabric. You do not need to repeat the steps for the other switches in the fabric.
SNMPv3	All Cisco switches with the same credentials are discovered in the fabric. If you have switches with different credentials, repeat the previous steps for those switches.

Migrating Cisco Switches from SMI-S to SNMP Discovery

You can convert Cisco switches from SMI-S to discovery for SNMP. Performance statistics, custom name, asset information, custom topology layouts, membership in an organization, and other historical data is removed when the Cisco switch is converted from SMI-S to SNMP discovery. There are slight differences in the information collected from Cisco switches through SMI-S and SNMP. For example, the Port Channel property is not available through SNMP.

To change the discovery method of Cisco switches from SMI-S to SNMP:

1. Delete existing Cisco SMI-S access points from either Step 2 (Topology) or Step 3 (Details). See Deleting Elements from the Product on page 157.
   Historical data about the Cisco switches is lost when you delete the existing access points; however, it is recommended you delete the access points to avoid confusion between the outdated access points and the new access points that will be created when you discover the Cisco switch through SNMP.
2. Change the discovery.exclude.CiscoSNMPProvider property to `false`, and set the same community string set for each of the Cisco SNMP switches in the fabric, as described in Pre-Discovery Steps for Cisco Switches Using SNMPv1 or SNMPv2 on page 91. The community string is not set by default on Cisco switches.
3. If you are switching to SNMPv3 discovery, do the following additional steps:
   a. Change the cimom.cisco.useSNMPv3 property to `true`.
   b. Set the cimom.cisco.snmp.authenticationProtocol to `MD5` or `SHA`.
   c. Set the cimom.cisco.snmp.privacyProtocol property to `DES`, `AES` or `None`

   See Pre-Discovery Steps for Cisco Switches Using SNMPv3 on page 92 for more information about these properties.
4. (*SNMPv1 or SNMPv2*) Change one Step 1 device entry per SAN to conform to SNMPv1 or

SNMPv2 discovery.

5. (*SNMPv1 or SNMPv2*) Change the username to the community string and remove the password. For information about modifying a discovery entry, see Modifying a Single IP Address Entry for Discovery on page 80.
6. Run Step 1 discovery only on one Cisco switch per SAN. For details, see Discovering Cisco Switches on page 90.
7. Run Step 3 discovery on the Cisco switch. The Cisco switch appears in the Default discovery group initially.
8. Repeat Steps 4 through 7 for one Cisco switch per SAN. All Cisco switches with the same credentials in a SAN will be discovered.

Migrating Cisco Switches using SNMP Discovery to SMI-S Discovery

You can convert your Cisco switches using SNMP Discovery to SMI-S discovery. Historical data, such as performance statistics, custom name, asset information, custom topology layouts, membership in an organization, is removed when the Cisco switch is converted from SNMP to SMI-S discovery. There are slight differences in the information collected from Cisco switches through SMI-S and SNMP. For example, the Port Channel property is available through SMI-S, unlike SNMP.

Cisco Switch Ports Discovered through SNMP	Statistics Migrated to SMI-S?
F ports	Yes
TE ports that correspond to SMI-S discovered E ports	Yes
TE ports without a corresponding SMI-S E port	No

To change the discovery method of Cisco switches from SNMP to SMI-S:

1. Delete existing Cisco SMI-S access points from either Step 2 (Topology) or Step 3 (Details) (see Deleting Elements from the Product on page 157).
   Historical data about the Cisco switches is lost when you delete the existing access points. It is, however, recommended that you delete the access points to avoid confusion between the outdated access points and the new access points that will be created when you discover the Cisco switch through SMI-S.
2. Change the cimom.cisco.useSNMPv3 property to *false*.
3. Change the discovery.exclude.CiscoSNMPProvider property to `true`, and follow the steps, as described in Pre-Discovery Steps for Cisco SMI-S Discovery on page 91.
4. Change one Step 1 device entry per SAN to conform to SMI-S discovery. Change the username to the user name for the switch and the password for the switch instead of the

community string. For information about modifying a discovery entry, see Modifying a Single IP Address Entry for Discovery on page 80.

5. Run Step 1 discovery only on one Cisco switch per SAN. For details, see Discovering Cisco Switches on page 90. HP Storage Essentials detects the rest of the switches in the Storage Area Network.
6. Run Step 3 discovery on the Cisco switch. The Cisco switch appears in the Default discovery group initially.
7. Repeat Steps 4 through 6 for each switch in the fabric.

Increasing the Time-out Period and Number of Retries for Cisco Switches in Progress

If you are having difficulty obtaining information from Cisco switches with SNMP connections during Get Details, you might need to increase the time-out period and the number of retries. By default, the management server gives a switch 5 seconds to respond to its requests for information during Get Details. If the switch does not respond the first time, the management server tries again. If it does not receive a response from the switch a second time, the management server says it cannot contact the switch.

To change the time-out period and number of retries for Cisco switches, modify the following properties:

1. Access the management server.
2. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the command for the time out, such as the following for Cisco switches:

   ```
   cimom.Cisco.Snmp.Timeout
   ```

5. Return to the Advanced page.
6. Paste the copied text into the Custom Properties box.
7. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
8. To modify the time-out period, set the corresponding property for your switch in the following table to the number of millisecond you want. The default is 5000 ms.
9. To modify the number of retries, repeat steps 4 through 6 by copying and pasting the cimom.Cisco.Snmp.Retries property. Set the property to the number of retries you want. The default is two retries. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

Discovering QLogic and HP StorageWorks M-Series Switches

The management server discovers QLogic and HP M-Series switches through SMI-S. See the support matrix for your edition for details on supported switch models and firmware revisions.

Keep in mind the following when discovering these switches with SNMP:

- When you discover these switches, you do not need to provide a password.
- The management server does not support provisioning for QLogic and HP M-Series switches. Only the active zone set and its zone members are reported.
- To manage a fabric of QLogic and HP M-Series switches, every switch in the fabric must be included in the discovery list. If a switch is not included in the discovery list, it might show up as a generic host system.
- The product displays Sun Qlogic adapters as two separate adapters.
- No ports are reported for uninstalled blades or GBICs.
- You must perform Get Details to obtain all available information from QLogic SMI-S switches—otherwise, attributes such as vendor, fabric, and port information will be missing for the QLogic SMI-S switches.

Keep in mind the following when discovering these switches with SMI-S:

- Before you can discover these switches with SMI-S, you must download and install the cimserver software. For more information, see the *HP StorageWorks M-Series for p-Class BladeSystems* documentation at http://www.hp.com/go/hpsim/providers.
- A user name and password are required to discover any SMI-S switch.
- You might see an error replicating the switch fabric name for QLogic-based switches. This error can be ignored.

To discover the switches:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the DNS name or primary IP address of the switch you want to discover.
6. In the User Name box, enter the user name for this switch. All SMI-S switches require a user name and password.
7. In the Password box, enter the password for this switch.
8. In the Verify Password box, enter the password of the switch again.

Discovering McDATA Switches

The management server supports the discovery of McDATA switches through SMI-S. The management server can discover multiple instances of Enterprise Fabric Connectivity (EFC) Manager.

The SMI-S setting lets you activate a zone set, in addition to creating, editing, and deleting zones and zone sets. You cannot manage or view information about zone aliases and nicknames are not supported.

Keep in mind the following:

- After an upgrade of the McDATA SMI-S provider to 2.5 from an earlier version, you must delete any existing McDATA switches that were previously discovered with the earlier McDATA provider and then run a new discovery before performing a Get Details.
- If you use EFC Manager, See the support matrix for your edition to verify the version requirements.
- Brocade 5000ni switches running in McDATA mode are managed by the Brocade SMI Agent and not by McDATA SMI-S. For more information, see Discovering Brocade Switches on page 88.
- After you discover a McDATA switch through a proxy, the IP address displayed next to the name of the switch is the IP address of the proxy for the switch in the Discovery, Topology, and Get Details screens. To find the IP address of the switch, click the link for the switch in the Topology or Get Details screen (**Discovery** > **Details**), and then click the **Properties** tab. The Properties tab can also be accessed by double-clicking the switch in System Manager.
- To add, remove, or replace McDATA switches after you discover the service processor, you must perform additional steps, see Managing McDATA Switches on page 103.
- All McDATA switches in a fabric must be managed by the same EFC Manager. Do not have more than one EFC Manager to a fabric for McDATA switches.
- If you want the management server to receive SNMP traps from McDATA switches, do one of the following:
  - If you discovered EFC Manager, enable SNMP trap forwarding to the management server on the EFC Manager, not on the individual switches.

    Or

  - If you discovered McDATA switches directly, enable SNMP trap forwarding on the switches, not in any other management software.

Before you can discover McDATA switches with SMI-S, you must download and install the McDATA SMI-S provider software. For instructions, see the *HP StorageWorks M-Series* documentation at http://www.hp.com/go/hpsim/providers. Check this site periodically to verify that you are running a current version of the SMI-S provider.

Caution: Do not install any providers on the management server.

Note the following when discovering these switches with SMI-S:

- Make sure that EFC Manager is installed and configured or add your switches to the SMI-S provider.
- A McDATA switch cannot be managed by more than one SMI-S provider.
- When you install the SMI-S provider, there are two modes:

  - In coexist mode, the SMI-S provider communicates with EFC Manager and adds all the switches in the managed list of EFC Manager.
  - In direct mode, you must add each switch to the SMI-S provider with its IP address, credentials and switch type. You can use a McDATA's `manageswitch.bat` file to manage the addition and deletion of switches.

- If you selected direct mode during the SMI-S provider installation, when you add switches, you must enter the switch type based on the McDATA model number even if your switch is an OEM model. For more information about the switch type, see your McDATA documentation.
- The SMI-S provider can be installed on the same server as EFC Manager.
- If you selected coexist mode during the SMI-S provider installation you can have only one EFC Manager server.
- If you are using EFC Manager you cannot add managed switches in direct mode. To add switches in direct mode you must remove them from EFC Manager first.
- If the SMI-S provider is installed on a machine other than the HP Storage Essentialsmanagement server, network links between them must pass http traffic on port 5988 (default) or https on port 5989. The port used by the SMI-S provider can be configured. See your switch documentation for more information.

To discover the proxy:

1. Select **Discovery**, and then click **Setup** in the upper-right pane of the HP Storage Essentials window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the proxy you want to discover.
6. In the User Name box, enter the user name.
7. In the Password box, enter the password.

   The user name and password are defined during the SMI-S provider installation. These credentials might be different from the EFC Manager credentials.
8. Re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.

12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovery is complete when the software displays the DISCOVERY COMPLETED message in the Log Messages box.

To obtain more information about the switch, you need to map the topology and obtain element details. See Building the Topology View on page 150 and About Get Details on page 151.

Excluding McDATA Switches from Discovery

Specific McDATA switches can be excluded from discovery by using system properties.

To exclude one or more switches from discovery, modify the cimom.mcdata.exclude property. Set the property cimom.mcdata.exclude to a comma-separated list of Worldwide Names (WWN) of the McDATA switches you want excluded, as shown in the following example:

```
cimom.mcdata.exclude=1000080088A07024,1000080088A0D0B6
```

The management server excludes the switches with the following WWNs: 1000080088A07024 and 1000080088A0D0B6

If the cimom.mcdata.exclude property is not modified, the management server discovers and obtains details from all McDATA switches.

Note: The IP addresses of excluded elements appear in the discovery lists (**Discovery** > **Setup**), topology (**Discovery** > **Topology**), or Get Details lists (**Discovery** > **Details**). The management server does not display additional information about excluded elements in the user interface. The management server, however, does mention in the logs (**Discovery** > **View Logs**) when a provider instance has been created for an excluded element. You can ignore this log message.

To modify the cimom.mcdata.exclude property:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the cimom.mcdata.exclude property.
4. Return to the Advanced page by going to **Configuration** > **Product Health**, and then clicking **Advanced** in the Disk Space tree.
5. Paste the copied text into the Custom Properties box.
6. Make your changes to the text in the Custom Properties box. Remove the pound (#) symbol in front of the property to make sure it is not commented out.
7. Add the WWNs corresponding to the switches you want to exclude from discovery. Separate additional WWNs with a comma; for example:

   ```
   cimom.mcdata.exclude=1000080088A07024,1000080088A0D0B6
   ```

   In this instance, 1000080088A07024 and 1000080088A0D0B6 are the WWNs for McDATA switches.
8. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

Managing McDATA Switches

Whenever you add, remove or replace McDATA switches in an already-discovered service processor, you must make the management server aware of those changes by performing Get Details to obtain information about the new switches from the service processor. For more information about adding switches, see Adding McDATA Switches below.

When you remove switches from the service processor, you must remove them from the management server. For more information about removing switches, see Removing McDATA Switches below.

When you replace McDATA switches, you add and remove the switches. For more information, see Replacing McDATA Switches on next page.

Adding McDATA Switches

After you add switches to an existing service processor, you must perform Get Details, as described in the following steps. If you are adding switches to a service processor that has not been discovered yet, see Discovering McDATA Switches on page 99.

Obtaining details takes some time. You might want to perform this process when the network and the managed elements are not busy.

To run Get Details:

1. Select **Discovery** > **Details**.
2. Click **Get Details**.

During Get Details, the software status light changes from green to red. You can view the progress of gathering details by accessing the logs. For more information, see Viewing Discovery Logs on page 162.

Removing McDATA Switches

After removing switches from a service processor, follow these steps to remove the switches from the management server database:

1. Delete the switches from the user interface by doing the following. These should be the same switches you removed from the service processor.
    a. Click **System Manager** in the left pane.
    b. Right-click the switch you want to delete.
    c. Select **Delete Element** from the menu.
    d. Select the following option:

```
Just delete Switch  [switch_name]. It may reappear the next
time you get topology information or element details.
```

e. Repeat Steps a through d for each switch you want to delete.

2. Verify that the switches were removed from the element list in Discovery Steps 2 and 3 as follows:

   a. To verify that the switches were removed from the element list in Discovery Step 3, select **Discovery** > **Details**.

   b. To verify that the switches were removed from the element list in Discovery Step 2, select **Discovery** > **Topology**.

Replacing McDATA Switches

After replacing switches in the service processor, you must make the management server aware of your changes by removing the old switches from the user interface and then performing Get Details so the management server can discover the new switches. If you are adding switches to a service processor that has not been discovered yet, see Discovering McDATA Switches on page 99.

To swap the switches, follow these steps on the management server:

1. Delete the switches that you removed from the service processor from the user interface:

   a. Click **System Manager** in the left pane.

   b. Right-click the switch you want to delete.

   c. Select **Delete Element** from the menu.

   d. Select the following option:

   ```
   Just delete Switch  [switch_name]. It may reappear the next
   time you get topology information or element details.
   ```

   e. Repeat Steps a through d for each switch you want to delete.

2. Verify that the switches were removed from the element list in Discovery Steps 2 and 3:

   a. To verify that the switches were removed from the element list in Discovery Step 2, select **Discovery** > **Topology**.

   b. To verify that the switches were removed from the element list in Discovery Step 3, select **Discovery** > **Details**.

   c. Select **Discovery** > **Details**.

   d. Click **Get Details**.

During Get Details, the software changes its status light from green to red. You can view the progress of gathering details by selecting **Discovery** > **View Logs**.

When the software finishes Get Details, it displays a message saying Get Details is complete on the **View Logs** page.

Discover Storage Systems, NAS Devices, and Tape Libraries

The following table lists the discovery requirements for storage systems, NAS devices, and tape libraries.

Discovery Requirements for Storage Systems, Tape Libraries, and NAS Devices

Element	Discovery Requirements	Additional Information
3PAR storage systems	Discover the 3PAR storage system directly.	Discovering 3PAR Storage Systems on next page
EMC CLARiiON storage systems	The EMC Navisphere Secure CLI is required for the management server to communicate with the CLARiiON storage system.	Discovering EMC CLARiiON Storage Systems on page 113
EMC Symmetrix storage system (Including EMC Symmetrix DMX storage systems)	Discover the server running the EMC Solutions Enabler.	Discovering EMC Solutions Enabler on page 107
Discovering HP StorageWorks EVA Arrays	Discover the Command View server.	Discovering HP StorageWorks EVA Arrays on page 118
Discovering HP StorageWorks MSA 1000 and 1500 Arrays	Discover the system (proxy) running the MSA 1000/1500 SMI-S provider.	Discovering HP StorageWorks MSA 1000 and 1500 Arrays on page 122
Discovering HP StorageWorks MSA P2000 G2 (2312fc/2324fc) Arrays	Discover the system (proxy) or DNS name of the system (proxy) running the P2000 G2 SMI-S provider.	Discovering HP StorageWorks MSA P2000 G2 (2312fc/2324fc) Arrays on page 123
Discovering HP StorageWorks SVSP	Discover an SVSP environment and the Virtualization Services Manager (VSM).	Discovering HP StorageWorks SVSP on page 125

Element	Discovery Requirements	Additional Information
Discovering HP StorageWorks XP Arrays	Discover the Command View Advanced Edition (AE) or the XP provider.	Discovering HP StorageWorks XP Arrays on page 127
HP and IBM Tape Libraries	Discover the server running the SMI-S provider for the tape library.	Discovering HP and IBM Tape Libraries on page 140

Discovering 3PAR Storage Systems

To discover a 3PAR storage system, the SMI-S server for the 3PAR storage system must be running. By default, the 3PAR SMI-S server is not started on the array. To start the SMI-S server, start the InForm CLI and run the following command:

```
startcim
```

This command starts the SMI-S server within a minute or so.

You do not need to provide the interop namespace because the management server includes the interop namespace for 3PAR storage systems in its default list.

To discover a 3PAR storage system:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the following for the 3PAR storage system you want to discover:

   ```
   <host>
   ```

   In this instance, <host> is the IP address or DNS name of the 3PAR storage system you want to discover.
6. Enter the user name of the storage system. The default username is `3paradm`
7. Enter the password of the storage system. The default password is `3pardata`
8. Re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.

12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.
13. Run Discovery Step 3 to collect array data.

Discovering EMC Solutions Enabler

If you are using a nethost file, edit it to allow the management server to discover the Solutions Enabler and the EMC Symmetrix storage systems it manages. For details, see the EMC documentation.

To discover and collect data from EMC Symmetrix arrays via an EMC Solutions Enabler server, make sure that port 2707 is open between the HP Storage Essentials management server and the EMC Solutions Enabler server. HP Storage Essentials communicates with EMC Solutions Enabler's service/daemon, storsrvd, which listens on port 2707.

To discover EMC Symmetrix storage systems, you must create and configure a VCM volume on the storage system. You must also configure the VCM database on the EMC Solutions Enabler host. See the *EMC Solutions Enabler Symmetrix CLI Command Reference* for details.

If error 214 is present in the discovery log or cimom.log during discovery, the SymAPI server is not licensed for remote connections. You must acquire and install the license before discovery can occur.

Required Licenses

To use all of the features of the management server, such as provisioning, with an EMC Symmetrix storage system, you must have licenses for the following products:

- Base
- DeltaMark
- SYMAPI Server
- Device Masking
- Configuration Manager
- Mapping Solution

Using Only One Subnet

To allow EMC Solutions Enabler to respond correctly, limit the management server to a single subnet. If your management server is on two or more subnets, discovering a storage array through EMC Solutions Enabler might not work. Limiting the management server to a single subnet allows EMC Solutions Enabler to respond correctly.

Using Multiple Solution Enablers to Discover EMC Arrays

HP Storage Essentials does not support the discovery of multiple instances of the EMC Solution Enabler software. If you have multiple instances of EMC Solution Enabler software installed, use the cimom.symmetrix.exclude property to exclude the discovery of all EMC Solution Enabler instances except for one. If you run into an issue with the discovered instance of EMC Solutions Enabler, you can easily modify the `cimom.symmetrix.exclude` property so that a second instance of EMC Solutions Enabler can be discovered. For information on how to modify the cimom.symmetrix.exclude property, see Excluding EMC Symmetrix Storage Systems from Discovery below.

Excluding EMC Symmetrix Storage Systems from Discovery

When multiple EMC Symmetrix storage systems are managed through a single Solutions Enabler, specific storage systems can be excluded from discovery by using system properties.

To exclude one or more Symmetrix storage systems from discovery, modify the cimom.symmetrix.exclude property. Set the property cimom.symmetrix.exclude to a comma-separated list of serial numbers of the storage systems you want excluded, as shown in the following example:

```
cimom.symmetrix.exclude=000183500570,000183610580
```

The management server excludes the storage systems with the following serial numbers: 000183500570 and 000183610580.

If the cimom.symmetrix.exclude property, the management server discovers and obtains details from all EMC Symmetrix Storage Systems managed by discovered Solutions Enablers.

The IP addresses of excluded elements appear in the discovery (**Discovery** > **Setup**), topology (**Discovery** > **Topology**) and Get Details lists (**Discovery** > **Details**). The management server does not display additional information about excluded elements in the user interface. The management server, however, does mention in the logs when a provider instance has been created for an excluded element. You can ignore the message that appears in the logs.

To modify the cimom.symmetrix.exclude property:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command:

   ```
   #cimom.symmetrix.exclude=000183500570,000183500575
   ```

4. Click **Close** to return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Remove the pound (#) symbol in front of the property to make sure it is not commented out. Add the serial numbers corresponding to the Symmetrix storage systems you want to exclude from discovery. Separate additional serial numbers with a comma, as shown in the

following example:

```
cimom.symmetrix.exclude=000183500570,000183500575
```

In this instance, 000183500570 and 000183500575 are serial numbers for Symmetrix storage systems.

7. When you are done, click **Save**.
8. The product notifies you if a restart of the AppStorManager service is required.

Excluding EMC Symmetrix Storage Systems from Forced Device Manager Refresh

The management server obtains most of its information about Symmetrix storage systems from the EMC Solutions Enabler (proxy server) it discovered. If the EMC Solutions Enabler does not have the latest information, the management server also displays the outdated information.

To make the management server aware of any changes, make sure the Solutions Enabler it discovered has the latest information. This can be done by forcing the Solutions Enabler to refresh its data. The management server is then made aware of these changes.

When the Force Device Manager Refresh option is selected, the management server refreshes the discovered EMC Solutions Enabler (proxy server), unless specified. If you do not want an EMC Solutions Enabler to be refreshed, you must assign the Symmetrix storage systems that use the Solutions Enabler to the cimom.emc.skipRefresh property.

To exclude EMC Symmetrix storage systems from a forced refresh:

1. Select **Configuration** > **Product Health** > **Advanced**.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command:

   ```
   #cimom.emc.skipRefresh=000183500570,000183500575
   ```

4. Click **Close** to return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Remove the pound (#) symbol in front of the property to make sure it is not commented out. Add the serial numbers corresponding to the Symmetrix storage systems you want the refresh to skip. Separate additional serial numbers with a comma, as the following example shows:

   ```
   cimom.emc.skipRefresh=000183500570,000183500575
   ```

   In this instance, 000183500570 and 000183500575 are serial numbers for Symmetrix storage systems. One of the ways to find the serial number is to double-click the storage system in System Manager, and then click the **Properties** tab.

7. When you are done, click **Save**.

   The product notifies you if a restart of the AppStorManager service is required.

8. To force a refresh for elements that are not configured to skip the refresh, select the **Force Device Manager Refresh** option on the Get Details page.
9. Click **Get Details**.

EMC Symmetrix Array User Authorization

The Array Authorization Access Control feature allows a Solutions Enabler storage admin to set up Symmetrix user authorization. All information regarding Symmetrix user authorization is stored within the Symmetrix array.

When this feature is enabled for a Symmetrix array, HP Storage Essentials is only able to discover the array or collect data for the array if the user is added to the list of authorized users (see the SYM CLI symauth command). In addition, the user must be assigned a Storage Admin or Admin role. If the user is assigned a lesser role, for example, Monitor—HP Storage Essentials is able to discover the array but will fail to collect certain data such as VMAX masking data. If HP Storage Essentials encounters an authorization error, an Event for the corresponding Symmetrix array is posted with text similar to the following:

```
WARNING: It appears that Access Control is enabled on the Symmetrix
Array 000123456789 and HP Storage Essentials was not authorized to
perform the requested operation(s). Please configure the Array so
that the HP Storage Essentials Server/User is in the Symmetrix
Authorization Users list and is assigned a StorageAdmin or Admin
role. Discovery and Data Collection may fail if user is not in
authorized list. Some data may be missing (i.e. masking data) if the
role is not StorageAdmin or higher. More details on this failure can
be seen in the symapi log on Solutions Enabler 192.168.0.130 server.
The current Authorization Users list can be checked by running the
SYM CLI command "symauth list -user"
```

See the SYM CLI guide or the SYM CLI manpage "symauth.1" in the subdirectory EMC\SYMCLI\Man\Man1 on the Solutions Enabler server for information on viewing and configuring Symmetrix array user authorization data.

Firewall Considerations

By default, HP Storage Essentials communicates with the EMC storsrvd daemon/service running on the Solutions Enabler server using RPC port 2707. This port must be open between the HP Storage Essentials server and the Solutions Enabler server in order for HP Storage Essentials to successfully discover Symmetrix arrays and gather corresponding data.

EMC Symmetrix SSL Certificate Verification

EMC Solutions Enabler APIs began enforcing SSL (Secure Sockets Layer) certificate verification starting with version 6.4. Previous versions of HP Storage Essentials used a pre-6.4 version of the EMC Symmetrix client APIs that was not subject to SSL certificate verification by the Solutions Enabler server (not even with newer versions of Solutions Enabler, for example, 7.0). HP Storage Essentials has updated its EMC Symmetrix client APIs to version 7.1 to enable new features such as thin provisioning and disk tiering. This version of the APIs is subject to SSL certificate verification by the Solutions Enabler server. HP Storage Essentials and EMC administrators need to be aware of the new security features and how to update the default configuration if necessary so that secure communication between HP Storage Essentials and the EMC Solutions Enabler server can be successfully established.

By default, EMC Solutions Enabler 7.0 (and newer) enforces SSL certificate verification during an SSL handshake between the Solutions Enabler server and a Solutions Enabler client (HP Storage Essentials). For HP Storage Essentials (the client) to successfully communicate with an EMC Solutions Enabler server (the server), an SSL handshake must be successfully completed. See the "Client/server Security" section of the *EMC Solutions Enabler Installation Guide* for information on configuring SSL and resolving common issues.

EMC SSL Certificates

EMC SSL certificates are required on both the Solutions Enabler server and the HP Storage Essentialsclient machines. The EMC Solutions Enabler server automatically creates its SSL certificates during installation. HP Storage Essentials automatically creates the required client side EMC SSL certificates during installation. On both the Solutions Enabler and HP Storage Essentials machines, these EMC SSL certificates are located in the following directory:

- Windows:

  `\Program Files\EMC\SYMAPI\config\cert`

- Linux:

  `/var/symapi/config/cert`

This location is a requirement of the EMC APIs and is not configurable on the HP Storage Essentialsmachine. For HP Storage Essentials installed on a 64-bit Windows OS, a directory link is created from `\Program Files (x86)\EMC\SYMAPI\config\cert` to `\Program Files\EMC\SYMAPI\config\cert`.

By default, the SSL certificates contain the fully qualified host name of the machine they were created on. The EMC certificate verification process is sensitive to DNS name resolution. The most common reason for SSL handshake errors between HP Storage Essentials and Solutions Enabler is due to DNS lookup errors on the host name and corresponding IP address of the host name stored in the certificate; for example:

- The EMC SSL certificate of the HP Storage Essentials host contains `mgmtsvrHouston01.datacenterAbc.hp.com`. The IP address is 192.168.0.20.
- The EMC SSL certificate of the Solutions Enabler host contains

`EmcHouston09.datatcenterAbc.hp.com`. The IP address is 192.168.0.130.

During the SSL handshake between the HP Storage Essentials client and the Solutions Enabler server, the Solutions Enabler server receives the HP Storage Essentials SSL client certificate, pulls out the host name, and then tries to verify the certificate by:

- `nslookup mgmtsvrHouston01.datacenterAbc.hp.com`, which returns 192.168.0.20 as expected
- `nslookup 192.168.0.20`, which returns internalHost.datacenterAbc.hp.com, which does not match what was in the certificate (mgmtsvrHouston01.datacenterAbc.hp.com)

The handshake, therefore, fails because nslookup on 192.168.0.20 fails to return the host name specified in the certificate.

The same type of verification occurs on the HP Storage Essentials host, where it attempts to verify the certificate sent by the Solutions Enabler server. In the event of a SSL handshake error, an error is logged in the HP Storage Essentials cimom log. The error message in the HP Storage Essentials cimom log looks similar to the following:

```
SymInitialize() failed with error code 512 (The remote client/server
handshake failed. Please consult symapi and storsrvd log files.
```

On the Solutions Enabler server, a log entry is made in the current storsrvd log that contains additional details about the reason for the SSL handshake failure.

If HP Storage Essentials encounters an SSL handshake failure, an event is posted with text similar to the following:

```
ERROR: EMC Provider SSL handshake error with EMC Solution Enabler
server at 192.168.0.130. HP Storage Essentials is not able to
communicate with the EMC Solutions Enabler server. The most common
reason for this error is DNS issues between the EMC Solutions Enabler
host and HP Storage Essentials host. Each host must be able to (A)
successfully get the IP of the other via nslookup, AND (B) be able to
get back the correct fully qualified host name via a reverse nslookup
on the IP returned from (A). Refer to the HP Storage Essentials
User's Guide for information on EMC security features, common issues,
and workarounds. More details about this SSL handshake error can be
found in the storsrvd log on the Solutions Enabler server at
192.168.0.130.
```

Other common configuration considerations can result in an SSL handshake error when using the default certificates, such as the Solutions Enabler or HP Storage Essentials host being multi-homed or belonging to a cluster. To resolve or work around the SSL handshake issues due to DNS errors or special configurations (multi-homed, clustered, and so forth), there are two basic approaches.

Resolution/Workaround 1: Update the SSL Certificate Using the manage_server_cert Script

The manage_server_cert script resides in the same directory as the certificates on the HP Storage Essentialshost and in the \Program Files\EMC\SYMCLI\bin directory on the Solutions Enabler host. To use the manage_server_cert script on the Solutions Enabler host, you must be in the certificate directory and specify the fully qualified name of the script because the script and the certificates are different directories; for example:

```
C:\Program Files\EMC\SYMAPI\config\cert> "C:\Program
Files\EMC\SYMCLI\bin\manage_server_cert.bat" list
```

In the previous example where the SSL handshake failed due to na nslookup error, the issue could be resolved by updating the SSL certificate on the HP Storage Essentials host by issuing the following command:

```
manage_server_cert.bat create mgmtsvrHouston01.datacenterAbc.hp.com
*.datacenterAbc.hp.com
```

This puts two host entries in the certificate. When the Solutions Enabler server receives this certificate from the HP Storage Essentials client, it does an nslookup on mgmtsvrHouston01.datacenterAbc.hp.com, which returns 192.168.0.20. It then does an nslookup on 192.168.0.20, which returns internalHost.datacenterAbc.hp.com. This matches on the second entry in the certificate and allows the reverse lookup verification to succeed.

If your HP Storage Essentials host cannot successfully resolve the Solutions Enabler server IP or host name using nslookup but can ping it, you must add the Solutions Enabler IP and hostname to the /etc/hosts file. You might also be able to fix the name resolution by adding the Solutions Enabler domain suffix to the /etc/resolv.conf file.

The Client/server Security section of the *EMC Solutions Enabler Installation Guide* provides details on SSL certificates and how to use the manage_server_cert script to manage the certificates for various configurations/scenarios.

Resolution/Workaround 2: Disable Client Certificate Verification on the Solutions Enabler Server

1. Set the storsrvd:security_clt_secure_lvl = NOVERIFY property in the EMC\SYMAPI\config\daemon_options file.
2. Restart the storsrvd daemon by rebooting the Solutions Enabler server or executing the following commands:

   ```
   stordaemon shutdown -immediate storsrvd
   ```

   ```
   stordaemon start storsrvd
   ```

The Solutions Enabler host will accept the HP Storage Essentials SSL certificate without executing the verification step that attempts to verify the host name in the certificate by nslookup and reverse lookup.

Discovering EMC CLARiiON Storage Systems

The EMC Navisphere Secure Command Line Interface must be installed on the management server for the management server to communicate with the CLARiiON storage system. EMC distributes the Navisphere Secure CLI as part of the EMC Navisphere Software Suite.

Contact your EMC representative for more information about obtaining the Navisphere Secure CLI. Distribution rights for the Navisphere Secure CLI belonging to EMC. After you install the Navisphere Secure CLI, restart the AppStorManager service.

When you use Navisphere Secure CLI, the management server is only able to discover CLARiiON arrays using the default port.

Before you discover a CLARiiON storage system, you must have already installed all required software components for that CLARiiON storage system. For more information, see the documentation for your storage system.

CLARiiON storage systems have two controllers called SPa and SPb with IP addresses. To use the provisioning feature in HP Storage Essentials with CLARiiON storage systems, you must discover both controllers. Make sure both controllers are kept in the same discovery group. If you are not planning to use the provisioning feature in HP Storage Essentials, you only need to discover one of the controllers.

In Navisphere Manager, add one of the following to the privilege user section:

- **Windows management server**:

  SYSTEM@<name_of_my_management_server>

  SYSTEM@<IP_of_my_management_server>

- **Linux management server**:

  ROOT@<name_of_my_management_server>

  ROOT@<IP_of_my_management_server>

The variables have the following meaning:

- `<name_of_my_management_server>` is the DNS name of the computer running the management server software
- `<IP_of_my_management_server>` is the IP address of the computer running the management server software

When you use the management server to discover the CLARiiON storage system, provide the IP address for the CLARiiON storage system and the user name and password used to log on to Navisphere.

Discovering LSI Storage Systems

When you discover LSI storage systems and IBM DS3xxx, DS4xxx, or DS5xxx arrays, keep in mind the following:

- Refer to the support matrix for a detailed listing of the models that are supported.
- Discover all controllers on an LSI storage system by entering the IP address of each controller. The management server discovers these controllers as one single storage system.

- The management server must have the User Name box populated to discover the LSI storage system. Even if your LSI storage system does not have a user name set, you must enter something in the User Name box.
- To obtain drive-related statistics, install a proxy host. Ensure that the proxy host has at least one LUN rendered by each controller of the array.
- A license key is required for each storage system and that the key is obtained from the Web site specified on the Activation Card that shipped with your storage system.
- LSI storage systems do not require a password for Get Details. If you do not want to use the management server for provisioning on LSI storage systems, select the **Do Not Authenticate** option. The management server will monitor the LSI storage system, but you will not be able to do provisioning tasks.
- LSI storage systems have two controllers with IP addresses. To use the provisioning feature in HP Storage Essentials with LSI storage systems, you must discover both controllers. Make sure both controllers are kept in the same discovery group. If you are not planning to use the provisioning feature in HP Storage Essentials, you only need to discover one of the controllers.

To discover LSI storage systems:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the controller or proxy you want to discover.
6. Enter the user name in the User Name box. If your LSI storage system does not have a user name, you must enter something in the User Name box, even though the storage system has no user name.
7. Leave the Password box blank if you do not want to do provisioning on the LSI storage system. To do provisioning, enter the corresponding password for controller or proxy.
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. If you do not plan to use provisioning in the product, select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering HDS Storage Systems

HiCommand Device Manager is required for the management server to communicate with an HDS storage system. To discover an HDS storage system, enter the IP address, user name, and password for the server running HiCommand Device Manager. Do not point to the disk array for the storage system.

The management server must be able to access the port that HiCommand Device Manager uses to listen. By default, HiCommand Device Manager listens on port 2001. The management server assumes this configuration at discovery time. If HiCommand Device Manager uses a different port, specify this other port when you discover HiCommand Device Manager.

You cannot scan an IP range to discover an instance of HiCommand Device Manager that listens on a port other than port 2001. The management server does not allow port numbers in the scanning of IP ranges, so you are not able to specify the port.

The management server communicates with HiCommand Device Manager through a nonsecure connection. If you want the management server to communicate with HiCommand Device Manager through a secure sockets layer (SSL) connection, you must modify an internal property or use HTTPS when you discover HiCommand Device Manager. See Communicating with HiCommand Device Manager over SSL on page 892.

To discover an HDS storage system that listens on a port other than 2001:

1. Access the Discovery Setup page (**Discovery** > **Setup**).
2. Click **Add Address**.
3. In the IP Address/DNS Name box, enter the name of the server and the port HiCommand Device Manager uses to listen separated by a colon, as in the following example:

   ```
   proxy2:1234
   ```

   In this instance:

   - proxy2 is the name of the server running HiCommand Device Manager
   - 1234 is the port HiCommand Device Manager uses to listen
4. In the User Name box, enter the user name for accessing HiCommand Device Manager. The default user name for HiCommand Device Manager is the following: system
5. In the Password box, enter the password for accessing HiCommand Device Manager. The default password for HiCommand Device Manager is the following: password
6. In the Verify Password box, re-enter the password for accessing HiCommand Device Manager.
7. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
8. Do not select the Do Not Authenticate option.
9. Click **OK**.

Excluding HDS Storage Systems from Discovery

When multiple HDS storage systems are managed through a single HiCommand Device Manager, specific storage systems can be excluded from discovery by using system properties.

To exclude one or more HDS storage systems from discovery, you must modify the cimom.hds.exclude property. Set the property cimom.hds.exclude to a comma-separated list of serial numbers of the storage systems you want excluded, as shown in the following example:

```
cimom.hds.exclude=61038,61037
```

The management server excludes the storage systems with one of the following serial numbers: 61038 and 61037.

If the cimom.hds.exclude property is not specified, the management server discovers and obtains details from all HDS storage systems managed by the discovered HiCommand Device Manager.

The IP addresses of excluded elements appear in the discovery (**Discovery** > **Setup**), topology (**Discovery** > **Topology**) or Get Details list (**Discovery** > **Details**). The management server does not display additional information about excluded elements in the user interface. The management server, however, does mention in the logs (**Discovery** > **View Logs**) when a provider instance has been created for an excluded element. You can ignore this message that appears in the logs.

To modify the cimom.hds.exclude property:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command:

   ```
   #cimom.hds.exclude=61038,61037
   ```

4. Click **Close** to return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Remove the pound (#) symbol in front of the property to make sure it is not commented out. Add the serial numbers corresponding to the HDS storage systems you want to exclude form discovery. Separate additional serial numbers with a comma, as follows:

   ```
   cimom.hds.exclude=61038,61037
   ```

   In this instance, 61038 and 61037 are serial numbers for HDS storage systems.
7. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

Excluding HDS Storage Systems from Forced Device Manager Refresh

The management server obtains most of its information about the HDS storage systems from the HiCommand Device Manager (proxy server) it discovered. If HiCommand Device Manager, does not have the latest information, the management server also displays the outdated information.

To make the management server aware of any changes, make sure the HiCommand Device Manager it discovered has the latest information. This can be done by forcing the HiCommand Device Manager to refresh its data.

When the Force Device Manager Refresh option is selected, the management server refreshes discovered HiCommand Device Manager (proxy server), unless specified. If you do not want a HiCommand Device Manager to be refreshed, you must assign the HDS storage systems that use HiCommand Device Manager to the cimom.HdsSkipRefresh property.

Before performing any provisioning operations, perform a forced refresh.

To exclude HDS storage systems from a forced refresh:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the following command:

   ```
   # cimom.HdsSkipRefresh=61038,61037
   ```

4. Click **Close** to return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Remove the pound (#) symbol in front of the property to make sure it is not commented out. Add the serial numbers corresponding to the HDS storage systems you want the refresh to skip. Separate additional serial numbers with a comma, as follows:

   ```
   cimom.HdsSkipRefresh=61038,61037
   ```

   In this instance, 61038 and 61037 are serial numbers for HDS storage systems.

   To find the serial number, double-click the storage system in System Manager, and then click the **Properties** tab.
7. When you are done, click **Save**.
   The product notifies you if a restart of the AppStorManager service is required.
8. To force a refresh for elements that are not configured to skip the refresh, select the **Force Device Manager Refresh** option on the Get Details page.
9. Click **Get Details**.

Discovering HP StorageWorks EVA Arrays

The management server supports the following Command View (CV) EVA array discovery options:

- Discovering EVA arrays using Command View 9.x and its SMI-S provider
- Discovering EVA arrays using Command View 8.x, or 9.0.x and the built-in EVA provider

If you upgrade to Command View EVA 9.1 from an earlier version of Command View you must perform a Discovery Step 1, and then Get Details. After performing the discovery, data from previous discoveries using earlier versions of Command View EVA is retained.

If you uninstall Command View EVA 9.1 and install an earlier supported version of CV EVA, you must perform a Discovery Step 1, and then Get Details for the change to take effect.

You can optionally use both Command View EVA 9.0.x (and earlier supported versions of CV EVA) and CV EVA 9.1 concurrently.

Before discovering EVA arrays, note the following:

- HP StorageWorks Command View (CV) EVA must be installed on a server that is not running HP Storage Essentials before you can discover an HP EVA storage system.
- If Command View EVA 9.x and the SMI-S provider are being used, SNMP traps are not used to convey events. You must install and configure the latest version of HP Insight Remote Support (IRS), as described in HP Insight Remote Support Required with CommandView EVA 9.x and the SMI-S Provider on page 534.
- If you have both active and standby Command View (CV) EVA proxy machines, you can discover both the proxy machine that is actively managing the array *and* the proxy machine that is not actively managing the array.

  To discover an EVA, the CV EVA server that is actively managing the EVA must be discovered. The EVA will not be discovered if only the CV EVA server that is passively managing the array is discovered. To continue collecting EVA data when an EVA fails over to the passive Command View EVA server, both the active and passive CV EVA servers must be discovered by HP Storage Essentials. If the passive CV EVA server does not have active management of any EVAs at the time discovery is run, no EVA will be listed for the discovered passive CV EVA server. If at some time an EVA becomes managed by the passive CV EVA server, a Get Details will detect the change and associate the EVA with the CV EVA server.
- If both proxy machines are discovered, keep them in the same discovery group. They can be moved to other discovery groups, but they must be moved together to the same group at the same time. When discovering the proxy machines separately, the machine that has already been discovered must be in the Default discovery group. For more information about discovery groups, see Managing Discovery Groups on page 155.
- If you run Discovery Get Details immediately after moving the EVA to a different Command View (CV) station in an Active/Passive setup, the EVA Volume, HSG, and Pool information under the Properties tab will be missing . To view these properties, wait until the Storage Abstraction Layer (SAL) refreshes and then re-run Get Details.

To determine provisioning support for HP StorageWorks Arrays, see About Setting Up Storage Partitioning on page 694 and About Setting Up Storage Partitioning on page 694.

EVA arrays can only be provisioned if they are actively managed by the Command View server through which they are discovered. When an EVA is discovered by the built-in EVA provider, a cache is created and populated with the current array configuration. Each subsequent cache refresh starts 30 minutes after completion of the previous cache refresh. The cache refresh time depends on the EVA configuration, model, and SAN traffic.

When you perform a provisioning operation (creating, deleting, or modifying a pool or volume), the cache information about provisioning is immediately updated. If you provision an EVA using Command View EVA or a different management station, the cached information about the EVA will not be accurate until the cache is refreshed.

When the EVA firmware and Command View EVA support RAID6, the management server by default creates RAID6 (enhanced) capable storage pools (disk groups) that are capable of RAID 0, 1, 5, and 6 volumes. Basic disk groups continue to be created for configurations that are not RAID6 capable, such as RAID 0, 1, and 5.

When HP EVA volumes are created, the volume name is given a suffix: Vol.Date-'<'*current_date*'>'.'<'*random_numbers*'>' for unique identification.

If the account used to discover Command View EVA has read-only permissions within Command View EVA, you will not be able to subscribe to events, nor will you be able to provision the array.

Discovering EVA Arrays Using Command View EVA

To discover an EVA array, follow these steps on the management server:

1. Select **Discovery** > **Setup** in the upper-right pane of the management server's home page window.
2. Click the **Add Address** button.
3. In the IP Address/DNS Name box, enter the IP address of the Command View server.
4. Enter the user name used to access the Command View server.
5. Enter the password used to access the Command View server.
6. If you entered a password in the previous step, re-enter the password in the Verify Password box.
7. (*Optional*) In the Comment box, enter any additional information. The information entered in this box is displayed in the Comment column in the Addresses to Discover list.
8. Do not select the Do Not Authenticate option.
9. Click **OK**.
10. To start discovering elements on the network, check the check box next to the elements you want to discover, and click **Start Discovery** on the IP Addresses tab.

Obtaining SNMP Traps Using Command View EVA

You must configure Command View EVA so it can send SNMP traps from the EVA to the management server. When the management server receives these SNMP traps, it converts them to WBEM Indications for display in its Event Manager.

Community String Requirements

If you are using the default community strings for Command View EVA and HP Storage Essentials, no changes to the community strings are needed. If the community strings are changed to non-default values, they must be a case-sensitive match.

Caution: Other applications might be using the default community strings to communicate with Command View EVA. If you change the community string in Command View EVA, you might break Command View EVA's connection to other applications. If a change is needed, you should change the community string in HP Storage Essentials to match the string in Command View EVA.

Obtaining SNMP traps from Command View

To obtain SNMP traps from Command View EVA:

1. Verify that the community strings follow the rules in Community String Requirements on previous page. For information on viewing or changing community strings, see one of the following:
   - Viewing or Changing the Community String in HP Storage Essentials below
   - Viewing or Changing the Community String in Command View EVA below.
2. Configure event and host notification. For instructions, see Configuring Event and Host Notification in Command View EVA on next page.

Viewing or Changing the Community String in HP Storage Essentials

To view or change the community string:

1. Select **Configuration** > **Product Health**.
2. Click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the `cimom.snmpTrapListenerCommunityString` variable. The management server uses the value that is listed last, so make sure to search to the end of the page to locate the latest version.
5. Click **Close** to return to the Advanced page.
6. Paste the copied text into the Custom Properties box.
7. Change the value by entering `cimom.snmpTrapListenerCommunityString=<value>`. In this instance, <value> is the desired community string value.
8. Click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

Viewing or Changing the Community String in Command View EVA

To view or change the community string:

1. Open the `C:\Program Files\Hewlett-Packard\Sanworks\Element Manager for StorageWorks HSV\config\cveva.cfg` file in a text editor on the Command View EVA server.
2. Find the following command lines:

```
# Authority. Default = Public
```

```
authority Public
```

3. Change the community string to the desired value. For example, to change the community string to public, enter `authority public`.
4. Restart the service for Command View EVA.

Configuring Event and Host Notification in Command View EVA

See the *HP StorageWorks Command View EVA User Guide* for instructions on configuring Command View EVA event notification.

Discovering HP StorageWorks MSA 1000 and 1500 Arrays

Before you can discover MSA arrays, you must download and install the HP MSA SMI-S Provider software. See the *HP StorageWorks Modular Storage Array* documentation at **http://www.hp.com/go/hpsim/providers** for more information. Check this web site periodically to verify that you are running a current version of the SMI-S provider.

Keep in mind the following:

- To determine provisioning support for HP StorageWorks Arrays, see About Setting Up Storage Partitioning on page 694.
- The Array Configuration Utility (ACU) application should not be running when HP Storage Essentials is using the MSA provider.
- The management URL on the Properties page for the MSA can be used only if the ACU is installed on the same host as the SMI-S provider and the Execution Mode is set to Remote Service. See the ACU Readme file for information about execution modes and how to change them.
- Selective Storage Presentation (SSP) for the array must be enabled for provisioning to work.
- Volumes on MSA 1000/1500 Arrays must be deleted in the reverse order of their creation. For example, if you have six volumes, and you want to delete the second one you created, you must delete the volumes one at a time, starting with the volume created sixth and continuing with the fifth, fourth, third, and then the second.
- The MSA SMI-S provider updates its cache every 4 minutes. If the array is managed by an application other than HP Storage Essentials, changes to the array configuration might not be reflected by a Get Details task that ran before the cache update.

To discover HP MSA storage systems:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentialshome page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.

5. In the IP Address/DNS Name box, enter the IP address or DNS name of the system (proxy) running the MSA 1000/1500 SMI-S provider.
6. Enter the user name used to access the MSA SMI-S provider. The default username and password is administrator.
7. Enter the password used to access the MSA SMI-S provider.
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering HP StorageWorks MSA P2000 G2 (2312fc/2324fc) Arrays

Before you can discover the HP StorageWorks MSA 2000 G2 storage system, you must download and install the HP MSA SMI-S Provider software. See the HP StorageWorks Modular Storage Array documentation at **http://www.hp.com/go/hpsim/providers** for more information. Check this web site periodically to verify that you are running a current version of the SMI-S provider.

Provisioning is not supported for HP MSA P2000 G2 (2312fc/2324fc) storage systems.

To discover HP MSA P2000 G2 (2312fc/2324fc) storage systems:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentialshome page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the system (proxy) or DNS name of the system (proxy) running the P2000 G2 SMI-S provider.
6. Enter the user name used to access the MSA P2000 G2 SMI-S provider. The default user name is **manage**.
7. Enter the password used to access the MSA P2000 G2 SMI-S provider. The default password is **!manage**
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.

9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

In the Host Security Groups page you may notice entries in the Initiators column with value `FF:FF:FF:FF:FF:FF:FF:FF`. Volumes shown are LUNs on the HP MSA P2000 G2 array that were configured with Default Mapping (see the product documentation for the HP MSA P2000 G2 web-based interface).

Discovering HP StorageWorks P2000 G3 Fibre Channel Modular Smart Arrays

Provisioning is not supported for the P2000 G3 FC MSA.

To discover P2000 G3 FC storage systems:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentialshome page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the P2000 G3 FC array.
6. Enter the user name used to access the P2000 G3 FC array. The default user name is **manage**.
7. Enter the password used to access the P2000 G3 FC array. The default password is **!manage**
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

In the Host Security Groups page, you may notice entries in the Initiators column with value FF:FF:FF:FF:FF:FF:FF:FF. Volumes shown are LUNs on the P2000 G3 FC array that were configured with Default Mapping (see the product documentation for the P2000 G3 FC web based interface).

Discovering HP StorageWorks SVSP

The HP StorageWorks SAN Virtualization Services Platform (SVSP) is a centralized management solution for storage pooling and virtual volume provisioning of HP and non-HP storage resources. SVSP services include volume management, data migration, SAN storage-based local and remote replication capabilities, synchronous and asynchronous mirroring, and thin provisioning. The centralized Virtualization Services Manager (VSM), which you can monitor using HP Storage Essentials, enables you to manage virtual disks that span multiple arrays, providing a single view of data across your storage environment.

To discover an SVSP environment, follow the instructions for the specific SVSP configuration implemented on your site(s):

- HP StorageWorks EVA array – see Discovering HP StorageWorks EVA Arrays on page 118.
- HP StorageWorks MSA array – see Discovering HP StorageWorks MSA 1000 and 1500 Arrays on page 122.
- Brocade switches – see Discovering Brocade Switches on page 88.
- Cisco switches – see Steps for Discovering Cisco Switches on page 94.

For all SVSP configurations, use HP Storage Essentials to discover and monitor the HP and SAN devices that make up your SVSP storage infrastructure. When discovering SVSPs, please note the following:

- For SVSP versions earlier than version 3.0.4, the capacity of the SVSP Point-in-Time (PiT) is included in the Storage Volume – Consumed Storage in Blocks property. You cannot identify and display the SVSP PiT instances and their individual sizes.
- For SVSP versions earlier than version 3.0.4, if the error “CIM_ERR_ACCESS_DENIED” occurs on an active VSM when you shut down the passive VSM, stop the SVSP SMI-S server on the active VSM, wait a minute or more, and then restart the SVSP SMI-S server.
- All ports are associated to the main SVSP storage virtualizer, instead of to their respective Virtualization Services Manager (VSM) or Data Path Module (DPM).
- Port Speed and Link Technology is not available from the SVSP SMI-S provider for front-end ports. For certain switches connected to back-end ports, the port speed is not returned and displays as 0 Gb/s.
- To correctly display external back-end topology in HP Storage Essentials, you must complete discovery of back-end storage devices. HP has tested HP EVA arrays and HP MSA P2000 G2 (2312fc/2324fc) arrays. For HP MSA P2000 G2 arrays, configure the Host Security Groups to map the MSA volumes to specific SVSP initiator port WWNs, instead of using default mapping where mapping the MSA volumes only to the generic all hosts (FF:FF:FF:FF:FF:FF:FF:FF) configuration.

- If either of the virtual disks that participate in an SVSP replication pair, such as Sync Mirror groups, are deleted without deleting the replication pair, an error is displayed in HP Storage Essentials during Get Details data collection for that SVSP.

For information about SVSP, see the HP StorageWorks SVSP website at http://h18006.www1.hp.com/products/storage/software/sanvr/index.html?jumpid=reg_R1002_USEN.

For information about the arrays supported by SVSP, visit http://www.hp.com/storage/SPOCK.

For information about infrastructure configurations supported by SVSP, see the SAN Design Guide at http://www.hp.com/go/SANDesignGuide and Operating Systems specific Connectivity Streams at http://www.hp.com/storage/SPOCK.

Discovering an Active Virtualization Services Manager (VSM)

The Virtualization Services Manager (VSM) facilitates creation and management of SVSP virtual disks and data copying between source and destination sites. Each SVSP has at least one VSM server, and the typical installation includes a minimum of two.

A VSM server can be configured as active or passive. A VSM server is active if it is running the VSM service processes from an active server IP address. As a rule, you should discover only active VSM servers in the Step 1 discovery list. If you attempt to include a passive VSM server in the list, a discovery failure of the passive VSM server occurs.

You can only discover the main active VSM server address. Therefore, if SVSP fails over to the passive VSM server, there can be a period of time where the data for SVSP is not refreshed until you fail the SVSP back to the original active VSM server.

To discover an active VSM server:

1. Select **Discovery > Setup** in the upper-right pane of the HP Storage Essentials home page.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or fully qualified domain name (FQDN) of the active VSM.
6. Enter the user name for the SMI-S agent on the active VSM. The default user name for the SMI-S agent is admin.
7. Enter the password for the SMI-S agent on the active VSM. The default password for the SMI-S agent is admin.
8. Re-enter the password in the Verify Password field.
9. (*Optional*) In the Comment field, enter additional information to display in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the **Do Not Authenticate** option.

11. Click **OK**.
12. Click **Start Discovery** on the IP Addresses tab.

The discovery process (Step 1) starts. After it completes, the SVSP is ready for data collection or Get Details (Step 3).

Discovering HP StorageWorks XP Arrays

Discover HP StorageWorks XP Arrays by using one of the following methods:

- Direct Discovery Using the XP Service Processor (SVP) on next page

  Or

- Proxy Discovery Using Command View XP Advanced Edition below

To determine provisioning support for HP StorageWorks Arrays, see Setting Up Storage Partitioning on page 700 and About Setting Up Storage Partitioning on page 694.

Proxy Discovery Using Command View XP Advanced Edition

HP StorageWorks Command View Advanced Edition must be installed on a server that is not running HP Storage Essentials before you can discover an HP XP storage system.

To do a proxy discovery using Command View XP Advanced Edition:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentials home page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the server running Command View Advanced Edition. The default user name for Command View Advanced Edition is the following: system
6. Enter the password used to access Command View Advanced Edition. The default password for Command View Advanced Edition is the following: manager
7. Re-enter the password in the Verify Password box.
8. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
9. Do not select the Do Not Authenticate option.
10. Click **OK**.
11. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Direct Discovery Using the XP Service Processor (SVP)

To do a direct discovery using SVP:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentials home page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address of the XP Service Processor (SVP).
6. Enter the user name used to access the XP storage system.
7. Enter the password used to access the XP storage system.
   The account must be a Partition Storage Administrator account.
8. If you entered a password in the previous step, re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering IBM Storage Systems or IBM SAN Volume Controllers

To discover IBM DS3xxx, DS4xxx, or DS5xxx arrays, use the discovery instructions in Discovering LSI Storage Systems on page 114. Refer to the support matrix for a detailed listing of the models that are supported.

HP Storage Essentials discovers IBM DS6xxx, DS8xxx arrays and SVCs (SAN Volume Controllers) through the IBM CIM agent, which can be embedded or installed on the IBM management console (HMC), depending on the firmware of the array. For installation and configuration information for the IBM CIM agent, refer to the IBM configuration.

To discover an IBM storage system or an IBM SAN Volume Controller (SVC), follow these steps to discover the IBM CIM agent:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentials home page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.

5. In the IP Address/DNS Name box, enter the IP address or DNS name of the system running the IBM CIM agent for the IBM Storage System or SVC you want to discover. In some versions of the product the IBM CIM agent is embedded. If you are not sure whether your IBM CIM agent is embedded, refer to the documentation for your IBM storage system.
6. If a non-default port is used, you must specify the port. Refer to the documentation for your version of the IBM CIM agent to determine the default port.
7. Type a colon (:) after the IP address or DNS name you entered in the **IP Address/DNS Name** box and then, enter a port number; for example:

   ```
   DNSName.companyname.com:1234
   ```

   In this instance, 1234 is the port number.
8. Enter the user name of the IBM CIM agent user.
   - Versions 5.2.1 of the CIM agent – The user name was set when the CIM agent was installed. For additional information about creating a user, see the *DS Open Application Programming Interface Reference Guide*.
   - Versions earlier than CIM agent 5.2.1 – The IBM CIMOM user name and password are defined with the setuser command.
9. Enter the password of the IBM CIM agent user.
10. Re-enter the password in the Verify Password box.
11. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
12. Do not select the Do Not Authenticate option.
13. Click **OK**.
14. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering IBM XIV Arrays

To use HP Storage Essentials to manage and monitor an IBM XIV array, you must discover the array's CIM Agent. The CIM Agent supports only the XIV Array on which the administrative module is located. You must discover a different CIM Agent for each IBM XIV array.

To discover the CIM agent for an IBM XIV array:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentials home page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.

5. In the IP Address/DNS Name box, enter the IP address or DNS name of the system running the administrative module. The IBM CIM agent is installed on the administrative module.
6. Type a colon (:) after the IP address or DNS name you entered in the IP Address/DNS Name box, and then enter a port number; for example:

   ```
   DNSName.companyname.com:5989
   ```

   In this instance, 5989 is the port number.
7. Enter the user name of the SMI-S Agent.

   The CIM client requires a SMI-S Agent user name and password to authenticate its requests. The XIV system administrator must use the IBM XIV Storage System GUI or the IBM XIV command-line interface (XCLI) to create the SMI-S Agent user name and password. To add a user for the SMIS Agent in the System, the XIV system administrator must enter the following in the XCLI (The following would be entered on one line.):

   ```
   smis_add_user user=UserName password=Password password_
   verify=Password [ current_password=Password ]
   ```

   In this instance:

   - UserName is the name of the new user account for the SMI-S agent.
   - Password is the password for the new user account for the SMI-S agent.
8. Enter the password of the SMI-S agent user.
9. Re-enter the password in the Verify Password box.
10. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
11. Do not select the Do Not Authenticate option.
12. Click **OK**.
13. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering Xiotech Storage Systems

You must have Xiotech's Intelligent Control (ICON) software installed. If you do not have the software, contact your Xiotech representative.

To discover a Xiotech storage system:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentials home page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.

5. In the IP Address/DNS Name box, enter the IP address or DNS name for the storage system and its namespace; for example:
   `<IP address/DNS name>/root/cimv2`

   In this instance:

   - <IP address/DNS name> is the IP address or DNS name of the storage system.
   - /root/cimv2 is its namespace.
6. A user name and password are required. Enter anything for the user name and password.
7. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
8. Select the **Do Not Authenticate** option.
9. Click **OK**.
10. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering HP NAS Devices on Windows

To discover an HP NAS device on Windows, you must first install a CIM extension on the device and then modify one of its properties files. For information on installing the CIM extension, see the *Installation Guide*.

To enable NAS support:

1. Connect to the NAS device on which you have installed the CIM extension.
2. Browse to the installation directory and open the APPQCime/conf directory.
3. Copy the nas.properties-sample file and paste a copy into the same directory.
4. Rename the copied file to nas.properties.
5. Open the file and locate the following line:

   ```
   # Set to true to enable NAS data collection; "false" is the
   default

   nas=false
   ```
6. Change the value to true to enable NAS support, as shown in the following example:

   ```
   nas=true
   ```
7. Save your changes and close the file.
8. Restart the CIM extension. See the *Installation Guide* for information about starting CIM extensions.

To discover an HP NAS device on Windows:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentials home page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the HP NAS device you want to discover.
6. Enter the user name of the HP NAS device. You must provide a privileged login.
7. Enter the password used to access the HP NAS device.
8. Re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering HP NAS Devices on Linux

To discover an HP NAS device on Linux, you must first install a CIM extension on the device and then modify one of its properties files. See the *Installation Guide* for information about starting CIM extensions.

To enable NAS support:

1. Connect to the NAS device on which you have installed the CIM extension.
2. Browse to the installation directory and open the /opt/APPQCime/conf directory.
3. Copy the nas.properties-sample file and paste a copy into the same directory.
4. Rename the copied file to nas.properties.
5. Open the file and locate the following line:

   ```
   # Set to true to enable NAS data collection; "false" is the
   default

   nas=false
   ```

6. Change the value to true to enable NAS support, as shown in the following example:

   ```
   nas=true
   ```

7. Save your changes, and then close the file.

8. Restart the CIM extension. For information on starting CIM extensions, see the *Installation Guide*.

To discover an HP NAS device on Linux:

1. Select **Discovery** > **Setup** in the upper-right pane of the HP Storage Essentials home page window.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the HP NAS device you want to discover.
6. Enter the user name of the HP NAS device. You must provide a privileged login.
7. Enter the password used to access the HP NAS device.
8. Re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering NetApp NAS Devices

Keep in mind the following:

- To communicate with the NetApp NAS device through SSL you have the flexibility to set the cimom.providers.netapp.useSSL property to true. This is a global setting and will cause all NetApp NAS devices to communicate using SSL. For more information, see Enabling SSL Communication with a NetApp NAS Device on next page.
- If you want the management server to be able to receive events from a NetApp NAS device, SNMP Event Traps must be enabled on the NetApp NAS device and you must add the IP address of the management server to the NetApp configuration.
- Use Get Details to update element information for NetAPP volumes created in an aggregate. If you use Update Element Data, the information for aggregate volumes is not collected and updated; use the more complete Get Details functionality instead.
- NetAPP devices appear as hosts when initially discovered using Step 1 Discovery. When a Get Details is run, the management server displays them properly.
- You must provide a privileged login, which is one of the following:

  - The root user
  - A user belonging to the Administrators group. This is a predefined group by NetApp.
  - A user belonging to a group that has the following roles: api-*, cli-*, login-http-admin, and at least one of the following: login-console, login-telnet, login-rsh, or login-ssh.
- Administrative HTTP access to the device can be restricted through the httpd.access and httpd.admin.access options. If you are restricting Administrative HTTP access, the management server needs to be registered with the device. This is done by adding the IP addresses of the management server to the httpd.admin.access option. For more information, see the NetApp NAS device documentation.

To discover a NetApp NAS device:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the NetApp NAS device you want to discover.
6. Enter the **User Name** of the NetApp NAS device. You must provide a privileged login.
7. Enter the **Password** used to access the NetApp NAS device.
8. Re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovery Information for NetApp Virtual Filers

To discover a NetApp virtual filer, provide the hostname/IP address of the physical filer along with the credentials of a user with administrator privileges to the NetApp physical filer in Step 1 discovery.

A virtual filer cannot be discovered if the hostname/IP address of the virtual filer is supplied in Step 1 or Step 3 discovery.

Enabling SSL Communication with a NetApp NAS Device

The configuration of the NetApp discovery address is flexible to allow individual filers to be contacted through https, rather than being contacted through an all or nothing approach.

To discover an individual NetApp device using SSL, enter a complete URL in the Step 1 Discovery address field, e.g., https://10.0.1.10:443. In this URL example, doing this will use SSL to contact the filer at 10.0.1.10 on port 443, which is the default NetApp SSL admin port.

If ALL the managed NetApp devices are configured for SSL communications, the cimom.netapp.useSSL custom property might be set to true, as shown in the following example. Doing this will then allow only the IP address to be entered in the Step 1 Discovery addresses field, and the connection will be attempted ONLY using SSL.

The following is an example for configuring to enable SSL communication with ALL of the managed NetApp NAS devices:

1. Select **Configuration** > **Product Health.**
2. Click **Advanced** in the **Disk Space** tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the following property:

   ```
   #cimom.providers.netapp.useSSL=true
   ```

5. Click **Close** to return to the Advanced page.
6. Paste the copied text into the Custom Properties box.
7. Uncomment the cimom.providers.netapp.useSSL property by removing the pound symbol (#) in front of cimom.providers.netapp.useSSL.
8. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

Discovering EMC Celerra

The management server communicates with the EMC Celerra device using the default SSL port (port number 443) configured on the device. If a non-default SSL port is configured on the device, you must specify the port along with the IP address or DNS name separated by a colon when you discover EMC Celerra devices.

You must provide the credentials of a user belonging to the nasadmin group and having the "XML API v2 allowed" Client Access role.

To enable the management server to receive events from the EMC Celerra device, you must enable SNMP traps on the device. You must add the IP address of the management server as an SNMP trap destination with proper community name. For more information on how to configure SNMP trap destination, refer to the EMC Celerra documentation.

To discover EMC Celerra:

1. Modify the discovery.exclude.CelerraProvider property so EMC Celerra can be discovered:
   a. Select **Configuration** > **Product Health**.
   b. Click **Advanced** in the Disk Space tree.

c. Paste the following into the Custom Properties field:

discovery.exclude.CelerraProvider=false

d. When you are done, click **Save**. The product notifies you if a restart of the AppStorManager service is required.

2. Select **Discovery** > **Setup** .
3. Select **Step 1** at the top of the page.
4. Click **Add Address** from the **IP Address** tab.
5. In the IP Address/DNS Name box, specify the IP address or the DNS name of the Control Station of the EMC Celerra device you want to discover.
6. Type the User Name and Password of a Celerra user, which is a part of the **nasadmin** group and has the "XML API v2 allowed" Client access role. By default, EMC Celerra has a user called **nasadmin** with password **nasadmin** that satisfies this criterion.
7. Re-enter the password in the Verify Password box.
8. (*Optional*) In the Comment box, enter any additional information. The information entered in this box appears in the Comment column in the Address to Discovery List (**Discovery > Setup**).
9. Do not select the Do Not Authenticate option.
10. Click **OK**.
11. Click **Start Discovery** on the IP address tab.

Note: After discovery when you run a GAED for Celerra, the values displayed for the – Total Physical Memory and Number of Processors properties for the Celerra Data mover on the Navigation page may differ from the actual values shown in physical hardware.

Discovering EMC Centera

Keep in mind the following:

- To communicate with the Centera device, the management server must be able to access the Centera TCP/UDP port (port number 3218). This port is used for the Application Server Access of the Centera Access node. You might not be able to discover the Centera device using a different port.
- The management server communicates with the Centera Access nodes to get information on the Centera device. However, a Centera Cluster could have more than one Centera Access node. You can provide information on the multiple access nodes during the discovery process by separating them with a semicolon. This enables the management server to communicate with the Centera cluster in case of Centera Access node failure.
- For the management server to be able to receive events from the EMC Centera device, SNMP traps must be enabled on the device. You must add the IP address of the management server as an SNMP trap destination with proper community name. For more information on how to configure SNMP trap destination, see the EMC Centera documentation.

Pre-Discovery Steps for EMC Centera Discovery

Before you can discover an EMC Centera device, you must install an EMC Centera SDK. Contact your EMC representative for more information about obtaining EMC Centera SDK. For information on installation, see Installing EMC Centera SDK on next page

By default, discovery of Centera is disabled.

To enable discovery:

1. Select **Configuration** > **Product Health**.
2. Click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. To enable the management server to accept and display the events generated by the device, ensure the value of the property `cimom.Centera.showEvents` is set to **true**. Setting this property value to **false** blocks the events generated by the device.
5. (*Optional*) To enable the management server to accept and display the events with severity information generated by the EMC Centera device, set the value of the `cimom.Centera.showEvents.showInformationSeverity` property to **true**. By default, this property is set to **false**. Retaining or setting the value of this property to **false** blocks the events with severity information.
6. To enable discovery, copy the following property:

   ```
   discovery.exclude.CenteraProvider=true
   ```

7. Click **Close** to return to the Advanced page.
8. Paste the copied text into the Custom Properties box.
9. Replace **true** with **false** so that the property and its value are displayed as follows:

   ```
   discovery.exclude.CenteraProvider=false
   ```

10. When you are done, click **Save**.
11. Restart the AppStorManager service.

Discovery Steps for EMC Centera

To discover an EMC Centera device:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or the DNS name of the EMC Centera access node, which is a part of the Centera cluster you want to discover.

6. Enter the User Name of the Centera device. You must provide a Centera profile with "Accesscontrol" and "Monitor Cluster" Management Roles.
7. Enter the Password used to access the Centera device.
8. Re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click **Start Discover** on the IP address tab to start discovering elements on the network.

Installing EMC Centera SDK

To install Centera SDK:

Windows management server

1. Extract the contents of the Centera SDK zip file to a folder.
2. Copy all .dll files from the lib32 folder to `%MGR_DIST%\Cimom\lib-native`.
3. Copy the FPLibrary.jar file from the lib folder to `%MGR_DIST%\Cimom\lib\ext`.

Linux management server

1. Extract the contents of the Centera SDK tar file to a folder.
2. Install Centera SDK by running the install script from the extracted folder.
3. Copy the FPLibrary.jar file from the lib folder to `$MGR_DIST/Cimom/lib/ext`.
4. Back up the runcim.sh file in `$MGR_DIST/Cimom/bin` so that you can revert to a previous version if necessary.
5. Open `$MGR_DIST/Cimom/bin/runcim.sh` in a text editor, and edit the LD_LIBRARY_PATH parameter so it resembles the following:

   ```
   LD_LIBRARY_PATH=<SDK_Dir>/lib/32/:$BASE_DIR/lib-native:$LD_LIBRARY_PATH
   ```

   In this instance, <SDK_DIR> is the location where the Centera SDK is installed. By default, the Centera SDK installer script installs the SDK in `/usr/local/Centera_SDK`.

   The example for the LD_LIBRARY_PATH parameter should appear on one line in the runcim.sh file.

   In this instance, /usr/local/Centera_SDK is the location where the Centera SDK is installed.

   Make sure that the text “export LD_LIBRARY_PATH” is still present in the next line in the runcim.sh file.

For example, if the SDK installation directory is /usr/local/Centera_SDK, then LD_
LIBRARY _PATH= /usr/local/Centera_SDK/lib/32/:$BASE_DIR/lib-native:$LD_LIBRARY_
PATH
export LD_LIBRARY_PATH

Discovering Sun NAS Devices

You do not need to provide the interop namespace because it is included in the management servers list of default namespaces.

To discover a Sun NAS Device:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the server running the SMI-S provider for the Sun NAS Devices you want to discover.
6. Enter the user name of the CIMOM/provider for the Sun NAS Devices you want to discover. You must provide a privileged login.
7. Enter the password used to access the CIMOM/provider for the Sun NAS Devices you want to discover.
8. Re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the Do Not Authenticate option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering HP X9000 Network Storage

HP Storage Essentials does not display the following information for some of the discovered X9000 systems:

- Some of the shares that are otherwise shown for a file system in the Fusion Manager
- Network adapter and network port details for the file server nodes
- Details of the dependent client hosts
- Dependent X9000 NAS system for a discovered NAS client

To discover a HP X9000 Network Storage system:

1. Select **Discovery > Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Addresses**.
5. In the **IP address/DNS Name** box, type the IP address or the DNS name of the HP X9000 Network Storage System's Fusion Manager you want to discover.

Note: If the X9000 device has an agile management console configuration, you must use the Cluster VIF or the IP address for discovering the X9000 device. The management server communicates with the X9000 device using the SSL port configured for the Fusion Manager on the device. If the Fusion Manager listens on a port other than 12443, you must specify the port number.

1. To specify the port number, type a colon (:) after the IP address or the DNS name provided in the previous step, and then enter the port number.
2. In the **User Name** box, type the user name of the device. The default user name is ibrix.
3. In the Password box, type the password that was assigned to this user.
4. Re-enter the password in the **Verify Password** box.
5. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery > Setup**).
6. Do not select the **Do Not Authenticate** option.
7. Click **OK**.
8. Click **Start Discovery** on the IP Addresses tab to start discovering elements on the network.

Discovering HP and IBM Tape Libraries

Before you can discover an HP or IBM tape library, you must download and install the corresponding SMI-S provider software.

- **IBM Tape Libraries.** See your IBM documentation and the support matrix for your edition for information about the SMI-S provider for IBM tape libraries.
- **HP Tape Libraries.** Download HP StorageWorks Command View for Tape Libraries (TL) Software from http://www.hp.com/go/support. Custom install the HP StorageWorks Command View TL Software, so you can select the SMI-S provider for HP tape libraries during the installation. All the libraries that Command View TL manages are discoverable when the SMI-S provider for HP Tape Libraries service is running. Refer to http://www.hp.com/go/hpsim/providers for more details.HP Storage EssentialsBackup Manager can also discover HP tape libraries through the supported backup software.

To discover an HP or IBM tape library:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.
3. Click the **IP Addresses** tab.
4. Click **Add Address**.
5. In the IP Address/DNS Name box, enter the IP address or DNS name of the SMI-S provider for the tape library.
6. Enter the user name and password of the provider running the tape library. The user name and password are the provider's user name and password, not the credentials for the operating system's user name. The default user name/password for IBM is cimuser/cimpass and for HP it's administrator/administrator unless you've made changes.
7. Enter the **Password** of the system running the tape library.
8. Re-enter the password in the Verify Password box.
9. (*Optional*) In the Comment box, enter any additional information. The information entered into this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Do not select the **Do Not Authenticate** option.
11. Click **OK**.
12. Click the **Start Discovery** button on the IP Addresses tab to start discovering elements on the network.

Discovering HP P4000 Devices

To discover an HP P4000 cluster device:

1. Click **Discovery**, and then click Setup in the upper-right pane of the HP Storage Essentials window.
2. Under Discovery Setup, select Step 1 at the top of the screen.
3. On the IP Addresses tab, click **Add Address**.
4. Enter the virtual IP, VIP, of the cluster.

The device should appear in the details screen with a device name consisting of the management group name and name of the cluster; for example, ManagementGroup0:Cluster0.

Related Topic:

HP P4000 iSCSI Information on page 147

HP P4000 System and Device Topology

The iSCSI cluster is linked to hosts through direct IP connections. HP Storage Essentials does not discover or display end-to-end IP topology through switches. IP links are shown as links on the system topology directly to the consuming device.

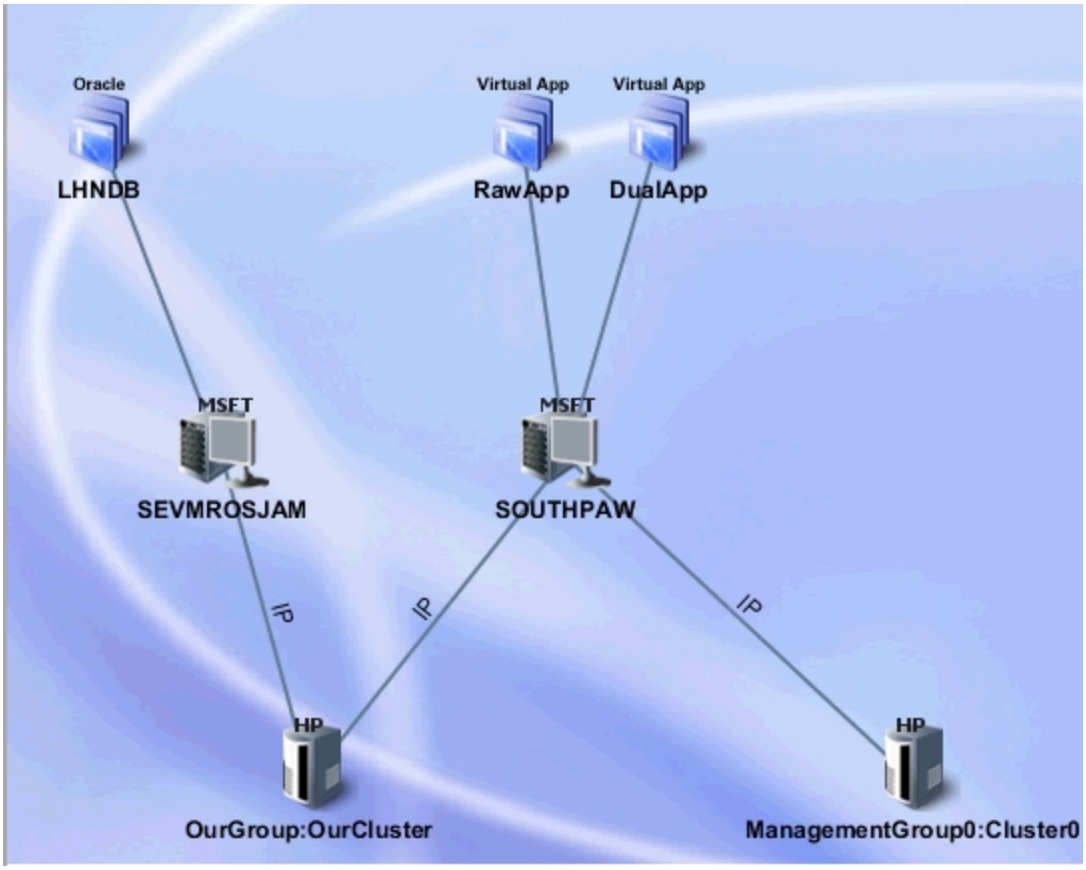

A more detailed graphical view of end-to-end application stitching can be viewed through the device topology page. The following illustration shows how an application, either mounted on a logical drive or raw partition on a host, is linked to an IP network through a particular host network port to an HP P4000.

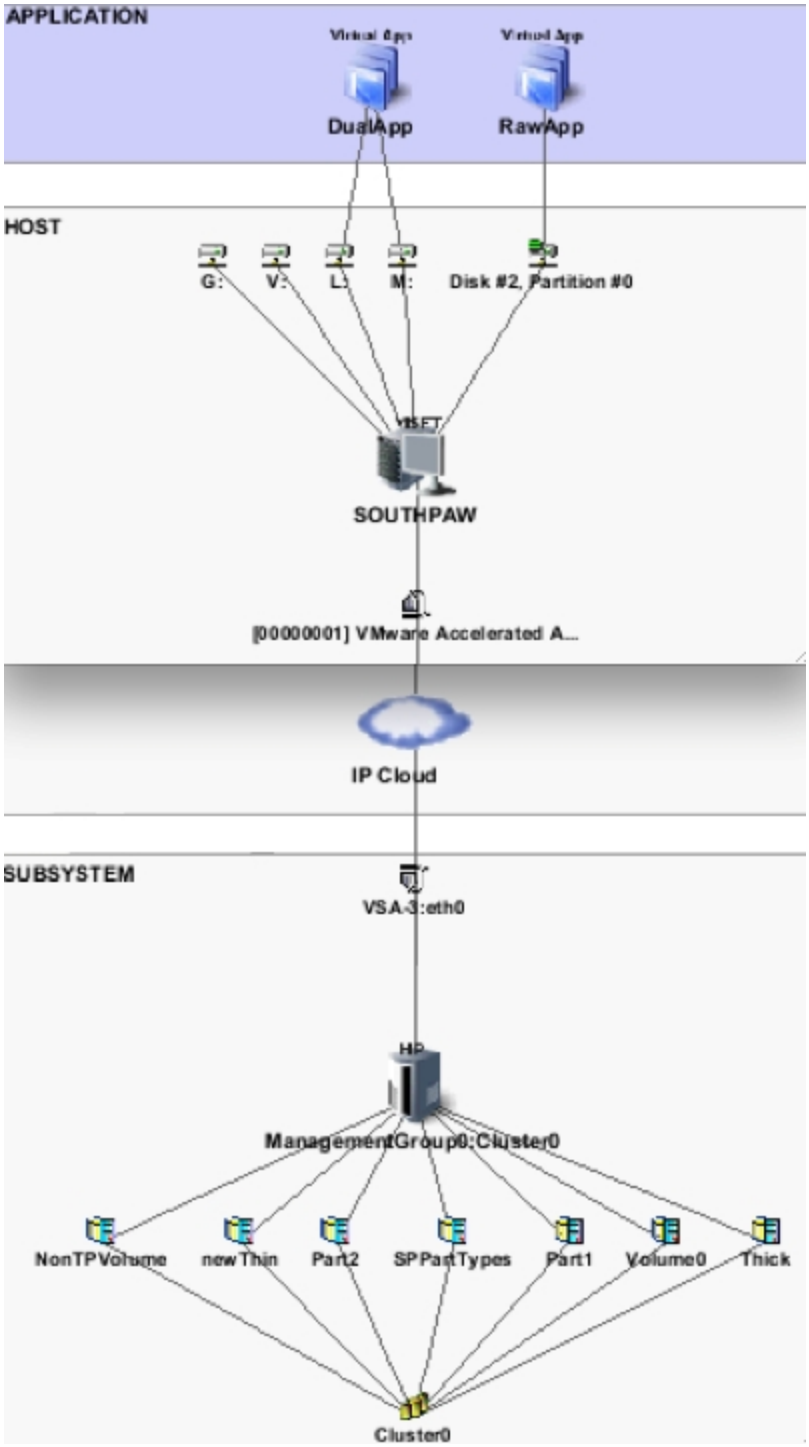

HP P4000 Device Navigation

The device navigation page is the central location to access information about the HP P4000. The navigation panel is broken into slices of the device: Front Physical, Logical, and Dependencies.

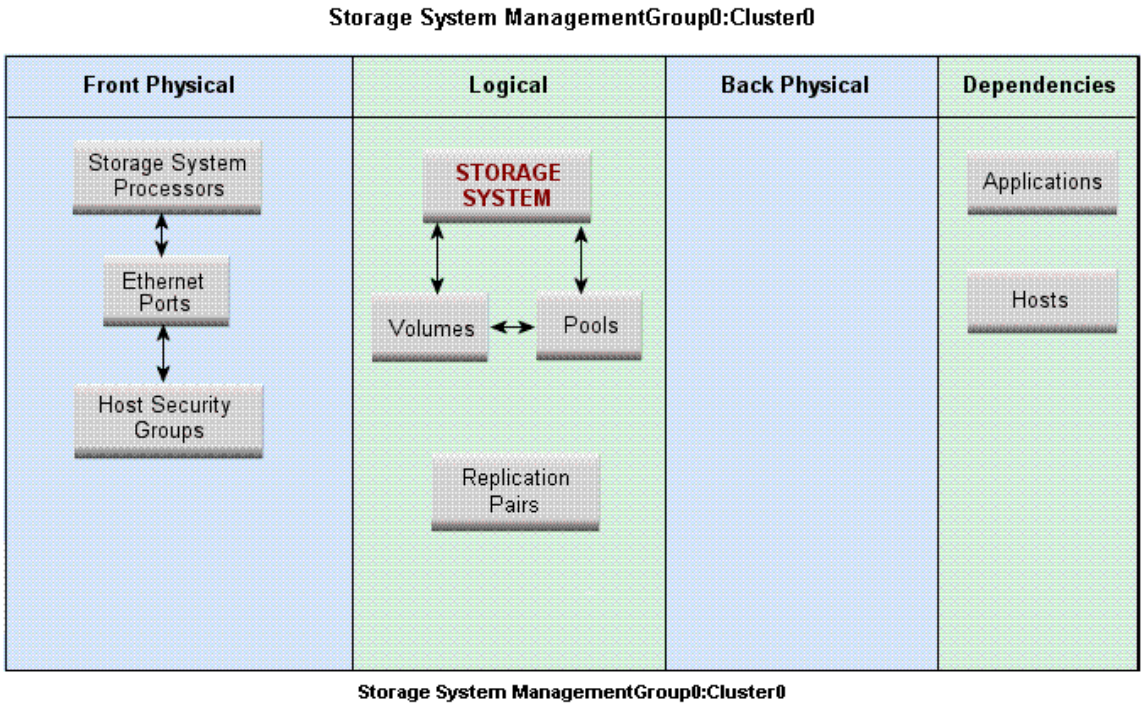

Primary Owner	
Description	HP P4000 VSA server 00:50:56:B5:53:F4

Front Physical

The presentation of iSCSI storage is through the front end of the device. This section provides detailed configuration and connection information from cluster nodes (Storage System Processors), ports (Ethernet Ports), and assigned servers (Host Security Groups).

The Storage System Processors contain a list of nodes in the cluster and provide access to detailed information for each node, including ports on the node, status, and software version.

Storage System Processors

Name	Description
VSA-1	HP P4000 VSA server 00:50:56:B5:4F:7B.00:50:56:B5:53:F4
VSA-2	HP P4000 VSA server 00:50:56:B5:4F:7B.00:50:56:B5:11:22
VSA-3	HP P4000 VSA server 00:50:56:B5:4F:7B.00:50:56:B5:4F:7B

Selecting a storage processor reveals the detailed properties for that node.

Storage System Processor VSA-1

Description	HP P4000 VSA server 00:50:56:B5:4F:7B.00:50:56:B5:53:F4	**Model**	VSA
Contacted	2010-04-12 21:50	**Record Created**	2010-04-11 21:41
Status	up	**Identifying Description**	[eth0]
Other Identifying Information	[16.118.234.223]	**Discovery Status**	Contacted
Version	SANiQ 8.1.00.0047	**Storage System**	ManagementGroup0:Cluster0

IP Ports

VSA-1:eth0

Ethernet Ports list all the ports on the cluster, together with the cluster node they are connected to. The name of the cluster node is pre-appended to the port name.

IP Ports

Name	Storage System Processor	MAC Address	IP Addresses	Network Card	Port Speed	Link Technology
VSA-2:eth1	VSA-2	00:50:56:B5:11:22:00	0.0.0.0	VirtualAdapter	1000 Mb/s	Ethernet
VSA-1:eth0	VSA-1	00:50:56:B5:53:F4	16.118.234.223, 16.118.234.219	VirtualAdapter	1000 Mb/s	Ethernet
VSA-2:eth0	VSA-2	00:50:56:B5:11:22	16.118.234.224	VirtualAdapter	1000 Mb/s	Ethernet
VSA-3:eth0	VSA-3	00:50:56:B5:4F:7B	16.118.234.225	VirtualAdapter	1000 Mb/s	Ethernet

When looking at a host with iSCSI bindings, the Port Speed column might be blank if the host is running Windows 2003.

Host Security Groups contains a list of assigned servers with their Host IQN, or if discovered, a link to the server, followed by the list of volumes assigned to that server.

Host Security Groups

Filter

Page 1 of 2 Showing 1-10 out of 11 Total (0 Selected) Display: 10 rows

Select All Pages | Unselect All Pages

Name↑	Initiators	Volumes
iqn.1987-05.com.cisco:01.f2cf5b667936	iqn.1987-05.com.cisco:01.f2cf5b667936	t2(LUN 0)
iqn.1991-05.com.microsoft:erittphilip1.cup.hp.com	iqn.1991-05.com.microsoft:erittphilip1.cup.hp.com	+ Volumes(LUNs)
iqn.1991-05.com.microsoft:sedev010	iqn.1991-05.com.microsoft:sedev010	+ Volumes(LUNs)
iqn.1991-05.com.microsoft:southpaw.selab.usa.hp.com	SOUTHPAW::[00000001] VMware Accelerated AMD PCNet Adapter	+ Volumes(LUNs)
iqn.1994-05.com.redhat:2e3337a4faa7	iqn.1994-05.com.redhat:2e3337a4faa7	rhelTest(LUN 0)
iqn.1994-05.com.redhat:eab6a4577c68	iqn.1994-05.com.redhat:eab6a4577c68	t2(LUN 0)
iqn.1998-01.com.vmware:cc3srv1-4699da59	iqn.1998-01.com.vmware:cc3srv1-4699da59	+ Volumes(LUNs)
iqn.1998-01.com.vmware:cc3srv2-3d2480d0	iqn.1998-01.com.vmware:cc3srv2-3d2480d0	+ Volumes(LUNs)
iqn.1998-01.com.vmware:cc4srv3-299bbd30	iqn.1998-01.com.vmware:cc4srv3-299bbd30	+ Volumes(LUNs)
iqn.1998-01.com.vmware:cc4srv4-7abdca9b	cc4srv4.selab.usa.hp.com::vmk0	+ Volumes(LUNs)

Logical

Logical refers to the inventory of all volumes and snapshots, pools summarizing total cluster capacity, and replication pairs.

The Volumes panel lists all volumes and allows one to be selected in order to show the detailed properties page.

Storage Volume HugeThin

Thinly Provisioned	true	**Contacted**	2010-04-19 10:38
Record Created	2010-04-19 10:38	**Replication Level**	2
Block Size	1,024	**Status Information**	Enabled
Raw Storage	1,024 MB	**Availability**	
Volume Type	Normal	**Snapshot**	false
Composition		**Discovery Status**	Contacted
Data Organization		**Consumable Blocks**	20,971,520
Device ID	iqn.2003-10.com.lefthandnetworks:managementgroup0:11506:hugethin	**Description**	HugeThin
Raid Type	Network RAID-10	**Composite Volume**	false
Consumed Storage In Blocks	524,288	**No Single Point Of Failure**	
Number Of Blocks	20,971,520	**Purpose**	
Access		**Storage Pool**	Cluster0
Storage System	ManagementGroup0:Cluster0		

Keep in mind the following:

- Raid Type indicates the type of data protection level provided by the volume RAID.
- Thin Provisioning (ThP) information is shown through the “Thinly provisioned” flag, as well as showing the exact storage consumed on the device “Consumed Storage.” The illustration shows that the 20Gb volume (Number of Blocks) is only consuming 512Mb of the carved space, and 1Gb if considering replicas (Raw Storage).
- Replication Pairs contains the volume-to-snapshot relationships, including the time the snapshots were last updated. The “when synced” property is the only property that is collected from the internal WBEM provider running on the cluster node.

Dependencies

The Dependencies column of the navigation page reveals the applications and client hosts that are using storage presented by this cluster.

Dependent Applications

Application	Host	Mount Point	HBA Port	Storage System Port	Storage Volume	LUN	Composition
DualApp (created)	SOUTHPAW	L:	[00000001] VMware Accelerated AMD PCNet Adapter	VSA-2:eth0	Volume0	0	
DualApp (created)	SOUTHPAW	M:	[00000001] VMware Accelerated AMD PCNet Adapter	VSA-2:eth0	Part1	0	
DualApp (created)	SOUTHPAW	M:	[00000001] VMware Accelerated AMD PCNet Adapter	VSA-2:eth0	Part2	0	
RawApp (created)	SOUTHPAW		[00000001] VMware Accelerated AMD PCNet Adapter	VSA-2:eth0	newThin	0	

For each application and the mount point it uses, the dependent application table lists the connection path from the host to the storage array volume that provides the storage.

Dependent Hosts

Host Name	Operating System	Mount Point	Storage Volume
SOUTHPAW	Windows XP		Thick
SOUTHPAW	Windows XP	G:	SPPartTypes
SOUTHPAW	Windows XP		newThin
SOUTHPAW	Windows XP	L:	Volume0
SOUTHPAW	Windows XP	M:	Part2
SOUTHPAW	Windows XP	M:	Part1
SOUTHPAW	Windows XP	V:	NonTPVolume
cc4srv4.selab.usa.hp.com	ESX Server	iSCSI Static LUN	cc4srv4_vol
cc4srv4.selab.usa.hp.com	ESX Server		RawESX2

HP P4000 iSCSI Information

If you access the Navigation tab for a host that has an iSCSI port connected to an ISCSI disk on an HP P4000 array, you will see an iSCSI Physical column.

The iSCSI Physical column provides the following buttons:

- Ethernet Card
- Ethernet Ports
- Bindings

If you select the Ethernet Card button, you will see the vendor model and serial number of the Ethernet card.

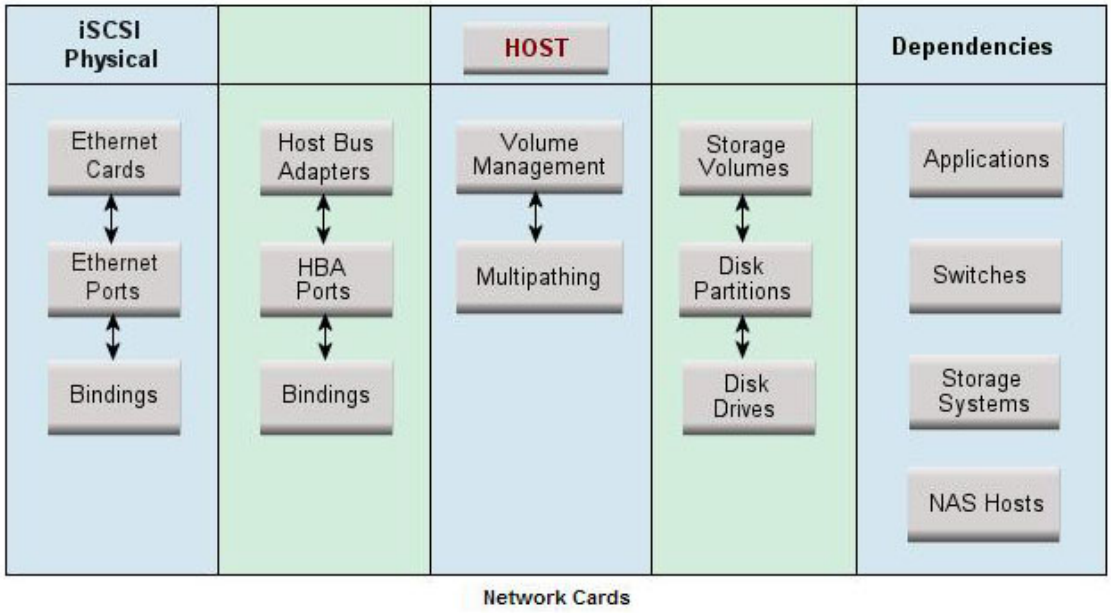

Name	Vendor	Model	Serial Number
iSCSI Initiator Root\SCSIADAPTER\0000_0	Microsoft Corporation	iSCSI Initiator	MSFT-05-1991

If you select the Ethernet Ports button, you will see the MAC address and the IP addresses on the host that is used to connect to the P4000 array. Each NIC card has its own unique IP address and MAC address.

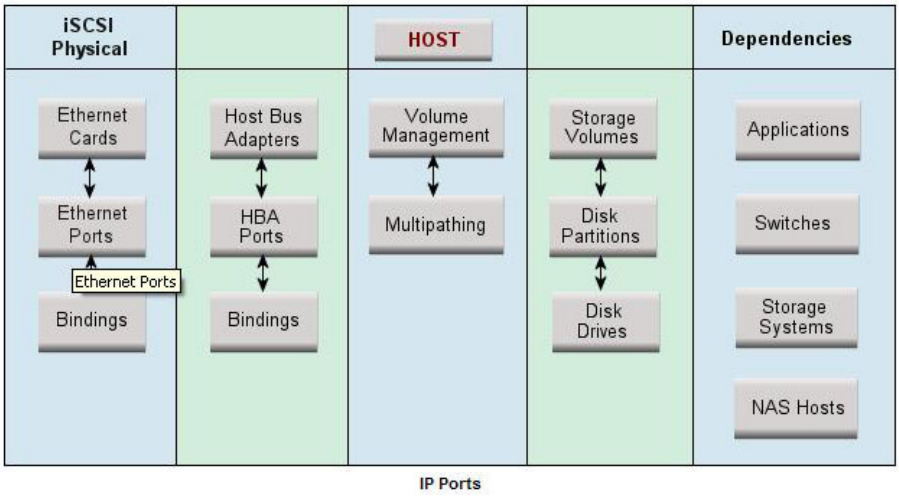

Name	MAC Address	IP Addresses	Network Card	Port Speed
[00000001] VMware Accelerated AMD PCNet Adapter	00:50:56:B5:63:EA	16.118.234.226, 0.0.0.0	iSCSI Initiator Root\SCSIADAPTER\0000_0	

If you select the Bindings button, you will see the following:

- Port: Name of the port.
- IP address: IP address of the port on the host.
- Target IP address: IPaddress of the port on the storage system.
- Target LUN: Name of the LUN on the storage array.
- Disk: Name of the disk on the host.

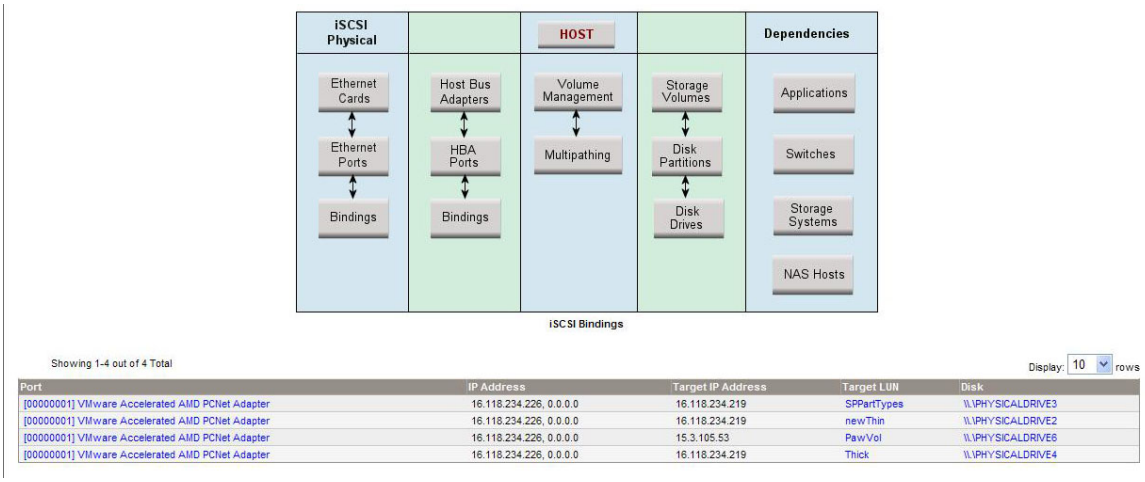

See HP P4000 Device Navigation on page 143.

Building the Topology View

After you discover elements, the management server requires you to build a topology view, which is a graphical representation of port-level connectivity information.

If a switch has more than one connection to an element, the number of connections is displayed above the line linking the switch and the element. For example, if the number two is shown between a switch and a storage system, it means that the elements have two connections to each other. To view the port details for the connection, right-click the element and select **Show Port Details** from the menu.

If the topology changes, you can update how the element is viewed in the topology by selecting the element and clicking the **Get Topology for Selected** button in the Get Topology for discovered elements page (select **Discovery** > **Topology** in the upper-right pane of the HP Storage Essentials home page). The management server obtains enough information about where the element is connected in the topology; for example, showing where a switch connected to a host.

If the management server detects an element but it cannot obtain additional information about it, it marks the element with a question mark in the topology. To learn more about fixing detected and/or disconnected elements, see Troubleshooting Topology Issues on page 882.

The user interface in HP Storage Essentialsmight load slowly while the topology is being recalculated. It might also take more time to log on to the management server during a topology recalculation. For more information, see Recalculating the Topology on page 894.

To obtain enough information to display the topology in System Manager:

1. Click the **Discovery** menu in the upper-right corner of theHP Storage Essentials home page.
2. Click **Topology** in the upper-right corner. The discovered elements are selected.
3. Select **All Discovery Groups** or click **Specified Discovery Groups** to specify a customized list. If you are obtaining the topology for the first time, select **All Discovery Groups**.

   For information on selecting a custom discovery list, see Creating Custom Discovery Lists on page 154.
4. Click **Get Topology**.

   The management server obtains the topology for selected elements and displays the Log Message page. After the management server builds the topology, a link appears to take you to System Manager so you can verify the topology view.

   You can also access System Manager by clicking **System Manager** in the left pane.
5. Review the topology for errors or changes.

If you see errors in the topology, look at the log messages, which can provide an indication of what went wrong. Look at Event Manager for additional information. Access Event Manager by clicking the **Event Manager** button in the left pane. For more information, see Viewing Discovery Logs on page 162 and Troubleshooting Topology Issues on page 882.

If the topology for an element in your network changes, select the element and click **Get Topology** (**Discovery** > **Topology**) to update the information.

Modifying the Properties of a Discovered Address

You can modify the user name and password the management server uses to access a device. However, whenever a user name and/or password changes on a device the management server monitors, the management server must be made aware of the change. For example, if the password for a host is changed, you must update the management server database with the new password. For more information, see Modifying a Single IP Address Entry for Discovery on page 80.

If you use this window to change the user name and password stored in the management server's database, it does not change the device's user name and password.

To change the discovery properties of an element:

1. Select **Discovery** > **Topology** or **Discovery** > **Details**cin the upper-right pane of the HP Storage Essentials home page window.
2. Click the **Edit** ( ) button corresponding with the element you want to modify.
3. To move an element to another discovery group, select its new discovery group from the **Discovery Group** menu.
4. Click **OK** in the Edit Discovered Element window.

Get Details

About Get Details

Get Details is required to obtain detailed information from discovered elements. Get Details must be performed before you can do provisioning and/or obtain provisioning information, such as data about zone sets and LUN numbers.

Keep in mind the following:

- Running Get Details takes time. You might want to perform this process when the network and the managed elements are not busy. To obtain a picture of device connectivity quickly, click **Get Topology** on the Topology tab.

Reports show data from the last successful Get Details and report cache update. When a scheduled Get Details finishes, the report cache refresh does not automatically run after Get Details. The default interval for report cache refresh is six hours. For information about refreshing the report cache, see Refreshing the Report Cache on page 365.

- Make sure you have created schedules for Get Details, so it occurs periodically. See the online help for **Configuration** > **Details** for more information.
- During Get Details the data you see in the user interface is not updated until the data collection is finished.
- During Get Details the topology in System Manager is recalculated. While the topology is being recalculated, the loading of the user interface might be slow. It might also take more time to log on to the management server during a topology recalculation.
- You can use discovery groups to break up Get Details. For example, instead of running Get Details for all elements, you could specify only the elements in Discovery Group 1. For more information, see Using Discovery Groups on the facing page.
- When an element in a discovery group is updated, its dependent elements are also updated.
- You can quarantine elements to exclude them from Get Details. For example, to get information about all the elements in a discovery group except for one, you can quarantine that element. For more information, see Placing an Element in Quarantine on page 160.
- If a problem occurs with a host or SMI-S element during Get Details, the host or element is automatically quarantined. To remove the element from quarantine, see Removing an Element from Quarantine on page 160.
- To receive status reports about Get Details, see Configuring E-mail Notification for Get Details on page 863 for information about how to configure this option.
- If an element changes and you run Get Details while the provider cache is updating, an error might occur or the gathered details might be inconsistent with the actual element status.
- CLARiiON and LSI storage systems have two controllers with IP addresses. If you want to use the provisioning feature in HP Storage Essentials with these storage systems, you must discover both controllers. Make sure both controllers are kept in the same discovery group. If you are not planning to use the provisioning feature in HP Storage Essentials, you only need to discover one of the controllers.

Running Get Details

To obtain details about the elements on the network:

1. Select **Discovery** > **Details**.
2. Select **Include infrastructure details**, which gathers the latest information about SAN details. You do not need to select **Include backup details** unless you already discovered hosts running backup applications and installed CIM extensions on those hosts. For information about discovering master backup servers, see Step 1 – Discovering Your Hosts and Backup Manager Hosts on page 179.
3. Select **Force Device Manager Refresh** if you want the management server to tell the device managers for your storage systems to obtain the latest information. If you do not select **Force Device Manager Refresh**, the management server gathers information from the external databases such as HP, HDS, and EMC storage systems with the assumption that the information in the external database is up to date. See the following topics for more

information: Excluding EMC Symmetrix Storage Systems from Forced Device Manager Refresh on page 109 and Excluding HDS Storage Systems from Forced Device Manager Refresh on page 117.

4. Select **All Discovery Groups** or click **Specified Discovery Groups** to specify a customized list. If you are running Get Details for the first time, select **All Discovery Groups**.

   For information on selecting a custom discovery list, see Creating Custom Discovery Lists on next page.

5. Click **Get Details**.

   During Get Details, the software changes its status light from green to red and the HP Storage Essentialslog opens and shows the progress of Get Details.

   When the software finishes getting all element details, it displays GETTING ALL DETAILS COMPLETED on the View Logs page and the status light turns green.

6. See Adding a Discovery Schedule on page 338 for information about automating the gathering of all element details.

Stopping the Gathering of Details

Obtaining details takes some time. If the network and managed elements are busy, you might need to stop the gathering of details and reschedule it for another time.

Note: If you stop the gathering of details, you should reschedule it. This type of collection obtains detailed information about elements in the network.

To stop the gathering of details:

1. Select **Discovery** > **View Logs**.
2. On the **View Logs** page, click the "Click here" portion of the following message:

   ```
   Click here if you wish to stop getting details.
   ```

3. When you are asked if you are sure you want to stop Get Details, click **OK**.

   The management server stops gathering details.

   Existing operations will finish before the management server stops gathering details.

4. Schedule a time to resume getting details.

Using Discovery Groups

The discovery groups feature is sometimes called *segmented replication* because it allows you to run Get Details for a segment of elements. Because HP Storage Essentials runs more slowly when Get Details is in progress, it is helpful to break the process into segments which can then be run at night or on multiple days. For example, if Get Details for all elements takes twelve hours, you could break the elements into several small groups and schedule Get Details to run at night on multiple days.

Note: For more about data collection, see About Get Details on page 151.

When planning discovery groups, consider the following requirements and capabilities:

- By default, HP Storage Essentials is configured with a default discovery group plus four additional groups.
- Discovery groups affect the amount of memory needed for HP Storage Essentials. Before configuring discovery groups, check the support matrix and verify that your system meets the memory requirements for using discovery groups.
- Do not move elements between discovery groups when Get Details is running. If you do this, an error will occur when Get Details tries to discover elements that were moved.
- An element can be a member of one discovery group at a time.
- Elements discovered through SMI-S and hosts discovered with CIM extensions from version 5.1 and later of HP Storage Essentials cannot be added to discovery groups. These elements can, however, be placed independently into scheduled Get Details tasks without being part of a discovery group. This allows you greater flexibility when gathering discovery data. For more information, see Creating Custom Discovery Lists below.
- When an element in a discovery group is updated, its dependent elements are also updated.
- Each discovery group communicates over a specific port.

The defaults are:

Discovery Group Ports

Default	5986
Discovery Group 1	5984
Discovery Group 2	5982
Discovery Group 3	5980
Discovery Group 4	5978

Creating Custom Discovery Lists

You can create a discovery list for Get Details or Get Topology that will allow you to select a set of discovery groups to use the next time Get Details runs.

1. Select **Discovery** > **Details or Discovery** > **Topology**.
2. Click the **Specified Discovery Groups** link.
3. Select the check box next to each item you want to add to the discovery list.

   Elements discovered through SMI-S and hosts discovered with CIM extensions from version 5.1 and later of the product appear in the list individually. You can add individual elements, discovery groups, or both to the same discovery list.

The Specify Discovery List page offers a set of filters to help you find discovery groups quickly. For more information, see Filters on the Specify Discovery List Page below.

4. Click **Add Selected Discovery Groups to Discovery List** to move them into the Discovery List.

   Do not run Get Details for all discovery groups simultaneously.

5. Click **OK** to save and return to the previous window. The elements are selected in the elements table.

6. Click **Get Details** or **Get Topology**.

Filters on the Specify Discovery List Page

The filter area is collapsed by default. To expand the filter area, click the **+** symbol. The following filters are supported:

- Discovery Group Name Contains – Use this filter to retrieve all the discovery groups whose name contains the specified string.
- Element Name Contains – Use this filter to retrieve all discovery groups containing an element with the specified substring in its name.
- Discovery Group Type – Use this filter to see only discovery groups of the specified type.
- Element Type – Use this filter to see only discovery groups that contain the specified element type.

To apply the filter settings, click **Filter** to refresh the content of the page. To restore the filters to their default settings, click **Reset**, and refresh the page.

Managing Discovery Groups

To manage discovery groups from the Discovery Setup page:

The Default discovery group cannot be edited.

1. Select **Discovery** > **Details or Discovery** > **Topology**.
2. Click **Manage Discovery Groups**.

   The Discovery Groups page shows a list of your discovery groups, including the name, Port Number, and included elements.

3. Click **Edit**.
4. To rename the group, enter a new name in the Name box.
5. To add a member, select the member from the Potential Members section, and then click the **Add Selected Items to Discovery Group** button to move it into the Discovery Group Members section.

   The Edit Discovery Group page offers a set of filters to help you find potential members quickly. For more information, see Filters on the Edit Discovery Group Page on next page.

6. To remove a member, select the member from the Discovery Group Members section, and then click the **Remove Selected Items from Discovery Group** button to move it into the Potential Members section.

   The path to the log file for the discovery group is listed at the top of the page.

7. Click **OK**.
8. Click **Back to Discovery Page**.

Filters on the Edit Discovery Group Page

The filter area is collapsed by default. To expand the filter area, click the **+** symbol. The following filters are supported:

- Access Point Contains – Use this filter to retrieve all the access points whose name contains the specified string.
- Element Name Contains – Use this filter to retrieve all discovery groups containing an element with the specified substring in its name.
- Element Type – Use this filter to see only potential members that contain the specified element type.
- Discovery Group Name Contains – Use this filter to retrieve all the discovery groups whose name contains the specified string.

To apply the filter settings, click **Filter** to refresh the content of the page. To restore the filters to their default settings, click **Reset**, and refresh the page.

Moving Elements Between Discovery Groups

All elements are initially placed in the Default discovery group. You can move elements between discovery groups.

Do not move elements between discovery groups when Get Details is running. If you do this, an error will occur when Get Details tries to discover elements that were moved.

Method 1: Select Discovery Group

To select a new discovery group for an element:

1. Select Discovery Setup (**Discovery** > **Details**). The Get Details page appears.
2. Select the check box for the element you want to move.
3. Click **Move to Discovery Group**. The Select Discovery Group window appears.
4. Select the new discovery group for the selected element.
5. Click **OK**. HP Storage Essentials notifies you that it can take a few minutes to move an element.
6. Click **OK**. The elements are moved to the new discovery group.

Method 2: Edit a Discovered Element

To edit a discovered element:

1. Select Discovery Setup (**Discovery** > **Details**). The Get Details page appears.
2. Click the **Edit** () button next to the element you want to modify.
3. Select a new discovery group in the **Discovery Group** menu.
4. Click **OK**. HP Storage Essentials notifies you that it can take a few minutes to move an element.
5. Click **OK**. The elements are moved to the new discovery group.

Deleting Elements from the Product

When you delete an element, all of its information, except for its statistics are removed . The product saves the statistics for a deleted device for three days by default. See Restoring Statistics from Deleted Elements on page 159.

To completely delete an element from the management server you must remove the elements, such as a switch or proxy, that were used to discover the element. If you do not delete all switches and proxies that were used to discover the element, the element might reappear the next time you Get Details.

For example, assume you want to delete Switch_A. Switch_B and Switch_C were used to discover Switch_A. If you delete only Switch_B and Switch_A, Switch_A will most likely reappear when you Get Details because it is still accessible by Switch_C.

You can delete an element within the following tools:

- **System Manager or Chargeback Manager** – Gives you the option of deleting just the element or deleting the element and the elements that use the same switches and proxies for access.
- **Discovery Step 2 (Topology) or or Step 3 (Details)** – Gives you the option of deleting multiple elements at a time. You are not given a detailed list of other elements you must delete; however, you can use the table on the Discovery screen to determine which switches and proxies provided access.

Deleting an Element Using System Manager or Chargeback Manager

To delete an element using System Manager or Chargeback Manager:

1. Do one of the following:
   - **In System Manager** – Right-click an element and select **Delete Element** from the menu.

If you are blocking pop-ups and you use the right-click menu to delete an element from System Manager, the Delete window is blocked and you are unable to delete the element. You must disable the popup blocker before you can delete the element.

Or

- **In Chargeback Manager** – Click the **Delete** ( ) button for the element you want to delete.

2. If the element has multiple access points, you are asked which to delete. Do one of the following:
   - **Delete the element and its access points.** This option lists not only the switch you want to delete, but also the other elements that use the same switches and proxies as the element you want to delete. For example, assume you want to delete Switch_A. Switch_B was used to discover Switch_A. Let's assume Switch_B is also the only path to Switch_D. If you delete Switch_B, you will no longer have access to Switch_D. This option would list Switch_D as one of the other elements that need to be deleted.

     An access point is the intersection of the IP address and the provider that discovered the IP address. A provider is software that is used to gather information about an element.

     Or

   - **Delete the element.** The element might reappear the next time you obtain element details. This is because not all switches and proxies connected to the element have not been removed. For example, assume you want to delete Switch_A. Switch_B is connected to Switch_A. If you do not delete Switch_B, the next time you obtain element details Switch_B will most likely find Switch_A again.
3. Click **OK**.

Deleting Elements Using Discovery Step 2 (Topology) or Step 3 (Details)

To delete multiple elements using Discovery Step 2 (Topology):

1. Select **Discovery** > **Topology** or **Discovery** > **Details** in the upper-right pane of the HP Storage Essentials home page.
2. Determine the access points for the element you want to delete. In the following figure, QBrocade2 is accessed by two switches: 192.168.10.25 and 198.168.10.22. You must delete both access points to completely remove the element. As a result, the QBrocade5 switch will also be removed because it has the same access points as QBrocade2.

☐	192.168.10.25	Switch	QBrocade2, QBrocade5	admin		
☐	192.168.10.21	Switch	QBrocade1	admin		
☐	192.168.10.22	Switch	QBrocade2, QBrocade5	admin		
☐	192.168.10.24	Switch	QBrocade3, QBrocade4	admin		

3. Select all of the access points for the element you want to delete, and then click the **Delete**

button just above the table.

For example, assume you want to delete QBrocade2 in the previous figure. You would select the two listings for QBrocade2 on the Discovered Elements tab and click the **Delete** button in the **Get Topology for Discovered Elements** table. If you delete only one of the listings, QBrocade2 and QBrocade5 still appear in the topology, since they are still accessible from one of the switches.

When you are asked if you want to remove the access points and its associated elements, keep in mind these elements will not be deleted if they are accessible from an access point not listed in the Delete Access Points window. For example, assume you selected access point 192.168.10.25 to be deleted. You are then told that switch1 will be deleted along with the access point. Assume also that switch1 is accessible from another access point, 192.168.10.29. When you remove access point 192.168.10.25, switch1 will still be accessible because it can be accessed from another access point that has not been removed.

4. Click **OK** to remove the access points listed in the Delete Access Points window.

The access points are removed. If the elements listed have no other access points, they are no longer accessible from the management server.

Restoring Statistics from Deleted Elements

The product saves the statistics for a deleted elements for three days by default. If you rediscover the element within three days, its statistics become available again. Only the statistics are restored. No other configuration data that might have been associated with the element, such as its license or proxy host information, is restored.

The save and restore functionality should not be used as a long term method to retain data for deleted elements. Although the information is stored internally in the database, it is only accessible once the deleted element is rediscovered. You should only modify the custom property if you cannot rediscover the device within three days and need more time to do the rediscovery.

To change how long the element is saved, modify the `rb_deletetime` property by setting a new value (`rb_deletetime=5 for example`) on the Advanced page (**Configuration** < **Product Health** < **Advanced** page).

Working with Quarantined Elements

When an element is quarantined, it is not included in the Get Details process until it is removed from quarantine. For more information, see Removing an Element from Quarantine on next page. If a problem occurs with a host or SMI-S element during Get Details, the host or element is automatically quarantined.

Placing an Element in Quarantine

When you click the **Get Details** button on the Get Details page, the management server automatically obtains details for the elements in the selected discovery group. Assume you want to discover all the elements in a discovery group, except for one, which is being taken off of the network for maintenance. You can use the quarantine feature to exclude this element from discovery.

After you perform Get Details for the discovery group containing the quarantined elements, the quarantined elements appear as missing throughout the product. The management server marks the quarantined elements as missing because it cannot obtain details from the quarantined element.

To quarantine an element:

1. Select the check boxes for the elements you want to quarantine on the Get Details page.
2. Click **Set Quarantine**.
3. When you are asked if you want to quarantine the selected elements, click **OK**.

The elements you quarantine appear with a flag ( ) in the Quarantined column on the Get Details page.

The elements are excluded from discovery until you clear them from quarantine.

Removing an Element from Quarantine

To remove an element from quarantine:

1. Select the check boxes for the elements you want to remove from quarantine on the Get Details page.

   Quarantined elements appear with a flag ( ) in the Quarantined column on the Get Details page.
2. Click **Clear Quarantine**.
3. When asked if you want to remove the selected elements from quarantine, click **OK**.

   The next time you perform Get Details for the element, the management server gathers data from it.

Updating the Database with Element Changes

After you initially discover the elements, information about them might change. To update database with these changes, perform the following steps.

Keep in mind the following:

- 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.

- If you are adding, removing or replacing McDATA switches, you must use a different procedure. For more information, see Managing McDATA Switches on page 103.
- Running Get Details takes time. You might want to perform this process when the network and the managed elements are not busy. To obtain a picture of device connectivity quickly, click the **Get Topology** button on the Topology tab.

To update the database:

1. Select **Discovery** > **Details**.
2. Select **Include infrastructure details**, which gathers information about SAN details.

   Include backup details is used for gathering information for Backup Manager. You do not need to select it unless you have already discovered hosts running backup applications and installed CIM extensions on those hosts. For more information about discovering master backup servers, see Step 1 – Discovering Your Hosts and Backup Manager Hosts on page 179.
3. The management server obtains most of its information from device managers for storage systems with external databases, such as HP, HDS, and EMC storage systems. Select **Force Device Manager Refresh** if you want the management server to tell the device managers for your storage systems to obtain the latest information. If you do not select Force Device Manager Refresh, the management server gathers information from the external databases based on the assumption the information in the external database is up-to-date.

   For more information, see: Excluding EMC Symmetrix Storage Systems from Forced Device Manager Refresh on page 109 and Excluding EMC Symmetrix Storage Systems from Forced Device Manager Refresh on page 109.
4. Click the **Get Details** button on the Get Details page.
5. View the status of the gathering of element details by looking in the **View Logs** page. For more information about the messsages viewed in this tab, see Viewing Discovery Logs on next page.
6. Verify the topology is displayed correctly by accessing System Manager. Access System Manager by clicking its button in the left pane.

Notifying the Software of New Elements

When you add a new element to the network, such as a host, perform discovery to make the management server aware of the new element.

Keep in mind the following:

- 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.

- If you started a CIM Extension on a Sun Solaris host with the ./start -users command, in the command, you must provide a user name to be used to discover the host. For example, if you use `./start -users <myname:yourname>` (in this instance, myname and yourname are valid UNIX accounts) to start the CIM Extension, `myname` or `yourname` and its password must be used to discover the host.
- If this is a new installation of the management server and you have Brocade switches, download and install the Brocade SMI Agent software as described in the *HP StorageWorks B-Series* document at http://www.hp.com/go/hpsim/providers.
- Additional steps are required for discovering McDATA switches; the steps vary according to your network configuration. For more information, see Discovering McDATA Switches on page 99.
- EMC CLARiiON storage systems require additional steps for discovery. For more information, see Discovering EMC CLARiiON Storage Systems on page 113 for more information.
- After you discover a McDATA switch, the IP address displayed next to the name of the switch is actually the IP address of the service processor for the switch in the Get Details screens. To find the IP address of the switch, click the link for the switch in the Topology screen (**Discovery** > **Topology**) or Get Details screen (**Discovery** > **Details**) and then click the **Properties** tab. You can also access the Properties tab by double-clicking the switch in System Manager.

Viewing Discovery Logs

Use the View Logs page to obtain the status of the following:

- Discovery
- Building the Topology
- Backup details

During these operations, the management server displays its status at regular intervals.

During a Step 1 Discovery the log messages shown on the View Logs page sometimes appear out of order. You might see log messages with a timestamp of 11:12 followed by log messages with an 11:11 timestamp.

To view logs for these operations:

1. Select **Discovery** > **View Logs**.
2. To view the progress of Get Details, click the **Infrastructure** tab.
3. To view the progress of Backup Details, click the **Backup** tab.
4. To obtain the latest status, click **Get Latest Messages**.

If the software is unable to discover or obtain information about a device, the log messages might provide some information as to where the problem occurred.

For example, if a host was not discovered, the log messages might indicate that the provider configuration for that device was never created. This could mean the software was given the wrong user name and/or password for that host. As a result, the software logged onto the host with a guest account, which does not have enough permissions to start Windows Management Instrumentation (WMI).

The logs show data from the most recent discovery, test, or data collection task.

Look at Event Manager for additional information. See About Event Manager on page 529

Viewing the Status of System Tasks

The Task Dashboard allows you to view the status of the tasks running on the management server. The dashboard provides the name of each task, its latest status, and the time the status was last reported.

To view the status of system tasks:

1. Select **Discovery** > **System Tasks**.
2. To obtain the latest status, click **Get the Latest Status**.

The following task statuses are provided by the Task Dashboard.

Task Status Descriptions

Status	Description
Not Found	This task cannot be found on this server.
Completed	This task was completed successfully.
Failed	This task failed with an error.
Aborted	This task was aborted by the user or other automated actions.
In Progress	This task is in progress. CPU and disk activities are active on this server.
Queued	This task is scheduled to be executed in the future.
Rejected	This task was rejected by this server.

Device-Specific Replication Information

HP Storage Essentials presents replication-state information using SMI-S terminology. Some SMI-S terms do not have an obvious device-specific equivalent.

See the following topics to find the vendor-specific terms and how HP Storage Essentials maps them with SMI-S.

- EMC Clariion Array Replication on next page
- EMC Symmetrix Array Replication on page 166

EMC Clariion Array Replication

HP Storage Essentials presents replication-state information using SMI-S terminology. Some SMI-S terms do not have an obvious device-specific equivalent. This topic explains how HP Storage Essentials maps EMC terminology with SMI-S.

Clariion

HP Storage Essentials supports SnapView Clone (Mirror - Local), MirrorView (Mirror - Remote), and SnapView Snapshot (Snapshot - Local). It does not collect data about SanCopy (Clone - Local and Clone - Remote).

SnapView Clone

SnapView Clone is a Local Mirror (a synchronized copy of the source element). The replica type is "Full Copy" and the copy type is "Sync." The "when synched" field is not exposed via NaviCLI commands and is not populated within HP Storage Essentials.

Mirror View

SnapView Clone is a Remote Mirror (a synchronized remote copy of the source element). The replica type is "Full Copy" and the copy type is "Sync." The "when synced" field is not exposed via NaviCLI commands and is not populated within HP Storage Essentials.

Snapview Snapshot

SnapView Snapshot is a Point-in-Time, associated virtual copy of the source element. The target element enables visibility into a session where Snapview Session is the Point-In-Time representation of the source element. The replica type is "Full Copy" and the copy type is "UnSyncAssoc." The "when synced" field is not exposed via NaviCLI commands and is not populated within HP Storage Essentials. Sync state is also not exposed via NaviCLI command for snapshots.

Here is a mapping for sync state and sync maintained based on the sync state value from relevant NaviCLI command output.

Sync State from NaviCLI	Sync State	Sync Maintained
Synchronizing	"ResyncInProgress"	True

Sync State from NaviCLI	Sync State	Sync Maintained
Synchronized	"Synchronized"	True
Consistent Out-Of-Sync	"Consistent"	False
	"Out of Sync"	False
	"State Unknown"	False

HP Storage Essentials must have Navisphere installed to discover replication information. It communicates with the Navisphere Host agent through the Navisphere CLI via port 6389. The following example illustrates how HP Storage Essentials CMS, NaviCLI, and two Clariion arrays could communicate with each other.

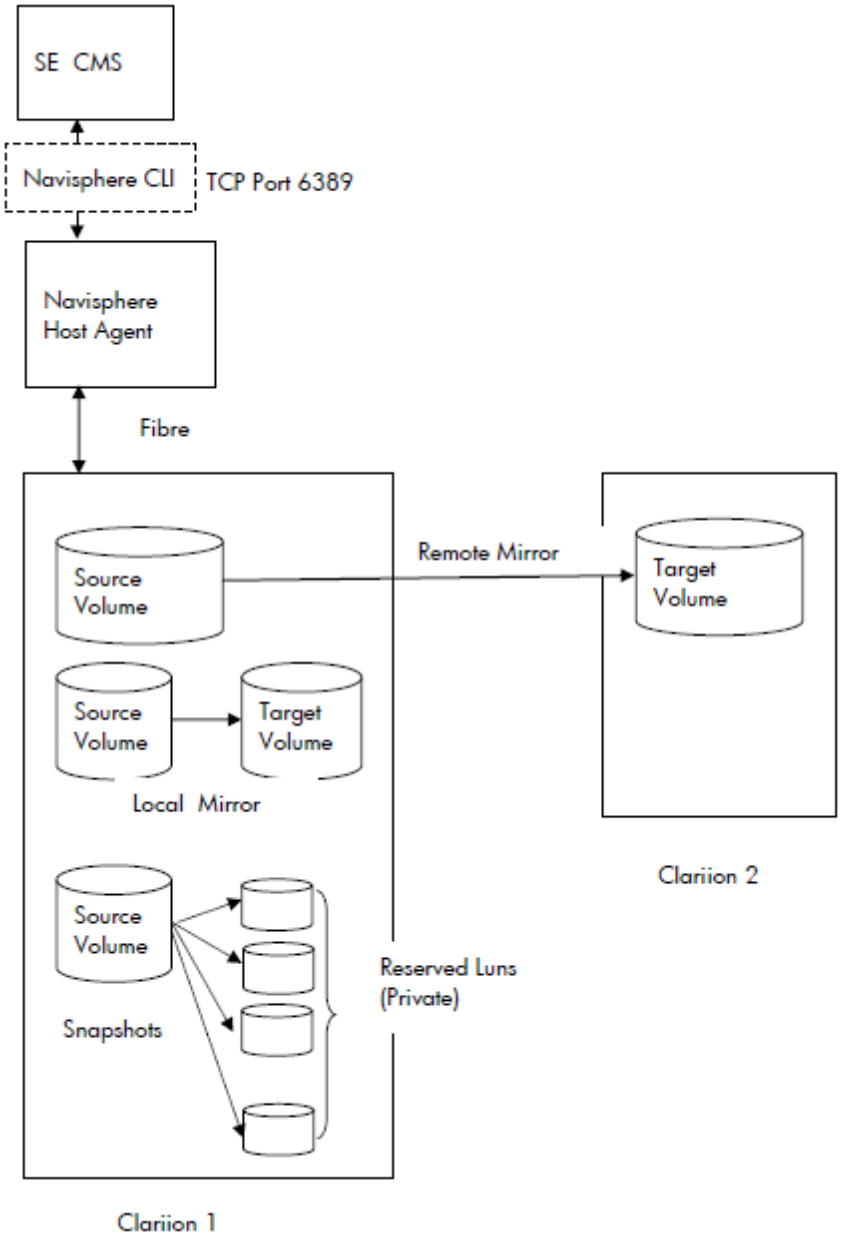

EMC Symmetrix Array Replication

HP Storage Essentials presents replication-state information using SMI-S terminology. Some SMI-S terms do not have an obvious device-specific equivalent. This topic explains how HP Storage Essentials maps EMC terminology with SMI-S.

Symmetrix

HP Storage Essentials supports local replication via business continuance volume (BCV) and TimeFinder Snap and Clone. Remote replication is supported via remote data facility (RDF).

BCV

Replication pairs are only recognized for BCV volumes that are paired with a standard volume. BCV volumes that have never been paired are not shown because there is no replication pair. BCV replica pairs always have a copy type of "sync" and a replica type of "full copy."

The following table maps the BCV pair states into the remaining SMI-S fields: sync state and sync maintained. The "when synced" field is not exposed via EMC APIs and is not populated within HP Storage Essentials.

BCV Pair State	Sync State	Sync Maintained
Sync in progress	ResyncInProgress	True
Synchronized	Synchronized	True
Split in progress	Fracture in progress	False
Split	Fractured	False
Restore in progress	Restore in progress	False
Split no incremental	"DTMF reserved" *EMC_SYNCSTATE_SPLIT_NO_INCREMENTAL* *Proprietary value 32761 == Short.MAX_VALUE-6*	False
Restored	"DTMF reserved" *EMC_SYNCSTATE_RESTORED* *Proprietary value 32760 == Short.MAX_VALUE-7*	False
Split before sync	"DTMF reserved" *EMC_SYNCSTATE_SPLIT_BEFORE_SYNC* *Proprietary value 32759 == Short.MAX_VALUE-8*	False

BCV Pair State	Sync State	Sync Maintained
Split before restore	"DTMF reserved" *EMC_SYNCSTATE_SPLIT_BEFORE_RESTORE* *Proprietary value 32758 == Short.MAX_VALUE-9*	False
Broken	"Broken"	False

RDF

HP Storage Essentials shows all RDF volume pairings.

Here is the mapping for copy type and replica type based on the RDF's current mode:

EMC RDF Mode for Replica Pair	Copy Type	Replica Type
Synchronous	Sync	Full copy
Asynchronous	Async	Full copy
Adaptive copy	Async	Full copy
Semi-synchronous	Async	Full copy

Here is the mapping for sync state and sync maintained based on the RDF's pair state or status:

RDF Pair State	Sync State	Sync Maintained
Sync in progress	"ResyncInProgress"	True
Synchronized	"Synchronized"	True
Split	"Fractured"	False
Failed over	"DMTF reserved" *EMC_RDF_STATE_FAILED_OVER* *Proprietary value 32766 == Short.MAX_VALUE - 1*	False
R1 updated	"DMTF reserved" *EMC_SYNCSTATE_R1_UPDATED* *Proprietary value 32765 == Short.MAX_VALUE - 2*	True

RDF Pair State	Sync State	Sync Maintained
R1 update in progress	"DMTF reserved" *EMC_SYNCSTATE_R1_UPDINPROG* *Proprietary value 32764 == Short.MAX_VALUE - 3*	True
Suspended	"DMTF reserved" *EMC_SYNCSTATE_RDF_SUSPENDED* *Proprietary value 32763 == Short.MAX_VALUE - 4*	False
Partitioned	"Broken"	False
Mixed	"DMTF reserved" *EMC_SYNCSTATE_RDF_MIXED* *Proprietary value 32762 == Short.MAX_VALUE - 5*	False
Invalid	"DMTF reserved" *EMC_SYNCSTATE_RDF_INVALID* *Proprietary value 32757 == Short.MAX_VALUE - 10*	False
Consistent	"Idle"	True

TimeFinder Snap and Clone

EMC TimeFinder Snap and Clone always have a sync maintained value of false, and a Replica type of Full Copy. Their copy type is UnSyncAssoc for Snap and UnSyncUnAssoc for clones. The following list compares EMC terminology with HP Storage Essentials terminology.

EMC Term	HP Storage EssentialsTerm
NA	Not Available
Copy in Progress	ResyncInProgress
Copied	Synchronized
Copy On Access	Copy On Access
Invalid	State Unknown

EMC Term	HP Storage EssentialsTerm
Create In Progress	PrepareInProgress
Created	Prepared
Copy On Write	Copy On Write
Restored	Restored
Terminate In Progress	Terminate In Progress
Restore In Progress	Restore In Progress
Failed	Failed
Recreated	Recreated
PreCopy	PreCopy
Split	Fractured
Unknown	State Unknown

HP Storage Essentials must have access to the EMC Solutions Enabler software in order to discover replication information. It communicates with Solutions Enabler via port 2707. The following example illustrates how HP Storage Essentials CMS, Solutions Enabler, and two EMC arrays could communicate with each other.

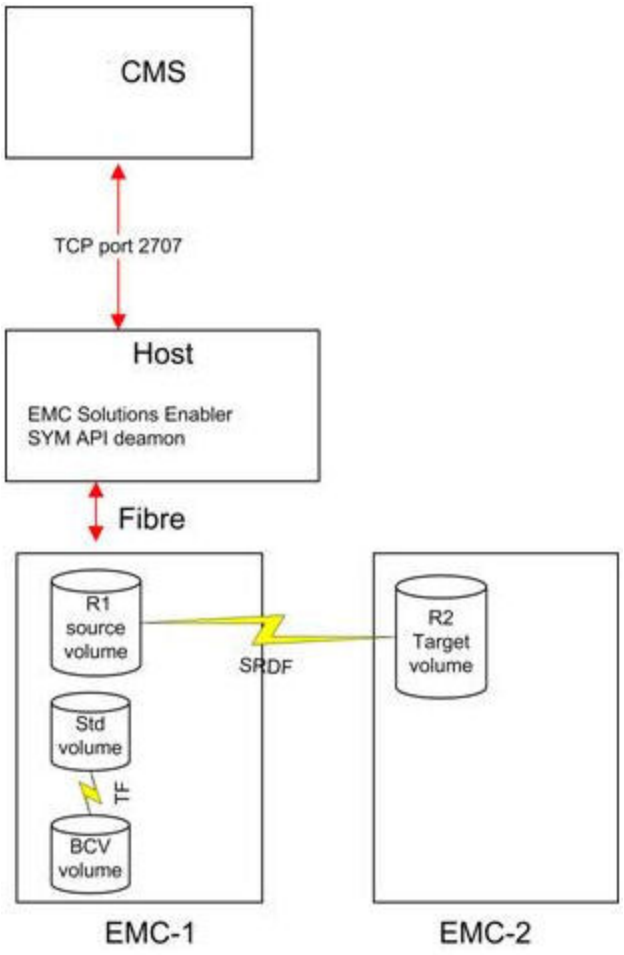

HDS Array Replication

HP Storage Essentials presents replication-state information using SMI-S terminology. Some SMI-S terms do not have an obvious device-specific equivalent.

Here are the HDS terms and how HP Storage Essentials maps them with SMI-S:

	TrueCopy (Sync & Async)	Universal Replicator	Shadow Image	C.O.W. Snapshot
Locality	Remote pair	Remote pair	Local pair	Local pair
Replica Type	Full copy	Full copy	Full copy	After delta
Copy type	Sync/Async depending on cache journaling in use	Async	Sync	UnSyncAssoc
Sync State	Paired, idle, failed, suspended	Active, halted, stopped	Copy, pair, PSUS	Idle or pair

HP EVA Array Replication

HP Storage Essentials presents replication-state information using SMI-S terminology. Some SMI-S terms do not have an obvious device-specific equivalent.

This topics explains how HP Storage Essentials maps HP EVA terms with SMI-S.

HP Storage Essentials communicates with Command View EVA to obtain replication information. By default, communication is done on TCP port 5989 over SSL.Command View. EVA communicates with the actual device over a fiber channel connection.

Local Replication via HP Business Copy EVA

HP Business Copy EVA makes local copies of virtual disks using snapclones, snapshots, mirrorclones, and pre-allocated containers. Replicated virtual disks are located on the same storage system as the source. The following features are built into HP Command View EVA.

- Snapclones – independent point-in-time copies
- Snapshots – dependent point-in-time copies
- Mirrorclones – ongoing copy

Snapclones

HP Storage Essentials does not support EVA snapclones because they are independent copies. Once the source volume data is copied to the target snapclone, there is no longer any replication relationship between the source and target, and the target becomes a standalone vdisk like any other. HP Storage Essentials can detect a snapclone if the creation (aka normalizing) is in progress while HP Storage Essentials is in the process of a Get Details task.

If this occurs, HP Storage Essentials will show the details of the snapclone at the time the data was queried, and that data will not change until the next Get Details task. (There would be no progress updates syncstate, when synced, sync maintained, and so forth.) On the next Get Details, the snapclone will probably disappear from HP Storage Essentials because it will be done normalizing, and will be seen by HP Storage Essentials as an independent volume with no replication relationship.

	Snapshots	Mirrorclones
Locality	Local pair	Local pair
Copy type	UnSyncAssoc	Sync when synchronized, Async when fractured
Replica type	After delta	Full copy
Sync state	Idle or broken if there is an error in the DR group link	Synchronized or fractured

	Snapshots	Mirrorclones
Sync maintained	False	True while sychronized, false while fractured or detached
When synced	Date and time the replica was created	Date and time the replica was created

Remote Replication via HP Continuous Access EVA

HP Continuous Access EVA makes remote copies of virtual disks. Replicated virtual disks are located on a different storage system from the source; typically, at a geographically separate site. Remote replication requires HP StorageWorks Continuous Access EVA.

CV EVA terms "source" and "destination" are equivalent to HP Storage Essentials terms "source" and "target."

CV EVA write mode (synchronous/asynchronous writethrough of data) should not be confused with CopyType (Syn/Async) in HP Storage Essentials. CopyType refers to the replication pair's relationship. Sync means the source is always kept in sync with the target. Async means the target is disassociated from the source volume as in, for example, a point-in-time copy.

The CV EVA SMI-S provider uses a caching scheme to provide consistent data and better performance to client applications. This may cause a replica pair's properties to not appear (in HP Storage Essentials) to be in sync with what CV EVA shows. When the EVA SMI-S provider's per-EVA cache is refreshed (typically every 30 minutes) the replica pair's data is refreshed.

	Remote Replicas via HP Continuous Access (DR Groups)
Locality	Source/a target depending on which device is being viewed
Copy type	Sync or async when I/O is suspended
Replica type	Full copy
Sync state	Synchronized or fractured when I/O is suspended
Sync maintained	True, false when I/O is suspended
When synced	Date/time the replica was created

HP SAN Virtualization Services Platform (SVSP) Replication

	Snapshots	SnapClone Groups	Async Mirror Groups	Sync Mirror Groups
Locality	Local Pair	Local Pair	Local Pair	Local Pair
Replica type	Full copy	Full copy	Full copy	Full copy
CopyType	UnSyncAssoc	UnSyncUnAssoc	Async	Sync
Sync state	Idle Resync In Progress Restore In Progress Copy In Progress	Copy in progress Idle	Synchronized Fractured Broken	Synchronized Resync In Progress Fractured Broken

Remote replication pairs are not supported for HP SVSP devices.

CopyType defines the type of (copy) association between a source and target. The supported values are:

- "Async" – Create and maintain an asynchronous copy of the source.
- "Sync" – Create and maintain a synchronized copy of the source.
- "UnSyncAssoc" – Create an unsynchronized copy and maintain an association to the source.
- "UnSyncUnAssoc" – Create an unsynchronized copy with a temporary association that is deleted upon completion of the copy operation.

Because SnapClone CopyType is UnsyncUnAssoc, the replication pair association is transient. If you run a GEAD in HP Storage Essentials while the snap is being created, you might see the pair show up in the HP Storage Essentials GUI. But if you do not run a GAED while the transient association briefly exists (or you run a GAED later after it is gone) you will not see the replication pair for the SnapClone in the HP Storage Essentials GUI.

Sync State describes the state of the association with respect to replication activity. The supported values are:

- "Resync In Progress" – Synchronization or resynchronization is in progress. This may be the initial copy or subsequent changes being copied.
- "Synchronized" – An async or sync replication is currently synchronized.
- "Restore In Progress" – An operation is in progress to copy the synced object to the system object.

- "Idle" – The normal state for an UnSyncAssoc replica.
- "Broken" – The relationship is non-functional due to errors in the source, the target, the path between the two, or space constraints.
- "Fractured" – An async or sync replication is fractured.
- "Copy In Progress" – A deferred background copy operation is in progress to copy the source to the replica target for an UnSyncAssoc association.

HP XP Array Replication

HP Storage Essentials presents replication-state information using SMI-S terminology. Some SMI-S terms do not have an obvious device-specific equivalent.

HP Storage Essentials maps the following HP XP terms as follows:

	Continuous Access	HP Continuous Access Journal	HP Business Copy	HP XP Snapshot
Locality	Remote pair	Remote pair	Local pair	Local pair
Replica type	Full copy	Full copy	Full copy	After delta
Copy type	Sync/async depending on cache journaling in use	Async	Sync	UnSyncAssoc
Sync state	Paired, idle, failed, suspended	Active, halted, stopped	Copy, pair, psus	Idle, pair

Whenever the locality is a remote pair, the remote system serial number and volume ID are displayed. Volume ID is the devNum (CU:LDEV converted to decimal). If the remote system is also discovered by HP Storage Essentials, the replication table links directly to that volume on the remote system.

For Universal Replicator and Continuous Access Journal, HP Storage Essentials displays the individual journal groups containing the journal LDEVs and categorizes their storage capacity separately so that it is accounted for but not considered as available capacity.

NetApp Devices Replication

HP Storage Essentials discovers NetApp devices using the NetApp DATA ONTAP SDK over HTTP or HTTPS. Most DATA ONTAP 7.x devices are supported.

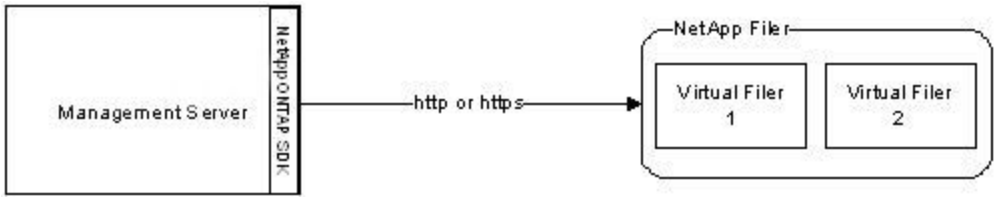

To discover a NetApp device, use FQDN, IP address, or HTTP(S) URL. If all NetApp filers are configured using HTTPS, you can set the internal custom property "cimom.netapp.useSSL=true" to enable users to enter just the FQDN or IP address instead of the full HTTPS://FQDN:443. The assumption is that the default port will be used for SSL communication.

NetApp virtual filers are discovered through the main physical filer's address. Once you perform initial identification, any devices configured through the NetApp Multistore license are shown alongside the main device in the discovery screen.

Snapshot

Snapshot replications are point-in-time, frozen deltas of the files since the last snapshot. These are taken periodically and after changes are made on the file system (after delta). These replicas are local to the filer only; hence, "local pair" for the locality.

Snapshot	
Locality	Local pair
Replica type	After delta
Copy type	UnsyncAssoc
Sync state	Frozen

SnapMirror

SnapMirror replications are full copy replicas of the source volume and are synchronized according to time periods that users configure. So that users can understand the location of these remote replicas, a Locality field describes whether the source or target resides on the local system.

	SnapMirror
Locality	Remote pair
Replica type	Full copy
Copy type	Async
Sync state	Target always synchronized as it is periodically updated to be a replica. Source is idle/busy depending on whether or not a SnapMirror update is in progress.

HP P4000 Device Replication

You can view snapshot copies that are configured on an HP P4000 cluster through the Replication Pairs panel.

The table in the panel follows the SMI-S Copy Services profile and is used to provide a common set of terms across all devices. Only local snapshots are collected from an HP P4000 cluster.

Select All Pages | Unselect All Pages

Source	Target	Copy Type	Replica Type	When Synced	Sync State	Sync Maintained	Locality	Remote System Id	Sync State Collection Time
NonTPVolume	NonTPVolume_Sch_RS_1_Pri.3573	UnSyncAssoc	After Delta	2010-04-09 22:17	Synchronized	true	Local Pair		2010-04-09 20:14
Part1	Part1_SS_1	UnSyncAssoc	After Delta	2009-12-16 18:24	Synchronized	true	Local Pair		2010-04-10 21:05
newTP	newTP_SS_1	UnSyncAssoc	After Delta	2009-11-18 22:17	Synchronized	true	Local Pair		2010-04-10 21:05
vol0_replica	vol0_replica_RS_1	UnSyncAssoc	After Delta	2009-11-10 21:16	Synchronized	true	Local Pair		2010-04-10 21:05
newAlert	newAlert_Sch_SS_1.389	UnSyncAssoc	After Delta	2010-04-09 22:53	Synchronized	true	Local Pair		2010-04-09 20:14
testRemote	Part1_Sch_RS_1_Rmt.498	UnSyncAssoc	After Delta	2010-04-09 22:54	Synchronized	true	Local Pair		2010-04-09 20:14
Part1	Part1_Sch_RS_1_Pri.496	UnSyncAssoc	After Delta	2010-04-09 21:54	Synchronized	true	Local Pair		2010-04-09 20:14
Part1	Part1_Sch_RS_1_Pri.497	UnSyncAssoc	After Delta	2010-04-09 22:24	Synchronized	true	Local Pair		2010-04-09 20:14
NonTPVolume	NonTPVolume_SS_1	UnSyncAssoc	After Delta	2009-11-10 21:16	Synchronized	true	Local Pair		2010-04-10 21:05
Part1	Part1_Sch_RS_1_Pri.498	UnSyncAssoc	After Delta	2010-04-09 22:54	Synchronized	true	Local Pair		2010-04-09 20:14

A collector can be configured to update the When Synced column information more frequently than each Get Details interval.

Properties include the source, destination, and state of the replication. The state can be collected at a user-defined time interval through an HP Storage Essentials collector.

Selecting a volume shows the volume and the replicas that are either the source or target of that volume. The full replica details can also be viewed as a property page, as follows:

Replication Pair Part1 - Part1_SS_1

Sync State Collection Time	2010-04-10 21:05	**Copy Type**	UnSyncAssoc
Sync Maintained	true	**Sync State**	Synchronized
Contacted	2010-04-09 20:14	**Record Created**	2010-04-07 12:00
Locality	Local Pair	**Discovery Status**	Contacted
Replica Type	After Delta	**Description**	
When Synced	2009-12-16 18:24	**Remote Element Identifier**	
Remote System Identifier		**Source Storage Volume**	Part1
Storage System	ManagementGroup0:Cluster0	**Target Storage Volume**	Part1_SS_1

3 Discovering Applications, Backup Hosts, and Hosts

This section contains the following topics:

Step 1 – Discovering Your Hosts and Backup Manager Hosts

Before you can discover your applications, you must discover their hosts. You discover hosts in the same way you discovered your switches and storage systems. You provide the host's IP address, user name and password. The user name and password must be from a valid account or you can enter credentials that were provided in the **cxws.default.login** file, as described in the section, "Creating Default Logins for Hosts," in the installation guide.

Unlike switches and storage systems, you must have installed a CIM extension on the host if you want to obtain detailed information about the host and its applications. See the support matrix for your edition for information about which backup applications the management server supports without a CIM extension installed.. For information about installing CIM extensions, see the "Deploying and Managing CIM Extensions" chapter of the installation guide.

For information about discovering clustered hosts, see Host and Application Clustering on page 267.

For information about discovering virtual machines, see Discovering Virtual Machines on page 184.

The management server automatically detects file servers on hosts through discovery. Before you map the topology (Step 2 in Discovery Setup), make sure the option for File System Viewer is selected, as described in Step 2 – Build the Topology on page 203.

The management server also detects the backup applications its supports, such as Veritas NetBackup, HP Data Protector,EMC NetWorker, and IBM Tivoli Storage Manager. If you are licensed for Backup Manager and you want to manage and monitor your backup applications, select **Include backup details** when you run Get Details, as described in Step 4 – Get Details on page 204.

Keep in mind the following:

- You must install a CIM extension on any virtual machines that will be participating as a cluster node.

- Direct iSCSI links to hosts are only displayed if a CIM extension is running on the host. For VMs discovered through the ESX or VC server, these direct iSCSI links will not be seen because they are not discovered through the ESX or VC server.
- After installing the CIM extension on a Data Protector system on Windows, check the Logon account for the DataProtector CRS service and verify that it matches the AppStorWin32Agent service. To determine the Logon account for the DataProtector CRS service, go to **Control Panel** > **Administrative Tools** > **Services**, select the DataProtector CRS service, access its Properties page, and select the **Logon** tab. To determine the Logon account for the AppStorWin32Agent service, go to **Control Panel** > **Administrative Tools** > **Services**, select the AppStorWin32Agent service, access its Properties page, and select the **Logon** tab.
- If you change the password of a host after you discover it, stop and restart the CIM extension running on the host, and change the host password in the discovery list.
- If your license lets you discover UNIX and/or Linux hosts, the Test button for discovery reports SUCCESS from any UNIX and/or Linux hosts on which the management server can detect a CIM extension. The CIM extension must be running. The management server reports "SUCCESS" even if your license restricts you from discovering certain types of hosts. For example, assume your license lets you discover Solaris hosts but not AIX hosts. If you click the **Test** button, the management server reports "SUCCESS" for the AIX hosts. You will not be able to discover the AIX hosts. The IP address is not discoverable, because of the license limitation.
- To receive status reports about Get Details, see Configuring E-mail Notification for Get Details on page 863 for information about how to configure this option.
- Depending on your license, you might not be able to access Backup Manager, File System Viewer and/or monitor certain applications might not be available. See the List of Features to determine if you have access toBackup Manager, File System Viewer and/or are able to monitor the other applications. The List of Features is accessible from the Documentation Center (**Help** > **Documentation Center** in HP Storage Essentials). To learn more about File System Viewer, see the File Servers Guide, which is also available from the Documentation Center.
- If you are unable to discover a UNIX host because of DNS or routing issues, see Unable to Discover a UNIX Host Because of DNS or Routing Issues on page 893.
- Get Details can hang if obtaining information from an AIX host where SAN storage was previously available is no longer visible to the operating system. You might need to reboot the management server to resolve this issue.
- When discovering a Linux host from the management server, the operating system/server type is not available.
- If you started a CIM extension on a Sun Solaris host using the `cim.extension.parameters` config file or the `./start -users` command, the user name provided in the command must be used to discover the host. For example, if you use `./start -users myname:yourname` (in this instance, myname and yourname are valid UNIX accounts) to start the CIM extension, you must use myname or yourname and its password to discover the host.

- If you try to discover a Solaris host with multiple IP address, the management server picks only one IP address for discovery.
- You can configure the management server to obtain information about your Backup Manager hosts at a set interval. See the topic "Scheduling Backup Collection for Backup Managers" in the User Guide for more information about collectors.
- The backup collection for Data Protector runs as follows:
  - By default, the backup collection does not run when you start the CIM extension. The backup collection is triggered once Get Details runs.
  - During the background collection, the following processes are involved:
    - **Session background collector** runs every 15 minutes.
    - **Media background collector** runs every 24 hours.

Discovery of hosts consists of the following tasks:

- **Setting Up** – Finding the elements on the network. See Step 1 – Set Up Discovery for Hosts below.
- **Topology** – Mapping the elements in the topology. See Step 2 – Build the Topology on page 203.
- (Optional) Step 3 – View the Topology on page 203
- **Details** – Obtaining detailed element information. See Step 4 – Get Details on page 204.

Step 1 – Set Up Discovery for Hosts

Some elements require additional steps before discovering hosts. If you are discovering:

- Virtual machines, see Discovering Virtual Machines on page 184 before starting the discovery process.
- Backup servers, see Discovering Backup Servers on page 201 before starting the discovery process.

To discover hosts:

1. Click **Discovery** > **Setup**.
2. If several of the elements in the same domain use the same name and password, click the **Set Default User Name and Password** link. Provide up to three user names and passwords.

The management server tries the default user names and passwords for elements during discovery. For example, if you have a several hosts using the same user name and password, add their user name and password to the list of default user names and passwords. If one of the hosts is connected to a storage system with another user name and password, you would also add this user name and password to the list. Do not specify the user name and password for the storage system in the individual range because that overrides the default user name and password.

To access a Windows-based device, prefix the user name with domain_name\, as shown in the following example. This is required by the Windows login mechanism.

```
domain_name\username
```

In this instance:

- domain_name is the domain name of the element
- username is the name of the account used to access that element

To add an IP address range to scan:

1. Click the **IP Ranges** tab.
2. Click the **Add Range** button.
3. In the **From IP Address** box, enter the lowest IP address in the range of the elements you want to discover.
4. In the To IP Address box, enter the highest IP address in the range of the elements you want to discover.
5. In the User Name (Optional) box, enter the user name.

   To access a Windows-based device, prefix the user name with the Windows domain name, as shown in the following example. It is required by the Windows login mechanism.

   ```
   domain_name\username
   ```

   In this instance:

   - domain_name is the domain name of the element
   - username is the name of the account used to access that element
6. In the **Password (Optional)** box, enter the password corresponding to the user name entered in the User Name box.
7. Enter the password from the previous step in the **Verify Password** box.
8. In the **Comment** box, enter a brief description of the servers. For example, Servers in Marketing.
9. Click **OK**.
10. Repeat steps b through i until all of the IP ranges have been entered.
11. Click the **Start Scanning** button.

The elements the management server detects during the scan are added to the Addresses to Discover list on the IP Addresses tab.

To add a single IP address or DNS name to discover:

1. Click the **IP Address** tab.
2. Click the **Add Address** button.
3. In the **IP Address/DNS Name** box, enter the IP address or DNS name of the device you want to discover.
4. In the **User Name (Optional)** box, enter the user name.

   This box can be left blank if one or more of the following conditions are fulfilled:

   - The element's user name and password are one of the default user names and passwords.
   - The element does not require authentication.

   To access a Windows-based device, prefix the user name with the Windows domain name; for example:

   ```
   domain_name\username
   ```

   In this instance:

   - domain_name is the domain name of the machine
   - username is the name of your network account
5. In the **Password (Optional)** box, enter the corresponding password for the user name entered in the previous step.

   This box can be left blank if one or more of the following conditions are fulfilled:

   - The element's user name and password are one of the default user names and passwords.
   - The element does not require authentication.
6. If you entered a password in the previous step, entered the password in the **Verify Password** box.
7. In the **Comment** box, enter a brief description of the server. For example, Server Used for Nightly Backups.
8. Click **OK**.

To start discovering elements on the network, click the **Start Discovery** button on the IP Addresses tab. The software discovers the IP addresses selected.

During discovery, the following take place:

- The status light changes from green to orange.
- You are shown the Log Messages page. To view the status of discovery, click **Discovery** > **View Logs**.

Discovery is complete when the DISCOVERY COMPLETED message is displayed in the Log Messages box.

Discovering Virtual Machines

The following topics provide instructions for discovering VMware virtual machines and Solaris virtual servers.

- Port Requirements for Discovering Virtual Servers on page 187
- Differences between Virtual Machines with a CIM Extension Installed and those Without on page 187
- Disabling Automatic Discovery of Virtual Machines on page 188
- Known Issues for ESX Servers on page 189

Discovering VMware Virtual Machines

You must install and run VMTools on each virtual machine. If VMTools is not running, the virtual machine will be unmanaged and only limited data will be available. For example, unmanaged virtual machines will not be displayed on the element topology for the associated discovered hosts.

Virtual machines are discovered in the same way as physical hosts, but there is an additional consideration for virtual machines. Virtual machines can be discovered through the VirtualCenter or through the individual ESX Servers. If you discover virtual machines through the VirtualCenter, you must provide the user name and password for a VirtualCenter account that can view or access the ESX Servers or virtual machines that you want to discover.

You can use any VirtualCenter account credentials, provided that the associated user's role has Datastore Browse privileges.

All ESX Servers and virtual machines that the VirtualCenter account can view or access are automatically discovered. For example, if a VirtualCenter has 15 ESX Servers and you provide the user name and password for a user account that can view or access just five ESX Servers, only those five ESX Servers are discovered. For this reason, discovering the VirtualCenter is the recommended process.

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

However, if you intend to use custom discovery lists, it is necessary to discover each ESX Server individually because discovering the VirtualCenter results in one access point for all the ESX Servers managed by that VirtualCenter. If you discover the ESX Servers individually, you will have an access point for each server, and all of the virtual machines are still discovered automatically. If you discover virtual machines through the individual ESX servers, you must use the ESX server's credentials.

To discover applications hosted on a virtual machine, or you want the virtual machine to participate as a cluster node, you must discover the virtual machine as described in Step 1 – Set Up Discovery for Hosts on page 181. In addition, you must install a CIM extension on the virtual machine. CIM extensions should not be installed on virtual servers. For information about installing CIM extensions, see the "Deploying and Managing CIM Extensions" chapter of the installation guide.

If you perform additional Get Details for a virtual machine, you must include the access points for both the virtual machine and its associated VirtualCenter or ESX Server. Performing Get Details for just the virtual machine will result in a lack of connectivity between the virtual machine and the ESX Server.

The management server discovers templates as powered off virtual machines. Templates are only discovered when you discover virtual machines through the VirtualCenter. If you discover individual ESX servers directly, the templates will not be found.

For ESX 4.x, the management server checks the status of VMTools on the virtual machine. If VMTools is not running on the virtual machine, then the management server cannot discover the virtual machine as a managed host. You can find the status of VMTools by looking at the VMTools field on the Properties tab for the virtual machine. If the VMTools field says "GuestToolsRunning," then VMTools is running on the virtual machine. There are multiple ways to access the Properties tab. One way is to double-click the virtual machine in System Manager and then click the Properties tab.

How Virtual Elements are Displayed

Virtual elements are displayed in Discovery Step 2 as follows:

In Discovery Step 1, if you discover the following	Discovery Step 2 displays the following
VirtualCenter	The VirtualCenter's access point with the associated virtual servers listed in the Elements colum; for example: • **IP address/DNS Name** (of the VirtualCenter) – https://192.168.1.1 • **Elements Column** – Names of the virtual servers managed by the VirtualCenter

In Discovery Step 1, if you discover the following	Discovery Step 2 displays the following
Virtual server	The virtual server's access point; for example: • **IP address/DNS Name** (of the virtual server) – https://192.168.1.1 • **Elements Column** – Virtual server name
Virtual machine with VMTools	The virtual server's or VirtualCenter's access point; for example: • **IP address/DNS Name** (of the virtual server or VirtualCenter) – https://192.168.1.1 • **Elements Column** – Virtual server or VirtualCenter name
Virtual machine with VMTools and a CIM extension	The virtual machine's access point; for example: • **IP address/DNS Name** (of the virtual machine) – cxws://192.168.1.1 • **Elements Column** – Virtual machine name

Virtual elements are displayed in Discovery Step 3 as follows:

If you get details for the following	Discovery Step 3 displays the following
VirtualCenter	The VirtualCenter's access point with the associated virtual servers listed in the Elements column; for example: • **IP address/DNS Name** (of the VirtualCenter) – https://192.168.1.1 • **Elements Column** – Names of the virtual servers managed by the VirtualCenter
Virtual server	The virtual server's access point; for example: • **IP address/DNS Name** (of the virtual server) – https://192.168.1.1 • **Elements Column** – Virtual server name
Virtual machine with VMTools	There is no access point for a virtual machine unless it has a CIM extension installed and is configured for discovery in Step 1.

If you get details for the following	Discovery Step 3 displays the following
Virtual machine with VMTools and a CIM extension	The virtual machine's access point. The virtual machines will also be listed in the Elements column of the associated virtual server. For example: • **IP address/DNS Name** (of the virtual machine) – cxws://192.168.1.1 • **Elements Column** – Virtual machine name

Excluding Virtual Machines from Discovery

To reduce the number of MAPs counted, exclude virtual machines from discovery by setting the cimom.discovery.exclude.vmware.vm property to true. When the cimom.discovery.exclude.vmware.vm property is set to true, data from ESX servers is collected but not data from virtual machines.

To exclude virtual machines from discovery:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Paste the following text into the Custom Properties box.

   ```
   cimom.discovery.exclude.vmware.vm=true
   ```

4. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

Port Requirements for Discovering Virtual Servers

Use the following default ports when discovering virtual servers or VirtualCenters:

- **HTTPS** – Port 443
- **HTTP** – Port 80

Non-standard ports can be specified; for example: https://192.168.1.1:444.

Differences between Virtual Machines with a CIM Extension Installed and those Without

The management server does not require that CIM extensions be installed on virtual machines, but additional functionality is provided for virtual machines with a CIM extension installed.

Feature	CIM Extension Not Installed	CIM Extension Installed
Application Discovery	No. Applications cannot be discovered.	Yes. All supported applications can be discovered.

Feature	CIM Extension Not Installed	CIM Extension Installed
File System Type	No. VMware does not provide enough information to know the file system type of the OS.	Yes. Behaves just like a physical host with a CIM extension installed.
File System Percentage Used	Yes. Capacity Manager and Report Optimizer will report the used, free, and total capacity of the virtual machine partitions.	Yes
Disk Partition Discovery	No. Disk level information is not available.	Yes
Connectivity to ESX Server (Topology)	Yes. Application level topology will be available.	Yes
Drive Type of Storage Volume	No	Yes
Storage Based Chargeback	No. Chargeback Manager requires application discovery which requires a CIM extension.	Yes
Raw Device Mapping (RDM)	Yes	Yes
Multipathing and Volume Management	No	Yes
FSRM Support	No	Yes
Host Performance	No	No

Disabling Automatic Discovery of Virtual Machines

In the current version of the management server, you can disable automatic discovery of virtual machines on ESX servers by changing a JBoss property. You might want to disable automatic discovery of virtual machines so that you do not exceed the total MAPs permitted by your licenses.

In previous releases, if you configured the management server to discover a virtual center or individual ESX servers, Step 2 and Step 3 discovery automatically discovered all of the virtual machines on ESX servers and counted each as a MAP.

Disable the automatic discovery of virtual machines, as described in Excluding Virtual Machines from Discovery on page 187.

If virtual machines were previously discovered, after changing the property, the virtual machines will no longer be discovered and will show up as missing. If the virtual machines were not deleted, they will continue to show up as missing in System Manager, but without any connectivity. They will not be counted as a MAP. Missing virtual machines will be restored if the property is changed back to false and Get Details is performed.

Known Issues for ESX Servers

A known third-party issue related to ESX Servers causes the management server to present incomplete or erroneous information. The issue occurs when a LUN is shared by more than one ESX Server. The following problems are a result of this issue:

- Some shared external storage volumes for a virtual machine are reported with drive types of local instead of external.
- A virtual machine's element topology will appear as having only local (to the ESX Server) storage instead of external storage.
- The Volumes column in the Multipathing Software table for a virtual machine is blank instead of containing the name of the external storage volume.
- In the End to End Connectivity Report, ESX Servers reporting back as not connected display "Not connected to external storage" in the Storage System column.
- iSCSI-based storage, such as storage from the HP StorageWorks P4000, mounted directly to the host with a drive letter shows as expected on the Dependent Storage Systems page for the host. If the mounted drive is created from an ESX datastore created using the iSCSI volume, the mount point will display the ESX datastore identification instead of the host-specified drive letter.

Discovering Solaris Containers

Solaris Containers is a server virtualization technology implemented by Sun for the Solaris operating system. Solaris Containers provide isolation between software applications or services using flexible software-defined boundaries.

Applications can be managed independently of each other, even while running in the same instance of the Solaris Operating System. Solaris Resource Manager and Solaris Zones software partitioning technology are both parts of the Solaris Container environment.

These components address different qualities the container can deliver and work together to create a complete container. A zone is a virtualized operating system environment created within a single instance of the Solaris Operating System.

When you create a zone, you produce an application execution environment in which processes are isolated from the rest of the system. This isolation prevents processes that are running in one zone from monitoring or affecting processes that are running in other zones. Solaris zones have been introduced in the Solaris 10 operating system. Solaris defines two types of Solaris zones:

- **Virtual server/physical host (Global Zone)**: The virtual server/physical host is the default zone for the system and the zone used for system-wide administrative control. All processes run on the virtual server/physical host if there are no virtual machines/Solaris Containers (non-global zones) that were created by the global administrator. Virtual machines/Solaris Containers (non-global zones) are also sometimes referred simply as zones.
- **Non-Global Zone (virtual machine/Solaris Container)**: The various instances of the virtual operating system environment, which are created to execute applications correspond to the virtual machine/Solaris Container. The virtual machines/Solaris Container are configured to have virtual network interface, one or more file systems and a virtual console.

HP Storage Essentials enables you to discover the zone portion of the Solaris Containers virtual infrastructure. The Solaris Containers virtual infrastructure in System Manager, Capacity Manager and element topology provides a comprehensive and convenient way to track storage.

The Solaris Containers infrastructure has two types of host:

- **The physical host or the Global Zone**: To maintain uniformity with other server virtualization support in HP Storage Essentials, the physical host or global zone is also referred to as the virtual server in HP Storage Essentials.
- **Solaris Containers or the Non Global Zone**: To maintain uniformity with other server virtualization support, Solaris Containers are referred to as virtual machines in HP Storage Essentials.

Each virtual server/physical host IP address corresponds to a single access point. The virtual servers/physical hosts can be distributed among available discovery groups for load balancing. All the functionality applicable to a Solaris managed host would be applicable to the virtual server/physical host.

For the agentless virtual machine/Solaris Container, HP Storage Essentials displays the connection between the file system of a virtual machine/Solaris Container and corresponding device (partition, host logical volume, file system) of the virtual server/physical host and onto a remote SAN Storage.

A virtual machine/Solaris Container is considered for discovery in all of its states. If the virtual machine/Solaris Container is in the running state when discovered, it is considered as a managed host and in all the other states it is considered as a unmanaged host.

During the building of the topology of virtual servers and virtual machines, virtual servers/physical hosts and virtual machines/Solaris Container are discovered along with few of their components.

During the Get Details of virtual servers and virtual machines, virtual servers and virtual machines are discovered, along with all of their components. Applications running on virtual servers and virtual machines are also discovered in this step.

Oracle configured on file systems is supported on Solaris virtual machines/Solaris Container. Oracle on raw device or on ASM is not supported in Solaris virtual machines/Solaris Container. CIM Extensions should not be installed on Solaris virtual machine/Solaris Container for Oracle discovery.

If you delete a Solaris Container and perform a Step 3 Detailed Discovery on the management server, the deleted Solaris Container still appears in the Policy Manager, Capacity Manager, and Report Data Collectors pages.

Steps for Discovering Solaris Containers

To discover Solaris Containers:

1. Install the CIM extension for Solaris on the virtual server/physical host (global zone).

   Never install a CIM extension on the virtual machine/Solaris Container (non-global zone). You might be tempted to install a CIM extension for Oracle, but Oracle configured on file systems is supported on virtual machines/Solaris Containers without a CIM extension. Oracle on raw device or on ASM is not supported on the virtual machine/Solaris Container.

2. Select **Discovery** > **Setup** and click the **Add Address** button.
3. Type the IP addresses of the Solaris host with the CIM extension in the IP Address/DNS Name field.
4. Type the password of the Solaris host with the CIM extension in the Password field.
5. Retype the password in the Verify Password field.
6. Click **OK**.
7. Build the topology as described in Step 2 – Build the Topology on page 203 (optional) and perform Get Details, as described in Step 4 – Get Details on page 204.

Discovering IBM VIO

The IBM Virtual IO infrastructure has two types of host:

- **The physical host or the VIO servers** - This is equivalent to the term virtual servers supported in HP Storage Essentials.
- **The virtual hosts or the VIO clients** - This is equivalent to the term virtual machines supported in HP Storage Essentials.

The discovery of IBM VIO requires the discovery of the virtual servers and all the virtual machines.

The management server can discover virtual machines on which CIM extensions have not been installed. To enable agentless discovery, the CIM extensions for AIX running on the VIO server uses the AIX CLIs through SSH to get various properties of each VIO client. To enable SSH communication, you must install the SSH service on each of the VIO client and the SSH client on the Virtual IO server. The AIX CIM extension uses the SSH channel to fetch VIO client details by using the IP address and the other credentials provided during Discovery Step 1.

To enable discovery of virtual machines, you must install CIM extensions on the selected virtual servers. You are not, however, required to install CIM extensions on each virtual machine.You are required to install the CIM extension on the virtual machines only if the virtual machine is attached to a host bus adapter connected to a SAN. Provide the IP address of the selected virtual server for discovery. VIO servers are discovered in the same way as physical hosts.

To complete the discovery of virtual machines, provide the IP address of each virtual machine hosted on a VIO server.

Steps for Discovering IBM VIO

Keep in mind the following:

- You must provide the IP addresses of all the VIO clients during discovery. This enables the CIM extensions installed on the virtual IO server to discover the VIO clients.
- You are not required to install the CIM extensions on the VIO clients.
- You must find the partition ID of the VIO clients in relation to the VIO server hosting it.
- Do not include the IP address of the VIO server while providing the IP addresses range in the **Add Range for Discovery** window. HP Storage Essentials does not support discovery, if the IP address of the VIO server forms a part of the IP address range.

Step 1 - Discovering Virtual Servers as Host:

Before you discover the virtual servers as host,make sure that a CIM extension is installed on the selected VIO server. For more information on installing the CIM extensions, see the *HP Storage Essentials Installation Guide*.

To discover a virtual server:

1. Select **Discovery** > **Setup**.
2. Select **Step 1** at the top of the page.
3. Click **IP Addresses** tab.
4. Click **Add Address**tab, the **Add Address for Discovery** window opens.
5. In the **IP Address/DNS Name** field, type the DNS Name/IP address of the VIO server with the CIM extension.
6. In the User Name box, type the user name of the VIO server with the CIM extension.
7. In the Password box, type the password of the VIO server with the CIM extension.
8. In the Verify Password box, re-type the password.
9. (Optional) In the Comment box, enter any additional information. The information entered in this box is displayed in the Comment column in the Addresses to Discover list (**Discovery** > **Setup**).
10. Select the **Is VIO Server** check box. This marks the specified hosts as a VIO server. The client discovery details appear only if you mark the host as a VIO server.

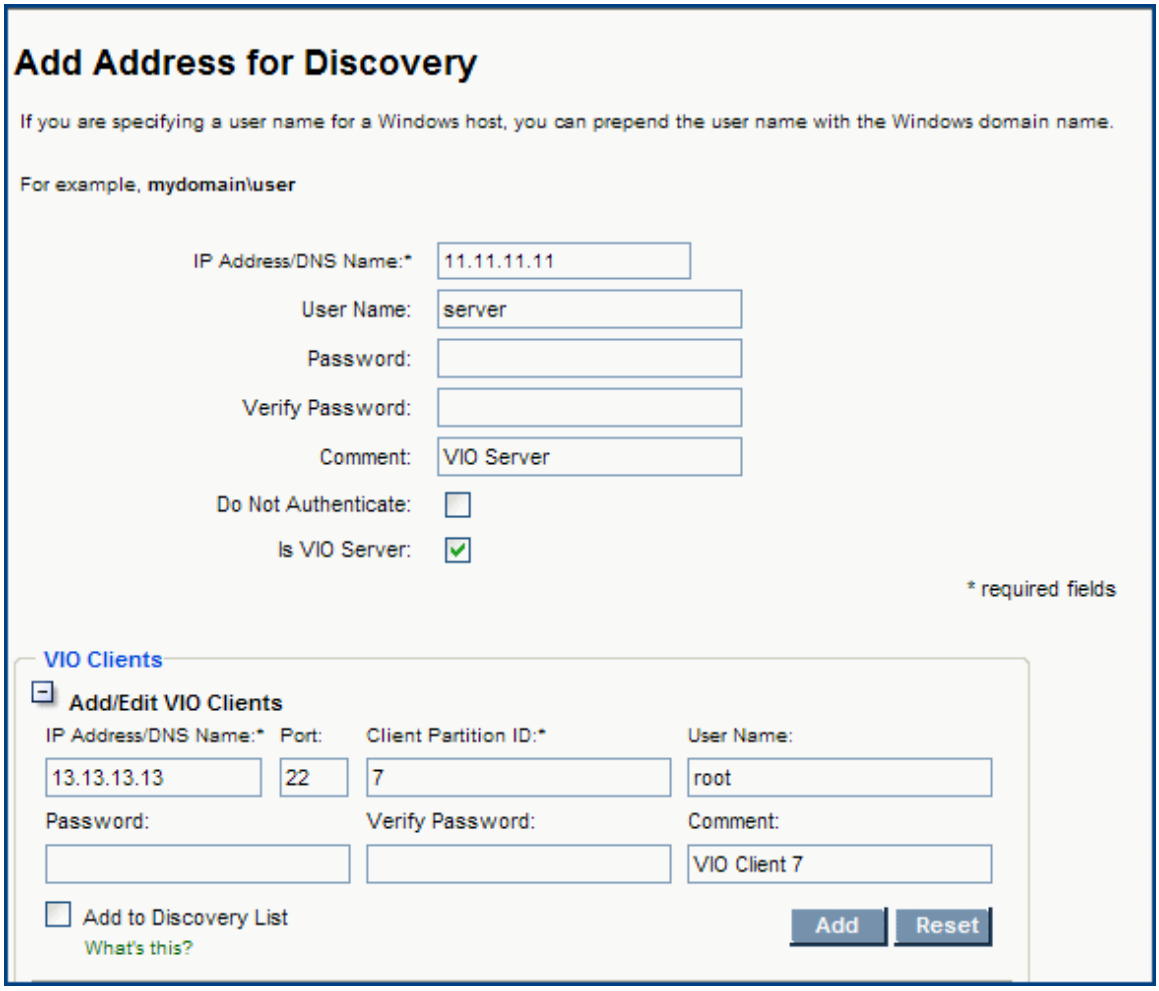

Step 2 - Discovering Virtual Client

To discover a virtual client:

1. In the IP Address/DNS Name field, type the DNS Name/IP address of the VIO clients.
2. By default, the Port box is populated with 22, but you can change the default port number.
3. In the Client Partition ID box, provide the client partition ID. To find the partition ID, log on to the host or the IBM Hardware Management Console. Or you can log on to the VIO client and run the command `uname -Ls` to find the partition ID.
4. In the User Name box, provide the user name of the VIO client.
5. In the Password box, type the password of the VIO client.
6. In the Verify Password box, re-type the password.
7. Click **Add**.

8. (*Optional*) Select the **Add to discovery list**check box. When you select this option, the VIO client information is added to the discovery list. Use this option only when CIM extensions are installed on the VIO client, or the VIO client is attached to a host bus adapter.

Note: It is not necessary to install CIM extensions on the VIO client. However, you must install CIM extension if the VIO client is attached to a host bus adapter connected to the SAN. If the VIO client is fetching the SAN resources through the VIO server, you need not install the CIM extensions or select **Add to discovery list** option.

Understanding IBM VIO Limitations in HP Storage Essentials

The following limitations are known for IBM VIO with this release of HP Storage Essentials:

- HP Storage Essentials currently does not recognize the physical layer of the machine. Therefore, it treats each VIO server as an individual machine. This is reflected in all the reports and navigation pages of the VIO server.
- A VIO client discovered through Secure Shell (SSH) is reported as an external storage, if the VIO client's disk is mapped to the VIO server's SAN disk. However, if the VIO client's disk is mapped directly to a host bus adapter SAN disk, it is reported as having local storage.
- A VIO client discovered with the CIM extensions is reported as having local storage, if the VIO client's disk is mapped to the VIO server's SAN disk. However, if the VIO client disk is mapped directly to host bus adapter SAN disk, it is reported as external storage.

  Note: You must use an ssh protocol version of 2.0 or above to enable the discovery of a VIO client.
- On a virtual client, if more than one virtual target device with multiple vhosts exists, it cannot fetch the respective vhost number for the different virtual target devices on virtual clients.

Prerequisites for Agentless Discovery of Data Protector

If you have a CIM extension installed, the product will automatically use the CIM extension to discover Data Protector.

Before you discover a Data Protector server that does not have a CIM extension installed, you must do the following:

1. Install the Data Protector Client on the management server. See Step 1 – Install the Data Protector Client on the facing page.
2. Create the DPREPORTER user group for Data Protector Reporter. See Step 2 – Create a User Group for Data Protector Reporter on page 198
3. Create a user in the DPREPORTER user group. See Step 3 – Create a User in the DPREPORTER User Group on page 199
4. Install the Data Protector 6.1 patches on top of the Data Protector 6.1 client or upgrade to the Data Protector 6.11 client. See Step 4 – Install the Data Protector Patch on page 200

Step 1 – Install the Data Protector Client

Install the Data Protector Client on the HP Storage Essentials management server as described in the following steps. These steps apply to Data Protector 6.11, 6.1 and 6.0.

- Linux Installation Steps below
- Windows Installation Steps below

Linux Installation Steps

To install the Data Protector Client:

1. Open the /etc/services file in a text editor, such as vi.
2. Search for 5555 in the text editor.
3. Comment the following two lines in the text editor as follows:

   ```
   #personal-agent 5555/tcp # Personal Agent
   #personal-agent 5555/udp # Personal Agent
   ```

4. Save the services file, and exit the text editor.
5. Copy the Data Protector tar file and extract the tar file.
6. Go to the LOCAL_INSTALL directory.
7. Run the Data Protector installation by entering the following command at the command prompt:

   ```
   ./omnisetup.sh
   ```

8. When asked which components to install, select only the following:
   - User Interface
   - Java GUI Interface

Windows Installation Steps

To install the Data Protector Client:

1. Select the **Client** option in the Setup Wizard and click **Next**.

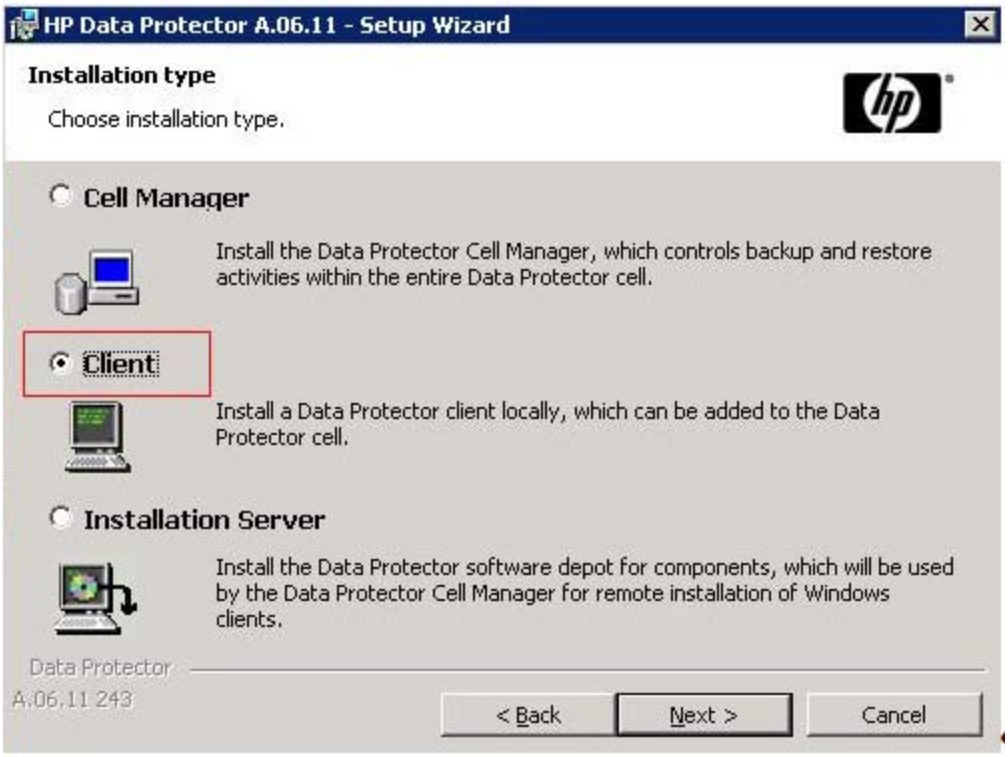

2. Leave the Cell Manager name field blank and click **Next**.

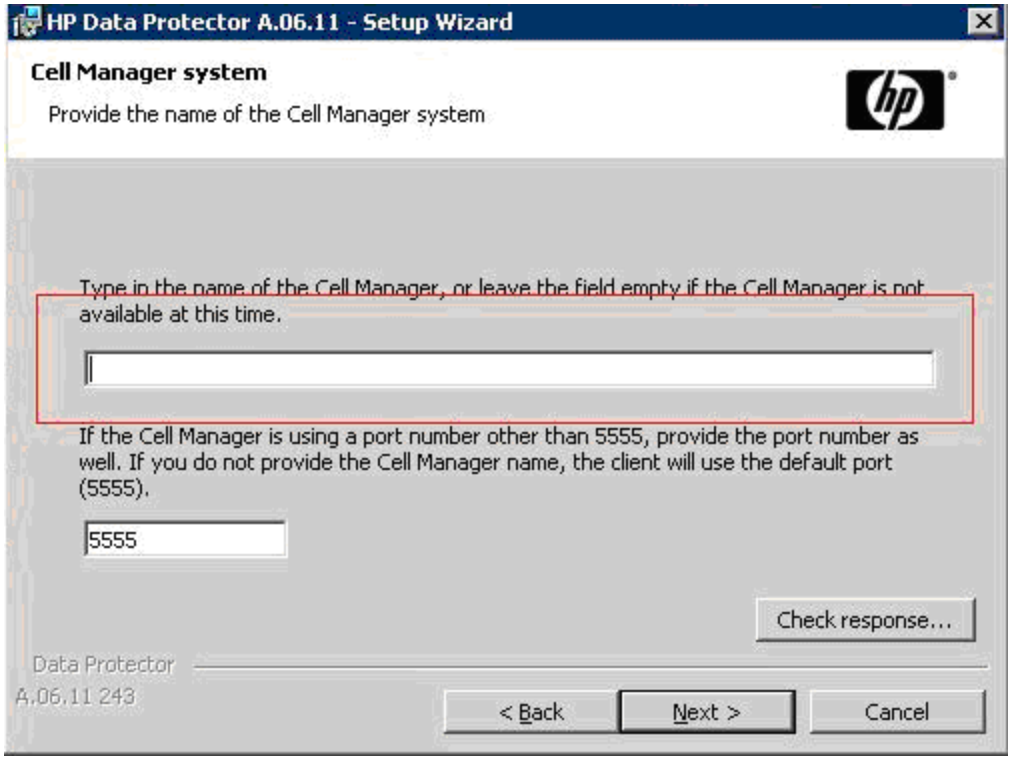

3. Deselect all options, except for the User Interface option, which is selected in the following figure. Click **Next** when done.

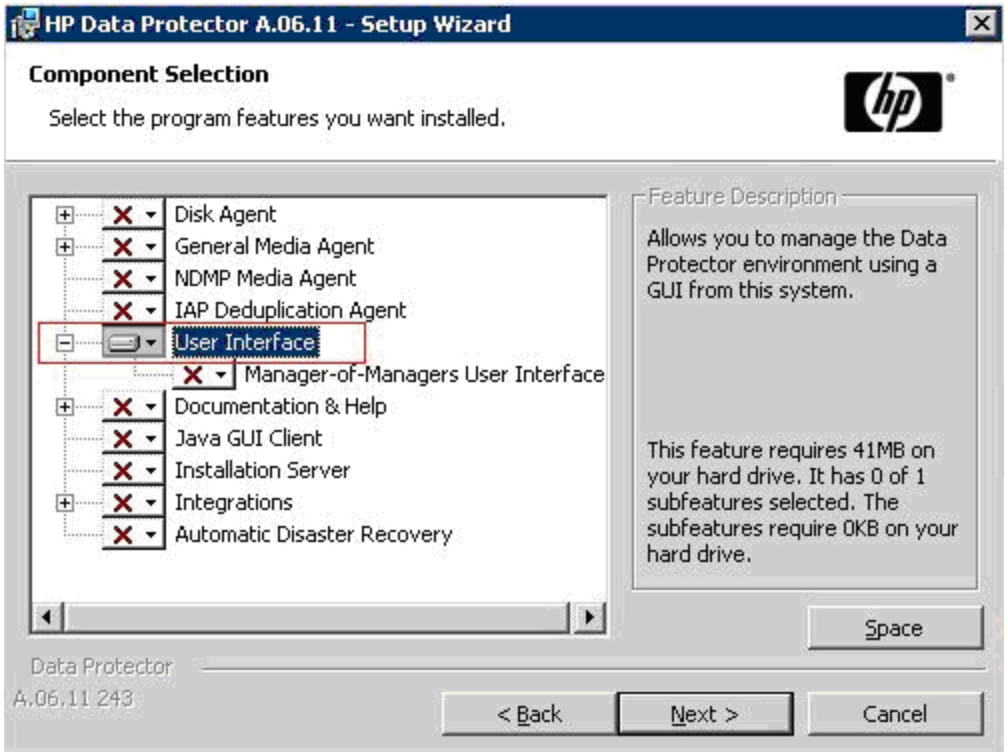

4. Complete the installation by following the steps in the Wizard.

Step 2 – Create a User Group for Data Protector Reporter

Ask your Data Protector Administrator to create a user group for Data Protector Reporter in the Data Protector Cell Manager Console Client as follows:

1. Open the Data Protector Cell Manager Console Client.
2. Go to **Users**. Right-click **Users**, and then click **Add User Group**.

3. Provide the user group name `DPREPORTER`.
4. Deselect the **Start restore** option in the Data Protector User Rights pane. This option is selected by default.
5. Select the following user rights in the Data Protector User Rights pane:
   - Device Configuration
   - Media Configuration
   - Reporting notifications

   The selections should resemble the following:

6. Click **Finish** to create the new user group.

Step 3 – Create a User in the DPREPORTER User Group

Ask your Data Protector Administrator to create a user within the DPREPORTER User Group as follows:

1. (*Windows only*) Before creating the user, make sure that the AppStorManager service, which is the service for HP Storage Essentials, is started on the Storage Essentials management server with the context of a Local Administrator user as the Log On User. You can check in the properties of the Service as follows:

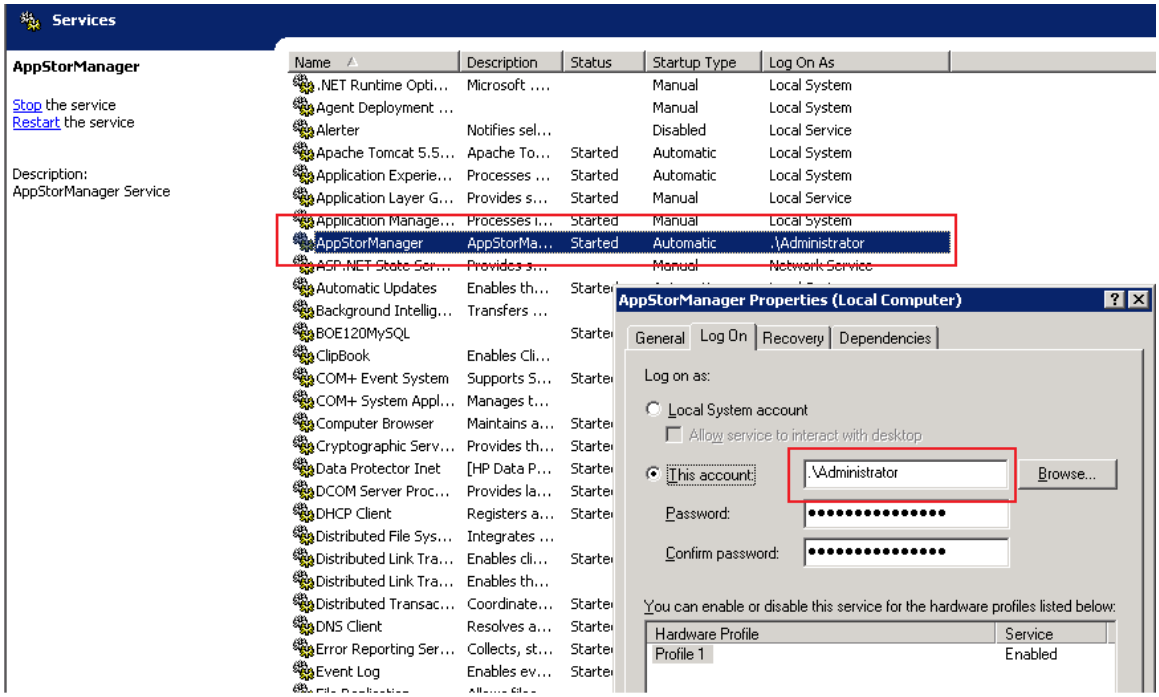

2. Right-click the DPREPORTER group and select **Add/Delete Users**.
3. In the Name field, provide one of the following:
   - **Linux**: The name of the user under which the HP Storage Essentials server process is running. By default, this name is the 'root' user.
   - **Windows**: The name of the user with which the HP Storage Essentials AppStorManager service is running. You can determine the user by looking for the account specified in the **This Account** field on the Log On tab. In this case, the user is Administrator.
4. In the Group/Domain field, provide one of the following:
   - **Linux**: The group information of the user under which the process is running. This can be verified by running the command 'id root' on the HP Storage Essentials management server.

- **Windows**: The host name of the HP Storage Essentials management server, since the AppStorManager service is started as the Local Administrator User.

5. In the Client field, select the DNS name or IP address of the HP Storage Essentials management server.
6. Click **>>** to apply your new user.
7. Click **Finish** to add your new user to the user group.

Step 4 – Install the Data Protector Patch

You need to install the following patch, depending the operating system of the HP Storage Essentials management server, on top of the Data Protector 6.1 client or upgrade to the Data Protector 6.11 client:

- **Linux**: DPLNX_00077
- **Windows**: DPWIN_00417

If you own a valid support contract, you can download patches from http://support.openview.hp.com/selfsolve/patches. You need an HP Passport Account for login. When you access the Patches Search page, select **All Products SPIs** and enter the name of patch, such as DPWIN_00417, in the Optional: Enter keyword(s) or phrases field. Click **Search**. The link to the patch appears under the Search button.

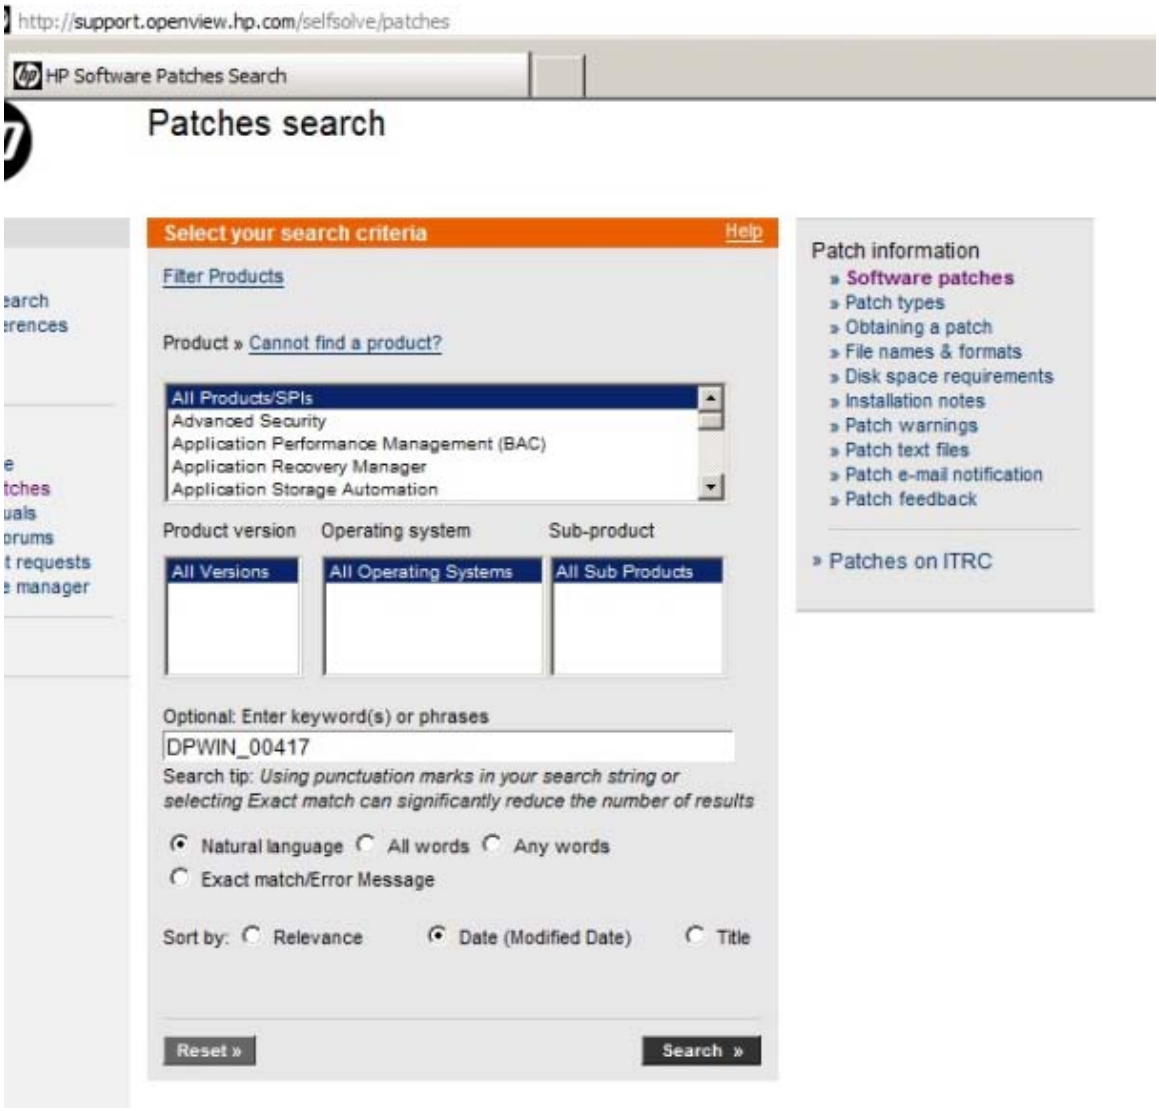

If you do not install the patch or do not upgrade to Data Protector 6.11, the following occurs in Backup Manager:

- Media and media pools details do not appear for discovered backup hosts.
- Policy Details for any session are not displayed in the Policy Detail tab.
- Schedule Details for any session are not displayed in the Schedule Detail tab.

Discovering Backup Servers

Backup Manager monitors your backup applications running on discovered hosts.

Complete the steps in this section if you want to discover backup applications, such as Veritas NetBackup , HP Data Protector, EMC Networker, and IBM Tivoli Storage Manager. See the support matrix for your edition for more information on supported platforms. See Prerequisites for Agentless Discovery of Data Protector on page 194 before you discover Data Protector servers.

1. Confirm that a CIM extension is installed on the server on which Veritas NetBackup or HP Data Protector or EMC Networker or IBM Tivoli Storage Manager is installed. See the Installation Guide for information about installing CIM extensions. Starting with HP Storage Essentials 9.4, agentless discovery for HP Data Protector is supported. You can now discover Data Protector on a host, that does not have any CIM extension installed.

   Note: The CIM extension only supports one backup solution on a host. If more than one backup applications are installed on the same host, only Data Protector is discovered by default and other applications are ignored by the CIM extensions. If Veritas NetBackup and EMC Networker are installed on the same host, only NetBackup is discovered by default. Networker is ignored by the CIM extension.

2. Discover the host that is the HP Data Protector, NetBackup ,EMC Networker or IBM Tivoli Storage Manager Master Server as described in Step 1 – Set Up Discovery for Hosts on page 181.

   Note: To discover IBM Tivoli Storage Manager, create an admin user on the IBM TSM providing the same user name and password used for host discovery.

3. If the server was previously discovered:

   a. Select **Discovery** > **Setup**.

   b. Delete the server .

   c. Select the Topology tab.

   d. Delete the server.

   e. Use the Test button to view the following information in View Logs:

      - Name of the backup application, such as NetBackup, Networker, DataProtector, and Tivoli Storage Manager.
      - Version of the backup application. Refer to the support matrix for your edition to determine if the version displayed is supported by HP Storage Essentials.

   The message “Backup Application Software not available.” will appear in View Logs if Backup application software is supported but not installed on the host or Backup Media server or the backup client is installed on the server.

4. You can configure the management server to obtain information about your Backup Manager hosts at a set interval.

Limitations with Discovering the Data Protector Server without a CIM Extension

You can discover the Data Protector server without a CIM extension; however, there are some limitations with this discovery method:

- Drive Utilization details are not shown in Drive Utilization tab in Backup Manager.
- Frequency and schedule window information is not populated for a session in the Schedule Detail tab.
- The MoM Server field is blank for a backup host where MoM is also configured along with Data Protector Cell Manager.
- Status, device and media pool details are not populated in the Policy Details tab for sessions.

Step 2 – Build the Topology

After you discover elements, the management server requires you build a topology view, which is a graphical representation of port-level connectivity information.

Note: The management server's user interface might load slowly while the topology is being recalculated. It might also take more time to log on to the management server during a topology recalculation.

To make the software aware of the devices on the network:

1. Click **Discovery** > **Topology**. The discovered elements are selected.
2. Click the **Get Topolog** button. The management server obtains the topology for selected elements.

   The Log Message page is displayed by the management server. After the management server builds the topology, a link appears to take you to System Manager so you can verify the topology view. You can also access System Manager by clicking **System Manager** in the left pane.
3. If you see errors in the topology, look at the log messages, which can provide an indication of what went wrong. Look at Event Manager for additional information. Access Event Manager by clicking the **Event Manager** button in the left pane. To obtain troubleshooting information, see the Troubleshooting Topology Issues on page 882.

   If the topology for an element in your network changes, select the element and click **Get Topology** in **Discovery** > **Topology** to update the information.

   The software obtains just enough information about where the element is connected in the topology; for example a switch connected to a host.

(*Optional*) Step 3 – View the Topology

Verify that the topology is displayed correctly by accessing System Manager.

To access System Manager:

1. Click the **System Manager** button in the left pane.
2. When asked if you want to trust the signed applet, click **Always**.

   The Always option prevents this message from being displayed every time you access System Manager, Capacity Manager, and Performance Manager.

The elements are shown connected to each other in the topology.

If you see a question mark above a host, the management server cannot obtain additional information about that element.

If a switch has more than one connection to an element, the number of connections is displayed above the line linking the switch and the element. For example, assume the number two is shown between a switch and a storage system. This means the elements have two connections to each other. To view the port details for the connection, right-click the element and select Show Port Details from the menu. If the topology changes, you can update how the element is viewed in the topology by selecting the element and clicking the Get Topology for Selected button in the Get Topology for discovered elements page (**Discovery** > **Topology**). The management server obtains just enough information about where the element is connected in the topology, for example a switch connected to a host.

The management server marks an element as “discovered” in the topology if the management server discovers an element but it cannot obtain addition information about it. To learn more about fixing discovered and/or disconnected elements, see Troubleshooting Topology Issues on page 882.

Step 4 – Get Details

After you obtain the topology of the network, you should obtain detailed information from the discovered elements. Get Details must be pe2rformed before you can do provisioning and/or obtain provisioning information, such as data about zone sets and LUN numbers. Clusters won’t be recognized until Get Details is completed. Get Details must be run on all of the participating nodes of application clusters.

Keep in mind the following:

- Unless you install CIM extensions and explicitly discover virtual machines using their own IP Address, they are not listed as access points on the Get Details page. Virtual machines can be viewed by looking at an ESX Server’s property page, or by clicking the Virtual Machines button on an ESX Server’s navigation page.
- Running Get Details takes time. You might want to perform this process when the network and the managed elements are not busy. To obtain a picture of device connectivity quickly, click **Get Topology** on the Topology tab.
- Reports show data from the last successful Get Details and report cache update. When a scheduled Get Details finishes, the report cache refresh does not automatically run after Get Details. The default interval for report cache refresh is six hours. For information about refreshing the report cache, see Refreshing the Report Cache on page 365.
- During Get Details the data you see in the user interface is not updated until the data collection is finished.
- During Get Details, the topology in System Manager is recalculated. While the topology is being recalculated, the loading of the user interface might be slow. It might also take more time to log on to the management server during a topology recalculation.

- You can use discovery groups to break up Get Details. For example, instead of running Get Details for all elements, you could specify only the elements in Discovery Group 1. For more information, see Using Discovery Groups on page 153.
- When an element in a discovery group is updated, its dependent elements are also updated.
- To monitor and manage backup servers, select **Include backup details**. If you also want to manage and monitor the host itself, select **Include infrastructure details**; otherwise, the host appears as a generic element in the topology in System Manager.
- If Get Details includes an AIX host, three SCSI errors (2 FSCSI error and 1 FCS error) per IBM adapter port are displayed in the system log. You can ignore these errors.
- You can quarantine elements to exclude them from Get Details. For example, to get information about all the elements in a discovery group except for one, you can quarantine that element. For more information, see Placing an Element in Quarantine on page 160.
- If a problem occurs with a host or SMI-S element during Get Details, the host or element is automatically quarantined. To remove the element from quarantine, see Removing an Element from Quarantine on page 160.
- To receive status reports about Get Details, see Configuring E-mail Notification for Get Details on page 863 for information about how to configure this option.

To obtain details:

1. Click **Discovery** > **Details** in the upper-right corner.
2. Verify that the **Include backup details** option is selected if you want to monitor and manage backup applications in Backup Manager.
3. Verify that the **Include infrastructure details** option is selected. This option is required to manage and monitor your elements not related to the backup infrastructure.
4. Click the **Get Details** button.

During Get Details, the status light changes from green to red. You can view the progress of gathering details by clicking **Discovery** > **View Logs**.

When the Get Details is finished, GETTING ALL DETAILS COMPLETED is displayed on the View Logs page.

Step 2 – Setting Up Discovery for Applications

Keep in mind the following when discovering applications:

- Make a list of the applications you want to monitor. Configure your applications first as described in this section and then run discovery.
- You should have already installed a CIM extension on the hosts that have the applications you want to discover. After you installed the CIM extension, you should have already discovered the host. See Step 1 – Discovering Your Hosts and Backup Manager Hosts on page 179.

You can configure the management server to monitor hosts and applications, such as Oracle, Microsoft Exchange server, Caché, and Sybase Adaptive Server Enterprise, in addition to Microsoft SQL servers and file servers. To obtain detailed information about the host and its applications, you must install a CIM extension on the host. See the "Deploying and Managing CIM Extensions" chapter of the *Installation Guide*.

The following is an overview of what you need to do. It is assumed you already discovered the hosts running your applications.

See Step 1 – Discovering Your Hosts and Backup Manager Hosts on page 179, and then set up the configurations for your applications on the management server. Some applications require you to provide additional discovery information about the application. Finally, perform discovery, map the elements in the topology, and then run Get Details. Get Details takes some time. Perform this step when the network is not busy.

See the following topics for more information:

- Creating Custom User Names and Passwords on Managed Database Instances below
- Monitoring Oracle on the facing page
- Monitoring Microsoft SQL Server on page 218
- Monitoring Sybase Adaptive Server Enterprise on page 227
- Monitoring Microsoft Exchange on page 231
- Monitoring Caché on page 233
- Monitoring IBM DB2 on page 239
- Monitoring IBM Informix on page 243
- Application Discovery Test on page 246

Creating Custom User Names and Passwords on Managed Database Instances

If user credentials managing more than one database instance are changed, make sure that the other database instances using those credentials are updated properly.

Keep in mind the following:

- Depending on the password policy, SQL Server 2005 might require that passwords be alphanumeric. For this reason, a managed SQL Server 2005 database instance might not accept the default managed database password (password) during user credential creation. A script is provided to input an alphanumeric password for SQL Server 2005. For all other applications, this script is optional.
- Do not use the SYS user or users having SYSDBA/SYSOPER privileges for discovering Oracle applications from HP Storage Essentials

The user credentials script names for each database type are as follows:

Database Type	Script Name
Oracle	CreateOracleActWithCustomPwd.sh (or .bat) or CUSTACCT.COM (for OpenVMS)
SQL Server	CreateSQLServerActCustomPwd.bat
Sybase	CreateSybaseActCustomPwd.bat
Caché 5.0.20	createCacheDB50UserCustomPwd.sh (or .bat)
Caché 5.2 and 2007.1	createCacheDBUserCustomPwd.sh (or .bat) or CUSTUSER.COM (for OpenVMS)

After changing the user credentials on a managed database instance, the user credentials must be changed on the HP Storage Essentials management server.

The following steps do not apply to DB2 and Informix databases.

To change the user credentials on the HP Storage Essentials management server:

1. Select **Discovery** > **Setup**.
2. Click the **Applications** tab.
3. In the Database User Credentials section, click **New**.
4. Enter the user name that was used for creating the account on the managed database instance.
5. Enter the password that was used for creating the account on the managed database instances.
6. Enter a description of the managed database instance.
7. Select the database type from the drop-down menu.
8. **SQL Server only**: Select the Authentication mode from the drop-down menu. If you select Windows Authentication, enter the domain controller.
9. Click **OK**.

The Manages column of the User Credentials table is not populated until the user credentials are assigned to an application instance.

Monitoring Oracle

For instructions on monitoring and managing Oracle, see the following:

1. Optional – Enable Autoscan on next page
2. Step A – Create the APPIQ_USER Account for Oracle on page 209
3. Step B – Provide the TNS Listener Port on page 212
4. Step C – Set Up Discovery for Oracle on page 212

After you complete these steps, you must discover Oracle and perform Get Details. See Step 3 – Discovering Applications on page 247.

Before you begin, make sure you purchased the module that lets you monitor Oracle. Contact customer support if you are unsure if you purchased this module.

Optional – Enable Autoscan

Autoscan allows Oracle instances to be discovered automatically without your having to enter the application setup information. By default, discovery of Oracle through autoscan is disabled.

To enable autoscan:

1. Select **Configuration > Product Health > Advanced.**
2. Add the following line to the Custom Properties section:

   ```
   oracleautoscan=true
   ```

3. Click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

Auto scans are supported for both Oracle standalone instances and RACs. However, Oracle instances configured as failover cluster resources should always be discovered by explicitly specifying the instance configuration as described in Discovering Single Instance Oracle Failover Clusters on page 216.

Autoscan for Oracle is supported on HP-UX, AIX, Solaris, and Linux platforms. Autoscan support for Oracle 11gR1 on these platforms requires the latest CIM extension to be installed on that managed host. Autoscan for Oracle is not supported for applications running on Solaris Containers. Auto scans for Oracle11gR2 are supported only for standalone instances. Discovering an Oracle11gR2 RAC using autoscan is not supported.

To discover Oracle on other platforms, enter the application information as described in Step C – Set Up Discovery for Oracle on page 212.

If you are discovering an Oracle 11g instance using autoscan, the LISTENER.ORA file must exist. It should be located in one of the following directories:

- `<Oracle_Home>/network/admin`
- `/etc`
- `/var/opt/oracle`

If LISTENER.ORA is not located in those directories, use the TNS_LOC parameter in the `cim.extension.parameters` file to specify where the file is stored. Restart the CIM extension for your changes to take effect.

If there are two LISTENER.ORA files specified in the TNS_LOC parameter, only those Oracle instances that are being serviced by listeners configured in any one of the LISTENER.ORA files are discovered by autoscan. To discover the other Oracle instances, enter the application information as described in Step C – Set Up Discovery for Oracle on page 212.

If a listener is configured with a non-default alias (a listener name other than LISTENER) in the LISTENER.ORA file, the listener must be started by entering the command `lsnrctl start <listenername>`. This allows the Oracle10g instances that are serviced by this listener to be discovered using autoscan.

Step A – Create the APPIQ_USER Account for Oracle

The management server accesses Oracle through the APPIQ_USER account. This account is created when you run the CreateOracleAct.bat script (on Microsoft Windows) or CreateOracleAct.sh (on UNIX platforms) or CRACCT.COM (on OpenVMS) on the computer running the Oracle database you want to monitor. The account has create session and select dictionary privileges to be used with the management server.

To create a user account with a custom user name or password, run CreateOracleActWithCustomPwd.bat (on Microsoft Windows) or CreateOracleActWithCustomPwd.sh (on UNIX platforms) or CUSTACCT.COM (on OpenVMS). For more information, see Creating Custom User Names and Passwords on Managed Database Instances on page 206.

Keep in mind the following:

- The `CreateOracleAct.bat` script must run under SYS user.
- Create the APPIQ_USER account on the Oracle Database you want to monitor, not on the management server.
- You should have already installed the database for the management server.
- Verify that the instance TNS (Transparent Name Substrate) listener is running so that the management server can find the Oracle installation and its instances. For example, on Microsoft Windows 2000, you can determine if the instance TNS listener is running by looking in the Services window for OracleOraHome10TNSListener for Oracle 10g and OracleOraHome11gR2TNSListener for Oracle 11g. The name of the TNS listener might vary according to your version of Oracle. See the Oracle documentation for information about verifying if the instance TNS listener is running. You can also verify the listener is running by entering the following at the command prompt:

  ```
  snrctl status
  ```

  If the listener is not running, you can start it by typing `lsnrctl start` on the command line.
- When creating the APPIQ_USER account on an Oracle Real Application Cluster (RAC) Database, this script should be run only once, on any one of the instances of the Oracle RAC Database. Since all the instances of an Oracle RAC access the same Database, it is sufficient to create the APPIQ_USER account on any one of the instances. However, for Oracle11gR2 RAC Database, you must run this script on the Oracle RAC database.
- To exclude instances from being autoscanned, do not create the APPIQ_USER account on those instances.
- Make sure you have all the necessary information and read through the following steps before you begin.

To create the Oracle user for the management server:

1. Log on as follows:

   IBM AIX, SGI IRIX, HP-UX, Linux or Sun Solaris:

   a. Log on to an account that has administrative privileges.

   b. Mount the StorageEssentials DVD (if not auto-mounted).

   c. Go to the `/CimExtensionCD1/DBIQ/oracle/unix` directory by typing the following:

      ```
      # cd /DVD/DBIQ/oracle/unix
      ```

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

   Microsoft Windows:

   Go to the `DBIQ\oracle\win` directory on the CIM extensions DVD.

   OpenVMS:

   a. Log on to an account that has administrative privileges.

   b. Mount the StorageEssentials DVD (if not auto-mounted) using the following command:

      ```
      $ MOUNT /MEDIA=CDROM
       /UNDEFINED_FAT=STREAM:32767/OVERRIDE=IDENTIFICATION
      DQB0
      ```

      In this instance, DQB0 is the CDROM drive.

   c. Go to the directory containing the Oracle agent creation script using the following command:

      ```
      $ SET DEF DQB0:[OVMS.DBIQ.ORACLE]
      ```

2. Make sure you have the password to the SYS user account because you cannot run the script without it.

3. Run `CreateOracleAct.sh` (on UNIX), or `CreateOracleAct.bat` (on Microsoft Windows), or `CRACCT.COM` (on OpenVMS) on the computer with the Oracle database. On OpenVMS, run CRACCT.COM on the host using the following command.

   ```
   $ @CRACCT.COM
   ```

   The script creates a user with create session and select dictionary privilege on a managed Oracle instance. You can use a remote Oracle client to run the script.

4. Specify the Oracle instance name, which must be visible to the client, as the first input when running the script. The script prompts you for the name of the Oracle instance on which to create the user for Oracle management packages and the password of the SYS account.

   You must be able to specify the default and temporary tablespaces for APPIQ_USER during the installation. You can enter `users` as default and `temp` as temporary if these tablespaces exist in the Oracle Instance.

5. Repeat the previous step for each Oracle instance you want to manage.

This script does the following in order:

- Creates the APPIQ_USER account.
- Grants create session and select on dictionary tables privileges to APPIQ_USER, enabling the management server to view statistics for the Oracle instances.

Removing the APPIQ_USER Account for Oracle

If you no longer want the management server to monitor an Oracle instance, you can remove the APPIQ_USER account for that Oracle instance by running the `UninstallOracleAct.bat` script (on Windows) or `UninstallOracleAct.sh` script (on UNIX) or RMACCT.COM (on OpenVMS).

Keep in mind the following:

- Before you remove the APPIQ_USER account for an Oracle instance, make sure no processes are running APPIQ_USER for that Oracle instance. The management server uses APPIQ_USER to obtain information about the Oracle database. For example, a process would be using APPIQ_USER if someone was using Performance Manager to view monitoring statistics about that Oracle instance. One of the ways to make sure APPIQ_USER is not being used is to temporarily remove the host running Oracle (**Discovery** > **Topology**). After you removed the APPIQ_USER account for Oracle, discover and perform Get Details for the host if you want to continue monitoring it.
- If you receive a message about not being able to drop a user that is currently connected while you are removing the APPIQ_USER account for Oracle, re-run the script for removing APPIQ_USER.
- When removing the APPIQ_USER account from an Oracle RAC Database, this script should be run only once, on any one of the instances of the Oracle RAC Database. Since all the instances of an Oracle RAC access the same Database, it is sufficient to remove the APPIQ_USER account from any one of the instances.

To remove the APPIQ_USER account:

1. Log on as follows:

   UNIX:

   a. Log on to an account that has administrative privileges.

   b. Mount the StorageEssentialsDVD (if not auto-mounted).

   c. Go to the /CimExtensionsCD1/DBIQ/oracle/unix directory by typing the following:

      ```
      # cd /DVD/CimExtensionsCD1/DBIQ/oracle/unix
      ```

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

   Windows:

   Go to the `CimExtensionsCD1\DBIQ\oracle\win` directory on the StorageEssentials DVD.

   OpenVMS:

a. Mount the StorageEssentialsDVD (if not auto-mounted) using the following command:

```
$ MOUNT /MEDIA=CDROM
UNDEFINED_FAT=STREAM:32767/OVERRIDE=IDENTIFICATION
DQB0
```

In this instance, DQB0 is the CDROM drive.

b. Go to the directory containing the Oracle agent creation script using the following command:

```
$ SET DEF DQB0:[CimExtensionsCD2.OVMS.DBIQ.ORACLE]
```

2. Verify that you have the password to the SYS user account.

   At the prompt, provide the password for this user account.

3. Run `UninstallOracleAct.bat` (on Windows) or `UninstallOracleAct.sh` or `RMACCT.COM` (on OpenVMS).

   The script removes the management software for the specified Oracle instance.

   You can use a remote Oracle client to run this script.

4. When asked for the Oracle instance name, enter the name of the Oracle instance you do not want the management server to monitor. The name must be visible to the client.

5. Provide the password for the SYS user account.

The APPIQ_USER account for the specified Oracle instance is removed. The management server can no longer monitor that Oracle instance.

Step B – Provide the TNS Listener Port

This step is required for discovering Oracle instances using autoscan.

If your Oracle instances use a different TNS Listener Port than 1521, follow these steps to change the port:

1. Select **Discovery** > **Setup**, and then click the **Applications** tab.

   The TNS Listener Port setting applies to all Oracle instances you monitor.

2. To assign a new port, click the **Create** button for the Oracle Information table.

3. Enter the new port number and click **OK**.

4. If necessary, click the  button to remove the old port number.

Monitoring Oracle clusters requires an additional step. If you are monitoring Oracle, see Step C – Set Up Discovery for Oracle. If you are discovering an Oracle cluster, see Discovering Single Instance Oracle Failover Clusters on page 216.

Step C – Set Up Discovery for Oracle

Keep in mind the following:

- If you are discovering an Oracle cluster, see Discovering Single Instance Oracle Failover Clusters on page 216.
- On Linux and Microsoft Windows operating systems, discovery of Oracle databases that are using Oracle Automatic Storage Management (ASM) requires the latest CIM extension to be installed on that managed host.

To discover Oracle instances without using autoscan:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. Click the **Create** button for the Database Information table.
3. In the **Host IP/DNS Name** box, enter the IP address or DNS name of the host running Oracle.
4. In the **Management IP/DNS Name** box, enter the IP address the listener is listening on for the Oracle instance. The IP address can be a virtual IP or a host IP. You can find the IP address in the `listener.ora` file for the monitored database. You can find the file in the following directory on the host of the monitored database. Do not look for it on the management server.

   Windows: `%ORA_HOME%\network\admin\listener.ora`

   UNIX: `$ORACLE_HOME/network/admin/listener.ora`
5. In the **Server Name** box, enter the Oracle System Identifier (SID) of the Oracle database you want to monitor.
6. In the **Port Number** box, enter the monitored port.

   If you are not sure of the monitored port, check the `listener.ora` file of the monitored database application. You can find the file in the following directory on the host of the monitored database. Do not look for it on the management server.

   `%ORA_HOME%\network\admin\listener.ora`

   The port can be found in the following code:

```
LISTENER =
 (DESCRIPTION_LIST =
 (DESCRIPTION =
 (ADDRESS_LIST =
 (ADDRESS = (PROTOCOL = TCP)(HOST = localhost)(PORT = 1521))
 (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC0))
 )
 )
 )
```

7. Select **ORACLE** from the Database Type menu.
8. Click **OK**.

Discovering Oracle Real Application Clusters (RAC)

Since Oracle RAC is an active-active application cluster, one RAC instance can provide information for the whole RAC. Regardless of the instance through which the database is accessed, the same sets of tables are accessed. This includes the data dictionary tables that are used to understand the logical and physical storage organization of the Oracle RAC application.

Discovery of Oracle RAC Instances Using One Instance

Because one RAC instance can provide information for the whole RAC, it is possible to identify and discover all the instances in the Oracle RAC cluster from any one of its instances. This means that the you can enter the application setup information for one instance of the Oracle RAC, and the management server will automatically discover the other instances, subject to certain conditions. The conditions to be satisfied for discovering all the instances of Oracle RAC using application setup information from one of its instances are the following:

- Only the Oracle RAC instances running on hosts already discovered and identified as part of the same cluster will be discovered as part of the Oracle RAC on the management server.
- The management server is able to contact the hosts running Oracle RAC instances using the short host name. The management server can be configured to access the hosts running Oracle RAC instances using the short name in the following ways:
  - On the management server, add entries for each host running an Oracle RAC instance in /etc/hosts (on UNIX platforms) or %WINDIR%\system32\drivers\etc\hosts (on Windows).
  - Add the domain of the host in the domain search list of the management server under the search option of /etc/resolv.conf (on UNIX platforms) or Append these DNS suffixes (in order) on the **Advanced TCP/IP Settings** > **DNS** tab (on Windows).
- The listener is configured on the same IP address that is used to discover the host. For example, on the Application Setup page, the management IP address for the application should be the same as the host IP address.
- Typically, all the instances of Oracle RAC will be listening on the same TNS port number. If this is not the case, the port numbers for the other instances should be specified in the default port list in the Application Setup page. For example, if SID1 is listening on TNS port LP1, and SID2 is listening on TNS port LP2, then it is possible to automatically discover SID2, provided that TNS port LP2 is part of the default port list in the Application Setup page.

To discover Oracle RAC:

1. Install the CIM extension on each node in the cluster.
2. If the cluster is not automatically discovered by the management server, create the cluster using Cluster Manager. For more information about Cluster Manager, see Host and Application Clustering on page 267.

3. Create the APPIQ_USER account on any one node in the cluster. See Step A – Create the APPIQ_USER Account for Oracle on page 209.
4. Click **Discovery** > **Setup** and discover the host for the first node by clicking the **Add Address** button and providing the appropriate information for discovering the host, as described in Adding an IP Range for Scanning on page 77.
5. Discover the first Oracle node as follows:
   a. Select **Discovery** > **Setup**, and then click the **Applications** tab.
   b. Click the **New** button in the Managed Databases section.
   c. In the **Host IP/DNS Name** box, enter the IP address or DNS name of the host running Oracle.

      In the **Management IP/DNS Name** box, enter the IP address the listener is listening on for the Oracle instance. The IP address can be a virtual IP or a host IP. You can find the IP address in the `listener.ora` file for the monitored database. Do not look for the file on the management server for this information. The file is located in the following directory on the host of the monitored database:

      `%ORA_HOME%\network\admin\listener.ora` (on Windows)

      `$ORACLE_HOME/network/admin/listener.ora` (on UNIX platforms)
   d. In the **Database Instance Name** box, enter the Oracle System Identifier (SID) of the Oracle database you want to monitor.
   e. In the **Port Number** box, enter the monitored port.

      If you are not sure of the monitored port, check the `listener.ora` file of the monitored database application. You can find the listener.ora file in the following directory on the host of the monitored database. Do not look for the listener.ora file on the management server for this information.

      Microsoft Windows:

      `%ORA_HOME%\network\admin\listener.ora`

      UNIX Platforms:

      `$ORACLE_HOME/network/admin/listener.ora`

      The port can be found in the following code:

```
LISTENER =
 (DESCRIPTION_LIST =
 (DESCRIPTION =
 (ADDRESS_LIST =
 (ADDRESS = (PROTOCOL = TCP)(HOST = localhost)(PORT = 1521))
 (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC0))
```

)

)

)

f. Select **ORACLE** from the Database Type menu.

g. If you created a custom user name as described in Creating Custom User Names and Passwords on Managed Database Instances on page 206, select the user name from the drop-down menu. If you used the custom password script to change the user name for the database instance, but you did not already add the custom user name to the management server, you can add it now by clicking **New User**.

h. Click **OK**.

Note: If all these conditions are satisfied, all the other instances in the Oracle RAC will be discovered, and the Oracle RAC application cluster will be constructed by the management server. If the other instances of the Oracle RAC are not discovered, repeat steps 4 and 5 for each node in the cluster.

About Discovery of an Oracle RAC Application Cluster on a Host Cluster Discovered Using Cluster Manager

When the underlying host cluster is not discovered, the management server will be "Oracle RAC safe," but not fully "Oracle RAC aware." Each instance will show up as a standalone Oracle application, and data will be collected for each instance separately (even though both instances will return identical capacity data). However, the management server does not explicitly identify and construct the Oracle RAC application cluster. Also, when the underlying host cluster is not discovered, other instances of the Oracle RAC cannot be discovered automatically as described in the Discovery of Oracle RAC Instances Using One Instance section.

However, if you create the host cluster at a later point in time, subsequent discovery of any instance in Oracle RAC will identify and construct the Oracle RAC application cluster. The management server will shift to "Oracle RAC aware" mode on top of the host cluster that you created.

Discovering Single Instance Oracle Failover Clusters

It is possible to operate a non-RAC Oracle instance as a clustered active/passive application. In this case, the single Oracle instance is configured as a cluster resource. The clustering software (such as VCS or Service Guard) is then responsible for monitoring the Oracle instance and failing it over to other operating nodes during a node failure.

In the case of a single instance failover cluster, the Oracle instance by itself will not be able to indicate that it is operating in clustered mode.

The conditions to be satisfied for discovering single instance Oracle failover clusters are as follows:

- All the hosts in the cluster configured to handle single instance Oracle failover should be discovered in the management server.
- The management server must be able to contact the hosts running the single instance Oracle failover instance using the short host name. The management server can be configured to access the hosts running a single instance Oracle failover instance using the short name in the following ways:
  - On the management server, add entries for each host configured for single instance Oracle failover instance in `/etc/hosts` (on UNIX) or `%WINDIR%\system32\drivers\etc\` hosts (on Windows).
  - Add the domain of the host in the domain search list of the management server under the search option of `/etc/resolv.conf` (on UNIX) or Append these DNS suffixes (in order) on the **Advanced TCP/IP Settings** > **DNS** tab (on Windows).

To discover a single instance Oracle failover application:

1. Install the CIM extension on each node in the cluster.
2. Create the APPIQ_USER account for the Oracle application from that node in the cluster in which it is currently running. See Step A – Create the APPIQ_USER Account for Oracle on page 209.
3. Click **Discovery** > **Setup** and discover the host for the first node by clicking the **Add Address** button and providing the appropriate information for discovering the host, as described in Adding an IP Range for Scanning on page 77.
   a. Discover the first Oracle node by selecting **Discovery** > **Setup**, and then clicking the Applications tab.
   b. Click the **Create** button for the Database Information table.
   c. In the Host IP/DNS Name box, enter the IP address of any one of the hosts in the cluster configured to handle the single instance Oracle failover in the application setup information. Be sure that the host with this IP address will be discovered in the management server.
   d. Enter the management IP for the single instance fail over Oracle application. Please note that the management IP configured for the single instance Oracle fail over cluster is dependent on underlying host cluster software.
   e. In the Server Name box, enter the Oracle System Identifier (SID) of the Oracle database you want to monitor.
   f. In the Port Number box, enter the monitored port. If you are not sure of the monitored port, check the listener.ora file of the monitored database application. You can find the listener.ora file in the following directory on the host of the monitored database. Do not look for the listener.ora file on the management server for this information.

      Microsoft Windows:

      `%ORA_HOME%\network\admin\listener.ora`

UNIX Platforms:

```
$ORACLE_HOME/network/admin/listener.ora
```

The port can be found in the following code:

```
LISTENER =
 (DESCRIPTION_LIST =
 (DESCRIPTION =
 (ADDRESS_LIST =
 (ADDRESS = (PROTOCOL = TCP)(HOST = localhost)(PORT = 1521))
 (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC0))
 )
 )
 )
```

g. Select **ORACLE** from the Database Type menu.

h. Select the check box **Discover as failover cluster** for discovering the Oracle failover cluster.

i. Click **OK**.

Deleting Oracle Application Information

If you do not want the management server to monitor an Oracle instance, follow these steps to remove its information:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. In the Managed Databases table, click the checkbox for the Oracle Application instances you do not want the management server to monitor.
3. Click **Delete**.
4. Perform Get Details to make the management server aware of your changes.

If Oracle Autoscan is enabled, the above step is not applicable.

Monitoring Microsoft SQL Server

If you plan to monitor SQL Server clusters, see Monitoring SQL Server Clusters on page 223.

Managing and monitoring SQL Servers requires the following tasks.

Step A – Create the User Account for the SQL Server

SQL Server 2000:

The management server accesses SQL Server through the appiq_user account. This account is created when you run the `CreateSQLServerAct.bat` or `CreateSQLServerActCustom.bat` script on the computer running the SQL Server database you want to monitor. This account has create session and select dictionary privileges, which allow the management server to view statistics for the SQL Server.

For more information about creating a custom user account or adding Windows authenticated users, see Custom User Accounts and Windows Authentication on page 226.

Keep in mind the following:

- Obtain the SQL Server name before you run the script.
- The database for the management server must already be installed.
- Make sure you have all the necessary information and read through the following steps before you begin.

To create the appiq_user account for SQL Server:

1. The script must run under the SA user account. To verify that the SA account is enabled, launch SQL Server's Query Analyzer tool and attempt to connect to the database as SA with the SA user's password.
2. To run the script on Microsoft Windows, go to the `DBIQ\sqlserver\win` directory on the CIM Extensions DVD.
3. Verify that you have the password to the SA user account. You cannot run the script without the password.
4. In a new command window, run the `CreateSQLServerAct.bat` script on the computer with the SQL Server database. You can use a remote SQL Server isql to run this script.
5. The script prompts you for the name of the SQL Server on which to create the appiq_user account. If you are creating the account on a default instance, enter the host name if the instance is non-clustered and the SQLNetwork Name if the instance is clustered. If you are creating the account on a named instance, enter the host name and the instance name as follows:

   For a non-clustered instance:

   `<Host Name>\<Instance Name>`

   For a clustered instance:

   `<SQL Network Name>\<Instance Name>`
6. If you are running the CreateSQLServerActCustom.bat script, you must provide a user name and password for the user account. The password must meet the password policy criteria described in Creating Custom User Names and Passwords on Managed Database Instances on page 206. If you are running the CreateSQLServerAct.bat script, the default password (password) is automatically used.

To create Windows authenticated users to manage a specific SQL Server, see Custom User Accounts and Windows Authentication on page 226.

7. The script prompts you for the SA user password. Enter the password. The appiq_user account is created.

To determine if the appiq_user account was added correctly to your SQL Server:

1. Open SQL Server Enterprise Manager.
2. Expand the user interface for SQL Server Enterprise Manager, then expand the specific SQL Server and select **Security**.
3. Double-click **Logins** and view the list of users authorized to access the SQL Server.
4. Click the refresh button in SQL Server Enterprise Manager. If the appiq_user is not listed, the management server is not able to discover the database.

To determine if the SQL Server is ready to accept connections from the management server:

1. Connect to the SQL Server installation through Query Analyzer using the account appiq_user and the password password.
2. Create a sample ODBC datasource for the SQL Server installation using the appiq_user account.
3. Click the **Test** button to test the datasource.
4. Repeat these steps for each SQL Server 2000 instance you want to manage.

SQL Server 2005 or 2008

The management server accesses SQL Server through the appiq_user account. To create this account, run the `CreateSQLServerActCustomPwd.bat` script on the computer running the SQL Server database you want to monitor. This account has create session and select dictionary privileges, which allow the management server to view statistics for the SQL Server.

To monitor SQL Server 2008, use the appiq_user creation scripts from HP Storage Essentials 6.1 or later.

For more information about using the CreateSQLServerActCustomPwd.bat script, see Custom User Accounts and Windows Authentication on page 226.

To access the Microsoft SQL Server performance metrics as a database user, you must have read permissions to the master.dbo.sysperfinfo table. To gain these permissions, you must recreate the SQL Server database user by running the `CreateSQLServerActCustomPwd.bat` or `CreateSQLServerAct.bat` script.

Step B – Provide the SQL Server Configuration Details

You must provide the server name for the SQL Server and port number for managing a SQL database.

If you have name resolution issues, your server might be discovered but your applications will not be discovered. To avoid this, add entries within the hosts file on the management server for the systems in question.

If SQL Server is discovered using Dynamic Port and the port is changed, you must update the port number in the Port Number box.

When configuring the System Application Discovery Settings for SQL servers, you must specify the following:

- **Host IP/DNS Name**: <IP Address>
- **Database Instance Name**: <SQL Server Name>
- **Port Number**: <SQL Port #>
- **Database Type**: SQLSERVER
- **User Name**: <User Name>

  (available only for the SQLSERVER database type)
- **Service Principal Name**: <SPN>

  (available only when the selected user is configured to use Windows Authentication)

To add information for discovering a SQL server:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. Click **New** in the Managed Databases section.
3. In the **Host IP/DNS Name** box, enter the IP address or DNS name of the host running SQL Server. You must provide the host name. You cannot use localhost or parenthesis.
4. You can leave the Management IP/DNS Name box blank. This box is for Oracle clusters. When this box is blank, the management server automatically lists the DNS name or IP address of the host under the **Host IP/DNS Name** column and the **Management IP/DNS Name** column.
5. In the **Database Instance Name** box, enter the SQL database server name you want to monitor. The SQL Server name is either the Windows system name (default) or the name specified when the SQL server was installed. It is one of the following:
   - The name specified at the time the SQL server was installed
   - The Windows system name (Windows 2000)
   - The local name (Windows 2003)

     For example, if a Windows 2003 server called SQLTEST has an IP address of 192.168.2.10 with the default SQL port (1433) and shows the name of (local) within SQL Enterprise Manager/SQL Server Management Studio, the correct system application discovery settings on the management server are the following:

- **Host IP/DNS Name**: 192.168.2.10
- **Database Instance Name**: SQLTEST
- **Port Number**: 1433
- **Database Type**: SQLSERVER
- **User Name**: mydomain\testuser (Windows Authenticated user)
- **Service Principal Name**: MSSQLSvc/sqltest.mydomain.com:1433 (SPN registered in the Active Directory)

6. In the **Port Number** box, enter the port used by SQL.

   To determine the correct SQL port number:

   SQL Server 2000:

   a. Open SQL Server Enterprise Manager.

   b. Expand the user interface for SQL Server Enterprise Manager, and select the specific SQL server. Right-click and select **Properties** from the menu.

   c. Click the **Network Configurations** button. On the General Tab, select the TCP/IP entry under the Enabled Protocols section, and click the **Properties** button.

   d. The resulting window shows you the TCP/IP port your SQL server uses. Provide this port number in the **Port Number** box on the management server.

   SQL Server 2005 or 2008:

   a. Open SQL Server Configuration Manager.

   b. Select the specific SQL Server 2005 or 2008 Network Configuration entry for the SQL Server 2005 or 2008 instance.

   c. Select the TCP/IP entry on the right pane, and click the Properties right click menu.

   d. From the IP Addresses tab, obtain the Port Number configured for the instance. Provide this port number in the Port Number box on the management server.

7. Select **SQLSERVER** from the Database Type menu.
8. Select a user name from the drop-down menu, or click **Create New User** to create a new user. If the authentication type of the selected user is Windows Authentication, enter the Service Principal Name. Click **Populate SPN** to get a suggested value for the Service Principal Name. The suggested value might not be the actual value registered in the Active Directory/Kerberos database.
9. Click **OK**.

Perform Get Details for your inputs to take effect. See Step 3 – Discovering Applications on page 247.

Removing the appiq_user Account for SQL Server

Before removing the appiq_user account for the SQL Server databases on a host, make sure no processes are running appiq_user for that SQL Server database. The management server uses appiq_user to obtain information about a SQL Server database. One way to make sure appiq_user is not being used is to temporarily remove the host running SQL Server (**Discovery** > **Topology**). After you remove the appiq_user account for SQL Server, discover and perform Get Details for the host if you want to continue monitoring it.

To remove the appiq_user account from the SQL Server databases on a host:

1. To run the script on Microsoft Windows, go to the `DBIQ\sqlserver\win` directory on the CIM Extensions DVD.
2. Verify that you have the password to the server administrator user account. You cannot run the script without the password.
3. Run the `DropSQLServerAct.bat` script on Microsoft Windows on the computer with the SQL Server database.
4. Enter the name of the SQL Server server.
5. Enter the password for the server administrator account.

The account for appiq_user is removed. The management server can no longer monitor the SQL Server databases on this host.

Deleting SQL Server Information

If you do not want the management server to monitor a SQL Server instance, follow these steps to remove its information:

1. Select **Discovery** > **Setup**, and click the **Applications** tab.
2. In the Managed Databases table, click the checkbox for the SQL Server instances you do not want the management server to monitor.
3. Click **Delete**.
4. Perform Get Details to make the management server aware of your changes.

Monitoring SQL Server Clusters

To monitor and manage SQL Server clusters:

1. Install CIM Extensions on each of the participating nodes.
2. Create the appiq_user account as described in Step A – Create the User Account for the SQL Server on page 218.

   This step must be run on any one of the participating host nodes of the SQL Server cluster.
3. Enter the server name and port number as described in Provide the SQL Server Name and Port Number for a Cluster on next page.

Provide the SQL Server Name and Port Number for a Cluster

The server name for the SQL Server and port number for managing a SQL Server cluster database must be provided in the following steps.

If you have name resolutions issues, your server might be discovered but your applications will not be discovered. You can address the name resolution issues by adding entries within the hosts file on the management server for the systems in question.

When configuring the System Application Discovery Settings for SQL servers, the following must be specified:

- **Host IP/DNS Name**: <IP Address>
- **Database Instance Name**: <SQL Server Name>
- **Port Number**: <SQL Port #>
- **Database Type**: SQLSERVER
- **User Name**: <User Name>
- **Service Principal Name**: <SPN>

  (available only when the selected user is configured to use Windows Authentication)

To add information for discovering a SQL Server cluster:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. Click **New** in the Managed Databases section.
3. In the **Host IP/DNS Name** box, enter the IP address or DNS name of at least one of the participating host nodes running SQL Server cluster. You must provide the host name. You cannot use localhost or parenthesis.
4. You can leave the Management IP/DNS Name box blank. When it is blank, the management server automatically lists the DNS name or IP address of the host under the Host IP/DNS Name column and Management IP/DNS Name column.
5. In the **Database Instance Name** box, enter the SQL database server name you want to monitor.

   The SQL Server name is one of the following:

   - The name specified at the time the SQL server was installed
   - The Microsoft SQL Network Name (the default instance)

   For example, if a SQL Server cluster instance called SQLCLUSTER is running on a 2-node Windows 2003 cluster (individual host node IP address being 192.168.2.10 and 192.168.2.11) at the default SQL port (1433) and shows the name Microsoft SQL Network Name within SQL Enterprise Manager / SQL Server Management Studio, the correct system application discovery settings on the management server is either of the following:

- **Host IP/DNS Name**: 192.168.2.10
- **Database Instance Name**: SQLCLUSTER
- **Port Number**: 1433
- **Database Type**: SQLSERVER
- **User Name**: mydomain\testuser (Windows Authenticated user)
- **Service Principal Name**: MSSQLSvc/sqlcluster.mydomain.com:1433 (SPN registered in the Active Directory)

  Or

- **Host IP/DNS Name**: 192.168.2.11
- **Database Instance Name**: SQLCLUSTER
- **Port Number**: 1433
- **Database Type**: SQLSERVER
- **User Name**: mydomain\testuser (Windows Authenticated user)
- **Service Principal Name**: MSSQLSvc/sqlcluster.mydomain.com:1433 (SPN registered in the Active Directory)

6. In the **Port Number** box, enter the port used by SQL.

   To determine the correct SQL Port Number:

   SQL Server 2000 Cluster

   a. Open SQL Server Enterprise Manager.

   b. Expand the user interface for SQL Server Enterprise Manager and select the specific SQL server. Right-click and select **Properties** from the menu.

   c. Click the **Network Configurations** button. On the General tab, select the TCP/IP entry under the Enabled Protocols section and click the **Properties** button.

   d. The resulting window shows the TCP/IP port your SQL server uses. Provide this port number in the **Port Number** box on the management server.

   SQL Server 2005 or 2008 Cluster

   a. Open SQL Server Configuration Manager.

   b. Select the specific SQL Server 2005 or 2008 Network Configuration entry for the SQL Server 2005 or 2008 instance.

   c. Select the TCP/IP entry on the right pane, and click the Properties right- click menu.

   d. From the IP Addresses tab, obtain the Port Number configured for the instance. Provide this port number in the Port Number box on the management server. If Dynamic Ports are used, the Port Number is located under **IPAll** > **TCP Dynamic Ports**.

7. Select **SQLSERVER** from the Database Type menu.
8. Select a user name from the drop-down menu, or click **Create New User** to create a new user. If the authentication type of the selected user is Windows Authentication, enter the Service Principal Name. Click **Populate SPN** to get a suggested value for the Service Principal Name. The suggested value might not be the actual value registered in the Active Directory/Kerberos database.
9. Click **OK**.

Perform Get Details for your inputs to take effect. See Step 3 – Discovering Applications on page 247.

Custom User Accounts and Windows Authentication

To create a custom user account or to add a Windows authenticated user for managing SQLServer, use the `CreateSQLServerActCustomPwd.bat` file. An account added using this script has create session and select dictionary privileges, which allow the management server to view statistics for the SQL Server.

Keep in mind the following:

- To add Windows authenticated users, the script must run under a Windows user account that has permission to create new users. Log on as that Windows user to the remote machine running SQL Server and run the `CreateSQLServerActCustomPwd.bat` script.
- Obtain the SQL Server name before you run the script.
- Make sure that the Windows user account to be added is available in the Active Directory and is enabled.
- Make sure that the SQL Server is registered in the Active Directory and Kerberos tickets can be issued for that SQL Server.

  Only Kerberos-based authentication is supported. NTLM is not supported for SQL Server management.
- You must have the Service Principal Name of the SQL Server.
- The database for the management server must already be installed.

To create a custom SQL user account or to add a Windows user:

1. The script prompts you for the name of the SQL Server on which to add the Windows user account. If you are adding the account on a default instance, enter the host name if the instance is non-clustered and the SQL Network Name of the instance is clustered. If you are adding the account on a named instance, enter the host name and the instance name as follows:

   For a non-clustered instance:

   `<Host Name>\<Instance Name>`

   For a clustered instance:

```
<SQL Network Name>\<Instance Name>
```

2. The script prompts you for the authentication mode to be used for the user account that is being added. To add a Windows user, enter WINDOWS as the authentication mode. To create a custom SQL account, enter MIXED as the authentication mode.
3. When the authentication mode is Windows, the script prompts you for the name of the Windows user account to be added. You must enter the username in the format DomainName\UserName. When MIXED mode is entered, the script prompts you for the SQL user name to be created and a password for that user.
4. When the WINDOWS mode is entered, the script uses the currently logged-in user account to connect to SQL Server and add the Windows user account. The Windows user account is added.

   When MIXED mode authentication is entered, the script prompts you for the SA user password to connect to SQL Server and create the new user. The new SQL user account is created.
5. To determine if the new user was added correctly to your SQL Server:
   a. Open SQL Server Management Studio.
   b. Expand the user interface for SQL Server Management Studio, expand the specific SQL Server, and select **Security**.
   c. Double-click **Logins** and view the list of users authorized to access the SQL Server.
   d. Click the **Refresh** button in SQL Server Management Studio. If the user added previously is not listed, the management server is not able to discover the database.
6. To determine if the SQL Server is ready to accept connections from the management server:
   a. Connect to the SQL Server installation through SQL Server Management Studio using the user account added.
   b. Create a sample ODBC datasource for the SQL Server installation using the user account added.
   c. Click **Test** to test the datasource.
7. Repeat these steps for each SQL Server 2000, 2005, or 2008 instance you want to manage using Windows authentication.

   Enter the database configuration details as described in Step B – Provide the SQL Server Configuration Details on page 220.

Monitoring Sybase Adaptive Server Enterprise

To monitor Sybase Adaptive Server Enterprise, you must:

- Create an APPIQ_USER account on the database for Sybase.
- Provide the database server name and port number.
- Discover the application.

The required drivers for Sybase Adapter Server Enterprise were automatically installed along with the management server.

Make sure you purchased Sybase IQ, the module that enables you to monitor Sybase Adaptive Server Enterprise. Contact your customer support if you are unsure if you purchased this module.

Step A – Create the APPIQ_USER account for Sybase

The management server accesses Sybase through the APPIQ_USER account. This account is created when you run the CreateSybaseAct.bat script (on Microsoft Windows) or CreateSybaseAct.sh (on UNIX platforms) on the computer running the Sybase database you want to monitor. The account has create session and select dictionary privileges to be used with the management server.

Note: To create a user account with a custom user name or password, run `CreateSybaseActWithCustomPwd.bat` (on Microsoft Windows) or `CreateSybaseActWithCustomPwd.sh` (on UNIX). For more information, see Creating Custom User Names and Passwords on Managed Database Instances on page 206.

Keep in mind the following:

- The script must run under SA user.
- Obtain the Sybase server name before you run the script.
- Create APPIQ_USER account on Sybase Database you want to monitor.
- You should have already installed the database for the management server.
- Make sure you have all the necessary information before you begin the installation. Read through the following steps before you begin.

To create the APPIQ_USER account for the Sybase server:

1. Do one of the following:
   - **To run the script on IBM AIX, SGI IRIX, or Sun Solaris**, log on to an account that has administrative privileges, mount the StorageEssentialsDVD (if not auto-mounted), and go to the `/CimExtensionsCD1/DBIQ/sybase/unix` directory by typing the following:

     ```
     # cd /DVD/DVD0/DBIQ/sybase/unix
     ```

     In this instance, /DVD/DVD0 is the name of the DVD drive

     Or

   - **To run the script on Microsoft Windows**, go to the `\DBIQ\sybase\win` directory on the CIM Extensions DVD.
2. Verify that you have the password to the SA user account.

   You are prompted for the password for this user account when you run the script.
3. Run the CreateSybaseAct.bat script (on Microsoft Windows) or CreateSybaseAct.sh script

(on UNIX platforms) on the computer with the Sybase database.

The script creates a user with login to master and select privilege on data dictionary tables on a managed Sybase instance.

You can use a remote Sybase isql to run this script.

4. Enter the Sybase instance name, which must be visible to the client, as the first input when running the script. The script prompts you for the name of the sybase server on which to create user for Sybase management packages and the password of the SA account.
5. Repeat the previous step for each Sybase server you want to manage.

The script does the following:

- First, creates the APPIQ_USER account.
- Next, grants "create session" and "select on dictionary tables" privileges to the APPIQ_USER account, which enables the management server to view statistics for the Sybase server.

Removing the APPIQ_USER Account for Sybase

Before you remove the APPIQ_USER account for the Sybase databases on a host, make sure that no processes are running APPIQ_USER for that Sybase database. The management server uses APPIQ_USER to obtain information about a Sybase database. One of the ways to make sure APPIQ_USER is not being used is to temporarily remove the host running Sybase (**Discovery** > **Topology**). After you remove the APPIQ_USER account for Sybase, discover and perform Get Details for the host if you want to continue monitoring it.

To remove the APPIQ_USER account for the Sybase databases on a host:

1. Do one of the following:
   - **To run the script on IBM AIX, SGI IRIX, or Sun Solaris**, log on to an account that has administrative privileges, mount the StorageEssentialsDVD (if not auto-mounted), and go to the `/DBIQ/sybase/unix` directory by typing the following:

     ```
     # cd /DVD/DVD0/DBIQ/sybase/unix
     ```

     In this instance, /DVD/DVD0 is the name of the DVD drive.

     Or

   - **To run the script on Microsoft Windows**, go to the `\DBIQ\sybase\win` directory on the DVD.
2. Verify that you have the password to the SA user account.

   You are prompted for the password for this user account when you run the script.
3. Run `UninstallSybaseAct.bat` (on Windows) or `UninstallSybaseAct.sh` (on Unix platforms).
4. Enter the name of the Sybase server.
5. Enter the password for the SA account.

The account for APPIQ_USER is removed. The management server can no longer monitor the Sybase databases on this host.

Step B – Provide the Sybase Server Name and Port Number

You must provide the Sybase server name and port number for managing the Sybase database in the following steps.

To add information for discovering Sybase Adaptive Server Enterprise:

1. Select **Discovery** > **Setup**, and click the **Applications** tab.
2. Click **New** in the Managed Databases section.
3. In the **Host IP/DNS Name** box, enter the IP address or DNS name of the host running Sybase.
4. You can leave the Management IP/DNS Name box blank. This box is for Oracle clusters. When you leave it blank, the management server automatically lists the DNS name or IP address of the host under the **Host IP/DNS Name** column and the **Management IP/DNS Name** column.
5. In the **Database Instance Name** box, enter the Sybase database you want to monitor.
6. In the **Port Number** box, enter the port that Sybase is using.
7. Select **SYBASE** from the Database Type menu.
8. If you created a custom user name as described in Creating Custom User Names and Passwords on Managed Database Instances on page 206, select the user name from the drop-down menu. If you used the custom password script to change the user name for the database instance, but you did not already add the custom user name to the management server, you can add it now by clicking **New User**.
9. Click **OK**.

Perform Get Details for your inputs to take effect. See Step 3 – Discovering Applications on page 247.

Deleting Sybase Information

If you do not want the management server to monitor a Sybase instance:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. In the Managed Databases table, click the checkbox for the Sybase instances you do not want the management server to monitor.
3. Click **Delete**.
4. Perform Get Details to make the management server aware of your changes.

Monitoring Microsoft Exchange

If you plan to monitor Microsoft Exchange Clusters, see Monitoring Microsoft Exchange Failover Clusters on page 233.

To monitor Microsoft Exchange, you must make the management server aware of domain controller access. After information for controller access has been added, discover Microsoft Exchange, map the topology and perform Get Details. To save time, delay these steps until you add the configurations for your other applications and hosts.

Monitoring Microsoft Exchange requires the following:

- Adding information for Microsoft Exchange Domain Controller Access
- Discovering the application (Step 3 – Discovering Applications on page 247).

Adding Microsoft Exchange Domain Controller Access

Before adding a domain controller, note the following:

- The hosts should recognize the management server by name, because a reverse look-up is required by both operating system security and Microsoft Exchange. Make sure the domain controller, Exchange server host, and management server are accessible to each other using the host name and the fully-qualified domain name.
- The user name you provide could be either the Windows logon name or Common Name (CN) of the Active Directory User for accessing the Microsoft Exchange server. If the CN is provided, make sure that the user resides under the default **Users** Organization Unit (OU). The Windows logon name should be in the format `Domain\Username`, and the corresponding user could be in any OU.

To find the CN for a user on a domain controller server:

1. Install the ADSIEdit MMC snap-in if it is not yet installed.
2. Select **Start** > **Run** and enter `adsiedit.msc`.
3. When the snap-in opens, expand the DOMAIN directory and navigate to the **CN=Users** folder to see the CN for each user in the Active Directory.

To provide information about your domain controllers:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. In the Exchange Information section, click **Create**.
3. Click the **Add New Domain Controller** link.
    a. In the Domain box, enter the domain name.
    b. In the Domain Controller Name box, enter the fully qualified DNS name for the domain controller.

c. In the User Common Name box, enter the Windows logon name or the Common Name (CN) of the Active Directory User for accessing the Microsoft Exchange server.

d. In the Domain Password box, enter the corresponding password for accessing the Microsoft Exchange server.

e. In the Verify Password box, re-enter the password for verification.

4. Click **Add**. The domain controller is added to the table.
5. Click **OK**.
6. Repeat these steps for each domain controller.
7. When all of your domain controllers are added, run `wmiadap /f` on the Exchange Server to refresh the Exchange data.

You must discover the host running Microsoft Exchange. See Step 3 – Discovering Applications on page 247.

Editing a Microsoft Exchange Domain Controller

To provide information about your domain controllers:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. Click the **Edit** button next to the Exchange domain controller you want to edit.
3. Enter a new User Name or Domain Password.
4. Click **Edit**. The domain controller updates are added to the table.
5. Click **OK**.

Deleting a Microsoft Exchange Domain Controller

To delete all of the domain controllers of a particular domain:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. Click the **Delete** ( ) button corresponding to the domain you want to remove.
3. Run Get Details for your changes to take effect.

To delete a particular domain controller in a domain:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. Identify the domain for the domain controller you want to remove and click the **Edit** ( ) button corresponding to that domain.
3. In the Edit window, click the **Delete** ( ) button corresponding to the domain controller you want to remove.
4. Run Get Details for your changes to take effect.

Monitoring Microsoft Exchange Failover Clusters

To monitor and manage Microsoft Exchange Failover Clusters:

1. Install CIM Extensions on each of the participating nodes of Microsoft Exchange Failover Cluster.
2. Add information for Microsoft Exchange Domain Controller Access. See Adding Microsoft Exchange Domain Controller Access on page 231.
3. Perform Get Details on each of the participating nodes of the Exchange Cluster.

Monitoring Caché

After you complete the monitoring steps, you must discover Caché. See Step 3 – Discovering Applications on page 247.

The required drivers for Caché are automatically installed along with the management server.

Before beginning, make sure you purchased Caché IQ, which is the module that lets you monitor Caché. Contact your customer support if you are unsure if you purchased this module.

Step A – Import the Wrapper Class Definitions into the Caché Instance

For Caché 5.2 and later versions:

1. Launch the Caché System Management Portal by right-clicking the Caché Cube icon in the system tray area of the Windows toolbar and selecting **System Management Portal**.
2. Click the **Classes** link under Data Management.
3. On the Classes page, select the **Namespaces** radio button, and then select **%SYS**.
4. Click **Import**.
5. Browse the DVD, select the wrapper xml file, and click **Open**.

   IBM AIX, Linux, or HP-UX:

   Log on to an account that has administrative privileges, and mount the StorageEssentials DVD (if not auto-mounted).

   The wrapper file is `/DVD/CimExtensionsCD1/DBIQ/cachedb/unix/cachedb_sqlprojs.xml`. In this instance, DVD is the name of the directory where you mounted the DVD.

   Microsoft Windows:

   The wrapper file on the StorageEssentials DVD is `\DBIQ\CimExtensionsCD1\cachedb\win\cachedb_sqlprojs.xml`.

   OpenVMS:

a. Log on as system and mount the StorageEssentialsDVD.

b. Copy the wrapper file. Here are two examples:

   Copy `DQB0:[OVMS.DBIQ.CACHE] SQLPROJS.XML` (in this instance, DQB0 is the DVD drive) to any internal location on the OpenVMS host.

   Copy `$DQB0:[OVMS.DBIQ.CACHE]SQLPROJS.XML $DKA0:[000000]SQLPROJS.XML`. In this instance, DKA0 is a local drive on the OpenVMS host.

c. Browse to $DKA0 and specify `SQLPROJS.XML` within $DKA0 as the import file.

6. After the file is opened, click **Select All**.

7. Select **Check here to compile imported items** and click **Import**.

The wrapper class definitions are imported into the Caché %SYS namespace.

Importing Wrapper Class Definitions

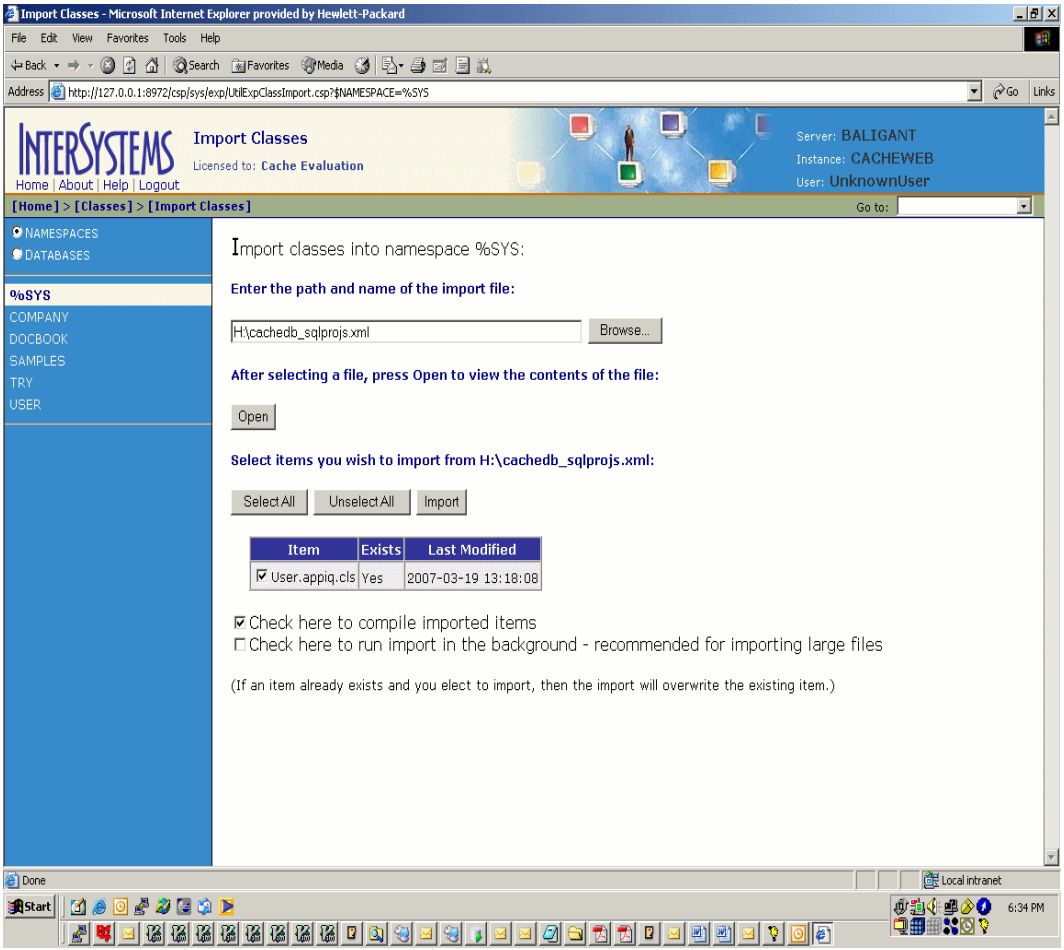

Step B – Create APPIQ_USER Account on the Caché Instance

The management server accesses Caché through the APPIQ_USER account. This account is created when you run the appropriate script (described below) on the computer running the Caché database you want to monitor. You can execute these scripts from the management server also.

This script creates APPIQROLE with execute permissions for the SQL projections imported into the Caché managed instance, creates an APPIQ_USER account, and assigns APPIQROLE to APPIQ_USER.

The script must run as the _SYSTEM user. Enter the Caché server name, the Super Server port number, and the password of the _SYSTEM user account as arguments for the script.

If you are running Caché 5.2 or later, and the Caché instance was installed using "Locked Down" security mode, see Normal and Locked Down Security Mode on next page before creating the APPIQ_USER account.

1. Create APPIQ_USER for the Caché instance either on the host or remotely, as follows:
   - **Create APPIQ_USER on the host**:
     - To run the script on IBM AIX, HP_UX, or Linux, log on to an account that has administrative privileges, mount the StorageEssentials DVD (if not auto-mounted) and go to the `/CimExtensionsCD1/DBIQ/cachedb/unix` directory by entering the following:

       ```
       # cd /DVD/CimExtensionsCD1/DBIQ/cachedb/unix
       ```

       In this instance, DVD is the name of the directory where you mounted the DVD.
     - To run the script on Microsoft Windows, go to the `DBIQ\cachedb\win` directory on the DVD.
     - To run the script on OpenVMS, log on as system, mount the DVD drive, and go to the `[OVMS.DBIQ.CACHE]` directory by entering the following:

       ```
       SET DEF DQB0:[OVMS.DBIQ.CACHE]
       ```

       In this instance, DQB0 is the name of the DVD drive.

   Or

   - **Remotely create APPIQ_USER from the management server:**
     - To run the script on Linux, go to the `/opt/<product name>/install/cachedb/unix` directory by entering the following:

       ```
       # cd opt/<product name>/install/cachedb/unix
       ```
     - To run the script on Windows, go to the `%MGR_DIST%\install\cachedb\win` directory.
2. Verify that you have the password to the _SYSTEM user account.

For later versions of Caché: run `createCacheDBUser.sh` (on UNIX), or `createCacheDBUser.bat` (on Windows), or or `CRUSER.COM` (on OpenVMS) on the computer with the CacheDatabase. To specify a custom user name or password, run `createCacheDBUserCustomPwd.sh` (on UNIX), or `createCacheDBUserCustomPwd.bat` (on Windows), or `CUSTUSER.COM` (on OpenVMS) on the computer with the CacheDatabase.

3. Enter the Caché server name, the Super Server port number, and the password of the _SYSTEM user account as arguments for the script. If you are running the custom user name and password creation script, enter the custom user name as the fourth argument and the custom password as the fifth argument.

   When invoking the scripts on OpenVMS, enclose the arguments in double quotes:

   ```
   $ @CRUSER.COM "<host name>" "<super server port>" "<password for _SYSTEM user>"
   ```

4. Repeat the previous step for each Caché instance you want to manage.

Normal and Locked Down Security Mode

For Caché 5.2 and later versions, if the Caché instance was installed using “Locked Down” security mode, follow these steps to create the APPIQ_USER account:

1. Launch the System Management Portal.
2. Click the **Security Management** link under System Administration.
3. On the Security Management page, click **Services**.
4. Click **%Service_Bindings** on the Services page.
5. On the Edit definition for Service %Service_Bindings page, do the following:
   a. Under Allowed Incoming Connections, click **Add** and enter the IP address of the management server in the Explorer User Prompt window.
   b. If the create APPIQ_USER scripts are being executed from the host on which Caché instance is running, add the IP address of the host.
   c. Click the **Service Enabled** check box on the Edit definition for Service %Service_Bindings page.
   d. Click **Save**.
6. Click the **Security Management** link under System Administration in the System Management portal.
7. On the Security Management page, click the **Users** link.
8. Click the **Edit** link for _SYSTEM user.
9. On the Edit Definition for User _SYSTEM page, click the **User Enabled** check box and enter a password for the _SYSTEM user in the Password and Confirm Password boxes.
10. Click the **Save** button.

Once the APPIQ_USER is created, the _SYSTEM user can be disabled from the System Management portal.

Removing the APPIQ_USER Account from the Caché Instance

If you no longer want the management server to monitor a Caché instance, you can remove the APPIQ_USER account and APPIQROLE for that Caché instance by running `dropCacheDBUser.bat` (on Windows), or `dropCacheDBUser.sh` (on UNIX platforms), or `DROPUSER.COM` (on OpenVMS).

Before you remove the APPIQ_USER account from the Caché instances on a host, make sure no processes are running APPIQ_USER for that Caché instance. The management server uses APPIQ_USER to obtain information about a Caché instance. One way to make sure APPIQ_USER is not being used is to temporarily remove the host running Caché (**Discovery** > **Topology**). After you remove the APPIQ_USER account for Caché, discover and perform Get Details for the host if you want to continue monitoring it.

For Caché 5.2 and later versions, if the Caché instance was installed using "Locked Down" security mode, make sure that the _SYSTEM user has been enabled before trying to remove the APPIQ_USER account.

To make sure that the _SYSTEM user has been enabled:

1. Launch the System Management Portal
2. Click the **Security Management** link under System Administration.
3. On the Security Management page, click the **Users** link.
4. Click the **Edit** link for _SYSTEM user.
5. On the Edit Definition for User _SYSTEM page, click the **User Enabled** check box and enter a password for the _SYSTEM user in the Password and Confirm Password fields.
6. Click **Save**.

Once the APPIQ_USER is removed, the _SYSTEM user can be disabled from the System Management portal. The %Service_Bindings service that was enabled before creating the APPIQ_USER can also be disabled.

1. Remove the APPIQ_USER account from the host either directly or remotely as follows:
   - **To remove the APPIQ_USER account from the host**:
     - To run the script on IBM AIX, HP_UX, or Linux, log on to an account that has administrative privileges, mount the DVD (if not auto-mounted), and go to the `CimExtensionsCD1/DBIQ/cachedb/unix` directory by entering the following:

       ```
       # cd /DVD/CimExtensionsCD1/DBIQ/cachedb/unix
       ```

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

- To run the script on Microsoft Windows, go to the `CimExtensionsCD1\DBIQ\cachedb\win` directory on the DVD.
- To run the script on OpenVMS, log on as system, mount the DVD drive, and go to the `[OVMS.DBIQ.CACHE]` directory by entering the following:

  ```
  SET DEF DQB0:[OVMS.DBIQ.CACHE]
  ```

  In this instance, DQB0 is the name of the DVD drive.

Or

- **To remotely remove the APPIQ_USER account from the Caché instance from the management server**:
  - To run the script on Linux, go to the `/opt/<product name>/install/cachedb/unix` directory by entering the following:

    ```
    # cd opt/<product name>/install/cachedb/unix
    ```

  - To run the script on Windows, go to the `%MGR_DIST%\install\cachedb\win` directory.

2. Verify that you have the password to the _SYSTEM user account.
3. Enter the Caché server name, the Super Server port number, and the password of the _SYSTEM user account as arguments for the script.

   When invoking the scripts on OpenVMS, enclose the arguments in double quotes:

   ```
   $ @DROPUSER.COM "<host name>" "<super server port>" "<password for _SYSTEM user>
   ```

4. Repeat the previous step for each Caché instance you want to manage.

After deleting the APPIQ_USER account from the Caché instance, follow these stepts to delete the wrapper class definitions:

For Caché 5.2 and later versions:

1. Launch the Caché System Management Portal.
2. Click the **Classes** link under Data Management.
3. On the Classes page, select the **Namespaces** radio button, and then click **%SYS**.
4. Click **Delete**.
5. Enter `User.appiq.cls` in the Enter search mask box, and click **Search**.
6. Select **User.appiq.cls** and click **Delete**.

Step C – Provide the Caché Instance Name and Port Number

To provide the Caché instance name and SuperServer port number for managing the Caché instance:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. Click **New** in the Managed Databases section.
3. In the Host IP/DNS Name box, enter the IP address or DNS name of the host running Caché.
4. You can leave the Management IP/DNS Name box blank. This box is for clusters. When you leave it blank, the management server automatically lists the DNS name or IP address of the host under the Host IP/DNS Name column and Management IP/DNS Name column.
5. In the **Database Instance Name** box, enter the Caché instance name you want to monitor.
6. In the Port Number box, enter the SuperServer port used by Caché.
7. Select **Cache** from the Database Type menu.
8. If you created a custom user name as described in Creating Custom User Names and Passwords on Managed Database Instances on page 206, select the user name from the drop-down menu. If you used the custom password script to change the user name for the database instance, but have not yet added the custom user name to the management server, add it now by clicking **New User**.
9. Click **OK**.

Perform Get Details for your changes to take effect. See Step 3 – Discovering Applications on page 247.

Deleting Caché Information

If you do not want the management server to monitor a Caché instance, follow these steps to remove its information:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. In the Managed Databases table, click the checkbox for the Caché instances you do not want the management server to monitor.
3. Click **Delete**.
4. Perform Get Details to make the management server aware of your changes.

Monitoring IBM DB2

After you complete the monitoring steps, you must discover the DB2 database and perform Get Details. See Step 3 – Discovering Applications on page 247.

Step A — Grant Privileges to the Specified User on the DB2 Database

The management server accesses DB2 through the system users that are used to manage the database. Use the `GrantDB2User` script to assign all of the necessary privileges to any database user who is a member of the SYSMON_GROUP.

Keep in mind the following:

- The script must be executed by a user who is a member of the DB2 administrator group; or example, the SYSADM_GROUP.
- Obtain the DB2 database name before you run the script.
- You should have already installed the database for the management server.
- Make sure you have all the necessary information before you begin the installation. Read through the following steps before you begin.

To grant privileges to the specified user on the DB2 database:

1. Do one of the following:
   - **To run the script on UNIX**:

     Log on to an account that has administrative privileges, mount the StorageEssentials DVD (if not auto-mounted), and go to the `/DBIQ/db2/unix` directory by entering the following:

     ```
     # cd /DVD/DVD0/DBIQ/db2/unix
     ```

     In this instance, DVD/DVD0 is the name of the DVD drive

   - **To run the script on Microsoft Windows**:

     Go to the `CimExtensionsCD1\DBIQ\db2\win` directory on the StorageEssentials DVD.

2. Run the `GrantDb2User.sh script` (on Unix) or the `GrantDb2User.bat` script (on Windows) on the computer with the DB2 database. The script assigns the necessary privileges to the specified user.

   Unix example:

   ```
   $ ./GrantDb2User.sh sample testusr /opt/ibm/db2/V9.5/bin

   Successfully granted LOAD authority to user "testusr" for database
   "sample"

   $
   ```

   Windows example:

   ```
   H:\DB2>GrantDb2User.bat sample testuser h:\DB2 "C:\Program
   Files\IBM\SQLLIB\BIN"

   "Successfully granted LOAD authority to user "testuser" for
   database "sample""

   H:\DB2>
   ```

Revoking Privileges

Before you revoke privileges for the user for the DB2 databases on a host, make sure that no processes are running for that DB2 database for that user. The management server uses the user to obtain information about a DB2 database. To ensure that the user is not being used, temporarily remove the host running DB2 (**Discovery** > **Topology**). After you revoke privileges for the user for the DB2 database, discover and perform Get Details for the host if you want to continue monitoring it.

To revoke privileges from the user for the DB2 databases on a host:

1. Do one of the following:
   - **To run the script on UNIX**:

     Log on to an account that has administrative privileges, mount the StorageEssentials DVD (if not auto-mounted), and go to the `/CimExtensionsCD2/DBIQ/db2/unix` directory by typing the following:

     ```
     # cd /DVD/DVD0/CimExtensionsCD2/DBIQ/db2/unix
     ```

     In this instance, /DVD/DVD0 is the name of the DVD drive
   - **To run the script on Microsoft Windows**:

     Go to the `\DBIQ\db2\win` directory on the DVD.
2. Run the RevokeDb2User script on the computer with the DB2 database.

   Unix example:

   ```
   $ ./RevokeDb2User.sh sample testusr /opt/ibm/db2/V9.5/bin

   Successfully revoked LOAD authority of user "testusr" for database
   "sample"

   $
   ```

   Windows example:

   ```
   H:\DB2>RevokeDb2User.bat sample testuser h:\DB2 "C:\Program
   Files\IBM\SQLLIB\BIN"

   "Successfully revoked LOAD authority of user "testuser" for
   database "sample""

   H:\DB2>
   ```

   The privileges are revoked from the user. The management server can no longer monitor the DB2 databases on this host.

Step B — Provide the Database Instance Name, Port Number, Database Name, and User Name

You must provide the DB2 instance name, port number, DB2 path, database name, and user name for managing the DB2 databases.

To add information for discovering DB2:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. Click **New** in the Managed Databases section.
3. In the Host IP/DNS Name box, enter the IP address or DNS name of the host running DB2.
4. You can leave the Management IP/DNS Name box blank. This box is for Oracle clusters. When you leave it blank, the management server automatically lists the DNS name or IP address of the host under the Host IP/DNS Name column and Management IP/DNS Name column.
5. In the Database Instance box, enter the DB2 instance name of the database you want to monitor.
6. In the Port Number box, enter the port used by DB2.
7. Select **DB2** from the Database Type menu.

   HP Storage Essentials displays additional fields when DB2 is selected.

   Provide the following information for the DB2 database:

   a. In the DB2 Path field, enter the absolute path to the DB2 executable. The DB2 path must be provided if the DB2 instance uses SMS tablespaces and capacity information for the same needs to be collected.

   b. In the Database Name field, enter the name of the DB2 database managed by the DB2 instance mentioned in step 5.

   c. Select one of the existing users who has privileges on the DB2 database from the User Name menu. You can also create a new user by clicking the **New User** button.

   d. Click the **Add to Table** button.

   e. Repeat steps b through d for all the databases that belong to the instance mentioned in step 5 and that must be monitored.
8. Click **OK**.

Perform Get Details for your changes to take effect. See Step 3 – Discovering Applications on page 247.

Deleting DB2 Information

If you do not want the management server to monitor a DB2 database, you can remove its information.

The **Delete** () button is disabled for DB2 instances with only one database record.

To remove DB2 information:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. In the Managed Databases table, click the checkbox for the DB2 instances you do not want the management server to monitor.
3. Click **Delete**.
4. Perform Get Details to make the management server aware of your changes.

Step C — Install the JDBC Driver for DB2 Databases

To install the JDBC driver:

1. Download the driver from: http://www-01.ibm.com/support/docview.wss?rs=4020&uid=swg21385217

   The driver is titled IBM Data Server Driver for JDBC and SQLJ (JCC Driver).
2. Place the driver jar files in the following location:

   Windows:

   `C:\hp\StorageEssentials\JBossandJetty\server\appiq\lib`

   Unix:

   `/opt/HP_Storage_Essentials/JBossandJetty/server/appiq/lib directory`
3. Restart the AppStorManager service.

Monitoring IBM Informix

After you complete the steps for monitoring IBM Informix, you must discover the Informix database and perform Get Details. See Step 3 – Discovering Applications on page 247.

Before you begin, make sure that you purchased Informix IQ, which is the module that lets you monitor Informix. Contact customer support if you are unsure if you purchased this module.

Step A — Create a Managed Database User Account for Informix

The management server accesses the Informix database through the managed database user account. For discovering and monitoring all Informix elements except sbspace and blobspace, the management server connects to the sysmaster database on the Informix database server using the managed database user account. For collecting sbspace and blobspace details, the management server connects to each database using the managed database user account and queries the necessary system catalogue tables. By default, any operating system user has SELECT privileges on the sysmaster database. In order to connect to each database and collect sbspace and blobspace information, the managed database user should have connect privileges on each database.

Keep in mind the following:

- The script must run under the root user.
- At least 250 KB free space should be available in the `/tmp` directory.

To grant permissions to the system user:

1. Log on as the root user, mount the StorageEssentialsDVD (if not auto-mounted), and go to the `CimExtensionsCD1/DBIQ/informix/unix directory` by entering the following:

   ```
   # cd /DVD/DVD0/DBIQ/informix/unix
   ```

   In this instance, DVD/DVD0 is the name of the DVD drive
2. Set the values for the following environment variables: INFORMIXDIR, INFORMIXSQLHOSTS and INFORMIXSERVER.
3. Run the `GrantInformixUser.sh` script on the computer where the Informix database is installed.
4. Enter the managed database user account. This is any operating system user and that was configured as a managed database user in HP Storage Essentials.

   Configuring “informix” and “root” as Managed Database User to discover and manage the Informix Dynamic Server is not recommended.
5. Enter the password for the Informix user. The database super user password is required to grant privileges to the managed database user for each database.
6. Repeat the previous steps for each Informix server you want to manage.

The script connects to the Informix database server with the user account informix, and grants privileges to the managed database user to allow it to connect to the individual databases and query system catalog tables.

Revoking Connect Privileges from the Managed Database User

To revoke connect privileges from the managed database user on Informix databases:

1. Log on as the root user, mount the StorageEssentialsDVD (if not auto-mounted), and go to the `CimExtensionsCD1/DBIQ/informix/unix` directory by entering the following:

   ```
   # cd /DVD/DVD0/DBIQ/informix/unix
   ```

   In this instance, DVD/DVD0 is the name of the DVD drive.
2. Set the values for the following environment variables: INFORMIXDIR, INFORMIXSQLHOSTS, and INFORMIXSERVER.
3. Run the `RevokeInformixUser.sh` script on the computer with the Informix database.
4. Enter the managed database user account.
5. Enter the password for the Informix user. The database super user password is required to revoke connect privileges from the managed database user.

The script revokes privileges from the operating system user so that they will not be able to connect to individual database.

Step B — Install the Informix JDBC Driver

HP Storage Essentials does not package and distribute the JDBC driver for Informix.

To install the JDBC driver for Informix:

1. Download the Informix JDBC driver 3.50.JC4 from IBM's portal at:

   http://www14.software.ibm.com/webapp/download/search.jsp?cat=&q0=&pf=&k=ALL&pn=Informix+JDBC&pid=&rs=&S_TACT=104CBW71&status=Active&S_CMP=&b=&sr=1&q=3.50&ibm-search=Search

2. Install the JDBC driver in a temporary location. For details about installing the JDBC driver, see the installation guide packaged with the JDBC driver installer.
3. Copy the `ifxjdbc.jar` file from the temporary location where the JDBC driver is installed and add it to the `$MGR_DIST/JBossandJetty/server/appiq/lib` directory. In this instance, $MGR_DIST is the location where HP Storage Essentials is installed.
4. Restart the AppStorManager server, which is the service for HP Storage Essentials.

Step C — Provide the Informix Server Name and Port Number

To provide the Informix server name and port number:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. Click **New** in the Managed Databases section.
3. In the Host IP/DNS Name box, enter the IP address or DNS name of the host running Informix.
4. You can leave the Management IP/DNS Name box blank. This box is for Oracle clusters. When you leave it blank, the management server automatically lists the DNS name or IP address of the host under the Host IP/DNS Name column and Management IP/DNS Name column.
5. In the Database Server text field, enter the name of Informix database server you want to monitor.
6. In the Port Number field, enter the port that Informix is using for client connection.
7. Select INFORMIX from the Database Type menu.
8. If you created a managed database user account as described in Creating Custom User Names and Passwords on Managed Database Instances on page 206, select that user name from the drop-down menu. If you have not yet created a managed database user account, you can add it now by clicking New User.
9. Click **OK**.

Deleting Informix Information

If you do not want the management server to monitor an Informix instance, you can remove its information as follows:

1. Select **Discovery** > **Setup** and click the **Applications** tab.
2. In the Managed Databases table, click the check box for the Informix instances you do not want the management server to monitor.
3. Click **Delete**.
4. Perform Get Details to make the management server aware of your changes.

Application Discovery Test

Application discovery allows you to test the configuration information entered during application setup. This allows you to verify the accuracy of the configuration information prior to running discovery.

Application discovery tests on unmanaged hosts are not supported.

A Test Discovery shows a number of repeated attempts by the same provider to access an element, but each attempt uses a different set of credentials. There can be at most three default credentials. This is normal behavior. The Test Discovery mechanism tries all available default credentials, as will Step 1 Discovery.

To run an application discovery test on Caché, Microsoft SQL, Oracle, Sybase, Informix, or DB2:

1. Select **Discovery** > **Setup**, and then click the **Applications** tab.
2. In the Managed Databases section, select the checkbox for the application on which you want to run a test discovery.

   You can only run a test discovery on one application at a time.
3. Click **Test**. The Log Messages windows displays with the results of the test discovery.

To run an application discovery test on Microsoft Exchange:

1. Select **Discovery** > **Setup**, and then click the **Applications** tab.
2. In the Exchange Information section, click the **Test** button in the row for the domain controller on which you want to run a test discovery. The Exchange Server Test Discovery dialog box appears.
3. To test all of the Exchange Servers, select the **All Exchange Servers** radio button. To select a subset of the Exchange Servers, enter the name of the Exchange Servers in a comma-separated list.

   The Exchange Server name can be the standalone Exchange instance name or the EVS name.
4. Click **OK**. The Log Messages windows displays with the results of the test discovery.

Step 3 – Discovering Applications

This step assumes you already discovered your hosts and provided discovery information for your applications. To discover an application:

- Detect the application (Step A – Detect Your Applications below)
- Obtain topology information about the application (Step B – Obtain the Topology on next page)
- Perform Get Details (Step C – Run Get Details on next page)

Keep in mind the following:

- This section assumes you already set up the discovery configurations for your applications as described in Step 2 – Setting Up Discovery for Applications on page 205.
- If you used a custom user name or password for the APPIQ_USER account, you must change the user name and password on the management server before performing Get Details. See Creating Custom User Names and Passwords on Managed Database Instances on page 206.
- If DNS records for your Microsoft Exchange servers are outdated or missing, the discovery of Microsoft Exchange can fail because Microsoft Exchange is dependant on Active Directory, which is dependant on DNS. Since Active Directory is dependant on DNS, Active Directory replication and Active Directory lookups can fail or contain errors if DNS records are not accurate.
- The management server is unable to discover Oracle on a Windows host if the host is on a private network behind a Windows proxy server. The management server can discover the Windows host through the Windows proxy server, but the management server is not able to detect Oracle.
- To run an application discovery test, see Application Discovery Test on previous page.

  Discovery consists of three steps:

  - **Setting up** – Finding the elements on the network.
  - **Topology** – Mapping the elements in the topology.
  - **Details** – Obtaining detailed element information.

Step A – Detect Your Applications

To make the software aware of the applications on the network:

1. Click **Discovery** > **Setup**.
2. To start discovering elements on the network, click the **Start Discovery** button on the IP Addresses tab.

The software discovers the IP addresses selected.

During discovery, the following occurs:

- The status light changes from green to orange.
- The Log Messages page opens. To view the status of discovery, click **Discovery** > **View Logs**.

When discovery is complete, the DISCOVERY COMPLETED message is displayed in the Log Messages box.

Keep in mind the following:

- If DNS records for your Microsoft Exchange Servers are outdated or missing, the discovery of Microsoft Exchange might fail because Microsoft Exchange is dependant on Active Directory, which is dependant on DNS. Since Active Directory is dependant on DNS, Active Directory replication and Active Directory lookups might fail or contain errors if DNS records are not accurate.
- If you are having problems discovering an element, see Troubleshooting Discovery and Get Details on page 860.

Step B – Obtain the Topology

The user interface can load slowly while the topology is being recalculated. It can also take more time to log on to the management server during a topology recalculation.

To obtain the topology:

1. Click **Discovery** > **Topology**. The discovered elements are selected.
2. Click the **Get Topology** button. The management server obtains the topology for selected elements.
3. Select the discovery group from which you want to obtain the topology. If you are obtaining the topology for hosts for the first time, make sure **All Discovery Groups** is selected.

   You can use discovery groups to break up getting the topology or getting details. For example, instead of obtaining the topology for all of the elements, you could specify that the management server gets the topology for only the elements in Discovery Group 1, thus saving you time. You add an element to a discovery group by modifying the properties used to discover the element. See Modifying the Properties of a Discovered Address on page 151.
4. If you see errors in the topology, look at the log messages, which can provide an indication of what went wrong. Look at Event Manager for additional information. Access Event Manager by clicking the Event Manager button in the left pane. To obtain troubleshooting information, see the Troubleshooting Topology Issues on page 882.

   If the topology for an element in your network changes, select the element and click **Get Topology** in **Discovery** > **Topology** to update the information.

   The software obtains just enough information about where the element is connected in the topology, for example a switch connected to a host.

Step C – Run Get Details

Obtain detailed information from the discovered applications as described in this section.

Keep in mind the following:

- Get Details takes some time. You might want to perform this process when the network and the managed elements are not busy.
- During Get Details the topology is recalculated. While the topology is being recalculated, the loading of the user interface might be slow. It might also take more time to log on to the management server during a topology recalculation.
- To obtain a picture of device connectivity quickly, click the **Get Topology** button on the Topology tab.
- When you do Get Details that includes an AIX host, three SCSI errors (2 FSCSI error and 1 FCS error) per IBM adapter port are displayed in the system log. You can ignore these errors.
- You can quarantine elements to exclude them from Get Details. See Placing an Element in Quarantine on page 160 for more information. Let us assume you want to discover all the elements in a discovery group, except for one. Perhaps the element you want to quarantine is being taken off the network for maintenance. You can use the quarantine feature to exclude one or more elements from discovery.
- If the management server is unable to obtain information from an element during Get Details as a result of a CIM extension failure, the management server places the access point where the CIM extension is located in quarantine. The management server then moves onto getting details for the next element in the Get Details table. These elements appear as missing until they are removed from quarantine. For information on how to remove an element from quarantine, see Removing an Element from Quarantine on page 160.

To obtain details:

1. Select **Discovery** > **Details**.
2. Select the discovery group from which you want to Get Details. If you are obtaining Get Details for hosts for the first time, make sure **All Discovery Groups** is selected.

   You can use discovery groups to break up getting the topology or Get Details. For example, instead of Get Details for all the elements, you could specify that the management server gets the element details for only the elements in Discovery Group 1, thus saving you time. You add an element to a discovery group by modifying the properties used to discover the element. See Modifying the Properties of a Discovered Address on page 151.
3. Click **Get Details**.

   During Get Details, the status light changes from green to red. You can view the progress of gathering details by clicking **Discovery** > **View Logs**.

   When discovery is complete, the DISCOVERY COMPLETED message is displayed in the Log Messages box.

   If the management server cannot communicate with an application, it labels the application as “Discovered.” The management server found the application, but could not obtain additional information about it.

4. See “Adding a Discovery Schedule” in the User Guide for information about automating the gathering of Get Details. If you run into problems with discovery, see Troubleshooting Discovery and Get Details on page 860.

Changing the Oracle TNS Listener Port

The software uses port 1521 by default to communicate with the TNS Listener service on the Oracle server. If your port is different or you use multiple ports, you can assign a new port number.

The hosts should recognize the management server by name, as a reverse look-up is required by operating system security as well as the Oracle Transparent Name Substrate (TNS).

To change this port number or to add ports:

1. Select **Discovery** > **Setup**, and click the **Applications** tab.
2. To assign a new port, click the **Create** button for the Oracle Information table.
3. Enter the new port number and click **OK**.
4. If necessary, click the  button to remove the old port number.
5. Verify that all elements were discovered by clicking the **Start Discovery** button.

For more information, see Troubleshooting Discovery and Get Details on page 860.

Known Issues about Applications

This section provides information about known issues with applications.

- Oracle ACFS shown with Drive Type "Local" even if the file system is on an External Drive. The Drive Type on the Storage Volumes page is shown as "Local" for Oracle Automatic Cluster File System (ACFS) file systems even if the ACFS file system is on an external disk.
- Unmounted Databases not shown on Properties Page for InterSystem Cache Databases. On the Properties Page for InterSystem Cache Database instances, unmounted databases are not shown under Logical Elements.
- sblobspace Reported for an Informix Server even if the sblobspace is Removed. The sblobspace reported for an Informix installation continues to be reported by the management server even if the sblobspace is removed.
- Usernames to Discover Applications must be Unique. In the Setup->Applications tab, user names are unique. A single user name with different passwords cannot be used to discover databases on multiple hosts; the user interface will show only one entry for a particular user name.
- Redo Groups on Raw Devices shown only for one RAC Instance. Redo groups appear in the topology for only one RAC instance in an Oracle RAC configuration with raw devices.

- Capacity Charts for Informix Databases show dbspaces. Although databases are listed on the Capacity pages, the Capacity Manager Charts display data for dbspaces for Informix databases.
- Cannot Create a Virtual Application on an Oracle RAC Shared Volume on Solaris x86. At this time it is not possible to create a virtual application on a shared Oracle RAC volume on Solaris x86. You will see the following message: “java.lang.NullPointerException.”
- Update Element Data (Single Element Refresh) does not Update all Oracle Failover Information. Performing a single element refresh does not update the Oracle Failover information about which node is active if there has been a failover. Get Details updates all the necessary information.
- Host Cluster Topology Does Not Show Oracle Database Instances as Shared. Oracle database instances on shared raw volumes in a cluster are not reported as shared on the Host Cluster Topology. The individual instances are shown as local to the host and not shared in the cluster. The Application Topology page shows the proper configuration.
- Status not Displayed for Oracle Database Instance Control Files. The status of the Oracle database instance's control files is not shown on the instance properties page.
- Exchange Services Statistics Chart Shows Raw Data. The Exchange Services Statistics Chart will report only the raw data available. It does not report on rolled-up data. This chart is being reconsidered, as a roll-up of a “service up” or “service down” value is not meaningful.

4 Agentless Discovery

Use agentless discovery to gather information about hosts based on host security groups, zones and zone aliases configured on storage systems and switches in the SAN. Hosts can be inferred based on specific search parameters and managed without installing a CIM extension.

The following functionality is not available for hosts inferred through agentless discovery:

- Automatic cluster membership detection
- Application support, such as Application Viewer, Backup Manager, and File System Viewer
- Host properties
- Full path calculations

If you set a system property, the product will guess the path calculations for inferred hosts based on host security group membership, but these calculations do not take into account the following:

- Account target mappings
- Logical drives
- Multipathing
- Volume Management

Host capacity information is available, but might not be accurate because it is based on the host security group. As a result, local disk capacity and all the mounted volume capacity are not displayed.

Creating Discovery Rules for Inferred Hosts

HP Storage Essentials treats the creation of inferred rules for hosts without a CIM extension as a two-step process. First you create the rule, as described in Step 1 – Create the Discovery Rule below, and then test the rule, as described in Step 2 – Test the Newly Created Rule on page 255.

Step 1 – Create the Discovery Rule

HP Storage Essentials can display and gather information from hosts without CIM extensions. You can create rules that effectively probe your switch and storage configurations to infer hostnames based on the World Wide Names of their HBA ports and correctly display them in System Manager.

Before creating rules, perform Step 1 and Step 3 discovery for the following elements:

- Switches and storage systems
- Hosts with CIM extensions installed

Agentless host discovery rules do not work for generic hosts that are grouped together in System Manager. You must ungroup generic hosts, as described in Ungrouping Discovered Hosts on page 471. If the host has a question mark above it and its name contains an underscore followed by several numbers, the host is considered a generic host since HP Storage Essentials could not obtain additional information about the host in Discovery step 3. If the host has a question mark and the word “inferred” after its name, the host was inferred through an agentless inference rule.

Virtual machines and iSCSI hosts also cannot be inferred using agentless rules. Agentless discovery is not supported for virtual machines.

Agentless rules can be imported and exported through the discovery lists. For more information about importing and exporting the discovery lists, see Importing Discovery Settings from a File on page 81 and Saving Discovery Settings to a File on page 83.

To create a rule for discovering agentless hosts:

1. Select **Discovery** > **Agentless Hosts**.
2. Click **Create Rule**.
3. Provide a name for the rule in the **Rule Name** field.
4. (*Optional*) Provide a description for the rule in the **Rule Description** field.
   Rule priority: Rules are run in a sequence from high to low priority. For example, a rule with a priority of 1 will run before a rule with a priority of 4.
5. (*Optional*) Select **Run this rule at completion of all Discovery Details** to discover new hosts and update information. If you select this option, the rule will run after every Discovery Step 3 (Get Details).

   It is recommended that you do not select this option because it will add a performance impact during each discovery. To update information for an inferred host, use the Update button on the host tab, as described in Viewing Agentless Hosts on page 262.
6. Select the type of information the rule will use to discover the hosts:
   - **Host Security Group** – HP Storage Essentials searches the host security group names on the storage systems for hosts. You must have storage systems discovered through Discovery Step 3.
   - **Zone** – HP Storage Essentials searches the zone name for hosts on the switches. You must have switches discovered through Discovery Step 3.
   - **Zone Alias** – HP Storage Essentials searches the zone alias name for hosts on the switches. You must have switches discovered through Discovery Step 3.

   Keep in mind the following when selecting Zone or Zone Alias as a scope:
   - You can run the rule from a management server where you have only discovered switches. You will be able to infer host names, but you will not obtain any storage details, since no storage has been discovered.
   - You do not need to discover the entire fabric.
   - Orphan zones and orphan zone aliases could return false inferences.

7. Provide an expression for agentless rules. These rules determine how the element will be discovered. See Creating Regular Expressions below for more information.
8. Click **Next**. The Test tab appears.
9. Continue with Step 2 – Test the Newly Created Rule below.

Step 2 – Test the Newly Created Rule

To use the Test tab to verify the rule you created:

1. Click **Start Test**.

   HP Storage Essentials displays the hosts it found with the expression you created.

   Agentless host discovery rules do not work for generic hosts that are grouped together in System Manager. You must ungroup generic hosts, as described in Ungrouping Discovered Hosts on page 471. Generic hosts are hosts discovered by HP Storage Essentials but additional information could not be obtained from them because they do not have a CIM extension installed. HP Storage Essentials designates generic hosts by a question mark in the topology.

   When you run an agentless host discovery rule in test mode, it reports on all zone/alias/HSG names that match the regular expression. If any of these are for hosts that already exist, such as host with a CIM extension, those hosts get reported with an empty HBA port column.

2. Click **Finish**. The inference rule is added to the Agentless Hosts Rules table.

You must run the rules for the hosts to be inferred through agentless discovery. For more information, see Running Rules on page 261.

Creating Regular Expressions

To infer agentless hosts, create a regular expression that meets the following criteria:

- Takes into account the naming convention of the zones, zone aliases, and host security groups in the environment so the host can be detected.
- Contains a capturing group that is used to display the host name. A capturing group is the characters within a set of parentheses.

For example, assume the agentless hosts you want to infer are prefixed with boston_, but you only want to display the host names without the boston_ prefix. You could use the following expression: boston_(.*)

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

If you wanted boston_ to be displayed in the host name and you still want only hosts with the prefix boston_ discovered, you could change the expression so that boston_ is included in the capturing group, as shown in the following expression: (boston_.*)

Note: You might need multiple rules for different naming conventions.

If you are not sure where to begin, consult the following examples to see if any match your environment. Try entering some of the basic expressions, such as .*_.*_.*, and see what is inferred. You can always add additional rules to narrow the range to detect a particular naming convention.

Examples of Regular Expressions

What is my environment?	What can I provide as an expression so HostName is displayed?	Result
`Boston_HostName_hba1`	`.*_(.*)_.*`	Strings that match the pattern of text_text_ text will be scanned. The text between the first and second underscores will be displayed as the host name.
`Boston-HostName-disk`	`.*-(.*)-.*`	Strings that match the pattern of text-text-text will be scanned. The text between the first and second dashes will be displayed as the host name.
`Boston-HostName_com`	`.*-(.*)_.*`	Strings that match the pattern of text-text_text will be scanned. The text between the first dash and second underscore will be displayed as the host name.
`Boston_storage_HostName`	`Boston_storage_(.*)`	Strings that match the pattern of Boston_ storage_text will be scanned. The text after the second underscore will be displayed as the host name.

What is my environment?	What can I provide as an expression so HostName is displayed?	Result
`Boston___HostName_disk`	`.*___(.*)_.*`	Strings that match the pattern of text____text_text will be scanned. The text between the third and fourth underscores will be displayed as the host name.
`uhcHostName` HostName is always the fourth character.	`...(.*)`	Strings that have four or more characters will be scanned and any characters after the third character spot will be displayed as the host name.
`HostName:hba`	`(.*):.*`	Strings that match the pattern of text:text will be scanned. Any text before the colon will be displayed as the host name.

What is my environment?	What can I provide as an expression so HostName is displayed?	Result

`boston_HostName_hba1`
`boise_HostName_hba1`
`marlborough_HostName_hba1`
but you do not want to discover `zebra_HostName_hba1`
	`[a-q]_(.*)_.*`	Strings that begin with any lowercase letter from a to q and matches the pattern of text_text_text will be scanned. Any text between the first and second underscore will be displayed as the host name.
		For uppercase letters use [A-Q].
		You can change the range to match your environment; for example, a-s or N-Z.

`boston1_HostName_hba1`
`boston3_HostName_hba1`
but you do not want to discover `boston9_HostName_hba1`
	`.*[1-3]_(.*)_.*`	Strings that have number 1, 2 or 3 before the first dash and that match the pattern.
		Any text between the first and second underscores will be displayed as the host name.
		You can change the range to match your environment; for example, 23 to 54.

What is my environment?	What can I provide as an expression so HostName is displayed?	Result
`HostName1_HostName2_HostName3`		Strings that have two underscores will be scanned. Text before, after, and between the underscores will be displayed as host names.
MRO_HostName_diskMy naming convention requires all zone names to begin with MRO, but I know a few have been created incorrectly and I want to capture those. For example, if I want to find any rogue zone names that do not start with "M" because my naming convention requires that all zones begin with "MRO," I would attempt to infer hosts with an expression like ([a-ln-zA-LN-Z]*).	`([a-ln-zA-LN-Z]*)`	This expression displays strings that begin with any letter except for the lowercase or uppercase letter M. The entire string would be displayed as the host name, so you could find the rogue zone names.

The notation used in the expressions are defined as follows.

Definition of Common Notation Used in Expressions

Expression	Definition
`()`	Capturing group. Any expression within a set of parenthesis is displayed for the host name. If you do not provide a capturing group, no host name will be displayed from the hosts that were detected from the expression.

Expression	Definition

.*	Any character zero or more times. Use this expression carefully. For example, the following expression matches any element that has the boston_ prefix:
	`boston_.*`
	If you want HP Storage Essentials to display any character after the boston_ prefix, add a capturing group as follows:
	`boston_(.*)`
	Assume though that you do not want to display all the characters after the boston_ prefix. If there is a character after .*, the wild card attribute will stop. For example, the following expression displays the characters that appear after boston_ and before _companyname:
	`boston_(.*)_companyname`
	Assume that all of your hosts do not end in _companyname. You can replace _companyname with _.* as follows:
	`boston_(.*)_.*`
	The expression matches all hosts with the prefix of boston_, and displays any character that is after boston_ but before the second underscore.

.	Any character. For example, assume the agentless hosts in your environment all have different naming conventions, but contain three characters before the host name. You could provide an expression as follows:
	`...(.*)`
	Hosts with the name BosHost1 or LasHostA would appear as follows in the topology:
	`Host1 and HostA`

[a-q]	Lowercase letter between a and q
[A-Q]	Uppercase letter between A and Q
[0-7]	Digits between 0 and 7

Expression	Definition

|	The OR operator. Use the OR operator when you have different naming conventions in your environment. For example, assume you want to match hosts prefixed with boston_ or boise_. You could use the following expression to match those hosts:
	`boston_(.*) | boise_(.*)`
	You could also use the OR operator to find hosts when the naming convention differs between host names. For example, assume you have some hosts that have underscores in their name and others that have dashes. You could use the following expression to match those hosts:
	`.*_(.*) | .*- (.*)`

For more information about regular expressions, go to:

http://java.sun.com/j2se/1.5.0/docs/api/java/util/regex/Pattern.html

Running Rules

You must run the rule for the host to be inferred through agentless discovery. When a host is inferred, the word (inferred) appears after the host name throughout the product, for example: HostName (inferred).

When you run a rule, an event is generated in Event Manager for each host inference. The event tells you the duration it took to run the rule and it also specifies the specific name of the rule that inferred each host.

The Run on Discovery column is cleared when a new discovery list is imported. Run the rules again to repopulate the column.

To run a report rule:

1. Select **Discovery** > **Agentless Hosts**.
2. Select a rule.
3. Click **Run Rule**.

HP Storage Essentials displays the hosts that are inference candidates based on the expression used. After the rule is executed, the inferred hosts are displayed in the System Manager topology.

A host detected through agentless discovery will have the word “Inferred” in parenthesis after its name on its properties page. In the topology, agentless hosts have a question mark above their icon. You can differentiate agentless hosts from generic hosts, which also have a question mark when displayed in the topology, because agentless hosts do not have an underscore followed by several numbers in their name.

Editing Rules

To edit a rule:

1. Select the rule in the Agentless Host table.
2. Click the **Edit** ( ) button.
3. Modify the rule as necessary.
4. Click **Next** and then click the **Start Test** button. HP Storage Essentials displays the hosts it found with the expression you modified.
5. Click **Finish**.

Deleting Rules

To delete a rule, select it from the Agentless Hosts Discovery Rules table and click **Delete** ( ) button.

Viewing Agentless Hosts

The Host tab displays hosts that have been inferred through agentless rules. A rule must have run at least once for the hosts associated with the rule to be displayed.

To access the Hosts tab:

1. Click **Discovery** > **Agentless**.
2. Click the **Hosts** tab.

You can modify the display so that you see only a subset of the agentless hosts discovered.

To filter the display on the Hosts tab:

1. Click the **Filter** link.
2. To filter by the name of the host, provide the name, or a portion of the name of the host, in the Host Name Contains text box.
3. Select one of the following from the Host Type box:
   - **All Agentless Hosts** - All agentless hosts are displayed.
   - **Rule-Discovered Hosts** - All agentless hosts that were discovered through agentless rules and not named are displayed.
   - **Named Generic Hosts** - Agentless hosts that have since been named are displayed.
4. Select one of the following from the Rule box:
   - **<All Rules>** - Any agentless host that was discovered through an agentless rule is displayed.
   - **Agentless Rule** - Select an agentless rule to display only the hosts that were discovered through that rule.

5. Click **Filter** to display the agentless hosts according to the filter. To reset the filter, click the **Reset** button.

You can remove hosts from the list. The hosts reappear in the list when the rule that was used to infer the deleted host runs again after Discovery Step 3.

Use the **Update** button to recalculate the changes in the host topology for inferred hosts and custom-named generic hosts.

An update calculates the mappings for a host. For example, if you added or deleted a new LUN or initiator port for an HBA in a host security group because you configured multipathing, you would not see the change in the topology for the inferred host until you run an update. The storage calculations displayed on the Presented Storage tab can also change to account for new configurations.

An update looks at the WWNs from the host as they are presented to the storage array through the host security group on the storage array. Inference is only as good as the configuration of the zoning and host security groups and how well your inference rules are created to capture that data.

When you run an update, for inferred or custom generic hosts, the update recalculates any changes that occurred with the addition or deletion of new host security group information. You also receive event notification for the following:

- Starting of the update process
- Ending of the update process
- Starting of resynthesis for each host. Resynthesis is the recalculation of the host, such as its topology, presented storage, and mappings to the inferred host.
- Completion of resynthesis for each host and how long it took

For examples of the messages displayed during an update of inferred hosts and discovered hosts, see Events Displayed in Event Manager when an Update for an Inferred or Discovered Host Occurs on next page.

To update agentless hosts:

1. Select the checkboxes for the hosts you want to update.
2. Click **Update**.

The Hosts tab displays the following information about the agentless hosts it inferred:

- **Host Name** – The name of the host.
- **Host Type** – HP Storage Essentials displays two host types:
- **Inferred** – An agentless host that was inferred through an agentless rule.
- **Discovered** – An agentless host that was given a generic custom name, as described in Assigning a Custom Name on page 500.

- **Rule Name** – The name of the rule that was used to infer the agentless host. This column is empty for custom-named generic hosts because they are not inferred by any rule.
- **Rule Scope** – The type of elements the rule used to find the inferred host
- **Host Security Group** – HP Storage Essentials searches the host security group names on the storage systems for hosts. You must have storage systems discovered through Discovery Step 3.
- **Zone** – HP Storage Essentials searches the zone name for hosts on the switches. You must have switches discovered through Discovery Step 3.
- **Zone Alias** – HP Storage Essentials searches the zone alias name for hosts on the switches. You must have switches discovered through Discovery Step 3.
- This column is empty for custom-named generic hosts.

Events Displayed in Event Manager when an Update for an Inferred or Discovered Host Occurs

The following example shows events in Event Manager when an update for an inferred or discovered host occurs.

Installing a CIM Extension on an Inferred Host

Install a CIM extension on an inferred host to obtain additional information about the applications installed on that host, local drive information, and the devices connected to its HBA ports.

The following occurs when you install a CIM extension on an inferred host:

- The host appears twice in ElementManager after Discovery Step 1 but before Discovery Step 3. The redundant host disappears once all the HBA ports are discovered through the CIM extension during Discovery Step 3.
- The host is identified by its DNS name after you install the CIM extension on it and complete Discovery Step 1 and 3. The HBA ports that remain inferred are those that are not discovered by the CIM extension. If you have an inferred host with a CIM extension and WWNs after Discovery Step 3, verify that your zoning and host group information is correct. The remaining WWN could belong to belong to a different host and orphan zone or an orphan host security group. Possibly, an orphan zone/host security group/zone alias existed, or the HBA was there

in the past and replaced with a new one and the outdated zone/host security group information was not removed. When the host is discovered with a CIM extension, it can leave the inferred host entry with the piece that was not resolved.

5 Host and Application Clustering

This section contains the following topics:

About Clustering

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

- Clusters are recognized as managed elements.
- System Manager supports clusters in all areas.
- The element topology shows which shared resources an application instance uses.
- Cluster capacity utilization is accurately reported.
- Capacity utilization trending is provided for applications running on clusters.
- The management server supports automatic discovery of several popular cluster servers, and allows management of other clusters through Cluster Manager.

Discovering Clusters

The following cluster services support automatic discovery:

- HP Serviceguard Cluster on HP-UX
- IBM High Availability Cluster Multi-Processing (HACMP) on IBM AIX
- Microsoft Cluster Services (MSCS) on Windows 2003 and 2008
- Oracle Clusterware Cluster on Red Hat Enterprise Linux 5
- Oracle Clusterware Cluster on Red Hat Enterprise Linux 4
- Veritas Clusters on HP-UX and Solaris
- VMware Clusters

Cluster services that do not support automatic discovery can be discovered manually using Cluster Manager. See Manual Discovery of Host Clusters on page 280.

The following application clusters are supported:

- Oracle Real Application Clusters (RAC)
- Microsoft Exchange 2003 FailOver Clusters and 2007 Single Copy Cluster (SCC)
- Microsoft Exchange 2007 Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR)
- Microsoft SQL Server 2000, 2005 and 2008
- Oracle FailOver Clusters

The LCR mechanism uses a single exchange server to replicate a copy of the storage groups. The CCR mechanism replicates the database and transaction logs for each storage group from an active node to a passive node.

For information about discovering application clusters, see Discovering Applications, Backup Hosts, and Hosts on page 179.

For a complete list of supported configurations, see the support matrix for your edition. The support matrix is accessible from the Documentation Center (**Help** > **Documentation Center**).

Known Issues with Host Clustering

- Unmounted Volume Capacity is not updated for Linux CRS Clusters when Storage is Added. When new storage is added to a Linux CRS cluster and is not yet mounted or is being used as raw storage, the capacity for that storage is not reported.
- Cluster Host Instances are not Removed from the System Topology After Being Deleted from Discovery. After you discover a cluster in the management server, deleting the cluster hosts from the management server discovery will not remove all instances of the clustered elements from the System Topology. Select and delete the unwanted elements manually from the System Topology screen to clear them.
- Total Capacity Summary Chart shows only Local File Systems for a Cluster Node (SE-3459). In the Capacity Chart tab for a clustered host, the Total Capacity Summary chart data for Total, Used, and Free includes only the local file systems.
- Manual Cluster Builder does not Support Volume Manager Volumes. At this time the Manual Cluster Builder will allow the selection of volume manager volumes, but the cluster is built using logical disks or disk partitions. Support of volume manager volumes in a manually built cluster is not available at this time.
- Automatically Detected Clusters and Shared Resources. Although the management server can detect the shared resources in a cluster, there might be inconsistencies in what is presented by the management server. If the cluster configuration detected and shown by the management server is incorrect, delete the partially detected cluster and use the manual cluster builder to assemble the cluster within the management server user interface.
- Adding Application and File Servers to a Cluster, Removing from a Cluster can Lead to Double-Counting. Moving an existing application server or file server for which the management server has already collected data into a cluster can result in double-counting of information and the loss of all history information associated with the host and applications on the host. When you move a host into a cluster, it is best to remove all applications and file

information from that host first. It might be simplest to delete the host from the management server user interface and rediscover the host alone with none of its applications or file data, and then add the host to the cluster. From that point, you can begin to discover the applications and file data on the newly clustered host. Removing a host

- Storage Marked as Shared Through Cluster Builder Not Shown as Remote on Property Page. On the shared logical drive property page the Remote Storage property is set to “false” for external drives marked as shared through the Cluster Builder. Shared drives in discovered clusters report the property correctly.
- Use the Collectors Tab for a Host Cluster to Start/Stop Report Data Collectors. The general Report Data Collectors Tab “Action” button has an incorrect status for host clusters. Use the host cluster's Collectors Tab to start and stop Report Data Collectors.
- Cluster Shared Resource Information for Disk Partitions and Disk Drives Limited for Manually Built Cluster. In a manually built cluster the information shown for shared disk partitions and disk drives is limited: the Cluster Host Capacity shows as zero; on the Cluster Host Properties page the Shared Resource Volume is blank; on the File Server Scan Page the Cluster Host Volume is blank, meaning it is not possible to do a File Scan on these types of resources.
- Cluster Builder Allows Selection of Logical Disk Elements. The Cluster Builder feature allows the selection of logical disk configurations such as Volume Manager volumes when building a cluster. Limit your Cluster Builder Shared Resource selections to the lowest level disk elements available in the list.
- Always Specify at Least One Shared Resource when Building a Cluster. When you build a cluster, always specify at least one Shared Resource in Step 3. If you do not, the built cluster will not appear as clustered in the topology.
- Deleting a Cluster Deletes Custom Commands on each Element in the Cluster. If you delete a manually built or discovered cluster from the management server, any user-defined custom commands will be removed from the elements the cluster comprises. These custom commands will need to be added back to the elements manually.

Automatic Discovery of Host Clusters

The following configurations support automatic discovery:

- HP ServiceGuard Cluster on HP-UX
- IBM High Availability Cluster Multi-Processing (HACMP)
- MSCS on Windows 2003 and 2008
- NetApp Clusters
- Oracle Clusterware Cluster on Red Hat Enterprise Linux 5
- Oracle Clusterware Cluster on Red Hat Enterprise Linux 4
- Veritas Clusters on HP-UX and Solaris
- VMware Clusters

Keep in mind the following:

- Additional steps are required for HACMP. Follow the steps in Requirements for Discovering IBM High Availability Cluster Multi-Processing on the facing page and Discovering HACMP Clusters on the facing page.
- NetApp devices do not share resources between cluster nodes.
- To enable automatic discovery of Oracle Cluster Ready Services (CRS) clusters on RHEL 5.5 when the /etc/init.d/init.crsd file has been deleted and the CRS service has been started using a custom script, set the ORACLE_CRS_HOME parameter in the cim.extension.parameters file so it points to the directory where the Cluster Ready Services were installed.
- VMware clusters must be discovered via the virtual center. If a cluster node is discovered separately using ESX server credentials, this node will not be shown as part of the cluster.
- On HACMP, a resource group should be configured for concurrent volume groups for HP Storage Essentials to show application-cluster topology and host-cluster shared resources and topology.
- For automatic discovery of Oracle Cluster Ready Services (CRS) clusters on RHEL 4 and RHEL 5, do one of the following:
  - Enable Oracle autoscan. See Optional – Enable Autoscan on page 208.

    Or
  - Provide the Oracle RAC details for Oracle RAC discovery in the Application Setup page, see Discovering Oracle Real Application Clusters (RAC) on page 214.

To discover hosts using any of these cluster services:

1. Discover your hosts as described in Discovering Applications, Backup Hosts, and Hosts on page 179. The clusters are automatically recognized by the management server.
2. The following optional steps enable you to select a preferred host from which shared resource capacity data will be collected.
   a. (*Optional*) Access Cluster Manager by right-clicking a cluster in System Manager and selecting Edit Cluster. The Cluster Manager Overview page is displayed.
   b. Click **Next**.
   c. (*Optional*) Cluster Manager Step 2 (Select Preferred Host for Cluster Shared Resources) is displayed. Select a preferred host for each of the cluster shared resources. Keeping the default selection of “None” will result in shared resource capacity data being collected from an available active host that shares the resource. Choosing a particular active host results in the specified host being used for data collection. If the specified host becomes unavailable, an available active host is used for data collection.
   d. Specify the preferred host for individual cluster shared resources. If a resource is not shared by the preferred host selection, the preferred host menu for that shared resource will continue to display the previous selection.
   e. When you finish specifying preferred hosts, click **Finish**.

Requirements for Discovering IBM High Availability Cluster Multi-Processing

You must set up the following before you can discover IBM High Availability Cluster Multi-Processing (HACMP):

- A CIM extension installed on every node.
- `bos.net.tcp.client`
- `Cldump`

Step 1 – Install a CIM Extension on Each Node of the Cluster

Install a CIM extension on each node of the cluster. Make sure that the CIM extension started.

Step 2 – Verify that the bos.net.tcp.client Package Meets the Version Requirement

Make sure that the `bos.net.tcp.client` package meets the version requirement according to the latest support matrix; otherwise, you will run into network issues with the host. If the bos.net.tcp.client package version requirement is not met, the discovery of HACMP methods for each node will be skipped. The nodes will be treated like a non-clustered AIX host.

Step 3 – Verify that Cldump Works Correctly

Make sure that the following commands work in each node of the clusters. The outputs from these commands should not be blank or contain any errors.

```
/usr/es/sbin/cluster/utilities/cldump
```

```
/usr/es/sbin/cluster/sbin/cl_lsvg
```

With earlier versions of AIX 6.1, cldump did not work unless the `/etc/snmpdv3.conf` file was modified. Check with the system administrators to make sure cldump works before proceeding.

Preferably for first time installations, make sure the cluster is in STABLE state from the cldump commands.

Discovering HACMP Clusters

HACMP supports two main methods of IP address tracking:

- **IP Alias.** Add the service IP address as an alias on a network interface in addition to the base IP address. This configuration is the default for HACMP 5.1 and later.
- **IP Replacement.** Replace the base (boot-time) IP address of an interface with the service IP address.

In both cases, there are individual node IPs and a cluster IP.

HP Storage Essentials supports the following types of discovery with HACMP:

- **Discovery via IP Alias**. Perform a Discovery Step 1 for all the nodes that have individual IP addresses that reside on the same subnet as the cluster IP. You do not need to discover the cluster IP. Then, perform a Discovery Step 3. There are no changes after failovers.
- **Discovery via IP Replacement where node IP is replaced**. On the node managing the cluster resources, that node's IP is replaced by the cluster IP. Perform a Discovery Step 1 of all the node IPs and cluster IP. Then, perform a Discovery Step 3.

  After any SAN file system failovers, the HACMP cluster resources are available in the other nodes. If you redo Discovery Step 3, the original node that was failed over is displayed as "missing." To avoid this, redo Discovery Step 1 for the cluster IP and the node IP that was previously not available and then redo Discovery Step 3.
- **Discovery via IP Replacement where there is a static NIC and IP**. When there is a network interface card or IP that will be static on the nodes regardless of the failover circumstances, it is best to discover the nodes via these interfaces.

Scenarios for Discovering HACMP Clusters

When discovering HACMP cluster nodes, choose the scenario that best fits your environment.

The following scenarios assume that `service_app.hpexample.com` is the (Service IP/Cluster IP) that is being failed over between the nodes. En is used in the typical AIX network interface.

Scenario 1: Discovery Through an IP Alias

Assume that Node_a and Node_b are always reachable through their fully qualified domain names (FQDN). Therefore, for discovery, the FQDN of the nodes should be used. In the following table, notice how `En0: Service_app.hpexample.com` (Service IP) is assigned to Node_a before the failover but to Node_b after the failover. Once En0: Service_app.hpexample.com (Service IP) is assigned to another node (Node_b), discovery Step 3 should be performed for Node_a and Node_b after a failover so thatHP Storage Essentials is aware of the new configuration.

Configuration Before and After a Failover (Scenario 1)

Before Failover	After Failover to Other Node

Node_a:
`En0: Node_a.hpexample.com`
`En0: Service_app.hpexample.com (Service IP)`
`En1: Heartbeat_a`
	Node_a:
	`En0: Node_a.hpexample.com`
	`En1: Heartbeat_a`

Before Failover	After Failover to Other Node
Node_b: En0: Node_b.hpexample.com En1: Heartbeat_b	Node_b: En0: Node_b.hpexample.com En0: Service_app.hpexample.com (Service IP) En1: Heartbeat_b

Initial Discovery Steps

To discover the nodes:

1. Perform discovery Step 1 to discover Node_a and Node_b (**Discovery** > **Setup**).
2. Perform discovery Step 3 (Get Details) to gather details for Node_a and Node_b (**Discovery** > **Details**).

After a Failover

You should always perform a discovery Step 3 (Get Details) for Node_a and Node_b after a failover so thatHP Storage Essentials is aware of the new configuration.

Scenario 2: IP Replacement Where the Main Interface Is Replaced at Startup

In this mode, the service IP is always reachable through the FQDN. However, one of the node's main interfaces is being replaced by the Service IP; therefore, the node is not reachable through its FQDN.

In the following table, notice how `En0: -` is assigned to Node_a before the failover but to Node_b after the failover. Once En0: - is assigned to another node (Node_b), discovery Steps 1 and 3 should be performed as described in "Discovery Steps After a Failover" after a failover so thatHP Storage Essentials is aware of the new configuration.

Configuration Before and After a Failover (Scenario 2)

Before Failover	After Failover to Other Node
Node_a: En0: - En0: Service_app.hpexample.com (Service IP) En1: Heartbeat_a	Node_a: En0: Node_a.hpexample.com En1: Heartbeat_a
Node_b: En0: Node_b.hpexample.com En1: Heartbeat_b	Node_b: En0: - En0: Service_app.hpexample.com (Service IP) En1: Heartbeat_b

Instead of trying to remember which node is the active node for Step 1 discovery, discover the FQDN for all the nodes and the service IP that replaces the main interface on a node. The node for which the main interface was replaced will be discovered automatically through the service IP and not through its FQDN.

Initial Discovery Steps

To discover the nodes:

1. Perform discovery Step 1 to discover Node_a and Node_b, in addition to `Service_app.hpexample.com` (**Discovery** > **Setup**).
2. Perform discovery Step 3 (Get Details) to gather details for Node_b and `Service_app.hpexample.com` (**Discovery** > **Details**).

Discovery Steps After a Failover

After a failover, HP Storage Essentials needs to be made aware of the new configuration. To discover the new configuration:

1. Perform discovery Step 1 to discover Node_a and Node_b, in addition to `Service_app.hpexample.com` (**Discovery** > **Setup**).
2. Perform discovery Step 3 (Get Details) to gather details for `Service_app.hpexample.com` and Node_a (**Discovery** > **Details**).

Scenario 3: IP Replacement Where the Main Interface Is Never Replaced and Instead Another Available Interface Is Replaced

In this mode, the Service IP is always reachable through the FQDN. One of the node's main interfaces is being replaced by the Service IP. However, each node has an extra interface (En2) that never changes. You can discover it as you did with Scenario 2. HP recommends that you follow this simpler method because it does not require a redo of discovery Step 1 after failovers.

In this mode, Node_a and Node_b are always reachable through their FQDNs. Therefore, for discovery, the FQDN of the nodes should be used. This mode does not require a redo of Step 1 post failover.

Notice in the following table how `En2: Service_app.hpexample.com` (Service IP) is moved from Node_a to Node_b during the failover and `En2: Node_b_temp.hpexample.com` is moved from Node_b to Node_a.

Configuration Before and After a Failover (Scenario 3)

Before Failover	After Failover to Other Node

Node_a:
`En0: Node_a.hpexample.com`
`En1: Heartbeat_a`
`En2: Service_app.hpexample.com (Service IP)`
	Node_a:
	`En0: Node_a.hpexample.com`
	`En1: Heartbeat_a`
	`En2: Node_a_temp.hpexample.com`

Before Failover	After Failover to Other Node
Node_b: En0: Node_b.hpexample.com En1: Heartbeat_b En2: Node_b_temp.hpexample.com	Node_b: En0: Node_b.hpexample.com En1: Heartbeat_b En2: Service_app.hpexample.com (Service IP)

Initial Discovery Steps

To discover the nodes:

1. Perform discovery Step 1 for Node_a and Node_b (**Discovery** > **Setup**).
2. Perform discovery Step 3 (Get Details) for Node_a and Node_b (**Discovery** > **Details**).

Discovery Steps After Failover

After a failover, perform a discovery Step 3 (Get Details) for Node_a and Node_b (**Discovery** > **Details**).

Scenario 4: IP Replacement Where the Main Interface Is Replaced and an Extra Network Interface Is Always Available

In this mode, the Service IP is always reachable through the FQDN. One of the node's main interface is being replaced by the Service IP. However, each node has an extra interface (En2) that never changes.

Configuration Before and After a Failover (Scenario 4)

Before Failover	After Failover to Other Node
Node_a: En0: - En0: Service_app.hpexample.com (Service IP) En1: Heartbeat_a En2: Node_a_perm.hpexample.com	Node_a: En0: Node_a.hpexample.com En1: Heartbeat_a En2: Node_a_perm.hpexample.com
Node_b: En0: Node_b.hpexample.com En1: Heartbeat_b En2: Node_b_perm.hpexample.com	Node_b: En0: - En0: Service_app.hpexample.com (Service IP) En1: Heartbeat_b En2: Node_b_perm.hpexample.com

Initial Discovery Steps

To discover the cluster:

1. Perform discovery Step 1 for Node_a_perm.hpexample.com and Node_b_perm.hpexample.com (**Discovery** > **Setup**).
2. Perform discovery Step 3 (Get Details) for Node_a_perm.hpexample.com and Node_b_perm.hpexample.com (**Discovery** > **Details**).

Discovery Steps After a Failover

After a failover, you must perform discovery Step 3 (Get Details) for Node_a_perm.hpexample.com and Node_b_perm.hpexample.com.

Scenario 5: IP Replacement Where Interfaces Fail Over in Multiple Steps

In this mode, the Service IP is always reachable through the FQDN. The node's main interface is being replaced by the Service IP. It fails over within the same node before failing over to the other node.

Configuration Before and After First Failover to Same Node (Scenario 5)

Before Failover	After First Failover to Same Node

Node_a:
`En0: -`
`En0: Service_app.hpexample.com (Service IP)`
`En1: Node_a2.hpexample.com`
`En2: Heartbeat_a`
	Node_a:
	`En0: Node_a1.hpexample.com`
	`En1: -`
	`En1: Service_app.hpexample.com (Service IP)`
	`En2: Heartbeat_a`

Node_b:
`En0: Node_b1.hpexample.com`
`En1: Node_b2.hpexample.com`
`En2: Heartbeat_b`
	Node_b:
	`En0: Node_b1.hpexample.com`
	`En1: Node_b2.hpexample.com`
	`En2: Heartbeat_b`

Initial Discovery Steps

To discover the cluster:

1. Perform a discovery Step 1 for Service_app.hpexample.com and Node_b2.hpexample.com (**Discovery** > **Setup**).
2. Perform a discovery Step 3 (Get Details) for Service_app.hpexample.com and Node_b2.hpexample.com (**Discovery** > **Details**).

Discovery Steps After First Failover to the Same Node

You must perform a discovery Step 3 (Get Details) for Service_app.hpexample.com and Node_b2.hpexample.com after the first failover to the same node (**Discovery** > **Details**).

Configuration Before and After Final Failover to Same Node (Scenario 5)

Second Failover to Other Node	Final Failover to Same Node

Node_a:
`En0: Node_a1.hpexample.com`
`En1: Node_a2.hpexample.com`
`En2: Heartbeat_a`
	Node_a:
	`En0: Node_a1.hpexample.com`
	`En1: Node_a2.hpexample.com`
	`En2: Heartbeat_a`

Node_b:
`En0: -`
`En0: Service_app.hpexample.com (Service IP)`
`En1: Node_b2.hpexample.com`
`En2: Heartbeat_b`
	Node_b:
	`En0: Node_b1.hpexample.com`
	`En1: -`
	`En1: Service_app.hpexample.com (Service IP)`
	`En2: Heartbeat_b`

Discovery Steps After Second Failover to Other Node

To discover the cluster after the second failover:

1. Perform a discovery Step 1 for Service_app.hpexample.com and Node_a2.hpexample.com (**Discovery** > **Setup**).
2. Perform a discovery Step 3 (Get Details) for Service_app.hpexample.com and Node_a2.hpexample.com (**Discovery** > **Details**).

Discovery Steps After Final Failover to the Other Node

After the final failover, perform a discovery Step 3 (Get Details) for Service_app.hpexample.com and Node_a2.hpexample.com (**Discovery** > **Details**).

Scenario 6: IP Alias Concurrent for Oracle and Other Databases

In this mode, Node_a and Node_b are always reachable through their FQDNs. All the database clustered resources are available at all times. Therefore, for discovery, the FQDN of the nodes should be used.

Configuration Before and After Failover (Scenario 6)

Before Failover	After Failover to Other Node
Node_a: En0: Node_a.hpexample.com En0: Service_app.hpexample.com (Service IP) En1: Heartbeat_a	Node_a: En0: Node_a.hpexample.com En1: Heartbeat_a
Node_b: En0: Node_b.hpexample.com En1: Heartbeat_b	Node_b: En0: Node_b.hpexample.com En0: Service_app.hpexample.com (Service IP) En1: Heartbeat_b

Initial Discovery

To discover the cluster before a failover:

1. Perform a discovery Step 1 for Node_a and Node_b (**Discovery** > **Setup**).
2. Perform a discovery Step 3 (Get Details) for Node_a and Node_b (**Discovery** > **Details**).

Scenario 7: Stacked IP with IP Aliases

In this mode, Node_a and Node_b are always reachable through their FQDNs. All the database clustered resources are available at all times. But each interface is stacked with multiple IPs.

Configuration Before and After Failover (Scenario 7)

Before Failover	After Failover to Other Node
Node_a: En0: Node_a1.hpexample.com Node_a2.hpexample.com Node_a3.hpexample.com Node_a4.hpexample.com En0: Service_app.hpexample.com (Service IP) En1: Heartbeat_a	Node_a: En0: Node_a1.hpexample.com Node_a2.hpexample.com Node_a3.hpexample.com Node_a4.hpexample.com En1: Heartbeat_a

Before Failover	After Failover to Other Node
Node_b: En0: Node_b1.hpexample.com Node_b2.hpexample.com Node_b3.hpexample.com Node_b4.hpexample.com En1: Heartbeat_a	Node_b: En0: Node_b1.hpexample.com Node_b2.hpexample.com Node_b3.hpexample.com Node_b4.hpexample.com En0: Service_app.hpexample.com (Service IP) En1: Heartbeat_a

Parameters to Control Host Agent Behavior for HACMP Cluster Nodes

The following parameters can be modified to change host agent behavior for HACMP Cluster nodes. Do not modify these parameters unless discovery problems exist.

socket.poll.interval Parameter

The `socket.poll.interval` parameter controls the time interval at which the host agent monitors changes in the IP address of the cluster node for IP replacement configuration. Do not modify this setting unless discovery problems exist.

To change this parameter:

1. If you do not already have the wrapper.user file, copy `wrapper.user-sample` to wrapper.user. If it was created, it can be found in the `/opt/APPQcime/conf` directory.
2. Open the wrapper.user file in a text editor such as Notepad.
3. If the socket.poll.interval parameter does not already exist in the file, add it to the file.
4. Specify the value in seconds for the socket.poll.interval parameter; for example:

   ```
   socket.poll.interval=50
   ```

   The default value is 30 seconds.
5. To turn off polling, set the parameter to 0.

hacmp.stabilization.interval Parameter

The `hacmp.stabilization.interval` parameter controls the time interval for which the host agent waits before restarting itself if the IP addresses configured on the cluster node change due to failover. This parameter is applicable only for IP Replacement configuration. Do not modify this setting unless discovery problems exist.

To change the hacmp.stabilization.interval parameter:

1. If you do not already have the wrapper.user file, copy `wrapper.user-sample` to wrapper.user. If it was created, it can be found in the `/opt/APPQcime/conf` directory.
2. Open the wrapper.user file in a text editor, such as Notepad.
3. If the hacmp.stabilization.interval parameter does not already exist in the file, add it to the file.
4. Specify the value in seconds for the hacmp.stabilization.interval parameter; for example:

   ```
   hacmp.stabilization.interval=150
   ```

   The default value is 120 seconds.

Manual Discovery of Host Clusters

If you are using a cluster service that does not support automatic discovery, you must manually create your clusters. For the list of cluster services that support automatic discovery, see Discovering Clusters on page 267.

To manually discover clusters:

1. Discover your hosts and applications as described in Discovering Applications, Backup Hosts, and Hosts on page 179.
2. Access Cluster Manager by right-clicking a host in System Manager and selecting **Build Cluster**. The Cluster Manager Overview page is displayed.
3. Click **Next**. Cluster Manager Step 2 (Specify Cluster Properties and Cluster Members) is displayed.

To specify the cluster properties and cluster members:

1. In the Cluster Properties section, specify the cluster name, cluster server type, and cluster virtual IP (if applicable).
2. In the Available Hosts section, select the hosts to add to the Cluster Members table. To use the filter to select the hosts, see Filtering Hosts on the facing page.
3. You can also use the Select Related Hosts button. Select a host in the table, and click **Select Related Hosts** to automatically select any related hosts.
4. After you select the hosts to add to the cluster, click **Add Selected Hosts to Cluster**. The selected hosts are added to the Cluster Members table.
5. Click **Next**.

   Cluster Manager Step 3 (Specify Cluster Shared Resources) is displayed.
6. Select **Automatic** or **Manual**.

If you select Automatic discovery:

1. Click **Display Cluster Shared Resources**. The table at the bottom of the page is automatically populated.
2. Click the **Edit** button for the first Cluster Shared Resource.

3. By default, only one node cluster node is specified. Specify the second node by unchecking the **None** checkbox, and selecting the correct resource from the drop-down menu.
4. Click **OK**.
5. Repeat these steps for each Cluster Shared Resource.

If you are building a DRS cluster for ESX Servers, only specify cluster shared resources for Shared Logical Disks. For Shared Volume Manager Volumes, set both of the nodes to None. This does not need to be done manually when ESX servers are discovered via the same Virtual Center. Automatic discovery will occur after the next Get Details.

If you select Manual discovery:

1. Enter a name in the Cluster Shared Resource Name box.
2. Select a resource type from the Resource Type menu. The menu includes the following resource types:
   - Logical Disk
   - Disk Partition
   - Volume Manager Volume
   - Disk Drive
3. If you are building a DRS cluster for ESX Servers, select **Logical Disk**. Selecting **Volume Manager Volume** results in problems with the cluster topology.
4. Select the relevant resource for each cluster host, and click **Save Selections as Cluster Shared Resource**. The selections are added to the Cluster Shared Resources table.
5. Repeat steps 1, 2 and 3 for each shared resource in the cluster.
6. Click **Next**.
7. Cluster Manager Step 4 (Select Preferred Hosts for Cluster Shared Resources) is displayed. Select a preferred host for each of the cluster shared resources. Shared resource capacity data will be collected from the specified node. Selecting “None” will result in no information being collected about the cluster shared resource.
8. Specify the preferred host for individual cluster shared resources. If a resource is not shared by the preferred host selection, the preferred host menu for that shared resource will continue to display the previous selection.
9. When you finish specifying preferred hosts, click **Finish**.

Once the manual discovery of a host cluster is done, you can discover applications on it as described in Discovering Applications, Backup Hosts, and Hosts on page 179.

Filtering Hosts

The Available Hosts table on Cluster Manager Step 2 (Specify Cluster Properties and Cluster Members) enables you to filter the list of hosts displayed.

To filter the list of hosts:

1. Click the **+ Filter** link to display the filtering options.
   If the volume filter is already displayed, the **– Filter** link is shown instead, which will collapse the filtering options.
2. Enter all or part of a volume name in the Name Contains box.
3. Select an operating system from the Operating System menu.
4. Enter all or part of a vendor name in the Vendor Contains box.
5. Enter a number in the Processors (>=) box.

   Hosts with at least as many processors as specified are displayed in the table.
6. Enter a number in the HBAs (>=) box.

   Hosts with at least as many HBAs as specified are displayed in the table.
7. Enter a number in the Ports (>=) box.

   Hosts with at least as many ports as specified are displayed in the table.
8. Click **Filter**.

   The table is updated to display only the elements that meet the filter criteria.
9. To reset the filter criteria, click **Reset**.

File Servers and Clusters

If you marked a host as a file server and move it into or out of a cluster, you must remove the file server data from the host and then re-mark it as a file server.

To remove the file server data from the host and re-mark it as a file server:

1. Select **Configuration** > **File System Viewer.**
2. Verify that the **File Servers** tab is displayed.
3. Select the file servers you want to remove, and then click **Delete**.
4. Click **Add File Server**.
5. Click the check boxes for the hosts you want marked as file servers.
6. Click **OK**. The hosts are marked as file servers, and you are returned to the **File Servers** tab.
7. Rescan the cluster member nodes and the cluster nodes, or incorrect data might be displayed.

Clustering in System Manager

System Manager seamlessly supports clusters in all areas. You can view connectivity information from all levels on a single canvas — from applications running on clusters, to the storage array spindles that share volumes for all the nodes of a cluster.

The following example shows how clusters are displayed in System Manager. The tree nodes on the List tab reflect the structure of the clusters.

The box on the left of the topology canvas shows a cluster with two hosts, and the box on the right shows a cluster with four hosts. Both clusters are in the expanded view mode, so all of the nodes are displayed. To minimize the view of a cluster, click the (-) button.

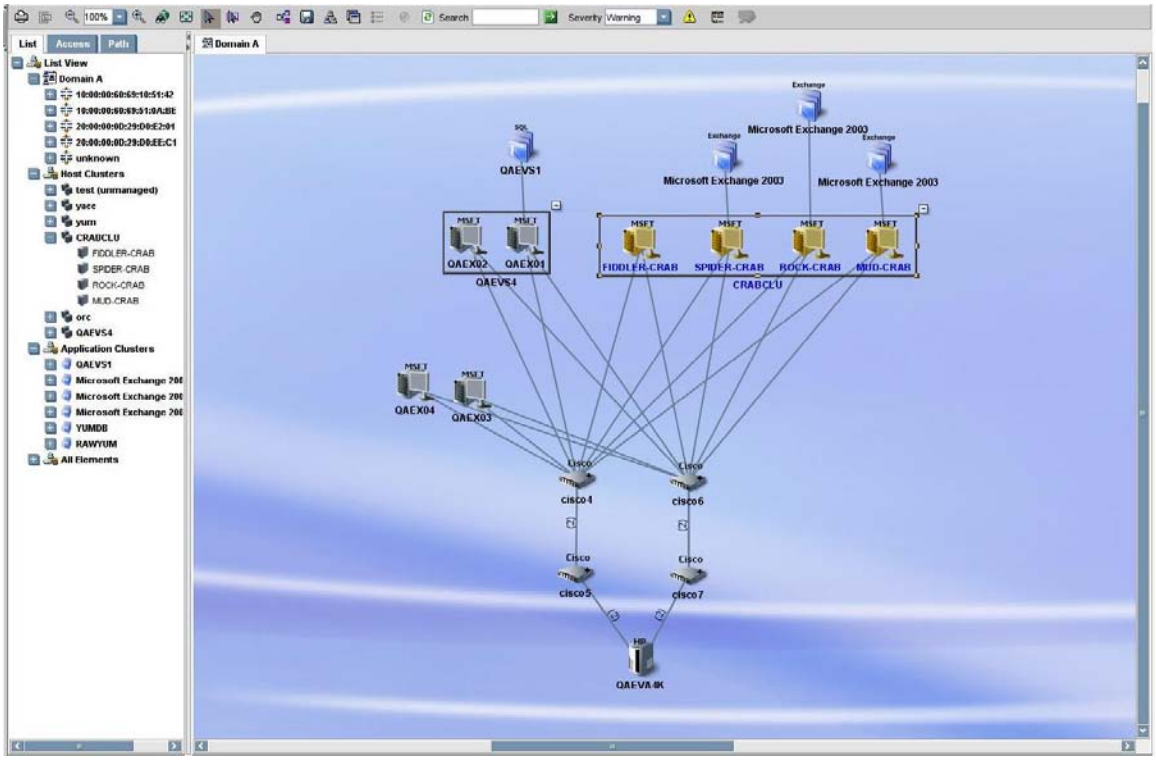

In the minimized view of a cluster, all of the nodes of the cluster are collapsed into a single box. To expand the display to show all of the nodes, click on the (+) button.

In the minimized view, a dotted line from an application to a cluster indicates that the application only runs on some of the clustered hosts. A solid line indicates that the application runs on all of the clustered hosts.

Double-click a cluster to open the Properties page for the cluster. Double-click an individual cluster node to open the Properties page for that node.

Clustering in Topology

Element topology expands System Manager's view to show exactly which shared resources a particular application instance uses. Individual paths from application nodes are listed in the path tree as well.

The following examples shows individual instances of Microsoft Exchange Server 2003 sharing HP EVA virtual disk array group shared resources.

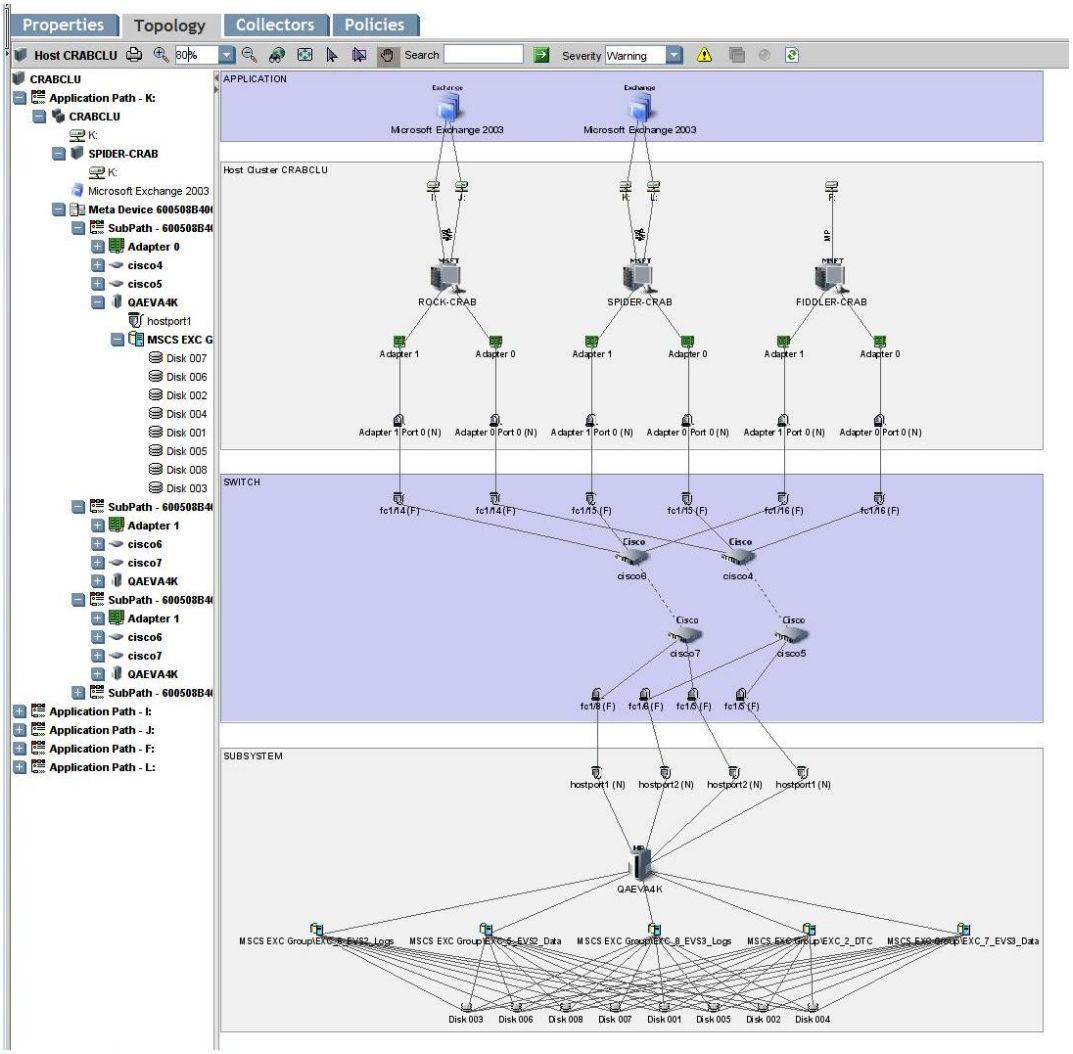

Clustering in Capacity Manager

Capacity Manager enables you to see the whole capacity utilization by the cluster. Clusters are represented as managed elements, and the capacity calculator intelligently avoids double counting of the capacity from individual nodes at the cluster level.

You can drill down to various levels to see the following details of cluster capacity utilization:

- Whole cluster capacity
- Individual application instance capacity
- Individual cluster node capacity
- Capacity trending over a period of time
- Shared resources of individual nodes

The following example shows how clusters are represented in Capacity Manager.

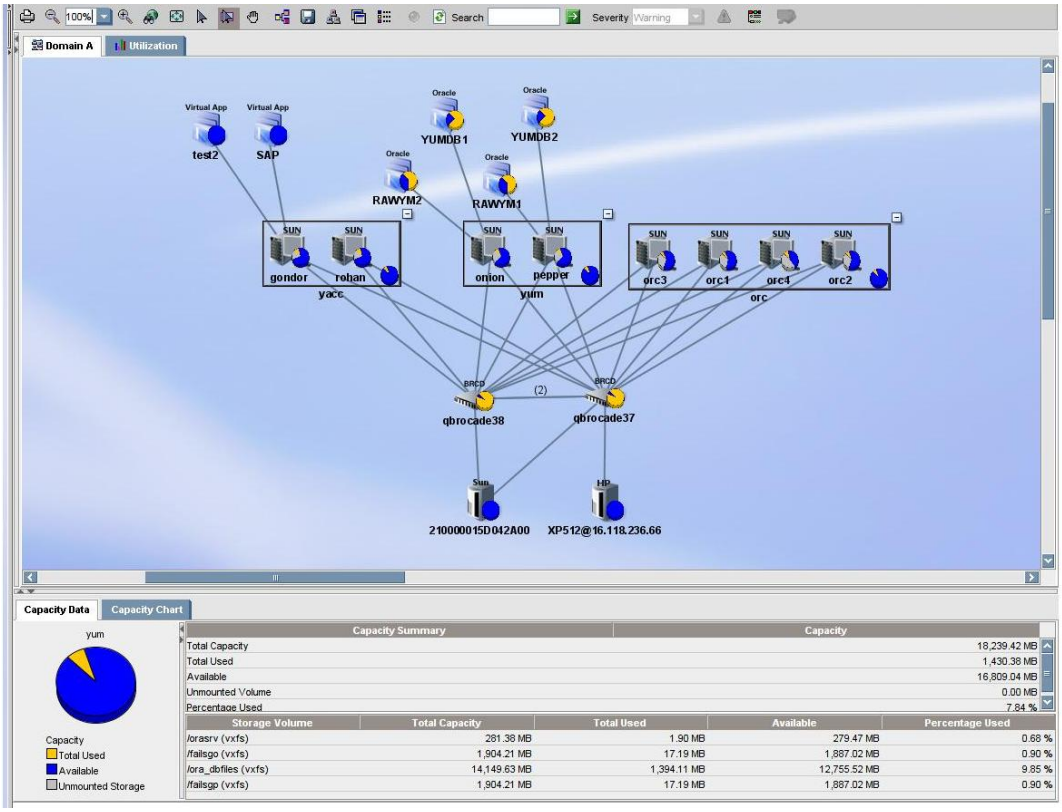
Search
Severity
Warning
Domain A
Utilization
Virtual App
test2
Virtual App
SAP
Oracle
YUMDB1
Oracle
YUMDB2
Oracle
RAWYM2
Oracle
RAWYM1
SUN
gondor
SUN
rohan
yacc
SUN
onion
SUN
pepper
yum
SUN
orc3
SUN
orc1
SUN
orc4
SUN
orc2
orc
BRCD
qbrocade38
(2)
BRCD
qbrocade37
Sun
210000015D042A00
HP
XP512@16.118.236.66
Capacity Data
Capacity Chart
yum
Capacity
Total Used
Available
Unmounted Storage
Capacity Summary
Capacity
Total Capacity 18,239.42 MB
Total Used 1,430.38 MB
Available 16,809.04 MB
Unmounted Volume 0.00 MB
Percentage Used 7.84 %
Storage Volume
Total Capacity
Total Used
Available
Percentage Used
/orasrv (vxfs) 281.38 MB 1.90 MB 279.47 MB 0.68 %
/failsgo (vxfs) 1,904.21 MB 17.19 MB 1,887.02 MB 0.90 %
/ora_dbfiles (vxfs) 14,149.63 MB 1,394.11 MB 12,755.52 MB 9.85 %
/failsgp (vxfs) 1,904.21 MB 17.19 MB 1,887.02 MB 0.90 %

6 Managing Security

Depending on your license, role-based security might not be available. See the List of Features to determine if you have access to role-based security. The list is accessible from the Documentation Center (**Help** > **Documentation Center**).

This section contains the following topics:

Security for the Management Server

The management server offers security that is based on the assignment of roles and organizations. Role-based security determines access to specific functionality according to the user account assigned to a role. Organization-based security determines if you can modify an element type, such as hosts. The management server ships with the Everything organization, which enables you to modify all element types.

See the following topics for more information:

About Roles

The management server ships with the following predefined roles. These roles determine which components of the software a user can access.

For example, users assigned to the Help Desk role have access to Application Viewer and Event Manager, but not to System Manager, Provisioning Manager, Backup Manager and Policy Manager. Likewise, users assigned to the domain administrator role have access to all of the features, as shown in the following table.

Default Role Privileges

Feature	CIO	Domain Administrator	Storage Administrator	Server Administrator	Application Administrator	Help Desk
Application Viewer	X	X			X	X
System Manager*	X	X	X	X	X	
Event Manager		X	X	X	X	X
Backup Manager	X	X	X	X	X	
Provisioning Manager		X	X			
Provisioning Administration		X	X			
Capacity Manager	X	X	X	X	X	
Policy Manager		X	X			
Chargeback Manager	X	X	X			
File System Viewer		X		X		
Performance Manager	X	X	X	X	X	
Access CLI		X	X			
Custom Commands		X	X			
System Configuration		X				

* Your account must belong to a role that has "System Manager" selected for you to be able to perform SAN zoning operations, such as creating zone aliases, zones, and zone sets.

Domain Administrator Role Privileges

Only users belonging to the Domain Administrators role can add, modify, and delete users, roles, and organizations. The Domain Administrator can only edit active organizations.

Domain Administrators can change the user names and roles of other domain administrators, but they cannot modify their own user name and roles while logged into the management server. Domain administrators can also edit their full name, e-mail, phone, and other details, as well assign and un-assign any organization.

System Configuration Option

If the System Configuration option is selected for a role, all users assigned to that role will have the following administration capabilities:

- Schedule discovery
- Find the CIM log level
- Save log files, e-mail log files
- Save the database, backup the database, and schedule a database backup
- Configure Event Manager, File System Viewer, and Performance Manager
- Configure reports and traps
- Set up the management server to send e-mail

If you do not want users belonging to that role to have those capabilities, do not assign the System Configuration option.

Roles Used to Restrict Access

Roles also restrict access to element properties, element records, and Provisioning Manager, as shown in the following table.

Default Role Privileges by Elements

Role	Application	Host	Switch	Storage System	Tape Library	Others
CIO	View	View	View	View	View	View
Domain Administrator	Full Control	Full Control	Full Control	Full Control	Full Control	Full Control
Storage Administrator	View	View	Full Control	Full Control	Full Control	Full Control
Server Administrator	View	Full Control	View	View	View	View
Application Administrator	Full Control	View	View	View	View	View
Help Desk	View	View	View	View	View	View

Options for Restricting a Role

You can assign one of the following options within a role to further allow or restrict access for a specific element:

- **Full Control** – Enables you view and modify the record for the element on the Asset Management tab, and perform provisioning if applicable.
- **Element Control** – Enables you view and modify the record for the element on the Asset Management tab. You cannot perform provisioning.
- **View** – Enables you only view element properties.

For example, if users belong to a role that only lets them view the element properties on storage systems, those users would not be allowed to perform provisioning on storage systems because their role does not have the Full Control option selected for storage systems. That same role could also have the Full Control option selected for switches, allowing the user to perform provisioning for switches. Thus, the user would not be able to provision storage systems, but would be able to provision switches.

You can modify roles and/or create new ones. For example, you can modify the Help Desk role so that the users assigned to this role can also view Provisioning Manager and modify servers.

About Organizations

You can use organizations to specify which elements users can access. For example, you can specify that some users have only access to certain switches and hosts. However, these users must already be assigned to roles that allow them to see switches and hosts.

Users assigned to an organization can see only the elements that belong to that organization. If users are assigned to more than one organization, they see all elements that belong to the organizations to which they are assigned. For example, assume you created two organizations: one called OnlyHosts that allowed access to only hosts and another called OnlySwitches that allowed access to only switches. A user assigned to OnlyHosts and OnlySwitches would have access to hosts and switches because those elements are listed in at least one of the organizations.

Organizations can also contain other organizations. An organization contained within another is called a child. The organization containing a child organization is called a parent. The figure below shows a parent-child hierarchy in which BostonWebHosts organization contains two child organizations, BostonWebHost_Windows and BostonWebHost_Solaris. BostonWebHosts is a parent because it contains two organizations.

Figure 1 Parent-Child Hierarchy for Organizations

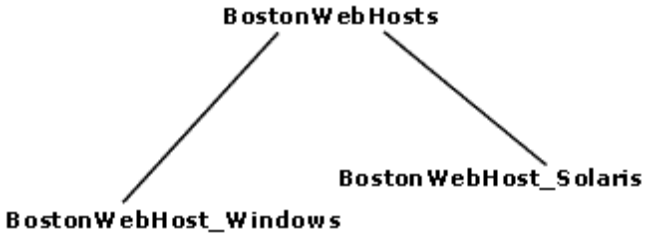

If a child contains organizations, it is also a parent. For example, if you add two organizations called BostonWebMarketing and BostonWebProduction to BostonWebHost_Windows. BostonWebHost_Windows would become a parent because it now contains two organizations. It would also be a child because it is contained in BostonWebHosts.

Parent organizations allow access to all elements listed in their child organizations. For example, users assigned to the organization BostonWebHosts can access not only the elements in BostonWebHost_Windows, but also those in BostonWebHost_Solaris. This is because BostonWebHosts is a parent of the two child organizations.

The parent-child hierarchy for organizations saves you time when you add new elements; for example, when you add a new element, you need to add it only once; the change ripples through the hierarchy. For example, if you add an element to BostonWebHost_Windows, not only users assigned to BostonWebHost_Windows would see this addition, but also users assigned to any of the parent organizations containing BostonWebHost_Windows. For example, users assigned to BostonWebHosts would also see the addition because it contains BostonWebHost_Windows; users assigned to only BostonWebHost_Solaris would not see the addition.

A child organization can be in multiple parent organizations. As shown in the following figure BostonWebHosts and NYWebHosts are not only children of the WebHosts organization, but they are also children of the US East Coast organization. For example, if you have a user that oversees all Web hosts in the company, you could assign that user to the WebHosts organization. Users managing hosts and storage systems on the East Coast would be assigned to the US East Coast organization, which is a parent of BostonWebHosts, NYWebHosts, and StorageSystems organizations. For example, if an element is added to NYWebHost_Solaris, users assigned to one or more of the following organizations would see the addition:

- NYWebHost_Solaris
- NYWebHosts
- WebHosts
- US East Coast
- Children in Multiple Organizations

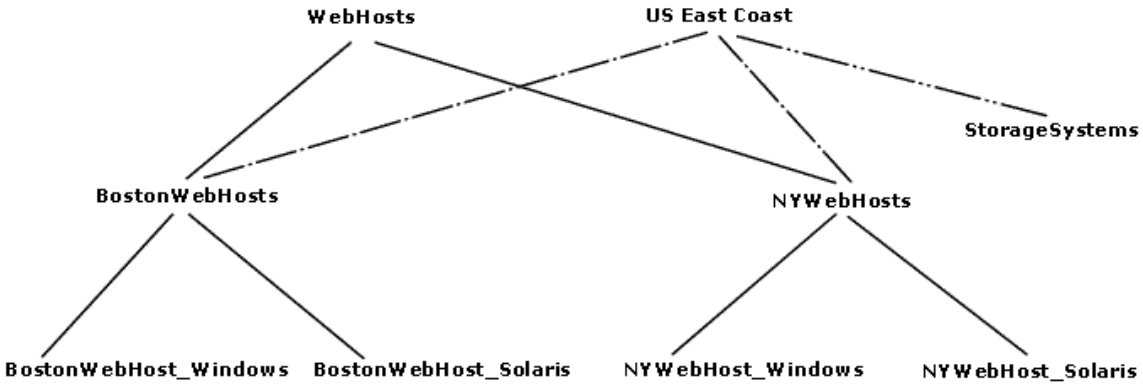

When you remove an element from an organization, users belonging to that organization or to one of its parents can no longer access that element if it is not a member of any other organization. For example, assume an element named MyHost was not only a member of BostonWebHost_ Solaris, but also had mistakenly became a member of BostonWebHost_Windows. If you remove MyHost from BostonWebHost_Solaris, users belonging to BostonWebHost_Solaris can no longer access the element. Users belonging to the following organizations would still see the element because the element is still a member of BostonWebHost_Windows.

- BostonWebHosts
- WebHosts
- US East Coast

Keep in mind the following:

- You cannot edit the Everything organization.
- A virtual machine cannot be moved to an organization that does not also contain its virtual server.
- Users can view all elements only in the Discovery pages. In all other pages, only the members of the active organization are available.
- Discovery lists (Discovery tab) are not filtered. Users can see all elements in the discovery lists regardless of their affiliation with an organization.
- Events from all elements regardless of the user's organization are displayed by Event Manager.

Planning Your Hierarchy

Before you begin creating organizations, plan your hierarchy. Do you want the hierarchy to be based on location, departments, hardware, software, or tasks? Or perhaps you want a combination of these options.

To help you with your task, create a table of users who manage elements on the network and the elements they must access to do their job. You might start seeing groups of users who oversee the same or similar elements. This table could help you in assigning users to the appropriate organizations.

Once you are done with planning your hierarchy, draw the hierarchy in a graphics illustration program, so you can keep track of which organizations are parents and children.

First create the child organizations and then their parents (see Adding an Organization on page 305).

Naming Organizations

When you create an organization, give it a name that reflects its members. You could use one or more of the following as a guideline:

- Type of elements that are members of the organization, such as switches, Sun Solaris hosts
- Location of the elements, such as San Jose
- Task, such as backup machines

You might find that it is easy to forget which containers are parents and which are children. When you name an organization, you could include a portion of the name of the dominant parent organization. For example, if you have two types of Web hosts in Boston, Microsoft Windows and Sun Solaris, you could name the two child organizations BostonWebHost_Windows and BostonWebHost_Solaris and their parent, BostonWebHosts.

About the SecurityProperties.properties File

The `SecurityProperties.properties` file contains several default properties. If this file is not present on your management server at **%MGR DRT%** > **Data** > **Configuration**, follow these steps:

1. Locate the sample file, `securityProperties.properties_sample`, rename it `securityProperties.properties`, and add it into the directory.
2. Restart the management server service.

Setting High-Strength SSL Cipher Suites

You can select a high-strength Secure Socket Layer (SSL) cipher suite by setting the useHighStrengthCipher property, Use this property to determine whether the SE CMS, as well as the CIM extensions, use high-strength (>112 bit length) ciphers during SSL communication.

If the useHighStrengthCipher property is set to "true", a high-strength cipher suite (>112-bit length) is used. If it is set to "false", the handshake process will use the strongest possible available cipher from a list of enabled ciphers (40-bit to 168-bit length).

By default, the useHighStrengthCipher property is set to "true", and therefore only >112-bit ciphers will be used until you change the property.

To set this property on the SE CMS, follow these steps:

1. In the custom properties window, set "useHighStrengthCipher=true" if you want only high-strength ciphers used. Set "useHighStrengthCipher=false" if you want the cipher selected from a list of enabled cipher suites that range from 40-bit to 168-bit length.
2. Restart the SE service.

To set this property on the CIM extensions, follow these steps:

1. Open the cim.extension.parameters file and set "useHighStrengthCipher=true" if you want only high-strength ciphers used. Set "useHighStrengthCipher=false" if you want the cipher selected from a list of enabled cipher suites that range from 40-bit to 168-bit length.
2. Restart the CIME.

Managing User Accounts

This section contains the following topics:

Adding Users

The following procedure explains how to add users and authorize privileges. You must belong to the Domain Administrator role to add or modify users.

Keep in mind the following:

- Windows – The user name and password must be alphanumeric and cannot exceed 256 characters. The user name cannot contain some special characters, see Using Active Directory/LDAP for Authentication on page 312 for more information.AD authentication for Windows LDAP server is not supported.
- Linux – The user name and password cannot exceed 256 characters.

To create an account:

1. Click **Security** > **Users**.
2. Click **New User** button.
3. Select a user type from the **User Type** list.

4. In the **Login Name** box, type a name for the user account; for example, jsmith.
   This name becomes the user name for the account.
5. (*Optional*) In the **Full Name** box, type a full name for the account.

   This information is used to provide a correlation between an account name and a user.

   The full name can contain spaces, but cannot be longer than 512 characters.

   Domain names and user names are case insensitive.
6. Assign the user account to a pre-existing role by selecting a role from the **Role** menu. Security for the Management Server on page 287 for more information about roles and organizations, including the parent-child hierarchy.
7. In the **Domain Controller Name** box, type the IP address or the fully qualified name of your primary Domain Controller server. You can also specify the secondary or additional controllers as a comma-separated list. This option is displayed only if you select the user type as Active Directory or LDAP. You should be able to ping the fully qualified name of the Domain Controller as well as its simplified name from the HP Storage Essentials management server.
8. In the **Distinguished Name** box, enter the distinguished name of the user; for example, `CN=NAME,CN=Users,DC=MyCompanyName,DC=Com`. This option is only applicable for LDAP users.
9. (*Optional*) In the **E-mail** box, enter the user's e-mail address.
10. (*Optional*) In the **Phone** box, enter the user's phone number.
11. (*Optional*) In the **Notes** box, provide additional information about the user.
12. (*Optional*) In the **Password** box, enter a password for the user account. This option is displayed only if you select the user type as Basic.
13. (*Optional*) In the **Verify Password** box, enter the password you entered previously.
14. Assign the user account to one or more organizations.

    The organizations determine which elements the user can manage. To assign a user account to an organization, select the organizations from the table.
15. Click **OK**.

Adding AD/LDAP Organizational Unit

The following procedure explains how to add AD or LDAP organizational unit details to the management server and assign a role to the organizational unit. You can also assign the organizational unit to one or more organizations.

Keep in mind the following:

- Any user belonging to AD/LDAP organization unit can log on to HP Storage Essentials management server using the appropriate password.

- You can add a nested organizational unit to the Management server. For adding, you must provide the entire hierarchy of the nested organizational unit in the AD/LDAP organizational unit box. For example, if there exists an organizational unit OU2 with a user say 'ouuser2' within an organizational unit OU1 with a user say 'ouuser1', you can add the nested organizational unit OU2 by entering OU1/OU2. In this case, only the user 'ouuser2' will be able to login into Management server.
- A user can be an individual user and can also be a part of an organizational unit added to the management server. In this case, the role and the organizational unit assigned to an individual user is applicable when the user logs in to the management server. For example, if there exists an organizational unit OU1 with a user say 'ouuser1'. Here, if the user 'ouuser1' is added to the management server through the Users page with its role set to 'Role1' and organization set to 'Org1' and if the organizational unit OU1 is added to the management server through the Organizational Unit page with its role set to 'Role2' and organization set to 'Org2', then when the 'ouuser1' logs in to the management server, it is logged in with role 'Role1' and organization 'Org1'.

To create an AD/LDAP organizational unit:

1. Click **Security** > **Users**.
2. Click **New AD/LDAP organization unit**.
3. In the **AD/LDAP organizational unit** box, type a name for the organization. This name must be present in the AD database.
4. Assign a pre-existing role to the organizational unit by selecting a role from the **Role** list. All users belonging to a specific organizational unit will have the same privileges as the organizational unit.
5. In the **Domain Controller Name** box, type the IP address or the fully qualified name of your Primary Domain Controller server to which the organizational unit belongs.
6. In the **OU Distinguished Name** box, type the distinguished name of the organizational unit. You must provide the distinguished name for an LDAP organizational unit.
7. Assign the organizational unit to one or more organizations. The organizations determine the elements that the users within the organizational unit can manage. To assign an organizational unit to an organization, select the organizations from the table.

Editing a User Account

Keep in mind the following:

- Only a user belonging to the Domain Administrator role is allowed to edit user accounts.
- The Admin account acts differently than the other accounts.
- You cannot add or remove organizations from the Admin account.
- You cannot remove the Everything organization from the Admin account.
- New organizations are automatically added to the Admin account when they are created.
- See Domain Administrator Role Privileges on page 288.

- User modifications take effect immediately even if the user is logged in to the management server.
- You cannot change the password for a user account that was authenticated against Active Directory/LDAP. To change the password for the user account, use Active Directory/LDAP. See Step 1 – Add Active Directory Users to the Management Server on page 313.
- A Super User can assign any other user belonging to the Domain Administrator role and Everything organization as a Super User. To be able to assign a user as the Super User, the user details must be present in theHP Storage Essentials database. The user must belong to the Domain Administrator role and to Everything organization.

  Everything organization is the default organization that enables users to access all current and future elements.
- Only a Super User can view the **Change Super User** tab.

To change your password, follow the steps in Changing Your Password on page 299.

To modify a user account:

1. Click **Security** > **Users**.
2. Click the **Edit** button () for the user account you want to modify.
3. To change the account name, enter a new name for the user account in the **Name** box; for example: jsmith. This name becomes the user name for the account. Domain names in user names must match the case of the domain name.
4. To change the name assigned to the user account, enter a new name for the account in the Full Name box. This provides a correlation between an account name and a user.
5. To change the role assigned to the user account, select a new role from the Role menu.
6. To change the e-mail address listed, enter a new e-mail address in the **E-mail** box.
7. To change the phone number listed, enter the user's new phone number in the **Phone** box.
8. Change or remove information from the **Notes** box if necessary.
9. To change the password:
   a. Select the Enabled option.
   b. Enter a new password in the **Password** box.
   c. Enter the password again In the **Verify Password** box.
   d. Click **OK**.
10. To change the organizations to which the user belongs, select or deselect the organizations from the table in the user interface.
11. Click **OK**. The user account is updated.

Editing a AD/LDAP Organizational Unit

To modify a AD/LDAP organizational unit:

1. Click **Security** > **User**.
2. Click the **Edit** button () for the AD/LDAP organizational unit you want to modify.
3. In the **AD/LDAP organizational unit**, type the new name for the organizational unit.
4. To change the role assigned to the organizational unit, select a new role from the **Role** list.
5. To change the Domain Controller Name, type the new IP address or the fully qualified name of your Primary Domain Controller server in **Domain Controller Name** box.
6. To change the distinguished name for an LDAP organizational unit, type the new distinguished name of the organizational unit in the **OU Distinguished Name** box.

   To change the organizations to which the AD/LDAP organizational unit belongs, select or deselect the organizations from the table.

   If you are logged on to the management server,you cannot modify the name and role of the organizational unit to which you belong.
7. Click **OK**. The AD/LDAP organizational unit is modified.

Assigning Super Users

Keep in mind the following:

- A Super User is any user who belongs to Domain Administrator role.
- A Super User can assign any other user belonging to the Domain Administrator role and everything organization as a Super User.
- To be able to assign a user as a Super User:
  - The user details must be present in HP Storage Essentialsdatabase.
  - The user must belong to Domain Administrator role.
  - The user must belong to Everything organization.
- Only a Super User can view the **Change Super User** tab.
- Any user assigned to roles having similar privileges as the Domain Administrator cannot be assigned as a Super User. These users are not listed in the **Select User** list in the **Change Super User** window to be chosen as Super User.

To change the Super User:

1. Click **Security** > **Users**.
2. Click **Change Super User** tab.
3. Select a user you want to assign as Super User from the list.
4. Click **OK**.

Changing the Password for a User Account

When changing the password for accessing the management server, keep the following in mind:

- Only a user belonging to the Domain Administrator role is allowed to change the password of another basic user.
- This change takes effect immediately, even if the user is logged into the management server.
- If a user account was authenticated against Active Directory/LDAP, you cannot use the management server to change that user's password. You must use Active Directory/LDAP to change the password.

To modify a password:

1. Click **Security** > **Users**.
2. Click **Users** from the menu.
3. Click the **Edit** button () corresponding to the user account you want to modify.
4. Click **Change Password**.
5. Enter a new password in the **New Password** box.
6. Enter the password again in the **Verify Password** box.
7. Click **OK**.

Changing Your Password

You cannot use the management server to change your password if your user name was authenticated against Active Directory/LDAP. For more information, see Step 1 – Add Active Directory Users to the Management Server on page 313.

To change the password you use to access the management server:

1. Click the name of your account in the upper-left corner.
2. On the **User Profile** tab, click the **Change Password** button.
3. Enter a new password in the **New Password** box.
4. Enter the password again In the **Verify Password** box.
5. Click **OK**.
6. Click the **Save Changes** button on the **User Profile** tab.

Your password change takes effect immediately.

Deleting Users

Keep in mind the following:

- You cannot delete the admin account.
- Only users belonging to the Domain Administrator role can delete users.
- You cannot delete a Super User account.

To delete a user account:

1. Click **Security** > **Users**.
2. Click the corresponding **Delete** button ( ). The user account is deleted.

Modifying Your User Profile

While you are logged into the management server, you are allowed to change the following information:

- E-mail address
- Full name
- Password
- Phone number

You are not allowed to modify the following:

- Login Name
- Organization affiliation
- Role

You must ask your Domain Administrator to make the changes.

To modify your user profile (other than name, role, and organization affiliation):

1. Click the name of your account in the upper-left corner.

2. On the User Profile tab, modify one or more of the following:
   - Full Name
   - E-mail address
   - Phone number
   - Password

     To change the password, click the **Change Password** button. See Changing Your Password on previous page.

     This feature is not available if your user name was authenticated against Active Directory or LDAP. Use Active Directory/LDAP instead.
3. When you are done, click **Save Changes**.

Modifying Your User Preferences

Use the User Preference tab to modify your user preferences for System Manager and Element Topology. The User Preference tab controls what is displayed for your user account.

To access the User Preferences tab:

1. Click the name of your account in the upper-left corner.
2. Click the **User Preferences** tab.

System Manager, Capacity Manager and Performance Manager Preferences

Select one of the following:

- **Load-on-Demand**: Does not populate the tree nodes or display elements in the topology when the page opens (Faster). Use this option for medium to large environments.
- (*Default*) **Automatic Loading**: Populates fabric tree nodes and displays all elements in the topology when the page opens (Slower).

System Manager and Element Topology Preferences

To change the severity icons you view in System Manager and in the element topology, select a severity level from the Display Severity icons with this severity level or higher menu.

To have events refreshed within a time period, select the **Refresh events automatically** box and enter in minutes how often you want the event information on the screen updated. If this option is set to every 5 minutes, the management server refreshes the severity icons displayed in System Manager and the element topology every 5 minutes.

Warnings for Slow Systems Operations

By default, the management server warns you when it encounters issues occurring when handling large amounts of data from storage systems, such as long load times.

If you do not want to be warned, clear the Warn about slow storage system operations option on the **User Preferences** tab. See Modifying Your User Preferences on previous page for information on how to access the User Preferences tab.

Viewing the Properties of a Role

If you are assigned the Domain Administrator role, you can determine which components a user can access by viewing the properties of the user's role.

To view the properties of a role:

1. Click **Security** > **Users**.
2. In the Role column, click the name of the role.

The following information for the selected role is displayed:

- Role Name – The name of the role. This name appears in the users table (**Security** > **Users**)
- Role Description – A description of the role.
- Access Level – How much access the user has to a type of element, such as hosts, storage systems, switches, and applications. See Security for the Management Server on page 287 for more information.

- Access to the <*product name*> – Components in the management server the user can access. In this instance, <*product name*> is the name of your product.

To learn how to edit a role, see Editing Roles on the facing page.

Viewing the Properties of an Organization

If you are assigned the Domain Administrator role, you can determine which elements a user can access by viewing the properties of the user's organization

To view the properties of an organization:

1. Click **Security** > **Users**.
2. In the Organization column, click the name of a organization.
3. Take one of the following actions:
   - To determine which elements are in a child organization, click the link of the child organization.
   - To learn more about an element, click the element's link to display the following information:

   Name – The name of the organization. This name appears in the users table (**Security** > **Users**)

   Description – A description of the organization

   Organization Members – Determines which elements the user can access. See Security for the Management Server on page 287 for more information.

To learn how to edit an organization, see Editing an Organization on page 307.

Managing Roles

This section contains the following topics:

Adding Roles

The management server ships with several roles. You can add roles to accommodate your organization; for example, a role for quality assurance. See Security for the Management Server on page 287 for more information about roles and organizations.

Keep in mind the following:

- The Role Name box does not accept special characters, except spaces and the following characters: $, -, ^, ., and _

- Only users belonging to the Domain Administrator role can add roles.

To add a role:

1. Click **Security** > **Roles**.
2. Click **New Role**.
3. In the Role Name box, enter a name for the role; for example, Quality Assurance.
4. The name can contain spaces, but cannot be longer than 100 characters.
5. In the Description box, enter a description for the role; for example: Role for those in quality assurance. The description cannot be more than 1024 characters long.
6. Select an access level for each element type:
   - Full Control – View and modify the record for the element (Asset Management tab) and perform provisioning.
   - Element Control – View and modify the record for the element (Asset Management tab).
   - View – View element properties (Options for Restricting a Role on page 290).
7. Select the features you want a user to be able to access.
8. Click **OK**.

Editing Roles

The software enables you to modify the default roles and the roles you created. See Security for the Management Server on page 287 for more information about roles and organizations.

Keep in mind the following:

- Only users belonging to the Domain Administrator role can modify roles.
- Domain administrators can change the user names and roles of other domain administrators, but cannot modify their own user name and roles while logged into the management server.
- After you click **OK** in the Edit Role window, any users assigned to the role you edited are logged out of the management server. Users see the changes when they log back into the management server.
- The Role Name box does not accept special characters, except spaces and the following characters: $, -, ^, ., and _

To edit a role:

1. Click **Security** > **Roles**.
2. Click the **Edit** () button.
3. Make your changes:
   - To edit the name of the role, change the name in the Role Name box. The name can contain spaces, but cannot be longer than 256 characters.

- To edit the description of the role, change the description in the Description box. The description cannot be more than 1024 characters.
- To change the access level, change the options selected in the table.
  - Full Control – View and modify the record for the element (Asset Management tab) and perform provisioning.
  - Element Control – View and modify the record for the element (Asset Management tab).
  - View – View element properties (see Options for Restricting a Role on page 290).

4. Select the features you want a user to be able to access.
5. Click **OK**.

Deleting Roles

Keep in mind the following:

- A role cannot be deleted if it contains a user.
- Only users belonging to the Domain Administrator role can delete roles.

To delete a role:

1. Click **Security** > **Roles**.
2. Select **Roles** from the menu.
3. Click the corresponding **Delete** button (). The role is deleted.

Managing Organizations

This section contains the following topics:

Adding an Organization

You can create new organizations to restrict access to certain elements. For example, if you do not want the help desk to have access to elements belonging to a certain group, you could create an organization that does not allow access to those elements. Once you assign users to that organization, they will only be able to access the elements you specified.

See Security for the Management Server on page 287 for more information about roles and organizations.

Keep in mind the following:

- Create child organizations first, and then their parents.
- Events from all elements regardless of the user's organization are displayed by Event Manager.
- Only users belonging to the Domain Administrator role can add organizations.
- Only active organizations can be edited.
- Moving a cluster from one organization to another moves all of the cluster's nodes to the target organization.
- File servers and their hosts must be in the same organization for File System Viewer to work properly.

To add an organization:

1. Click **Security** > **Organizations**.
2. Click the **New Organizations** button.
3. In the **Name** box, enter a name for the organization. The name of an organization has the following requirements:
   - Can contain spaces.
   - Can add digits to the beginning of an organization's name.
   - Cannot be longer than 256 characters.
   - Cannot contain the caret (^) symbol. The system allows the caret symbol to be entered, but the caret symbol should not be included in an organization's name.
4. In the **Description** box, enter a description for the organization. The Description box cannot have more than 1024 characters.

To add elements:

1. Expand the Element Types node and select the element type you want to add.
2. In the Potential Members pane, select the elements you want to add by clicking the appropriate check boxes.

3. Click **Add**. The selected elements are added to the Organization Members pane. To add storage volumes to the organization, see Adding Storage Volumes to an Organization below.

To add organizations:

1. Click the **Organizations** node.
2. In the Potential Members pane, select the elements you want to add by clicking the appropriate check boxes.
3. Click **Add**. The selected organizations are added to the Organization Members pane. The organizations in the Organization Members pane are listed as child organizations because they are now contained within the organization you are creating. See Security for the Management Server on page 287 for more information.
4. Click **OK** when you are done adding the elements and organizations.

Adding Storage Volumes to an Organization

Only users belonging to the Domain Administrator role can add storage volumes to an organization.

To add storage volumes to an organization:

1. Expand the Element Types node and select the Storage Systems node.
2. In the Potential Members pane, click the **Storage Volumes** tab and select a storage system from the Showing Volumes for Storage System menu.
3. To filter the list of volumes for a storage system, click the **Show Volume Filter** link, select the appropriate filter criteria, and click **Submit Query**.
4. Select the storage volumes you want to add to the organization. Click the **+Ports** link in the Ports column to see a list of the ports associated with a particular volume.
5. When you are finished selecting volumes, click the **Add** button located at the top of the pane.
6. Click **OK**. The selected volumes are added to the Organization Members pane.

Viewing Organizations

The Setup Organizations page lists the organizations with their descriptions. The page also shows the number of top-level elements, users, and child organizations assigned to each organization.

Only users belonging to the Domain Administrator role can view organizations.

The No. of Top Level Elements column provides the total number of elements assigned directly to an organization. This number does not include those within the child organization. A zero (0) in the Elements column indicates that the organization contains only child organizations; however, users assigned to that organization would have access to the elements assigned to its child organizations.

Assume an organization contains only two child organizations. As a result, 0 would be displayed under the No. of Top Level Elements column. Users assigned to that organization can access the elements assigned to the two child organizations.

Access the Setup Organizations page by clicking **Security** > **Organizations**.

To access information about a child organization, click its link in the Child Organization column.

Editing an Organization

When elements are removed from an organization, users belonging only to that organization are no longer able to access the removed elements.

See Security for the Management Server on page 287 for more information about roles and organizations.

Keep in mind the following:

- Depending on your license, role-based security might not be available. See the List of Features, which is accessible from the Documentation Center.
- Only users belonging to the Domain Administrator role can edit organizations.
- Only active organizations can be edited.
- You cannot edit the Everything organization.
- File servers and their hosts must be in the same organization for File System Viewer to work properly.

To edit an organization:

1. Click **Security** > **Organizations**.
2. Click the Edit ( ) button.
3. To change the name of the organization, enter a new name in the Name box.
   The name of an organization has the following requirements:
   - Can contain spaces.
   - Can add digits to the beginning of an organization's name.
   - Cannot be longer than 256 characters.
   - Cannot include special characters except spaces and the following characters: $, -, ., and _
   - Cannot contain the carot (^) symbol.
4. To change the description of the organization, enter a new description in the **Description** box.
   You cannot enter more than 1024 characters in the **Description** box.
5. Add or remove elements as described in Adding an Organization on page 305 and Removing Members from an Organization on next page.

6. When done adding or removing elements, click **OK** in the Add Organization or Remove Organization page.
7. In the Edit Organization page, click **OK**.

Removing an Organization

When an organization is removed, users assigned only to that organization are no longer able to access its elements. For example, assume you belong to two organizations, onlyHosts and onlySwitchesandHosts. The organization onlyHosts contains only hosts, and onlySwitchesandHosts contains switches and hosts. If you delete the onlySwitchesandHosts organization, you still have access to hosts because you still belong to the onlyHosts organization.

Keep in mind the following:

- You cannot remove the Everything organization, which is the default organization.
- Only users belonging to the Domain Administrator role can delete organizations.
- You cannot delete an organization that contains a user who belongs to no other organizations. For example, you could create an organization named Org1 that contains two users: User1 and User2. User1 belongs to two other organizations, and User2 belongs only to Org1. You would not be able to then delete Org1 because Org1 contains User2, and User2 does not belong to any other organizations.

To delete an organization:

1. Click **Security** > **Organizations**.
2. Click the Delete ( )button corresponding to the organization you want to remove. The software removes the organization.

Removing Members from an Organization

If you remove an element from an organization, users belonging to that organization or to one of its parents can no longer access that element if it is not a member of any other organization. For example, assume an element named MyHost is a member of BostonWebHost_Solaris, and also mistakenly becomes a member of BostonWebHost_Windows. If you remove MyHost from BostonWebHost_Solaris, users belonging to BostonWebHost_Solaris can no longer access the element. Users belonging to the BostonWebHost_Windows organization or to its parent can still see the element.

To remove elements from an organization:

1. Click **Security** > **Organizations**.
2. Click the Edit ( ) button for an organization, and then select the elements or child organizations you want to remove by clicking the appropriate check boxes in the Organization Members pane.
3. Click **Remove**.

Only users belonging to the Domain Administrator role can remove members from an organization.

Filtering Organizations

The management server provides a filtering feature that lets you designate which organizations are active in your view. For example, assume you belong to an organization named Hosts, and this organization contains two organizations: WindowsHosts and SolarisHosts. To view elements only in WindowsHosts and not in SolarisHosts organizations, use the filtering feature to activate only the WindowsHosts organization.

Keep in mind the following:

- Users assigned to the Admin account cannot filter organizations because the Admin account belongs to the Everything organization by default. As a result, these users do not have access to the filtering feature for organizations.
- If you do not want to view an element, deselect all child organizations containing that element. You must also deselect all parent organizations containing the child organization that has that element. For example, assume you do not want to view all Solaris hosts and all Solaris hosts are in the SolarisHosts organization. The SolarisHosts organization is contained in the Hosts organization. You must deselect the SolarisHosts organization and the Hosts organization if you do not want to see the Solaris hosts.
- The filter for organizations does not appear in Event Manager. Events from all elements regardless of the user's organization are displayed by Event Manager.
- Organization filtering does not affect the reports.

To filter an organization:

1. Click the  button at the top of the screen, or click the link listing the organizations you can view.
2. Deselect the organizations that contain the elements you do not want to obtain information about. For example, to view only the elements in the WindowsHosts organization, select only WindowsHosts. If you have a parent organization named Hosts that contains SolarisHosts and WindowsHosts, deselect SolarisHosts and Hosts. You must deselect Hosts because it contains organizations other than WindowsHosts.

   Keep in mind that you cannot deselect all organizations.

   If you belong to the Domain Administrator role, links are displayed for the organizations. To learn more about the contents of an organization, click its link.

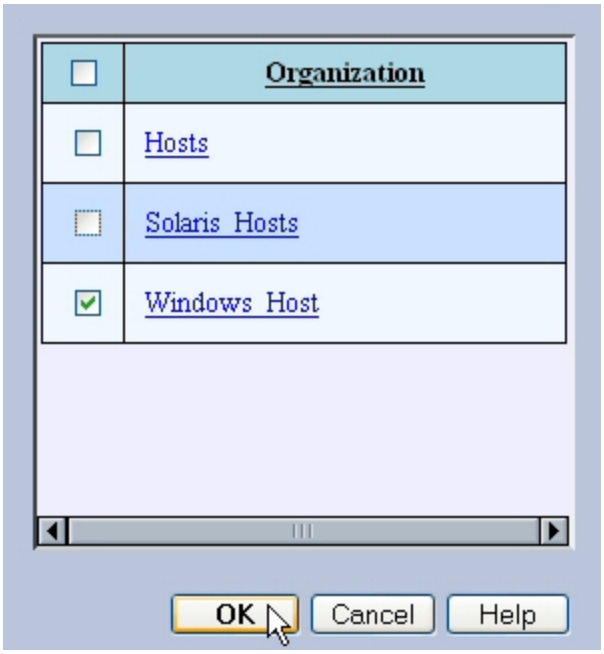

3. Click **OK**.

   You can now only obtain information about elements in the active organizations. These active organizations are listed in the link next to the filter button.

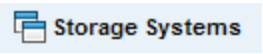

Changing the Password of System Accounts

Change the passwords to the following accounts to prevent unauthorized access.

- RMAN_USER - RMAN backup and restore; user has sys privilege; default password: backup
- DB_SYSTEM_USER - All database activity including establishing a connection to the management server database; default password: password

Use the Database Admin Utility to change the passwords of these accounts, so the management server is aware of the changes. Do not use Oracle to change the password for these accounts. Keep the new passwords in a safe location so that you can remember them.

The password requirements for the management server are:

- Must have a minimum of three characters.
- Must start with a letter.
- Can contain only letters, numbers, and underscores (_).
- Cannot start or end with an underscore (_).

To change the password of a system account:

1. Stop the AppStorManager service.
   - **Windows**:
     i. Go to the **Administrative Tools** > **Services** window.
     ii. Right-click **AppStorManager**.
     iii. Select **Stop** from the menu.
   - **Linux**:
     i. Open a command prompt window.
     ii. To stop the management server, enter the following:

     ```
     /etc/init.d/appstormanager stop
     ```

     iii. To see the status of the management server, enter the following:

     ```
     /etc/init.d/appstormanager status
     ```

2. Access the database utility by doing the following on the management server:
   - **Linux**:
     i. Set the display if you are accessing the Database Admin Utility remotely.

     To set Perl in your path, enter the following command at the command prompt:

     ```
     # eval `/opt/<SE Install Dir.>/install/uservars.sh`
     ```

     In this instance, /opt/<SE Install Dir.> is the directory containing the software. It is defined by `$APPIQ_DIST`.

     ii. Go to the `$APPIQ_DIST/Tools/dbAdmin` directory and enter the following at the command prompt:

     ```
     perl dbAdmin.pl
     ```

   - **Windows**:

     Go to the `%MGR_DIST%\Tools\dbAdmin` directory and double-click **dbAdmin.bat**.

3. Click **Change Passwords** in the left pane.
4. Select an account name from the User Name box.
5. Enter the current password in the Old Password box.
6. Enter the new password in the New Password box.
7. Re-enter the password in the Confirm Password box.
8. Click **Change**. The Database Admin Utility changes the password for the specified account.

Using Active Directory/LDAP for Authentication

The management server supports external authentication through Active Directory (AD) and Lightweight Directory Access Protocol (LDAP) directory services. When you configure the management server to use external authentication, user credentials are no longer stored in the management server database. This configuration centralizes all security related requirements to the enterprise AD/LDAP infrastructure, such as password expiration, resets, and complexity requirements.

When a user attempts to log on to the management server, the management server authenticates the user name and password against AD/LDAP for credential verification. If AD/LDAP verifies that this user has the correct credentials, the HP Storage Essentials management server checks if this user has been already added to HP Storage Essentials database. If both the conditions satisfy, it will allow this user access to the application.

Keep in mind the following:

- It is important to enable either AD or LDAP. You cannot enable both.
- To go back and forth between internal and external (AD/LDAP) authentication, change the `logintype` to `"activedirectory"` or `"ldap"` in the custom properties box.
- If you specify a Pre-Windows 2000 username on a Windows AD server, the Pre-Windows 2000 username must match the current AD username.
- It is possible to create two “admin” user accounts on the management server that differ by case when running with AD/LDAP authentication.
- Remote Active Directory login does not work if you have enabled SSL and have used a self-certification certificate. Active Directory functions properly if you use a Certification authority (CA) certificate and do not use the SSL option.
- Active Directory users with special characters in their name cannot login to HP Storage Essentials. Although Active Directory accepts special characters, HP Storage Essentials converts special characters, such as the following, to underscores (_) when they are entered in the Login Name field, and therefore the user names with special characters cannot be mapped to Active Directory:
  - semicolon (;)
  - open bracket ([)
  - close bracket (])
  - pipe (|)
  - equal sign (=)
  - plus sign (+)
  - asterisk (*)
  - question mark (?)

- less than sign (<)
- greater than sign (>)
- quote (")

To use AD/LDAP to authenticate your users:

- Step 1 – Add Active Directory Users to the Management Server below
- Step 2 – Configure the Management Server to Use AD or LDAP on next page

Step 1 – Add Active Directory Users to the Management Server

Before the management server is configured for Active Directory/LDAP, add active directory users to the management server. This step is required to prevent accidental access to the management server from other AD/LDAP users. Until the user is authenticated against AD/LDAP, the management server views the user as an internal user, whose password can be changed within the management server.

Once a user is authenticated against AD/LDAP, the user is tagged as an external user and the user's password must be managed through AD/LDAP.

To add a user to the management server:

1. Log on to the management server using the default admin user specified in Step 2 – Configure the Management Server to Use AD or LDAP on next page.
2. Create the users as described in Adding Users on page 294 observing the following rules:
   - domain\username format

     Prefix the user name with the domain name; for example, `domain\newuser`. The user name you create in HP Storage Essentials must match the user name in AD/LDAP. You can specify the user say user 1 belonging to a domain say domain1 in one of the following formats:

     i. domain1\user1
     ii. user1@domain1
     iii. user1

     If two users have the same user name and belong to different domain, you cannot use third format to specify the user name.You must use either the first or the second format to provide the user name.

     If the NETBIOS name is different from the domain controller name, only the following formats work:

     - Domain\username
     - username@domain

For example, assume you have a NETBIOS name of JAYLENO and you have a domain controller name of `win2k3r2x86.tonight.show.the.com`. The following user names work, but the username snehauser does not work:

- JAYLENO\snehauser
- snehauser@tonight.show.the.com

- Email format

  Provide the user name in email format; for example, user@domain.com.
  The user should be configured with the proper mail attribute in AD/LDAP.

It is not necessary to create a password, because the passwords used for login are those already configured on either the AD or LDAP server.

Step 2 – Configure the Management Server to Use AD or LDAP

To use AD/LDAP, you must specify the login type as Active Directory or LDAP.

The following sections contain instructions:

- To use AD, see Configuring the Management Server to Use Active Directory below
- To use LDAP, see Configuring the Management Server to Use LDAP on the facing page

Configuring the Management Server to Use Active Directory

You can configure HP Storage Essentials to authenticate users through Active Directory. You can use both email and domain\username for authentication.

You can provide details of a specific AD organizational unit and map it to the management server. The product can then gather user information from such an AD organizational unit. This enabled authentication privileges to any user belonging to that organizational unit.

To specify the management server to use Active Directory:

1. Select **Security** > **Users** to specify user data for AD users. For more information on creating an account, see Adding Users on page 294
2. Specify the login type as Active Directory. To specify the login type follow these steps:
   a. Select **Configuration** > **Product Health**.
   b. Click **Advanced** in the Disk Space tree.
   c. Type `logintype=activedirectory` in the Custom Properties box.
   d. Restart the AppStorManager service.

Creating User Accounts for Active Directory Authentication Through Email

HP Storage Essentials can authenticate email addresses through active directory. This feature enables users to log on with their email address for the user name and their Active Directory password for the password.

To authenticate through an email address:

1. Create a user in HP Storage Essentials (**Security** > **User**). Provide the user's email address for the user name, and set the user's email attribute in the domain controller. Do the following:
   a. Select the specified organization.
   b. Click **OK** when done. If you are not sure how to add a user, see Adding Users on page 294.
   c. Repeat this step for each user you want to add.
2. Specify logintype as AD in the Custom Properties box to enable Active Directory login, as described in Configuring the Management Server to Use Active Directory on previous page.

When users log on to HP Storage Essentials, they must provide the following information:

- Their email address in the username field.
- Their AD password for the password.

Configuring the Management Server to Use LDAP

The LDAP server requires a distinguished name (DN) and credentials. The DN can be configured, allowing name substitution and support for multiple DN configurations.

To configure the management server to use LDAP:

1. Select Security > Users to specify user data for LDAP users. For more information on creating an account, see Adding Users on page 294.
2. Specify the login type as LDAP, as follows:
   a. Select **Configuration** > **Product Health**.
   b. Click **Advanced** in the Disk Space tree.
   c. Type `logintype=ldap` in the Custom Properties box.
   d. Restart the AppStorManager service.

Optional Security Features

This section contains the following topics:

Prevent the Execution of Arbitrary Commands

Summary: Secure the management server by disabling areas of the user interface that allow execution of custom commands.

Follow these steps:

1. Browse to the file SecurityProperties.properties-sample located at:

   `<INSTALL_LOCATION>\Data\Configuration`

2. Save a copy as SecurityProperties.properties.
3. Open the new file with WordPad and comment in the following line:

   `security.disableCommandExecution=true`

4. Save the changes and close the file.
5. Restart the appstormanager service or reboot the appliance.

Expected Result: The right-click options for custom commands in System Manager are no longer available. Policy Manager no longer allows the creation/execution of custom commands.

Disable Provisioning at All Levels

Summary: Prevent element provisioning by removing the option from all areas of the user interface.

Follow these steps:

1. Verify that a provisioning license was installed.
2. Browse to the file SecurityProperties.properties-sample located at:

   `<INSTALL_LOCATION>\Data\Configuration`

3. Save a copy as SecurityProperties.properties.
4. Open the new file in a text editor such as WordPad and comment in the following line:

   `security.disableProvisioning=true`

5. Save the changes and close the file.

The product notifies you if a restart of the AppStorManager service is required.

Expected Results: The Provisioning Manager option is removed from the main menu. Provisioning as a right-click option in the System Manager user interface is no longer available.

Block CLI, Session Applets, and Secure API Invocations

Summary: Protect the management server against unauthorized access via external hosts and programs by configuring it to specify the transport protocols it will deny via API invocations. You can also block the execution of any local CLI session to protect the management server against unauthorized access.

Follow these steps:

1. Browse to file securityProperties.properties-sample located at:

   `<INSTALL_LOCATION>\Data\Configuration\`

2. Save a copy as `SecurityProperties.properties`.
3. Open the file in a text editor such as Notepad. The following list of configuration options can be denied:
   - **# local-rmi** – API invocations using rmi from localhost will be disallowed.
   - **# remote-rmi** – API invocations using rmi from remote hosts will be disallowed.
   - **# remote-http** – API invocations using http from remote hosts will be disallowed.
   - **# remote-https** – API invocations using https from remote hosts will be disallowed.
   - **# session-http** – API invocations using http from remote hosts and session id as authentication will be disallowed.
   - **# session-https** – API invocations using https from remote hosts and session id as authentication will be disallowed.
4. To deny any of these protocols, edit the line `security.deny.transport=` by specifying which transport protocols you want to deny (comma separated for multiple entries), and remove the #.
5. Save the changes and close the file.
6. Restart the appstormanager service or reboot the appliance.

In the following example, the modified syntax denies the execution of CLI from any remote host via all protocols, and denies session applets from remote hosts via http and https from their web browsers:

```
security.deny.transport=remote-rmi,remote-http,remote-https,session-http,session-https
```

Specifying “local-rmi” as a denied transport prevents CLI commands from being executed locally on the management server.

Expected Result: The execution of CLI commands can be blocked from all remote hosts using the RMI, http, or https protocols. Active screens (such as System Manager) can be blocked from view by remote hosts using http or http as a web browser protocol. If session applets are denied (session-http, session-https), the user on the remote host will receive a security transport error message when attempting to view any active screen, and be directed to contact an administrator.

Modify the Password Requirement

Summary: Enhance security by forcing users to create a password with a minimum amount of alpha-numeric characters.

Follow these steps:

1. Browse to the file SecurityProperties.properties-sample located at:

   `<INSTALL_LOCATION>\Data\Configuration`

2. Save a copy as SecurityProperties.properties.
3. Open the new file with WordPad and enter the following:

   `security.minUserPasswdLen=0`

4. Specify required amount of characters in place of “0” in the default statement.
5. Save the changes and close the file.
6. Restart the appstormanager service or reboot the appliance.

Expected Result: When new users are added to the management server, their password must meet the minimum length requirement as specified in the statement. If the password is too short, a message will indicate how many characters are required.

Note: Users who chose passwords before this feature was enabled will be not forced to change their passwords if they do not meet the length requirement.

Modify CIM Extensions on UNIX Hosts

Summary: The parameters file for CIM extensions can be modified to accept connections from specified management servers. Non-specified servers will be unable to discover UNIX hosts with specified parameters.

Follow these steps:

1. On the UNIX host where the CIM extension is installed, browse to the `cim.extension.parameters-sample` file located at:

   `<AGENT_INSTALL_DIR>\conf\`

2. Change the name of the file to `cim.extension.parameters`.
3. In the renamed file, modify the following line by removing the # and replacing the sample IP addresses with the IP addresses of the servers that are allowed to contact the CIMOM extension:

   `-mgmtServerIP 127.0.0.1,192.168.0.1`

   Multiple IP addresses must be comma separated.

4. Save the changes and close the file.
5. Restart the appstormanager service.

Expected Result: The UNIX host can only be discovered from the Management Servers specified by the allowed IP addresses.

7 Managing Licenses

This section contains the following topics:

About the License

The management server restricts the number of elements it manages through its license. It is important that you keep your license up to date with the requirements of your network. The management server has several different types of license restrictions, as shown in the following table.

License Restrictions

Type of Restriction	Description	Unit of Measurement
MAPs	The management software restricts the number of hardware elements it manages through the use of managed access points (MAPs) for hardware. A MAP is the sum of all storage access ports of all hardware elements that the management server manages. When a CIM extension is installed to discover a HP NAS system, this also counts as at least 1 MAP, or as many MAPs as there are FC ports. See related table information. (Cluster detection is not supported, however.) If the CIM extension is running on HP NAS, and if you use File System Viewer on the HP NAS, you must also take into account the number of terabytes (TB) for the File System Viewer, which would be the actual total size of the files scanned. When HP Storage Essentials discovers Brocade switches through SM-S, it discovers the switches in the fabric and adds the ports to the MAP count. To reduce MAP counts, restrict the number of Brocade switches discovered through SMI-S. See Excluding Brocade Switches from SMI-S Discovery on page 89. When HP Storage Essentials discovers HP Data Protector application running on a discovered host, it adds the Data Protector host to its MAP count. You can reduce the MAP counts by discovering the host as a backup server. To enable the discovery of the host as a backup server, select **Include backup details** option while running Get Details. You can also exclude additional devices to further reduce your MAP counts. For more information, see: • Virtual machines – Excluding Virtual Machines from Discovery on page 187. • HDS storage systems – Excluding HDS Storage Systems from Discovery on page 117. • McDATA switches – Excluding HDS Storage Systems from Discovery on page 117. • EMC Symmetrix storage systems – Excluding EMC Symmetrix Storage Systems from Discovery on page 108.	Number of MAPs

Type of Restriction	Description	Unit of Measurement
Data Protector Reporter Edition	Data Protector Reporter Edition by default does not come with MAPs, and therefore you cannot discover devices that have MAPs, such as switches, arrays and CIM extension, even though this functionality is displayed in the product and mentioned in the documentation. If you are, running Data Protector Reporter without MAPs, you can only discover your backup servers without a CIM extension installed as described inPrerequisites for Agentless Discovery of Data Protector on page 194.	
Backup Size	The management server determines licensing for Backup Manager through gigabytes (GB). The management server compares the number of GB for Backup Manager with what you are backing up. If you are backing up more than your license allows, you are warned the next time you log on to the management server.	Gigabytes (GB)
Raw NetApp Capacity	Raw NetApp Capacity is the total disk capacity (unformatted capacity) of all discovered NetApp filers.	Terabytes (TB)
Managed Exchange Instances	The management server determines licensing for Microsoft Exchange instances by counting the number of instances of Microsoft Exchange it manages.	Number of instances of Microsoft Exchange the software manages
Managed Database Instances	Total number of instances of the following databases managed by the software: • Microsoft SQL Server • Oracle • Sybase Adaptive Server Enterprise • InterSystems Caché The total is broken down by each type of database in the table.	Number of managed databases

Type of Restriction	Description	Unit of Measurement
File System Viewer	The management server determines licensing for File System Viewer through terabytes (TB). When you purchased File System Viewer, you were given a number of TB you were allowed by the management server to monitor. The management server detects the number of TB that are being monitored on file servers and verifies that number is at or below the purchased amount. You do not have to monitor everything associated with your file server. You can choose to manage only the mount points that are important to you. Only the files associated with these mount points are counted toward the file server TB. If you use File Server SRM to monitor NAS systems, the TB of the NAS systems must also be considered in the FIle Server total licensing TB count requirement.	Terabytes (TB)
NAS Manager	Licensing for NAS Manager is based on the number of raw NAS TBs managed. When a CIM extension is installed to discover a HP NAS system, this also counts as at least one MAP, or as many MAPs as there are FC ports. (Cluster detection is not supported.) If the CIM extension is running on HP NAS, and you use File System Viewer on the HP NAS, you must also take into account the number of TB for the File System Viewer, which would be the actual total size of the files scanned.	Terabytes (TB)
EVA Array Performance Packs	Each EVA Performance Pack license lets you monitor only one EVA array. To monitor multiple EVA arrays, you must purchase an EVA Performance Pack license for each EVA array.	EVA Array
XP and HDS Array Performance Packs	Each XP or HDS Array Performance Pack license lets you monitor only one XP or HDS array. To monitor multiple XP and/or HDS arrays, you must purchase an XP or HDS Array Performance Pack license for each array.	XP Array, HDS Array

Type of Restriction	Description	Unit of Measurement
VIO	• Each VIO server with no HBA port = 1 MAP • Each VIO server with 1 HBA port = 1 MAP • Each VIO server with X HBA ports = X MAP • Each VIO client with no HBA port = 1 MAP • Each VIO client with 1 HBA port = 1 MAP • Each VIO client with X HBA ports = X MAP	

The management server Current Usage Summary is first updated 6 hours after the management server (AppStorManager) starts. Updates occur every 24 hours thereafter. Elements that the management server has discovered before the update are not reflected in the Current Usage Summary table. The time for the update is determined when the management server is first started. For example, the first update of the Current Usage Summary table occurs 6 hours after the management server is first started. The following updates occur every 24 hours. If the management server is started for the first time at Noon, the first update of the Current Usage Summary table would occur at 6 pm. All following updates would always occur at 6 pm.

To update the Current Usage Summary table immediately, click the **Refresh License Usage** button on the Licenses page (see Refreshing the License Usage Table on page 327.

Element	Managed Access Point
Hosts	Each Fibre Channel port counts as one MAP. If a host has no Fibre Channel ports, the software assumes one MAP. The software does count direct attached storage, provided it is supported by the management server.
Virtual machines and servers	Virtual servers are treated like physical hosts. Each Fibre Channel port counts as one MAP. If a virtual server has no Fibre Channel ports, the software assumes one MAP. A virtual machine uses a MAP if it is running VMTools. It does not matter whether it is stored through internal or external storage, or whether it was discovered through the virtual server or through VirtualCenter. A virtual machine that is not running VMTools will be treated as unmanaged and will not use any MAPs. A virtual machine with CIM extensions installed will use one MAP regardless of whether or not VMTools is installed.
Switches	All ports on a switch are counted as MAPs.
Storage systems	The MAPs are the sum of all front-facing ports. Storage systems with FC ports that the software does not support, such as mainframe attached FICON, are still counted as MAPs. However, the management server does not count MAPs from storage systems that it does not support. See the release notes for information about supported storage systems.

The local Oracle database that HP Storage Essentials uses as its own database is not counted as a MAP.

Example 1:

Assume you have the following environment:

- Brocade (two switches of 12 ports each, one switch of 16 ports) – total 40 ports
- McDATA (one switch of 64 ports) – total 64 ports
- Windows 2000 and Solaris Hosts (10 hosts with two Fibre Channel connection each) – total 20 ports
- EMC Subsystem (one subsystem with 16 Fibre Channel ports) – total 16 ports

The software calculates 140 MAPs.

Example 2:

Assume you have the same configuration as the first example, and you add several devices to your network that the management server does not support. There are still 140 MAPs in this environment, because the management server does not count the ports from devices it does not support.

Example 3:

Assume you have the same configuration as the first example, with two Windows 2000 hosts that are directly attached to storage systems, no Fibre Channel connections, and no Fibre Channel ports, as shown in the following figure:

Figure 2 Example of Direct Attached Storage

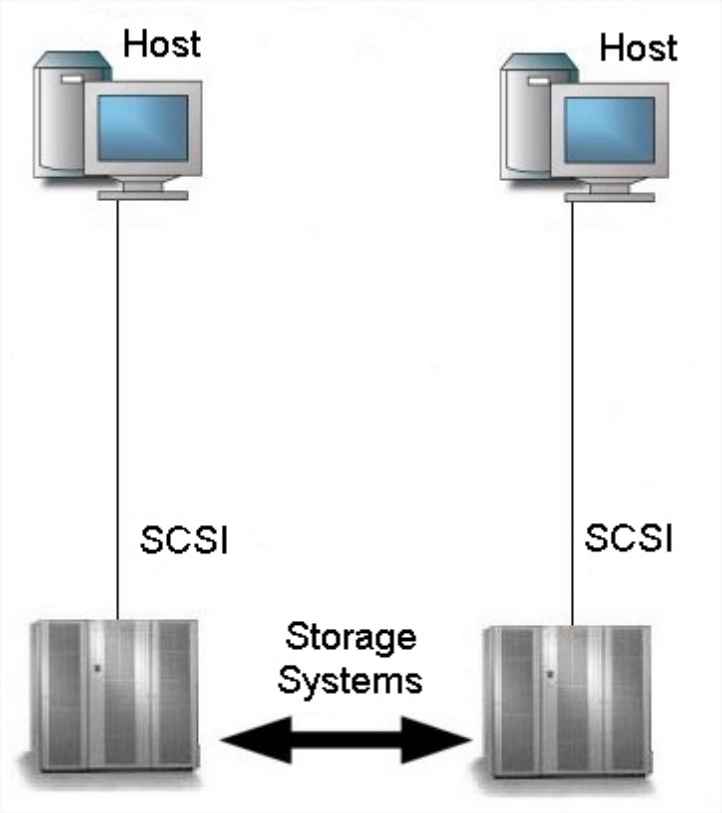

The software calculates four MAPs, because we assume one MAP for each host, even though it has no Fibre Channel ports. The storage systems are counted, because they are supported by the management server. If you include the MAPs from the first example (140 MAPs), your total would be 144 MAPs.

If you had a configuration that included a switch, two managed hosts, and several unmanaged hosts, the MAPs would not be used against the unmanaged hosts.

Some switches allow the user to turn off an unused GBIC (Gigabit Interface Converter). If a GBIC is turned off, the port is not counted. But if the GBIC is turned on, or if there is no GBIC, the port is counted.

Example 4:

Assume you wanted to order licensing to support a total of 850 MAPs of HP Storage Essentials, a total of 600 MAPs of HP Storage EssentialsChargeback Manager, a total of 25 TBs of HP Storage EssentialsFile System Viewer, a total of 20 MALs of HP Storage Essentials Exchange Viewer, a total of 5 HP Storage EssentialsReport Optimizer, one Concurrent User LTU (License To Use), a total of 10 HP Storage Essentials Performance Pack LTUs to monitor performance on a total of 10 HP EVA 8000 systems, and a total of 5 TBs of NAS Manager.

One EVA Performance Pack license allows you to manage only one EVA array. You would have to purchase multiple licenses to manage multiple EVA arrays. The same applies to the XP Performance Pack. Each license allows you to manage one XP array.

Your order would consist of the following:

- 17 HP Storage Essentials, 50 MAP LTU (17 X 50 MAPs = 850). This includes the anticipated related MAPs requirement for the NAS system.
- 12 HP Storage EssentialsChargeback Manager 50 MAP LTU (12 X 50 = 600)
- 25 HP Storage EssentialsFile System Viewer 1 TB LTU (25 X 1 = 25). This quantity includes the anticipated related TB usage for the NAS system.
- 20 HP Storage Essentials SRM Exchange Viewer 1 MAL LTU (20 X 1 = 20)
- 5 HP Storage Essentials Report Optimizer 1 Concurrent User LTU (5 X 1 = 5)
- 5 TB NAS Manager TB LTU (5 X 1 = 5)
- 10 HP Storage Essentials Performance Pack 1 Array LTU (10 X 1 = 10)

For more examples and information, refer to the product Quick Specs by selecting your product from the product links at the following web page:

http://h71028.www7.hp.com/enterprise/cache/123557-0-0-225-121.html

Importing a License File

If you cannot find the license file you want to import, or are interested in expanding your license for managing additional elements, contact your software or support representative for assistance.

When adding a license for a module that requires MAPs, first import the MAP license and then import the module add-on license.

The license agreement, which is in PDF format, is displayed the first time you access HP Storage Essentials. Install the latest version of a PDF reader, such as Adobe Acrobat Reader, on the client you plan to use to access HP Storage Essentials for the first time.

To import a license file:

1. Select **Security**.
2. Select **Licenses** from the menu.
3. Select **Import License File**.
4. Select **Browse**. The file system of the computer being used to access the management server appears.
5. Select the license file.
6. Select **OK**.

Viewing Cumulative Licenses

The View Cumulative License feature enables you to view the complete number of elements the management server supports at the current time. The software adds up the number of licensed components from the licenses and takes into account the expiration date. See About the License on page 319 for more information about the licensing capacities displayed.

You cannot modify the license file because it is encrypted. To increase the number of elements the management server is allowed to manage, follow your organization's procedures to contact your support representative.

To view cumulative licenses:

1. Select **Security**.
2. Select **Licenses** from the menu.
3. Select **View Cumulative Licenses**. The properties for the cumulative licenses are displayed.

In the **Cumulative License** window, each feature has a property that is set to either true or false. If a value for a property is set to true, you can access that feature. Likewise, if the value is set to false, you cannot access that feature.

You can determine how many elements your licenses supports by looking at the **Current Usage Summary** table at the bottom of the page. The cumulative number for each type of licensed capacity is displayed in this table.

To update the Current Usage Summary table immediately, click the **Refresh License Usage** button on the Licenses page (see Refreshing the License Usage Table below).

Refreshing the License Usage Table

To obtain the current license usage based on what is currently in the database, click the **Refresh License Usage** button on the Licenses page (**Security** > **Licenses**).

If you deleted several elements and want to obtain an up-to-date tally of the license usage in the Used Licenses column, you must click the **Refresh License Usage** button on the Licenses page (**Security** > **Licenses**). If you delete an element from the Discovery Step 3 (Get Details) page, such as a host, you could see more than one MAP freed up.

For example, if you delete a host running several applications that HP Storage Essentials monitored, you would most likely see several MAPs freed up if the host had several Fibre Channel ports or a virtual machine.

Viewing a Specific License

Do not manually edit the license. To increase the number of elements the management server is allowed to manage, contact technical support.

To view the content of an individual license:

1. Select **Security**.
2. Select **Licenses**from the menu.
3. Select the  button corresponding to the license you want to view. The license name and file name are listed, along with its properties.

You can determine how many MAPs and managed application licenses (MALs) this license supports by looking at the properties in the license file. However, that can be misleading if you have other licenses that also provide support for MAPs and MALs. To obtain a total of the MAPs and MALs that are supported, take a look at the cumulative licenses (see Viewing Cumulative Licenses on previous page).

The following properties are used for tracking MAPs and MALs:

- LICENSE_FSRM_SIZE_TB – The amount of space in Terabytes you are allowed for File System Viewer.
- LICENSE_MAL_DATABASE – The number of database application instances, such as Oracle and Sybase Adaptive Server Enterprise, the management server is allowed to monitor.
- LICENSE_MAL_EXCHANGE – The number of Microsoft Exchange instances the management server is allowed to monitor.
- LICENSE_MAPS – The number of MAPs the management software is allowed to manage.

Deleting a License

Before you delete a license, make a copy of it. If you delete the wrong license, you could lose access to certain features or access to the product. The management server saves the license files in the following folder:

`<drive where the management server is installed>\data\`

To delete a license:

1. Select **Security**.
2. Select **Licenses** from the menu.
3. Select the  button corresponding to the license you want to delete.

License Setup for Array Performance Pack

The HP Performance Pack license provides the ability to collect and report additional performance data for specified EVA, XP and HDS arrays, EMC Symmetrix, and NetApp systems. For more information, see the HP Storage Essentials Storage Performance Management Guide. It describes each of the HP Performance Pack products and explains how to set up licenses for them.

The number of required licenses depends on the number of arrays you want to include for additional collection and reporting. There is no license setup for NetApp devices.

Note: You must complete a Get Details for EVA, HDS, or XP arrays before importing the license for the EVA or XP, HDS Array Performance Pack. After importing the license, you can start the data collectors from the Performance Data Collection page (**Configuration** > **Performance** > **Data Collection**). Although EVA, HDS, and XP arrays are displayed after you run discovery, you must run a Get Details for the collectors to run properly.

As part of the license setup, a license page similar to the following one displays the used and maximum numbers of managed arrays.

If your license includes the Array Performance Pack capability, the current usage summary reports how many arrays can have this capability applied.

Current Usage Summary

Licensed Capacities	Used Licenses	Maximum Licensed
MAPs	415	500
Raw NAS Capacity	0.00 TB	9,999.00 TB
Managed Exchange Instances	0	1
Managed Database Instances	0	1
Managed Oracle Instances	0	
Managed SQL Server Instances	0	
Managed Sybase Instances	0	
Managed Caché Instances	0	
Managed File Server Storage	0.00 TB	9,999.00 TB
EVA Performance Pack Array-Instances *	0	1
XP, HDS Performance Pack Array-Instances *	1	2

* Use the Performance Licensing tab to apply licenses to storage systems.

After installing the licenses:

1. Click the **Performance Licensing** tab in License Manager and specify which EVA, XP, or HDS arrays you want to have the Array Performance Pack capability, as follows:

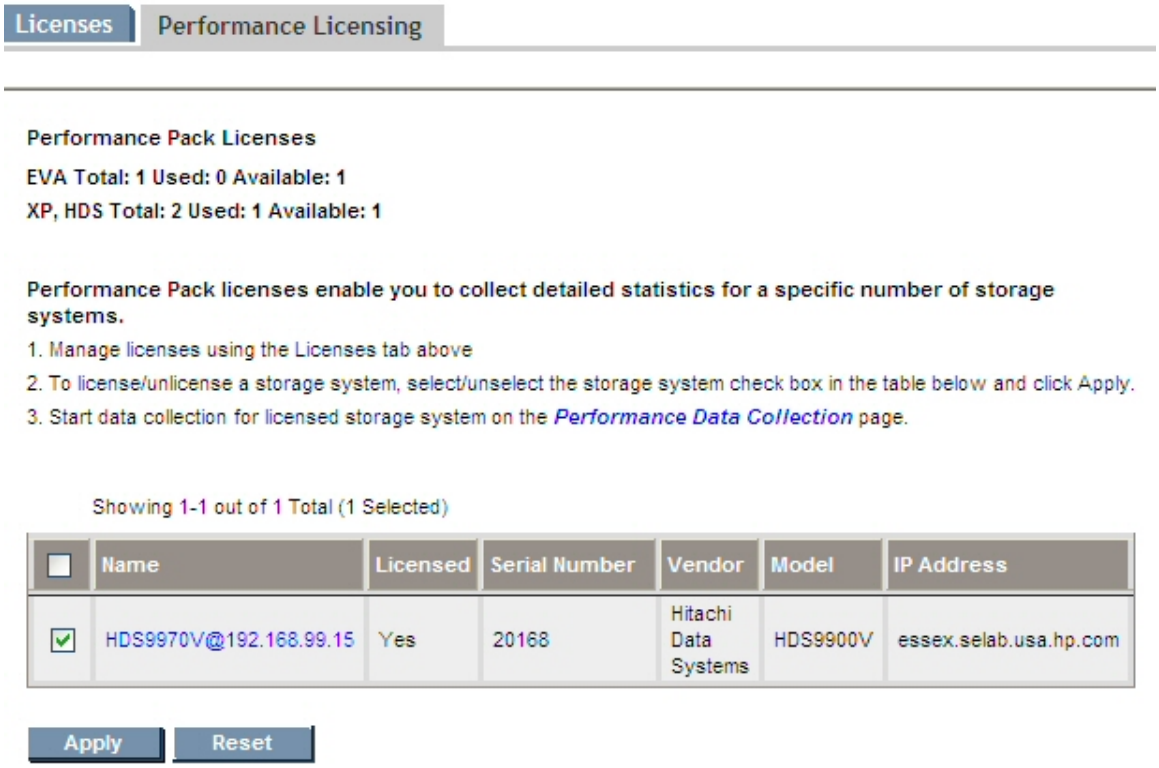

2. Click **Configuration** > **Performance** > **Data Collection**.
3. Start the data collectors for the licensed arrays, so that reporting data is obtained for the parameters specified.

XP P9500 Performance Pack Licensing with Command View Advanced Edition

XP P9500 storage arrays discovered using Command View Advanced Edition (CVAE) are not supported for performance in HP Storage Essentials and cannot be licensed for performance statistics.

Installation of Performance Pack 9.4.1 License

During installation, if you attempt to license P9500 storage arrays that use Command View Advanced Edition for discovery, you will receive an error message instructing you to remove the P9500 storage arrays from Command View Advanced Edition.

If you want to license the XP P9500 storage arrays for HP Storage Essentials performance metrics, you must rediscover the P9500 storage arrays directly through their Service Processors. After discovery, the P9500 storage arrays can be licensed for performance management.

Upgrades from 9.4.0 to 9.4.1

During upgrades, if you attempt to license P9500 storage arrays that use Command View Advanced Edition for discovery, you will receive an error message instructing you to remove the P9500 storage arrays from Command View Advanced Edition.

If you want to license the XP P9500 storage arrays for HP Storage Essentials performance metrics, you must rediscover the P9500 storage arrays directly through their Service Processors.

If you have previously discovered P9500 storage arrays (as well as other supported storage systems) using Command View Advanced Edition, you will need to complete the following steps to enable performance statistics for P9500 storage arrays.

To enable performance statistics for P9500 storage arrays:

1. Remove the P9500 storage arrays from Command View Advanced Edition
2. Remove the access point for the P9500 storage systems. This will also delete any other storage arrays managed by the access point.
3. Add the Service Processor access point for the P9500 storage arrays, then discover the P9500 arrays.
4. Enable performance licensing for the P9500 storage arrays.
5. Run Discovery Step 3 to get all element details for any storage arrays previously discovered using Command View Advanced Edition.
6. Enable licensing for these previously discovered storage arrays.

8 Configuring the Management Server

This section contains the following topics:

Trap Generation

You can configure the software so that events received by the system generate SNMP traps, which the software can send to another event-monitoring system. The software allows up to five SNMP trap destinations. The software can send either SNMPv1 or SNMPv2 traps. Either SNMP version you select will be used for all trap destinations. The default is SNMPv2. To change the default to SNMPv1, see Changing the Default to SNMPv1 on next page.

The software provides an SNMP MIB for each SNMP version that you can compile into your existing enterprise framework. This MIB contains trap definitions so your enterprise framework can understand the traps. The MIB can be found in AppIQ-Traps-v1.mib for SNMPv1 or AppIQ-Traps-v2.mib for SNMPv2 located in the UtilitiesCD/SNMP directory on the StorageEssentialsDVD. You should only compile one of the two MIBs into your enterprise framework. Choose the MIB file that corresponds to the SNMP version you are using.

The software does not have the capability to forward traps received from other devices. It can take events from Event Manager and create SNMP traps from them. These traps are generated using an SNMP MIB.

The software does receive SNMP traps from some devices. These traps are translated into events in Event Manager. When they are sent out as SNMP traps, the information in the trap will be the same as the original device trap, but the format of the trap will be different. For example, the trap will contain the original severity and description information, but the Trap OID, fields, and codes will be different.

To configure trap forwarding:

1. Access the management server, as described in Accessing the Management Server on page 60.
2. Select **Configuration** > **Traps**.
3. Select the **Enable Trap Generation** option.
4. In the SNMP Community String box, enter the SNMP community string, which is used for filtering.

   If you enter a value in the SNMP Community String box, the SNMP agent must know the SNMP community string entered in the box to receive the SNMP trap.
5. Click **Save**.
6. In the New SNMP Destination box, enter the IP address of the server running an SNMP agent.
7. Click **Add**.

Changing the Default to SNMPv1

To change the system to send SNMPv1 traps:

1. Select **Configuration** > **Product Health** and click **Advanced** in the **Disk Space** tree**.**
2. Click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the following command. How you copy the text depends on your Web browser.

   ```
   outgoingSnmpTrapVersion=v2
   ```
5. Return to the Advanced page (**Configuration** > **Product Health**. Then, click **Advanced** in the **Disk Space** tree.
6. Paste the copied text into the Custom Properties box. How you paste the text depends on your Web browser.
7. Modify the text as follows:

   ```
   outgoingSnmpTrapVersion=v1
   ```

8. To make sure the property is not commented out, remove the hash (#) symbol in front of the property.
9. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

Enabling Email Notification

Depending on your license, e-mail notification might not be available. See the List of Features to determine if you have access to e-mail notification. The List of Features is accessible from the Documentation Center (**Help** > **Documentation Center**).

The management server provides e-mail notification for reports and policies. For example, you can set up the management server to notify you by e-mail when the amount of free space on a host becomes too low.

You must assign an SMTP server from which the management server can send its e-mail notifications.

When you change the email configuration on the management server, the new values take effect immediately, but subsequent changes require the restart of the management server service.

To configure e-mail notification:

1. Access the management server, as described in Accessing the Management Server on page 60.
2. Select the **Configuration** > **E-mail Server** option in the upper-right corner.
3. *Required*: Select **Enabled** to enable e-mail notification.
4. *Required*: In the Name box, enter the DNS name or the IP address of the Simple Mail Transfer Protocol (SMTP) server you want to use to send the e-mail notification.
5. *Required*: In the Port box, enter the port of the SMTP server you want to use to send the e-mail notification.
6. In the User Name box, enter a user name for the SMTP server.
7. In the Password box, enter a password for the SMTP server.
8. In the Verify Password box, enter the password you entered previously.
9. *Required*: In the Sender box, enter the e-mail address of the sender. This address is displayed in the From box in the e-mail.
10. To have the replies go to an e-mail address other than the e-mail address specified in the Sender box, enter the e-mail address you want to receive replies in the Reply to box.
11. Click **Save**.
12. Send a test e-mail.

To send a test e-mail:

1. In the To box, enter an e-mail address. The software verifies that the address entered has a correct form.
2. In the Subject box, enter a subject to distinguish this e-mail from notification of a real event; for example:

   ```
   Testing E-mail Notification
   ```

3. In the Message box, enter a message; for example:

   ```
   I'm just testing e-mail notification.
   ```

4. Click **Send Test Message**.

Configuring Print Settings

To configure print settings:

1. Click the  button.
2. Use the fields on the Paper tab to modify the setup of the page. When you are done, click **Apply**. If you want the default settings, click **Default**.

   A preview of the printout is displayed in the right pane.

   Before you change the margins, decide on a unit of measurement.

   - **Paper format** – Select the paper size from the drop-down menu.
   - **Unit** – Select cm (centimeters) or inch for the margins.
   - **Paper width** – Displays the width of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
   - **Paper height –** Displays the height of the paper. You can modify the measurement in this field when you select the **Custom** option in the Paper format drop-down menu.
   - **Top margin** – Type a measurement for the top margin.
   - **Bottom margin** – Type a measurement for the bottom margin.
   - **Left margin** – Type a measurement for the left margin.
   - **Right margin** – Type a measurement for the right margin.
   - **Orientation** – Click an orientation for the printout.

3. Click the **View Selection** tab to modify how the printout will appear on the page. You can modify the following. When you are done, click **Apply**. If you want the default settings, click **Default**.

   A preview of the printout is displayed in the right pane.

   Before you change the margins, decide on a unit of measurement.

- **Start x** – Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. You can also enter negative numbers. Anything more than zero expands the printout to another page.
- **Start y** – Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom. You can also enter negative numbers.
- **Width** – Determines the width of the printout.
- **Height** – Determines the height of the printout.

To remove extra space around the topology, click the **Trimmed** button.

4. To change how many pages the printout will use, select one of the following. When you are done, click **Apply**. If you want the default settings, click **Default**.

   A preview of the printout is displayed in the right pane.

   Before you change the margins, decide on a unit of measurement.

   - **Unit** – Select cm (centimeters) or inch for the margins.
   - **Position/Size** – Lets you change the position and size of the printout so that it spans several pages:
     - **Start x** – Determines the horizontal placement of the printout on the page with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. Determines the horizontal placement of the printout. Anything more than zero expands the printout to another page.
     - **Start y** – Determines the vertical placement of the printout on the page with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom.
     - **Width** – Determines the width of the printout. If the width entered does not fit on the page, the printout wraps around to another page.
     - **Height** – Determines the height of the printout. If the height entered does not fit on the page, the printout wraps around to another page.
   - **Resolution (pixel/unit)** – Lets you change the resolution so that the printout spans several pages.
   - **Page** – Lets you expand the printout so it prints on several pages without modifying the graphic.

5. To preview your pages, click the **Preview** tab and then click the page you want to preview. The page appears in the right pane.
6. When you are ready to print, click **Print**.

7. Click **Close**.

To revert back to all of the original settings, click the **Default** button next to the **Print** button.

Setting the Date and Time for Scheduled Tasks

To set the date and time for a scheduled task:

1. Click the calendar icon,  .
2. In the Time box, enter the time in 24-hour format, with the hour and minutes separated by a colon; for example, 22:15. Click the date on which you want the task to run. Today's date is highlighted in pink. Click **Set**.

   The date and time appear in the Next Scheduled Run box in yyyy-mm-dd hh:min format.

   If you change the date in the box to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2010-11-31, the software assumes the date is 2010-12-01.
3. In the Repeat Interval box, enter an interval. Select one of the following units from the list:
   - **Second(s)**
   - **Minute(s)**
   - **Hour(s)**
   - **Day(s)**
   - **Week(s)**

Managing Discovery Schedules

You can schedule the management server to obtain discovery details at a specified interval. The management server provides the following types of discovery details:

- **Include infrastructure details** – The management server gathers detailed information about the SAN infrastructure. This process can be network intensive.
- **Include backup details** – For the latest backup information, schedule discovery to run nightly after you run backup sessions.

Note: HP recommends not scheduling the discovery of infrastructure details and backup details to run at the same time.

Adding a Discovery Schedule

Schedule getting details when the network is not busy.

All collectors are stopped during a discovery using the Include infrastructure details option. This means that information about the SAN, such as for Performance Manager, is not updated.

If you are creating multiple discovery schedules, avoid scheduling conflicts (concurrently scheduled discovery tasks) and make sure that each scheduled task has enough time to start and finish before the next discovery task is scheduled to start. See Do Not Run Overlapping Discovery Schedules on page 868.

Do not run Get Details for all discovery groups simultaneously.

Hosts discovered with CIM extensions cannot be added to discovery groups. These hosts can be placed independently into scheduled Get Details tasks without being part of a discovery group. This allows you greater flexibility when gathering discovery data.

If you are upgrading from a previous version of the product, and run Get Details against your hosts, they will be moved out of their existing discovery groups. After Get Details, each host is placed in its own discovery group. If the original discovery groups containing these hosts were included in scheduled tasks, the schedules would be modified to contain the new discovery groups for the hosts.

To schedule discovery details:

1. Access the Discovery page by selecting **Configuration** > **Discovery**.
2. Click the **Discovery Schedule** tab.
3. Click **New Schedule**.
4. In the Name box, enter a name for your discovery schedule.
5. In the Description box, enter a description for your discovery schedule.
6. Select one or more of the following:
   - **Include infrastructure details** – The management server gathers detailed information about the SAN infrastructure. This process can be network intensive.
   - **Include backup details** – Schedules discovery to run nightly after you run your backup sessions. It is recommended you do not schedule the discovery of infrastructure details and backup details to run at the same time.
   - **Force Device Manager Refresh** – Enables the device managers for HDS and EMC Symmetrix storage systems to obtain the latest information whenever getting discovery details. The management server obtains most of its information for HDS and EMC Symmetrix storage systems from their device managers. If the device managers do not have the latest information, the management server also displays the outdated information. For more information, see Excluding EMC Symmetrix Storage Systems from Forced Device Manager Refresh on page 109.
7. Select the **Enable** check box.
8. Set the date, time, and repeat interval for this task. For more information, see Setting the Date and Time for Scheduled Tasks on previous page.
9. Click **Next**.
10. Select the discovery groups you want included in the discovery:

    Only the elements in the discovery groups you select are included in discovery.

a. Select groups in the Available Discovery Groups section and click the **Add Selected Groups to Schedule** button to move them to the Discovery List section.

   The Create Discovery Schedule page offers a set of filters to help you find discovery groups quickly. For more information, see Filters for Discovery Schedules below.

b. Click **Finish**. The scheduled Get Details operation appears in the list of scheduled discoveries.

Filters for Discovery Schedules

The filter area is collapsed by default. To expand the filter area, click the **+** symbol. The following filters are supported:

- **Discovery Group Name Contains** – Use this filter to retrieve all the discovery groups whose name contains the specified string.
- **Element Name Contains** – Use this filter to retrieve all the elements whose name contains the specified string.
- **Discovery Group Type** – Use this filter to see only discovery groups of the specified type.
- **Element Type** – Use this filter to see only discovery groups that contain the specified element type.
- **Schedule Status** – Use this filter to see only discovery groups with the specified schedule status.

To apply the filter settings, click **Filter** to refresh the content of the page.

To restore the filters to their default settings, click **Reset** and refresh the page.

Disabling a Discovery Schedule

To disable a schedule for getting SAN topology details:

1. Access the Discovery page by selecting **Configuration** > **Discovery**.
2. Click the **Edit** ( ) button corresponding to the discovery schedule you want to disable.
3. Deselect the **Enable** option.
4. Click **Next**.
5. Click **Finish**. The schedule is disabled.

Editing a Discovery Schedule

To edit a schedule:

1. Access the Discovery page by selecting **Configuration** > **Discovery**.
2. Click the **Edit** ( ) button corresponding to the discovery schedule you want to modify.
3. If necessary, change the following properties:

- Name
- Description
- Type of discovery
- Schedule

4. Click **Next**.
5. If necessary, change the discovery groups you want assigned to the schedule.
6. Click **Finish**.

Removing a Discovery Schedule

To remove a schedule:

1. Access the Discovery page by selecting **Configuration** > **Discovery**.
2. Click the **Delete** ( ) button corresponding to the discovery schedule you want to remove.
3. When prompted to confirm, click **OK**. The schedule is removed.

Modifying Collector Settings for Newly Discovered Elements

The management server is capable of collecting many different types of data. Instead of using a single large process, these data are gathered using many different collectors. You can decide whether all or some of these collector schedules should start or stop when a new element is discovered.

To review the list of collectors that are available for each element type, access the Discovery page by selecting **Configuration** > **Discovery**, and click the **Collector Settings** tab.

The Default Collector Settings page provides a comprehensive listing of what collectors are available for each element type and what category each collector is classified as.

To help you find your collectors quickly, this page offers a set of filters. The filter area is collapsed by default. To expand the filter area, click the **+** symbol. The following filters are supported:

- **Element Type** – Retrieves the list of collectors designed to support the specified element types.
- **Collector Category** – Specifies whether you want to see a listing of all the collectors, or only collectors designed to gather performance statistics, report data, or capacity data.

To apply the filter settings, click **Filter** to refresh the content of the page. To restore the filters to their default settings, click **Reset**, and refresh the collector page.

You can modify the default collector settings for all future discovery elements by clicking the **Edit** ( ) button, or by selecting one or more collectors using the checkboxes and clicking **Edit Selected**.

The following properties are available when you modify default collector settings:

- **Enable selected collectors for newly discovered elements** – Each collector is designed to gather data for a specific type of elements. If you select this checkbox, the management server will start the selected collectors automatically whenever it discovers a system matching the element type supported by a selected collector. This means that the start time of the selected collectors for the newly discovered system is set to the same time as when the system is discovered.

  If you do not select this checkbox, the collectors are still available to the discovered elements, but the collector schedules are stopped by default. You can start any of the collectors after a system is discovered using one of the Data Collector Configuration pages. For example, **Configuration** > **Report, Configuration** > **Performance**, or **Configuration** > **Backup**.
- **Repeat Interval** – Specifies how often the collector should run on a recurring basis when the management server starts a collector schedule for an element type.

Managing Product Health

To obtain information from Product Health:

1. Add the management server to your discovery list:
   - Select **Discovery** > **Setup**.
   - Click the **Monitoring Product Health** link next to Step 1.
   - Click **Add**.
2. Discover the management server and include it in Get Details.

The Product Health menu option helps you monitor and manage the management server. At installation, a CIM extension is automatically installed on the management server so you can monitor the management server just as you would any other host.

Product Health does the following:

- **Disk space monitoring** – This feature keeps track of the management server's use of disk space. See Enabling Disk Space Monitoring on the facing page.
- **Database alert log** – The Database Alert Log scans the management server for critical errors at a specified interval and displays the information in its own chart. This setup frees up Event Manager for monitoring other elements. See Enabling the Scanning of Critical Events of the Management Server Database on page 352.
- **Log files** – You can view and download product logs, as described in Accessing the Log Files on page 345 and Downloading Logs Through the User Interface on page 348.
- **Scheduled RMAN backups** – This feature lets you schedule RMAN backups. See Scheduling RMAN Hot Backups on page 386. If the buttons on the RMAN Backup page are disabled, the product is set to No archive mode. See Changing the Archive Mode on page 395 for more information about changing the archive mode.

- **Advanced** – This feature lets you modify advanced settings so you can configure the product to run optimally in your environment. See Advanced Settings on next page and Customizing Properties on next page for more information.

Enabling Disk Space Monitoring

You can configure the management server to monitor itself just as it would any other element. This feature enables you to monitor the amount of free space the management server has left. The management server uses disk space for many of its operations, such as when it collects performance data, gathers element properties, generates events, and creates a backup.

To obtain information from Product Health, you must have already discovered the management server and obtained element details from it. For more information on how to discover a host, see Step 1 – Discovering Your Hosts and Backup Manager Hosts on page 179.

1. Make sure you performed Get Details for the management server. Discover the management server in the same way you would discover a host.
2. Select **Configuration** > **Product Health**.
3. Set the date, time, and repeat interval for this task. For more information, see Setting the Date and Time for Scheduled Tasks on page 338.
4. Click **Save Changes**.
5. To view the results of the monitoring, click the **Results** tab periodically.

Viewing the Results of Disk Space Monitoring

To make sure your management server does not run out of space, you should check the results for disk space monitoring.

The Results tab appears empty if the management server has not been included in Get Details. To obtain information from Product Health, you must have already discovered the management server and obtained element details from it. For information on how to discover a host, see Step 1 – Discovering Your Hosts and Backup Manager Hosts on page 179.

To access the results for disk space monitoring:

1. Select **Configuration** > **Product Health**.
2. Select **Disk Space**in the tree.
3. Click the **Results** tab in the Disk Space window.

The following information is displayed:

- Date/Time
- Disk Capacity
- Free Space
- Database Files
- Archive Files

- RMAN Files
- Temp Tablespace

Advanced Settings

The advanced page enables you to modify:

- **Product properties**. See Customizing Properties below for more information on how to modify product properties.
- **Java Memory Settings**. Do not modify the Java memory settings on the Advanced page (**Configuration** > **Product Health** > **Advanced**) unless instructed to do so by Technical Support. Incorrectly changing these settings could adversely affect the performance of the software.

Customizing Properties

The management server enables you to modify its properties. These properties control a variety of functionality, such as the ability to specify the number of time-outs for switches. You can always view the default setting of the properties by accessing the Default Properties page.

To customize properties:

1. Select **Configuration** > **Product Health**.
2. Click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the commands you want to modify. How you copy the text depends on your Web browser.
5. Return to the Advanced page.
6. Paste the copied text into the Custom Properties box. How you paste the text depends on your Web browser.
7. Make your changes in the Custom Properties box. To make sure the property is not commented out, remove the hash (#) symbol in front of the property.
8. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

Managing Logging

This section contains the following topics:

- Accessing the Log Files on the facing page
- About Log Files on the facing page
- Emailing Logs to Customer Support on page 347
- Downloading Logs Through the User Interface on page 348

Accessing the Log Files

You can obtain information about the software and CIMOM's transactions in the log files in the `%MGR_DIST%\logs` directory. CIMOM is a component in the CIM management infrastructure that handles the interaction between management applications and providers. There is a trace for the XML received from a CIMOM. The log files might contain information that is difficult to understand. HP recommends sending your log files to Customer Support, as suggested in Emailing Logs to Customer Support on page 347.

Some log files are more appropriate for your specific needs, and some are more useful for troubleshooting or other support needs. Some are for internal use only. The most useful ones for you are the following:

- `appstorm.<timestamp>.log` – Provides information about the transactions in the software, including web messages, EJB information, and general exceptions.
- `cimom.log` – Provides information about threads with CIMOM, such as provider transactions.

To view all log files, save them in a zip file, as described in Downloading Logs Through the User Interface on page 348.

About Log Files

On the management server, the following log files roll over on startup, at the start of a new day (midnight), and by size:

- `appstorm.<timestamp>.log`
- `AppstormProvisioning.log`
- `AppstormRemoteConsole.log`
- `Discovery.log`
- `GAEDSummary.log`
- `LicenseChanges.log`
- `userAudit.log`
- All CIMOM logfiles

- `GAEDDetails.csv` – The value of the custom property enableGaedDetailsLog can be set to true to enable, or false to disable, Discovery Details logging. By default, the enableGaedDetailsLog value is set true.
- `CIMOMDetails.csv` – To enable the CIMOM detail logging, set the priority to TRACE1 in cimomlog4jbase.properties. Otherwise, the priority can be set to DEBUG, WARN or ERROR. By defaul,t the priority value is TRACE1.

The following provide information about log file timestamps, sort criteria, configurable parameters, and adding a trace for the XML received from a CIMOM.

- **Log file timestamp** – A timestamp (YYMMDD-HHMMSS) is inserted into the filename at its creation, making its origin more quickly identified; for example, appstorm.20071012-122025.log.
- **Log file sort criteria** – Logfiles sort in order of their creation, based upon the timestamp in their filename.
- **Log file configurable parameters** – Configurable parameters for all log files are the following:
  - Maximum size of the logfile before it rolls over (MaxFileSize). This parameter resides in log4j.xml and is used to limit the size of an individual log file.
  - Maximum total amount of space the logfile can use (MaxTotalSize). This parameter resides in log4j.xml, and it is is used to limit the total size of a set of log files, such as appstorm logs.

Log file "appenders" manage the log file rollover when the MaxFileSize and MaxTotalSize parameters are reached. These parameters can be changed for any log file by using the appstorm.<timestamp>.log appenders at the following directory:

```
<management_server_install_
directory>/JBossandJetty/server/appiq/conf/log4j.xml
```

```
At the log4j.xmldirectory indicated above, change the appender values
to the new desired values:
```

```
<parem name="MaxFileSize" value="100MB"/> <!--Max size of a file
before it is rolled over -->
```

```
<parem name="MaxTotalSize" value="900MB"/> <!--Max size of a all
these log files, oldest is deleted when size is exceeded -->
```

Example 1 – Log File Rollover Based on Size

Assume the `appstorm.<timestamp>.log` file MaxFileSize=100MB and MaxTotalSize=900MB.

If the size of the current appstorm.<timestamp>.log file exceeds 100MB before the next day starts, a new appstorm.<timestamp>.log file is created.

If any rollover occurs, and the total size of all appstorm.<timestamp>.log files exceeds 900 MB, the oldest appstorm.<timestamp>.log files are deleted until the total size is below 900 MB.

Whenever a time-based or size-based rollover occurs, a footer is appended to the current file, and a header is placed on the new file. These headers and footers describe why the rollover occurred and the logfile to, or from, which it is being rolled.

Example 2 – Log File Rollover Based on Time

Assume a new day occurred. The current logfile (appstorm.20071012-154625.log) would receive this footer:

```
****Log File Rollover due to Time****

*****Next Log

File:<Installation_Directory>/logs/appstorm.20071013-000055.log****
```

The next logfile (appstorm.20071013-000055.log) would receive this header:

```
****Log Fle ROllover due to Time****

****Previousl Log

File:C:/hp/StorageEssentiaals/logs/appstorm.2007-154625.log****
```

Adding trace for XML received from CIMOM – Traces are normally very large files. The trace is, therefore, turned off by default. To add a trace, go into the properties file in the following directory:

```
%/JBOSS4_DIST%\server\appiq\conf
```

At the conf directory, uncomment the following line by deleting the pound sign (#):

```
#wbem.debug.sml=1
```

After uncommenting the line, set the level to at least 3 for the XML traces to be written. After they are written, you can go to the `/JBOSS4_DIST%\bin` directory to view them.

Emailing Logs to Customer Support

You can easily send logs to Customer Support with the e-mail logs feature, which saves the log files in a zip file and sends it to Technical Support.

Before you can use this feature, you must make the management server aware of your e-mail server. See Enabling Email Notification on page 335 for more information.

To send logs to Technical Support:

1. Select **Configuration** > **Product Health**.
2. Select **Log Files** in the Product Health tree.
3. Click **E-mail Logs**.

Downloading Logs Through the User Interface

If you run into problems with the management server, use the Download Logs feature to track the problem. This feature saves all the log files in a zip file, which is then stamped with the date and time (24-hour clock).

Some of the log files are generated only when you run certain features. For example, the reports.log file is only generated when you run reports.

To save all logs to a file:

1. Select **Configuration** > **Product Health**.
2. Select **Log Files** in the Product Health tree.
3. Click **Download Logs**.
4. When asked if you want to open or save the file, save the file.
5. Enter a name for the *.zip file, and select the directory to which you want to save the file.

   Make sure that the zip file is saved to a location other than the local disk drives of the management server.
6. Click **Save** in the Save As window.
7. If a blank page is displayed in the web browser, close this blank page.

Downloading Logs Through the Command Line

In addition to the Download Logs feature, there is also an automated process of gathering and downloading logs that is accessible from a command-line utility. This command-line utility is helpful in a situation when the user is unable to access the Download Logs feature from the user interface. For example, if HP Storage Essentials is unable to start.

(*Windows management servers only*) You can quickly gather system information and log files for troubleshooting by running the srmCapture.cmd program in `<installation directory>/tools`. The program provides a date and time-stamped zip file with this information.

The srmCapture.cmd program requires that zip.exe be in the same folder as srmCapture.cmd. If you are missing zip.exe, you can find it in the tools directory in both the ManagerCDLinux and ManagerCDWindows directories on the StorageEssentialsDVD.

To run the srmCapture.cmd program:

1. Open a command prompt window on the Windows management server, and go to the `<installation directory>/tools` directory.
2. Type the srmCapture command. The command has several parameters:

   ```
   srmCapture [/nowait] [/listmodules] [/?] [/help] [/usage]
   ```

- `/nowait`

  Non-interactive mode. The srmCapture command runs without prompting you with the message "press any key to continue."

- `/listmodules`

  Shows the dll files in use by each process (written to srmListProcesses.txt). If you use the /listmodules parameter, you must also include the /nowait parameter.

- `/?, /help or /usage`

  Provides information on how to use srmCapture.

The following are examples of srmCapture commands:

- `srmCapture`
- `srmCapture /?`
- `srmCapture /nowait`
- `srmCapture /nowait /listmodules`

The following information is gathered by srmCapture.cmd:

- List of environment variables, look for the `srmListEnvVar.txt` file.
- Results from running `ipconfig /all`, look for the `srmListIpconfigAll.txt` file.
- Results from running `netstat -noab`, look for the `srmListNetstatNoab.txt` file.
- Results from running `netstat -rte`, look for the `srmListNetstatRte.txt` file.
- Results from running `netsh diag show test`, look for the `srmListNetshDiagShowTest.txt` file.
- Install wizard log files (all files are located in `%systemdrive%\srmInstallLogs`).
- srmwiz.ini
- Oracle export log file
- File SRM log files
- File SRM configuration files
- Oracle log files
- Zero G registry content

If a message similar to `Current location, d:\Tools, is not writable` appears, the current working subdirectory is not writable. The srmCapture.cmd program goes through the following directories, in order, until it finds one that is writeable:

1. `%temp%`
2. `%tmp%`
3. `%systemdrive%`

Downloading the User Audit Log

You can determine who has been accessing your machine by viewing the user audit log.

To access the user audit log:

1. Select **Configuration** > **Product Health**.
2. Click **Download User Audit Log**.
3. Save the zip file.
4. Unzip the zip file.
5. Open the text file in a text editor. If you are using Notepad, you might need to select **Format** > **Word Wrap**.

Information is displayed as follows:

```
[2010-05-09 09:22:24] INFO [admin/1000/computername.companyname.com]
login succeeded
```

In this instance:

- [2010-05-09 09:22:24] is the time and date on which the action occurred.
- INFO is the level of warning.
- [admin/1000/computername.companyname.com] shows that the user logged in as admin from computername.companyname.com. 1000 is the DNS name.
- login succeeded is the action that occurred.

Downloading the Discovery Summary Log

To access status information from Get Detailsand view the summary logs:

1. Select **Configuration** > **Product Health**.
2. Click **Download Discovery Summary**.
3. Save the zip file.
4. Unzip the zip file.
5. Open the `GAEDSummary.log` file in a text editor. If you are using Notepad, you might need to select **Format** > **Word Wrap**.

Displaying a Log File in a Command Prompt Window

The software ships with tail.exe, which can display and update a log file in a command prompt window. This utility is extremely helpful if you do not want to enable the option that lets the management server service interact with the desktop.

To display a log file in a command prompt window:

1. Open a command prompt window.
2. Go to the `%MGR_DIST%\logs` directory by entering the following at the command prompt:

   ```
   c:\>cd %MGR_DIST%\logs
   ```

3. Enter the following at the command prompt window:

   ```
   %MGR_DIST%\logs>tail -f appstorm.<timestamp>.log
   ```

   In this instance, appstorm.<timestamp>.log is the log file you want displayed in the command prompt window, and <timestamp> is the timestamp for the log file.

The tail.exe utility checks the file for updates and appends them to text in the command prompt window.

Tail.exe is a program distributed under the GNU General Public License. For more information, see http://www.gnu.org.

Changing the Provider Log Level

The management server obtains information from its discovered elements through providers, which communicate with the hardware interface. These providers provide by default superficial logging to the `%JBOSS4_DIST%\server\appiq\logs\appstorm.<timestamp>.log` file. You can change the level of logging provided by selecting the new level of logging from the **Provider Log Setting** menu on the Product Health page (**Configuration** > **Product Health**), selecting **Log Files** and then clicking **Apply**. Only one logging level option can be selected.

Use the following table as a guideline for the different options. Several of the options mention providers. A provider is software that gathers information from an element, such as a switch.

Logging Levels

Log level Option	Description	Use When You
Default Logging	Provides superficial logging	Do not want additional logging.
Detailed tracing of Brocade provider	Provides detailed logging of the providers used to gather information from the Brocade switch.	Have difficulty obtaining information from a Brocade switch.
Detailed tracing of CIM Object Manager	Provides detailed logging of the infrastructure that handles the interaction between management applications and providers. The CIM Object Manager supports services such as event notification, remote access, and query processing.	Have difficulty obtaining information from the CIM Object Manager. You are having difficulty obtaining information from event notification, remote access, and query processing.

Log level Option	Description	Use When You
Detailed tracing of CLARiiON provider	Provides detailed logging of the providers used to gather information from CLARiiON storage systems.	Have difficulty obtaining information from CLARiiON storage systems.
Detailed tracing of HDS provider	Provides detailed logging of the providers used to gather information from HDS storage systems.	Have difficulty obtaining information from HDS storage systems.
Detailed tracing of HOST/SERVER provider	Provides detailed logging of the providers used to gather information from hosts and servers.	Have difficulty obtaining information from hosts or servers.
Detailed tracing of all providers	Provides detailed logging of the providers, which gather information from the elements.	Have difficulty obtaining information from more than one type of element.
Detailed tracing of SYMMETRIX provider	Provides detailed logging of the providers used to gather information from EMC Symmetrix storage systems	Have difficulty obtaining information from EMC Symmetrix storage systems.

Enabling the Scanning of Critical Events of the Management Server Database

You can configure the management server to scan for only critical events occurring with the database for the management server at a specified time interval. The management server displays the results of these scans under **Configuration** > **Product Health** > **Log Files**.

To enable this type of scanning:

1. Select **Configuration** > **Product Health**.
2. Select **Log Files** in the Product Health tree.
3. Set the date, time, and repeat interval for this task. For more information, see Setting the Date and Time for Scheduled Tasks on page 338.

   The minimal interval you can schedule is one day. If you select **Hour(s)**, **Minute(s)** or **Second(s)**, you must enter an interval that equals more than a day. For example, if you select **Hour(s)**, you must enter a value of 24 or more; if you select **Minute(s)**, you must enter a value of 1440 or more.
4. Click the **Enable** option.
5. Click **Save Schedule**.
6. To view the results of the scanning, click the **Results** tab periodically.

Viewing the Results of Logging

You can view when an error occurred and at what time it was discovered by accessing the Results tab for logging.

To access the Results tab:

1. Select **Configuration** > **Product Health**.
2. Select Log Files in the tree.
3. Click the **Results** tab in the Log Files pane. The following are displayed:
   - Scan Date/Time
   - Error Occurred Time
   - Error Description

Managing the Display of Events

This section contains the following topics:

Controlling the Display of Cleared and Deleted Events

You can control how the management server displays events by modifying one or more of the following tasks:

- **Clearing and deletion frequency** – The frequency table determines how often the user interface in Event Manager removes events and marks events as cleared. Events always display as they occur in the user interface.

  Events are only removed or marked cleared when their automatic delay time is completed. See Configuring the Clearing of Events on next page, and Configuring the Deletion of Events on page 355.

- **Clearing of events** – You can determine how often events are marked cleared. For more information, see Configuring the Clearing of Events on next page.
- **Deletion of events** – You can determine how often events are deleted. By default, events are deleted every two weeks. See Configuring the Deletion of Events on page 355.

Modifying the Clearing and Deletion Frequency

You can modify how often the user interface in Event Manager removes events and marks events as cleared. Events are still displayed as they occur in the user interface.

Events are removed or marked cleared only when their automatic delay time is completed. See the topics, Configuring the Clearing of Events below and Configuring the Deletion of Events on the facing page.

Assume you set the clearing and deletion frequency to every 15 minutes, with the initial time at 11:00 a.m., so that every 15 minutes the management server checks for events marked for deletion and for clearing, and it updates the user interface accordingly. Now assume Informational events are set to be cleared every hour, and an Informational event occurs at 12:20 p.m. Exactly one hour later, the management server marks this event as cleared, but the user interface is not updated, because the frequency update of the user interface is every 15 minutes. If you look at the Event Manager at 1:35 p.m., the event is be marked as cleared.

To modify the clearing and deletion frequency:

1. Access the Events page by selecting **Configuration** > **Events**.
2. Set the date, time, and repeat interval for this task. For more information, see Setting the Date and Time for Scheduled Tasks on page 338.

   The higher the interval, the more demand there is on the management server.
3. Select the **Enable** option.
4. Click **Save Changes**.

Configuring the Clearing of Events

Depending on the severity of an event, the management server might mark the event as cleared after 60 minutes. Events designated as Major and Critical are never marked as cleared. You can change the time delay in clearing an event, and you can specify that the management server never marks an event as cleared.

You might want to have unimportant events marked as cleared rather than automatically deleted. Depending on how you configured the deletion of events, you can view the cleared events at a later time.

The default settings for clearing events are the following.

Severity Level	Default Time Delay to Clear the Event (Hours)
Unknown	1
Informational	1
Warning	1
Minor	1
Major	Never
Critical	Never

To change the default time delay to clear an event:

1. Access the Events page by selecting **Configuration** > **Events**.
2. Perform one of the following tasks:
   - If you never want an event of the specified severity level marked as cleared, select the **Never** option next to the severity level.

     Or

   - To change the delay time in clearing an event, select one of the following units of measurement from the combo box and enter the number in the adjacent box:
     - **Minutes**
     - **Hours**
     - **Days**
     - **Weeks**
3. Click **Save Changes**.

Configuring the Deletion of Events

By default, the management server automatically deletes events after two weeks. However, you can specify for each severity level different time periods for deleting events. For example, you can modify the management server to delete events with the Information severity level every two days. You can also modify the management server to never delete events with the Critical severity level.

To change the default time delay to delete an event:

1. Access the Events page by selecting **Configuration** > **Events**.
2. Perform one of the following tasks:
   - If you never want an event of the specified severity level automatically deleted, select the **Never** option in the Automatic Delete Delay column.

     Or

   - If you want to change the delay time in deleting an event, select one of the following units of measurement from the list, and enter the number in the adjacent box:
     - **Minutes**
     - **Hours**
     - **Days**
     - **Weeks**

     For example, to have events that are a week old deleted, enter **1** and select **Weeks** in the combo box in the Automatic Delete Delay column.
3. Click **Save Changes**.

Managing File System Viewer

For information about configuring File System Viewer, see the “Overview of File System Viewer” chapter in the *File Servers Guide*.

Managing Backup Collection

This section contains the following topics:

Viewing Collectors for Backup Servers

The management server uses collectors to gather information for Backup Manager. Backup Manager provides information about backup protection, such as whether last night’s backup was successful.

Keep in mind the following for Data Protector:

- By default the backup collection does not run when you start the CIM extension. The backup collection is triggered once Get Details runs.
- During the background collection, the following processes are involved:
- **Session background collection.** It runs every 15 minutes.
- **Media background collection.** It runs every 24 hours.
- If you are running CIM extensions from 6.0 Service Pack 4 (6.0.4), the starting of a 6.0.4 CIM extension triggers the background collector schedule. Get Details does not trigger backup collection unlike later releases.

To manage collectors for Backup Manager:

1. Select **Configuration** > **BackupProtection**.
2. Select one of the collection tabs. The three collection tabs are Image Collection, Sessions Collection, and Media Collection.

The following are displayed on the collection tabs:

- **Backup Manager** – Displays the names of the backup servers.
- **Next Scheduled Run** – Displays the next time the management server is scheduled to obtain image details from the backup server.
- **Interval in Minutes** – Displays how often the management server is scheduled to obtain image details.
- **Running** – A check mark means the collector is running.
- **Edit Schedule –** Lets you modify the collection details
- **Start Collectors** – Click this button to start the collectors.
- **Stop Collectors** – Click this button to stop the collectors.

Scheduling Backup Collection for Backup Managers

You can configure the management server to obtain information about your master backup servers at a set interval. You can obtain image details, session details, or media details.

- Make sure these collectors run at least daily, so the latest backup information is displayed in Backup Manager.
- The process of Get Details takes time. It is recommended you perform this procedure when the network is not busy.
- It is not possible to run data collectors on quarantined elements. Attempting to do so will result in exceptions in the `appstorm.<timestamp>.log` file.
- Session collection should be enabled more frequently than a complete backup collection. The default interval for session collection is 60 minutes.

To obtain details for a backup server:

1. Select **Configuration** > **Backup**.
2. Select one of the collection tabs.
3. Select the management servers for which you want to obtain details.
4. Click the **Start Collectors** button.
5. Set the date, time, and repeat interval for this task. For more information, see Setting the Date and Time for Scheduled Tasks on page 338.
6. Click **OK**. The management server gathers information about image, session, or media details.

Image collection is not applicable to EMC NetWorker or IBM Tivoli Storage Manager (TSM).

Note: To discover IBM Tivoli Storage Manager, create an admin user on the IBM TSM providing the same user name and password used for host discovery.

Editing the Schedule of Backup Collection

To change when the management server obtains details for a backup server:

1. Select **Configuration** > **BackupProtection**.
2. Select one of the collection tabs.
3. Click the **Edit** () button corresponding to the Backup Manager you want to modify for its collection of details.
4. Click the **Start Collectors** button.
5. Set the date, time, and repeat interval for this task. For more information, see Setting the Date and Time for Scheduled Tasks on page 338.
6. Click **OK**. The management server gathers information about image, session, or media details.

Setting the Backup Sessions Retention Period

To set the retention values for sessions to be stored in the database:

1. Select **Configuration** > **Backup**.
2. Click the **Retention Configuration** tab.
3. Enter the number of days you would like sessions to be retained. The retention period must be a minimum of 30 days to a maximum of 1098 days.
4. Click **Submit**. The retention period is set.

Session Monitoring

To view the current running session details for Backup Manager:

1. Select **Configuration** > **Backup**.
2. Click the **Session Monitoring** tab.

The following are displayed:

- **Backup Manager** – Displays the names of the backup servers.
- **Next Scheduled Run** – Displays the next time the management server is scheduled to obtain current running session details for the backup server.
- **Interval (Minutes)** – Displays how often the management server is scheduled to obtain current running session details.
- **Running** – A check mark means the collector is running.
- **Edit –** Enables you to modify the collection details.
- **Start Collectors** – Click this button to start the collectors.
- **Stop Collectors** – Click this button to stop the collectors.

Drive Monitoring

To view drive monitoring details for Backup Manager:

1. Select **Configuration** > **Backup**.
2. Click the **Drive Monitoring** tab.

The following are displayed:

- **Backup Manager** – Displays the names of the backup servers.
- **Next Scheduled Run** – Displays the next time the management server is scheduled to obtain drive monitoring details for the backup server.
- **Interval (Minutes)** – Displays how often the management server is scheduled to obtain drive monitoring details.
- **Running** – A check mark means the collector is running.
- **Edit –** Enables you to modify the collection details.
- **Start Collectors** – Click this button to start the collectors.
- **Stop Collectors** – Click this button to stop the collectors.

Viewing the Status of Backup Collection

The management server keeps track of the collections it completes on the discovered Backup Managers on the Status tab. The Status tab displays the time when the collection started and ended on a Backup Manager, in addition to its status.

To access the Status tab:

1. Select **Configuration** > **Backup**.
2. Click the **Status** tab.

Managing Reports

This section contains the following topics:

Architectural Overview of Report Views and Report Cache Refresh

Management server reports are based on report views (also called materialized views). Report views are data snapshots of the management server schema at a certain time. The management server has a refresh process to refresh the report views. You can change the frequency of how often the report views are refreshed, as described in Refreshing the Report Cache on page 365.

Suggestions for Scheduling the Report Cache Refresh

The report cache (report views) is snapshot data of a management server schema up to a point in time. The report cache refresh is time- and resource-consuming.

To ensure that the report cache has the latest data from Get Details, schedule the report cache refresh after Get Details finishes.

Schedule the report cache refresh during off-peak hours.

For information on how to refresh the report cache, see Refreshing the Report Cache on page 365.

After an initial installation, make sure you already ran Get Details. Refresh the report cache to populate it with the data from Get Details. If you do not refresh the report cache, the cache is still in its initial state, which is empty. If the cache is empty, the reports will be empty because they obtain their information from the Report Database, which obtains its information from the report cache.

For information about when to schedule Get Details, report cache refresh, and global copies, so that the latest data appears, see the online help for the Report Database Admin Utility.

Report Refresh Status

The management server has two types of views for its reports. During a report cache refresh, these views are updated. You can check the status of the following views:

- MVIEWCORE_STATUS – This table keeps track of the refresh status of the core views. The core views are the views starting with mvc, mvca and mvcs, as described in Materialized Views on page 919.
- MVIEW_STATUS – This table keeps track of the refresh status of the regular views, which are views starting with mv_, as described in Materialized Views on page 919.

To query the MVIEW_STATUS table:

1. Enter the following at the command prompt:

   ```
   Sqlplus appiq_system/password
   ```

   In this instance, password is the password for the appiq_system account.

2. Enter the following at the command prompt:

   ```
   Sql>col lastRefresh format a30
   ```

   This command ensures that the data is displayed in a readable format.

3. Enter the following at the command prompt, on one line, with a space between the closing parenthesis and `lastRefresh`:

   ```
   Sql> select mviewname, to_char(last_refresh_time,'mm/dd/yyyy
   hh24:mi:ss') lastRefresh,
   ```

4. Enter the following at the command prompt:

   ```
   status
   ```

5. Enter the following at the command prompt:

   ```
   from mview_status
   ```

6. Enter the following at the command prompt:

   ```
   order by 2;
   ```

To query the MVIEWCORE_STATUS table:

1. Enter the following at the command prompt:

   ```
   Sqlplus appiq_system/password
   ```

   In this instance, password is the password for the appiq_system account.

2. Enter the following at the command prompt:

   ```
   Sql>col lastRefresh format a30
   ```

   This command ensures the data is displayed in a readable format.

3. Enter the following at the command prompt, on one line, with a space between the closing parenthesis and lastRefresh:

   ```
   Sql> select mviewname, to_char(last_refresh_time,'mm/dd/yyyy
   hh24:mi:ss') lastRefresh,
   ```

4. Enter the following at the command prompt:

   ```
   status
   ```

5. Enter the following at the command prompt:

   ```
   from mviewcore_status
   ```

6. Enter the following at the command prompt:

   ```
   order by 2;
   ```

Managing Collectors for Reports

The management server uses data collectors to gather information for reports. To view a report, you must have its corresponding collector running, and your report cache must be up-to-date. See Refreshing the Report Cache on page 365 for details.

To view collectors for reports, select **Configuration** > **Reports** and click the **Data Collection** tab.

For each element, the Report Data Collectors page displays all the collectors that the management server uses to gather report data. You can decide whether data should be collected for each element by starting or stopping the collector schedule. If a collector is gathering data when you stop it, that collector will complete the collection. You are stopping all future runs of the collector; not the currently active process. The Enabled column reflects whether you've started your collector schedule or not.

The Report Data Collector page offers a set of filters to help you find your collectors quickly. The filter area is collapsed by default. To expand the filter area, click the **+** symbol. The following filters are supported:

- Element Type – Use this filter to see the collectors for a specific element type or for all elements.
- Collector State – Use this filter to see all schedules or only the schedules that have been started or stopped.
- Element Name Contains – Use this filter to retrieve all the elements whose name contains the specified string.

To apply the filter settings, click **Filter** to refresh the content of the Report Data Collector page. To restore the filters to their default settings, click **Reset** and refresh the collector page.

In addition to changing the collector schedule after an element has been discovered, you might want to decide whether a collector schedule should be started or not for future discovery elements. You can access the default collector schedule settings by clicking the **Default Collector Settings** link above the Filter area, or select **Configuration** > **Discovery** and click the **Collector Settings** tab. For more information about the Collector Settings tab, see Modifying Collector Settings for Newly Discovered Elements on page 341.

Keep in mind the following:

- All collectors are stopped during Get Details. This means that during Get Details data is not updated.
- If a collector is scheduled to gather data while Get Details is occurring, the management server does not let that collector begin its task of gathering data. Collectors that were gathering data before Get Details can complete their task.
- You can set the interval for performance and report collectors to less than 15 minutes; however, setting the interval to less than 15 minutes might slow the performance of the product, increase memory usage and expand the size of the database on the management server.

It is not possible to run data collectors on quarantined elements. Doing so results in exceptions in the `appstorm.<timestamp>.log` file.

The Data Collection tab provides the following information.

Column Heading	Description
Element	Displays the name of the element from which this collector gathers information.
Element Type	Displays the type of element from which the collector gathers information.
Collector Type	Lists the type of collector that is responsible for providing information about an element.
Enabled	Displays the status of the collector. Collectors that are running display a check mark in this column.
Interval (Minutes)	Displays the interval in minutes between collector runs.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Edit	To edit the schedule for running a collector, click the **Edit** () button. For more information, see Editing a Collector Schedule on page 377.
Action	Displays one of the following buttons: • **Stop** – Stops the collector. The corresponding reports display only information gathered previously. See Stopping Report Collectors on next page. • **Start** – Starts the collector. When you start a collector, it begins gathering information for its corresponding reports. See Starting Stopping Collectors below.

Starting Stopping Collectors

After you click **OK**, if the date and time you set has not passed, the collector starts immediately. If the set time has passed, the collector starts 2 minutes after you click **OK**.

To start a collector:

1. Select **Configuration** > **Reports** and click the **Data Collection** tab.
2. Click the **Start** button corresponding to the collector you want to start.

   To start more than one collector at once, select more than one collector and then click **Start Collectors**.
3. Set the date, time, and repeat interval for this task. For more information, see Editing a Collector Schedule on page 377.
4. Click **OK**.

Stopping Report Collectors

When you stop a collector, the management server stops gathering the information for which the collector is responsible. For example, if a collector is not running, its corresponding reports are no longer receiving information to display. One of the following occurs:

- If there was originally no information gathered for the report, no data appears in the report.
- If information was previously gathered for the report, old data appears in the report.

To stop a collector:

1. Select **Configuration** > **Reports** and click the **Data Collection** tab.
2. Click the **Stop** button corresponding to the collector you want to stop. The collector stops gathering information for its corresponding reports.

   To stop more than one collector at once, select more than one collector and then click **Stop Collectors**.
3. Click **OK**.

Viewing Data Aging Statistics for Reports

Data aging includes Data Rollup and Garbage Collection. Data Rollup controls how often a set of data is summarized. For example, hourly data is rolled into the daily table periodically. Garbage Collection refers to how long a set of data is preserved before it is permanently removed from the database.

The settings on the Data Aging page control data aging for both reports and performance statistics.

Do not modify the data on the Data Aging page unless instructed by Customer Support. Changing them incorrectly can adversely affect the management server.

To view data aging statistics:

1. Select **Configuration** > **Reports**.
2. Click the **Data Aging** tab at the top of the screen. Data aging statistics are displayed in the table.

Perform the following steps *only* if customer support instructed you to modify one of the collectors on the page.

1. Click the **Edit** () button.
2. Set the date, time, and repeat interval for this task. For more information, see Setting the Date and Time for Scheduled Tasks on page 338.
3. Click the **Enable** option.

   If you are not allowed to disable the collector, the Enable option is unavailable.
4. *Garbage Collection only*: To change how long the data is preserved, enter an interval in the

Preserve box, and then select one of the following from the list to the right of the box:

- **Second(s)**
- **Minutes**
- **Hours**
- **Days**
- **Weeks**

5. Click **OK**.

Refreshing the Report Cache

The management server gathers information for reports from the database every 6 hours. The management server stores this information in its report cache and displays it when a report is requested. If you are seeing outdated information in the report, you can refresh the report using one of the following methods:

- **Refresh the report cache now** – See Refreshing the Report Cache Immediately on page 369.
- **Schedule the report cache to be refreshed** – See Scheduling a Report Cache Refresh on page 370.

For information about data gathered from the report cache, see About the Report Cache Tab below.

About the Report Cache Tab

The Report Cache tab displays detailed status of the report cache refresh activity. In case of failure during the report cache refresh, the Report Cache tab helps you identify the root cause.

The Report Cache tab lets you refresh the report cache manually and by scheduling, as described in Refreshing the Report Cache Immediately on page 369 and Scheduling a Report Cache Refresh on page 370.

You can view the details of the last successful report cache refresh, its status, details about the Path Snapshot and batch update, and detailed materialized views status.

The Report Cache tab provides the following information:

- Last Refresh Summary and Current Status information summarizing report cache refresh activity
- Collapsible Materialized Views Refresh Details section displaying details of the status of all the materialized views.

For information about obtaining a log of the report cache refresh, see Obtaining Report Cache Refresh Logs on page 371.

You access the Report Cache tab from **Configuration** > **Reports** > **Report Cache**.

The Last Refresh Summary section appears.

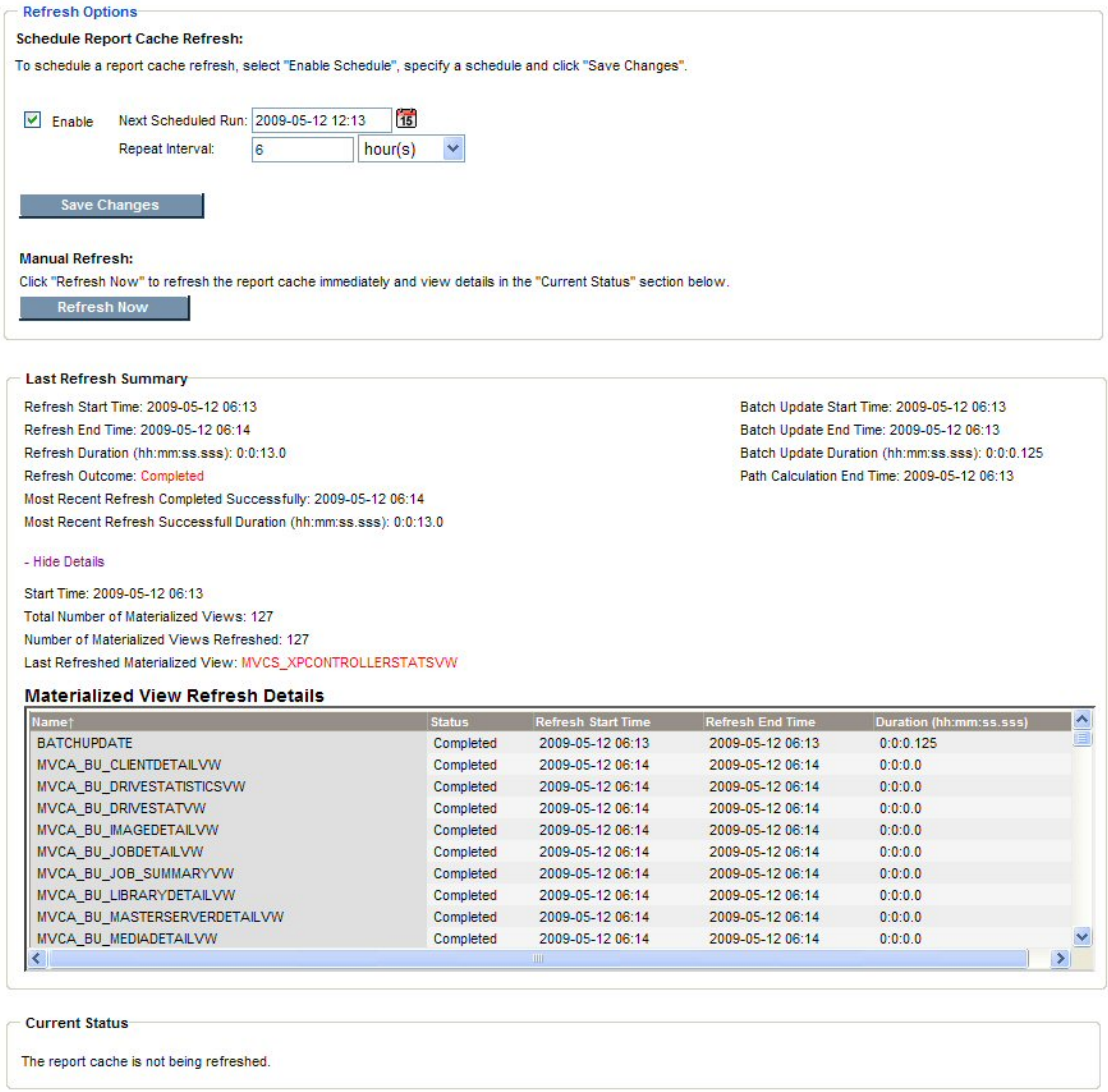

The collapsible details close up view also appaers.

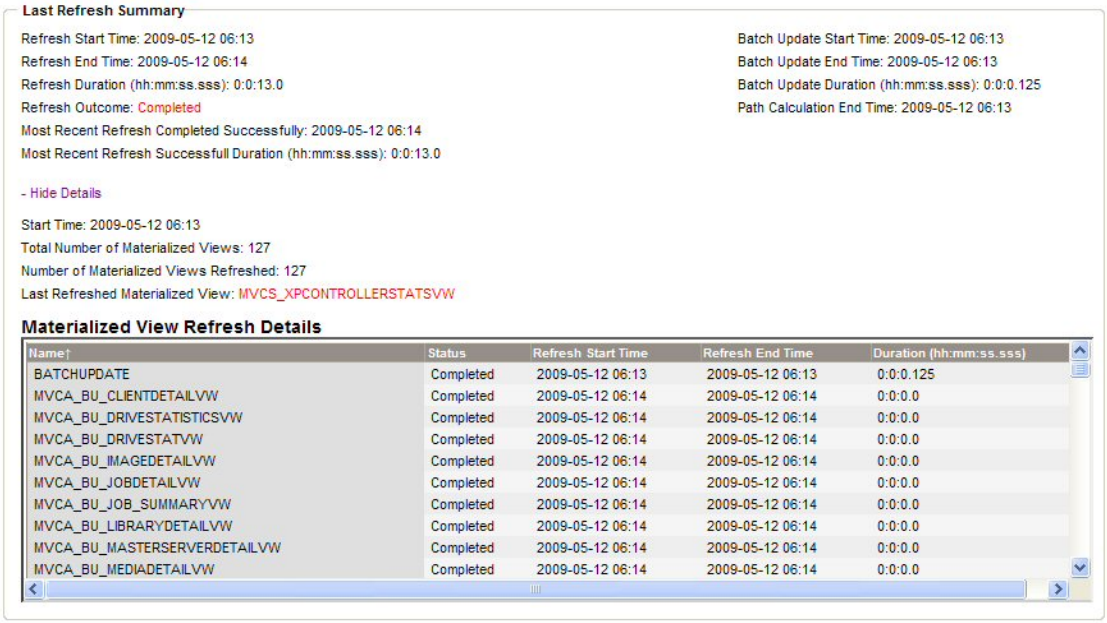

The Last Refresh Summary section displays the status of execution of the last report cache refresh task/process, and includes the following fields:

- **Refresh Start Time** – YYYY-MM-DD HH:MM format
- **Refresh End Time** – YYYY-MM-DD HH:MM format
- **Refresh Duration** – HH:MM:SS:SSS format
- **Refresh Outcome** – Successful or Failed
- **Most Recent Refresh Completed Successfully** – YYYY-MM-DD HH:MM format
- **Most Recent Refresh Successful Duration** – HH:MM:SS:SSS format
- **Batch Update Start time** – Batch update start time of refresh in YYYY-MM-DD HH:MM format
- **Batch Update End Time** – Batch update end time of refresh in YYYY-MM-DD HH:MM format
- **Batch Update Duration** – HH:MM:SS:SSS format
- **Path Calculation End Time**. Path Calculation end time of report cache refresh in YYYY-MM-DD HH:MM format
- **Start Time** – YYYY-MM-DD HH:MM format
- **Total Number of Materialized Views**
- **Total Number of Materialized Views Refreshed**
- **Last Refreshed Materialized View**
- **Materialized View Refresh Details**. Detailed information is displayed on the Report Cache

tab in a collapsible section, which is collapsed by default. Detailed information is presented in a table with the following columns:

- **Name** – Name of the materialized view
- **Status** – Not Yet Started (starting state after new install), Completed, or Failed
- **Refresh Start Time** – The start time of the last view refresh in YYYY-MM-DD HH:MM:SS format. This column will be blank in a new installation before the first execution of the report cache refresh process.
- **Refresh End Time** – The end time of the last view refresh in YYYY-MM-DD HH:MM:SS format. This column will be blank in a new installation before the first execution of the report cache refresh process.
- **Duration (in HH:MM:SS format)**. The time it took to refresh the materialized view during the last refresh. This column will be blank in a new installation before the first execution of the report cache refresh process.

From the Report Cache tab (**Configuration** > **Reports** > **Report Cache**), you can also access the Current Status section.

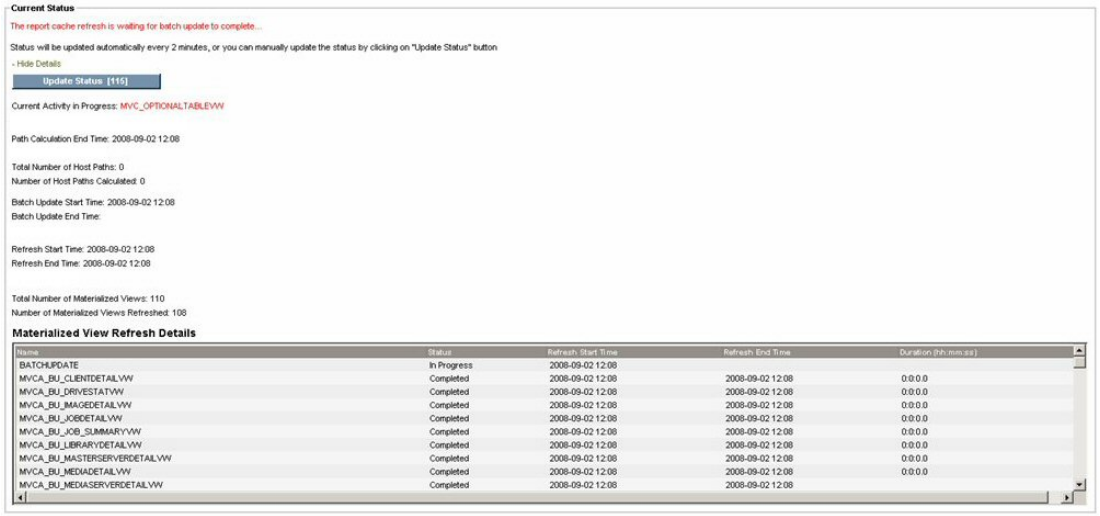

The Current Status section displays the status of currently executing processes of the report cache refresh task. The display contains a collapsible details section, and shows information only when a current report cache refresh is in progress; otherwise, a message saying that no report cache refresh activity is in progress is displayed. This section has the following sub-fields:

- **Status message** – Describes the progress of the report cache refresh activity.
- **Update Status button**. The button, which updates the status of each Materialized view, has a polling frequency of 120 seconds.
- **Current Activity in Progress** – Describes the current activity in progress.

- **Path Calculation End Time** – Path Calculation end time of report cache refresh in YYYY-MM-DD HH:MM format
- **Total Number of Host Paths (during Path Calculation)**
- **Number of Host Paths completed (during Path Calculation)**
- **Batch Update Start Time** – Batch Update Start time of report cache refresh in YYYY-MM-DD HH:MM format
- **Batch Update End Time** – Batch Update end time of report cache refresh in YYYY-MM-DD HH:MM format
- **Refresh Start time** – Start time of report cache refresh in YYYY-MM-DD HH:MM format
- **Refresh End Time** – End time of report cache refresh in YYYY-MM-DD HH:MM format
- **Total number of Materialized Views**
- **Number of Materialized Views Refreshed**
- **Materialized View Refresh Details** – Detailed information is displayed in a collapsible section which is expanded by default. The table contains the following columns:
  - **Name** – Name of the materialized view.
  - **Status** – Not Yet Refreshed (starting state after new install), In Progress (when it is being refreshed), or Completed (after it has finished refreshing).
  - **Refresh Start Time**. The start time of the last view refresh in YYYY-MM-DD HH:MM:SS format. This column will be blank in a new installation before the first execution of the report cache refresh process.
  - **Refresh End Time**. The end time of the last view refresh in YYYY-MM-DD HH:MM:SS format. This column will be blank in a new installation before the first execution of the report cache refresh process.
  - **Duration (in HH:MM:SS format)**. The time it took to refresh the materialized view during the last refresh. This column will be blank in a new installation before the first execution of the report cache refresh process.

If Get Details is occurring, wait for it to finish before clicking **Refresh Now**. This technique ensures the database is completely updated and thus your reports will be as accurate as possible. Get Details collects the latest data. When you refresh the report cache, the management server transfers the information collected from Get Details and transfers it to the report cache.

While the report cache is being refreshed, reports display no data.

Refreshing the Report Cache Immediately

To display the latest data in your reports, the report cache must be up to date. You can refresh the report cache immediately by clicking the **Refresh Now** button on Report Cache page (**Configuration** > **Reports** > **Report Cache**).

Refreshing the report cache frequently, such as every 10 minutes, could hurt the response time of the management server to perform other tasks.

If you find you are still viewing old information regarding elements on the network, you might need to perform Get Details. It is best to perform Get Details at regular intervals. See Adding a Discovery Schedule on page 338.

If you click the Refresh Now button while the report cache refresh process is in occurring, a message is shown in red in the status summary section stating, “The report cache is being refreshed.” Also, similar messages are displayed for Path Calculation and for Batch Update fields.

When you click the Refresh Now button, the following occurs:

- The Last Refresh Summary section is updated with the following information:
  - **Last Refresh Completed**. The time when the last report cache refresh completed, irrespective of the status. The time is in the YYYY-MM-DD HH:MM format.
  - **Refresh status**. Updates with report cache refresh status.
  - **Last Successful report cache refresh completed**. The time when the last report cache refresh completed. The time is in the YYYY-MM-DD HH:MM format. A report cache refresh is complete after all the materialized views have been refreshed.

  You can view the cache refresh of the materialized views by clicking **Show Details** in the Last Refresh Summary section.
- The Current Status section updates with the following additional information:
  - **Start Time**. The time (in YYYY-MM-DD HH:MM:SS format) when the report cache refresh process started.
  - **Number of materialized views refreshed**. Displays the number of views that are refreshed.
  - **Currently Refreshing Materialized View**. Name of the MVIEW that is currently being refreshed.

  The Current Status section displays “The report cache is not being refreshed.” once the report cache refresh completes.
- The **Update Status** button, which is a polling button in the details part of the status section is enabled and running. This button stops polling when the report cache refresh task completes.

For information about how to obtain a log of the report cache refresh, see Obtaining Report Cache Refresh Logs on the facing page.

Scheduling a Report Cache Refresh

When you schedule the refreshing of the report cache, keep in mind that the higher the frequency of the report cache interval, the more stress you put on the management server. A very frequent report cache interval, such as every 10 minutes, could hurt the response time of the management server to perform other tasks.

If you find you are still viewing old information regarding elements on the network, you might need to perform Get Details. It is best to perform Get Details at regular intervals. See Adding a Discovery Schedule on page 338.

To schedule a report cache refresh:

1. Select **Configuration** > **Reports**.
2. Click the **Report Cache** tab at the top of the screen.
3. Set the date, time, and repeat interval for this task. For more information, see Setting the Date and Time for Scheduled Tasks on page 338.
4. Click the **Enable** option.
5. Click **Save Changes**.

Obtaining Report Cache Refresh Logs

To obtain a log of the report cache refresh:

1. Click **Configuration** > **Product Health**. Expand the **Disk Space** node and click **Log Files**.
2. Click the **Download Report Cache Refresh Log** button.

   You are asked if you want to open or save a zip file. Click **Save**. The zip file contains the log files from the report cache refresh.

About Global and Local Reports

In releases previous to 6.2, global reports needed to be configured. This additional work is no longer required. The global reports in Report Optimizer contain information from all discovered management servers appear.

Report Optimizer also contains local reports. You can specify which management server you want to appear in those local reports by selecting a management server from the Local Report combo box. Refer to the online help for the Report Database Admin Utility.

Adding the Report Optimizer Server as a Trusted Site

If you are running Windows Server 2008 with the Internet Explorer Enhanced Security Configuration (IEESC) enabled, you must add the server running Report Optimizer as a trusted site.

When you access Report Optimizer directly, you are prompted to add the site as a trusted site.

When you access Report Optimizer from within HP Storage Essentials, you are not prompted to add the server as a trusted site and thus, you might run into difficulty with accessing Report Optimizer from within HP Storage Essentials.

To manually add Report Optimizer server as a trusted site:

1. In Internet Explorer, click **Tools** > **Internet Options** > **Security**.
2. Click **Trusted Sites** and then click **Sites**.

3. Add several variations of the server name. For example, if the server running Report Optimizer is named reportserver.usa.mycompany.com with IP address 192.168.1.1, you can enter the following variations of the site name:
   - The IP address of the server; in this example, http://192.168.1.1
   - The full name of the computer; in this example, http://reportserver.usa.mycompany.com
   - The computer name; in this example, http://reportserver

Managing Performance Collection

This section contains the following topics:

Managing Performance Collectors

The management server uses performance collectors to gather information for Capacity Manager and Performance Manager charts, as well as for monitoring.

To manage performance collectors, select **Configuration** > **Performance** and click the **Data Collection** tab.

For each element, the Performance Data Collectors page displays all the collectors that the management server uses to gather data. You can decide whether data should be collected for each element by starting or stopping the collector schedule. If a collector is gathering data when you stop it, that collector will complete the collection. You are stopping all future runs of the collector; not the currently active process. The Enabled column reflects whether you have started your collector schedule or not.

If you are not sure whether you want to start or stop a collector schedule, click **+Statistics** to see the list of performance statistics that the collector is designed to collect. Once the statistics are collected, you can use Performance Manager to review the data.

The Performance Data Collector page offers a set of filters to help you find your collectors quickly. The filter area is collapsed by default. To expand the filter area, click the **+** symbol. The following filters are supported:

- **Element Type** – Use this filter to see the collectors for a specific element type or for all elements. When you select a specific element type, the Collector Type filter is adjusted to show you the collector types that are meaningful to the selected element type.
- **Collector State** – Use this filter to see all schedules or only the schedules that have been started (Enabled) or stopped (Disabled).

- **Collector Type** – Use this filter to see all the collectors for a specific collector type or all collector types. This filter works in conjunction with the Element Type filter.
- **Element Name Contains** – Use this filter to retrieve all the elements whose name contains the specified string.

To apply the filter settings, click **Filter** to refresh the content of the Report Data Collector page. To restore the filters to their default settings, click **Reset** and refresh the collector page.

In addition to changing the collector schedule after an element has been discovered, you might want to decide whether a collector schedule should be started or not for future discovery elements. You can see the default collector schedule settings using the Default Collector Settings link above the Filter area, or select **Configuration** > **Discovery** and click the **Collector Settings** tab. For more information about the Collector Settings tab, see Modifying Collector Settings for Newly Discovered Elements on page 341.

Keep in mind the following:

- All collectors are stopped during Get Details. This means that during Get Details, data is not updated. Historical collectors, such as those available from the Configuration tab, are restarted when they are stopped during Get Details. Any charts that were active in Performance Manager when Get Details was started are not restarted.
- If a collector is scheduled to gather data while Get Details is occurring, the management server does not let that collector begin its task of gathering data. Collectors that were gathering data before Get Details can complete their task.
- You can set an interval for performance and report collector to less than 15 minutes; however, setting the interval to less than 15 minutes might slow the performance of the product, increase memory usage and expand the size of the database on the management server.

It is not possible to run data collectors on quarantined elements. Attempting to do so will result in exceptions in the `appstorm.<timestamp>.log` file.

The Data Collection tab provides the following information.

Column Heading	Description
Element	Displays the name of the element from which this collector gathers information.
Element Type	Displays the type of element from which the collector gathers information.
Collector Type	Displays the collector type. Click **+Statistics** to see the list of performance statistics that the collector is designed to collect.
Enabled	Displays the status of the collector. Collectors that are running display a check mark in this column.

Column Heading	Description
Interval (Minutes	Displays the interval in minutes between collector runs.
Next Scheduled Run	Displays the date and time when the collector is scheduled to run.
Edit	To edit the schedule for running a collector, click the **Edit** ( ) button. Then set the date and time. For more information, see Editing a Collector Schedule on page 377.
Action	Displays one of the following buttons: • **Stop** – Stops the collector. The corresponding reports display only information gathered previously. See Stopping Performance Collectors on the facing page. • **Start** – Starts the collector. When you start a collector, it begins gathering information for its corresponding reports. See Starting Performance Collectors below.

Starting Performance Collectors

To start a collector:

1. Access the page for performance collectors (**Configuration** > **Performance** > **Data Collection**).
2. Click the **Start** button corresponding to the collector you want to start.

   To start more than one collector at once, select more than one collector on the **Data Collection** tab and click **Start Collectors**.
3. Set the date, time, and repeat interval for this task. For more information, see Editing a Collector Schedule on page 377.

If you are asked to provide a proxy host:

1. Click **Browse**.
2. Select a proxy host from the menu, and then click **OK**.
3. Click **OK** again to set the time for starting the collector.
4. If you do not see any hosts displayed, verify that you have the latest CIM extension version installed and running on a host that can access the LSI storage system.
5. Click **OK**.

Stopping Performance Collectors

When you stop a collector, the management server stops gathering information for which the collector is responsible. For example, if a performance collector is not running, its corresponding statistics are no longer receiving information to display. One of the following occurs:

- If there was originally no information gathered for the statistic, no data appears for that statistic.
- If information was previously gathered for the statistic, old data appears for that statistic.

To stop a collector:

1. Access the page for performance collectors (**Configuration** > **Performance** > **Data Collection**).
2. Click the **Stop** button corresponding to the collector you want to stop. The collector stops gathering information for its corresponding reports.

   To stop more than one collector at once, select more than one collector and click **Stop Collectors**.

Viewing Data Aging Statistics for Performance

Data Aging includes Data Rollup and Garbage Collection. Data Rollup controls how often a set of data is summarized. For example, hourly data is rolled into the daily table periodically. Garbage Collection refers to how long a set of data is preserved before it is permanently removed from the database.

The settings on the Data Aging page control data aging for both reports and performance statistics.

Do not modify the data on the Data Aging page unless instructed by Customer Support. Changing them incorrectly can adversely affect the management server.

To view data aging statistics:

1. Select **Configuration** > **Performance** > **Data Collection**.
2. Click the **Data Aging** tab at the top of the screen.

   Perform the following steps *only* if Customer Support instructs you to modify one of the collectors on the page:

   a. Click the **Edit** () button.

   b. Set the date, time, and repeat interval for this task. For more information, see Setting the Date and Time for Scheduled Tasks on page 338.

   c. Click the **Enable** option.

      If you are not allowed to disable the collector, the Enable option is unavailable.
3. *Garbage Collection only*: To change how long the data is preserved, enter an interval in the

Preserve box, and select one of the following from the list to the right of the box:

- **Second(s)**
- **Minute(s)**
- **Hour(s)**
- **Day(s)**
- **Week(s)**

4. Click **OK**.

Editing the Locale and Currency Settings

The management server determines which languages and currency to display by looking at the language and currency settings for the operating system. You can override the management server's default locale and currency settings.

For example, assume the user interface for the management server is displayed in English, but you want to view it in Japanese. You could change the locale in the management server to Japanese without changing the locale setting of the computer running the management server.

When you change the locale and currency settings, the following occurs:

- The text in the product changes to the locale you set. For example, if you change the locale setting to French, the text in the product changes to French. Keep in mind the online help and PDFs stay in English regardless of the locale setting.
- The currency you selected is displayed in the product.

  The online help looks at the localization of the operating system instead of the localization settings of the management server to provide the corresponding language for some its components, such as the table of contents, index and search. If you set the locale to Korean and your client computer is set to Japanese, the help components listed above appear in Japanese.

Abbreviation	Full Name
CNY	Chinese Yuan
JPY	Japanese Yen
KRW	Korean Won
USD	U.S. Dollar

To change the locale and currency settings:

1. Select **Configuration** > **Locales**.
2. Choose a locale from the **Select Locale** menu.
3. Choose a currency from the **Select Currency** menu.

4. Click **Change Locale**.
5. Restart the management server to have the changes take effect.

Process Names

This section describes the process names on Windows and UNIX systems.

Process names on Windows

The following process names are displayed in the Windows Task Manager on the management server.

Process	Executable Name
Service Wrapper (launcher)	`AppStorMService.exe`
Application Server (JBoss)	`storApplicationServer.exe`
CIMOM for Default Discovery	`storCimomDefault.exe`
CIMOM for Discovery	`storCimom2.exe`
CIMOM for Discovery	`storCimom3.exe`

The executable names for the CIMOM's corresponding to the other discovery groups follow the same pattern as those described above.

Process Names on UNIX Systems

The following process names are displayed in UNIX systems by running the ps -ef command.

Process	Executable Name
Service Wrapper (launcher)	`AppStorMService.sh`
JBoss	`storApplicationServer.sh`
CIMOM for Default Discovery Group	`storCimomDefault.sh`
CIMOM for Discovery	`storCimom2.sh`
CIMOM for Discovery	`storCimom3.sh`

The executable names for the CIMOM's corresponding to the other discovery groups follow the same pattern as those described above.

If you are using the prstat utility, all of the processes are named java.exe.

Editing a Collector Schedule

There are several ways to modify a collector schedule:

- Locate the collector in the collector table and click the **Edit** () button in that collector's row.
- Locate the collector in the collector table and select the collector using the check box in the first column of the table. You can select one or more collectors this way. Click the **Edit Selected** button above the collector table to modify the selected collector's schedules at the same time.
- To modify all of the collectors displayed on the same page of the collector table, click the check box in the first column header, and then click the **Edit Selected** button above the table.
- To modify all the collectors displayed on all of the pages, click the **Select All Pages** link above the collector table, and then click the **Edit Selected** button above the table.

When you edit a collector schedule, you can change the following schedule properties:

- **Next Scheduled Run** – Specifies a start time for when the collector should run.
- **Repeat Interval** – Specifies how often the collector should run on a recurring basis.
- **Spread Start Time** – This checkbox can only be selected when more than one collector is being edited. If you select this checkbox, the management server automatically adjusts the start time of the selected collectors and spreads the load of the collectors across the specified repeat interval.

Creating Schedules Using Windows Task Scheduler

For Windows-based management server environments, you can create EVA Quarantine schedules using the Windows Task Scheduler.

To create EVA Quarantine schedules:

1. Create a new Windows scheduled task (**Start** > **Control Panel** > **Scheduled Tasks** > **Add scheduled task**).
2. Click **Next** and browse to the location of `AppIQSet.bat`.

   For Windows-based management servers, this path is typically from the sub-directory for the management server to `\cli\bin\AppIQSet.bat`.
3. Set the time you want to unquarantine the EVA access point.
4. Check the box to open **Advanced task settings**.
5. Add `-unquarantine <name(s) of access points>` as arguments after the batch file in the **Run** field; for example, from the management server sub-directory to `\cli\bin\AppIQSet.bat -domaindiscovery -unquarantine evahost1`.
6. In the management server, schedule a Get Details to occur after the EVA provider has had enough time to refresh its cache when it is unquarantined. Generally, 30 minutes is sufficient for this interval.
7. Follow the previous steps to create a second Windows task to quarantine the access point again, after the Get Details is completed.

9 Database Maintenance and Management

This chapter contains information about backing up and restoring the database using the Database Admin Utility.

This section contains the following topics:

Scheduling Database Maintenance Operations

Schedule a maintenance window of two to four hours weekly during off peak hours of operation to do the following database maintenance operation:

1. Stop the service for the management server. See Restarting the Service for the Management Server on page 68.
2. Access the Database Admin Utility. See Accessing the Database Admin Utility on page 389 for more information.
3. Using the Database Admin Utility, verify that the database is in an open state and the listener is running. See Checking the Database Status on page 390.

   If the database state is not in an open state and shows an error, obtain the following logs and contact technical support:

   Log file for the Database Admin Utility:

   `%MGR_DIST%/logs/dbAdmin.log`

   Database alert log, which can be found in one of the following locations:

   Windows:

   `\oracle\admin\APPIQ\bdump`

UNIX:

```
$ORACLE_BASE/admin/APPIQ/bdump
```

4. Reset the temporary tablespace, as described in Resetting the Temp and Undo Tablespace on page 393.
5. Export the database, as described in Exporting the Database on page 391.
6. If the database is running in archive mode, set the database to no-archive mode, as described in No-Archive Mode on page 383 and then clean the archive directory: `\oracle\oradata\APPIQ\archive` on Windows and `$ORACLE_HOME/oradata/APPIQ/archive` on UNIX systems.

   The archive directory only exists if you previously set the management server to archive mode.
7. Return the database to archive mode, as described in Changing the Archive Mode on page 395.
8. Start the service for the management server.
9. If the database is in archive mode, take a current RMAN backup by clicking the **Backup Now** button on the **Configuration** > **Product Health** > **Disk Space** > **RMAN Backup**. See Performing an RMAN Hot Backup on page 385 and Architectural Overview of RMAN Backups on page 384.

   On successful completion of RMAN backup, the backup is saved to the following directory:

   Windows:

   ```
   %ORA_HOME%\rman\current
   ```

   UNIX:

   ```
   $ORACLE_HOME/rman/current
   ```
10. Clean the following folders:

    Windows:

    ```
    %ORA_HOME%\rman\backup1
    %ORA_HOME%\rman\backup2
    ```

    UNIX:

    ```
    $ORACLE_HOME/rman/backup1
    $ORACLE_HOME/rman/backup2
    ```

Backup Directories in%ORA_HOME%\rman

Directory	Contains
current	The backup when the **Database Server Backup** button was clicked
log	A log of when the backup was done

Directory	Contains
backup1	Information from the automatic backup (alternating day)
backup2	Information from the automatic backup (alternating day)

Overview of Backups

The management server provides the following backups.

Export and RMAN backups should be done regularly and in combination.

Backup Type	Description	Files Backed Up	Database Mode
Support Database	Done through the Database Admin Utility. See Generating a Support Database on page 399 for more information about generating a support database.	Database Schema, Oracle network configuration files (tnsnames.ora, listener.ora), CIM repository, File System Viewer Passwords for all devices, managed servers and databases are removed from the backup.	Does not matter
Export backup	Done through the Database Admin Utility. See Exporting the Database on page 391 for more information about exporting the database.	Database Schema, Oracle network configuration files (tnsnames.ora, listener.ora), CIM repository, File System Viewer.	Does not matter
RMAN HOT backup	The backup is referred to as being “hot” because the management server is still running while the backup is occurring. You can configure the RMAN backup to run by default. See Scheduling RMAN Hot Backups on page 386.	Database files, Control files, Redo files, Archive files, Oracle network Configuration files (tnsnames.ora, listener.ora), CIM repository, File System Viewer.	To do an RMAN hot backup, the management server must be set to archive mode. See Changing the Archive Mode on page 395.

Backup Type	Description	Files Backed Up	Database Mode
RMAN Cold backup	Done through the Database Admin Utility.	Same files as an RMAN HOT backup.	If the management server is set to no-archive mode, users can perform an RMAN cold backup. See Changing the Archive Mode on page 395.

Backup Destination and Operation for RMAN

The management server has three backup points available for RMAN backups:

- **Windows**:

  `%ORA_HOME%\rman\current`

  `%ORA_HOME%\rman\backup1`

  `%ORA_HOME%\rman\backup2`

- **UNIX**:

  `$ORACLE_HOME/rman/current`

  `$ORACLE_HOME/rman/backup1`

  `$ORACLE_HOME/rman/backup2`

The scheduled backup writes in the backup1 and backup2 folders, in rotation. The Backup Now backup from the management server continues to overwrite in the current folder. See Architectural Overview of Report Views and Report Cache Refresh on page 359 for more information about RMAN backups.

Database Mode

The database can be set to archive mode or no-archive mode.

This section contains the following topics:

Archive Mode

To facilitate the HOT RMAN backup, you must change the database mode to archive mode, as described in Changing the Archive Mode on page 395. The default database archive destination is \oracle\oradata\APPIQ\archive on Microsoft Windows and $ORACLE_HOME/oradata/APPIQ/archive on UNIX systems. This destination can be modified as described in Changing the Archive Destination on page 397.

Depending on the input/output of the data, archiving can be in the range of 2.5 GB to 10 GB per day. The archive folder gets cleaned during on a scheduled RMAN backup.

When you change the database to archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and enable the RMAN backup scheduler. Take a current RMAN backup after switching to archive mode, as described in Performing an RMAN Hot Backup on page 385.

No-Archive Mode

You can change the database mode to run in no-archive mode with RMAN backup disabled, as described in Changing the Archive Mode on page 395. When you change the database to no-archive mode, you reset the logs SCN, set the archive parameter in the database parameter file and disable the RMAN backup scheduler.

Export the database after switching to no-archive mode. See Exporting the Database on page 391 for more information.

Keep in mind the following implications if you do decide to change the database to no-archive mode:

- If you set the management server to **no-archive mode**, it is up to you to back up the management server manually. If you forget to back up your management server and your management server fails, you will not have a database to import. To learn more about manual backups, see Restoring a Cold Backup on page 394. Export the database after you change the database mode to no-archive. See Exporting the Database on page 391.
- Scheduled RMAN backup sessions do not run. If you change the database mode to no-archive during an RMAN backup, the RMAN backup will error out.
- If the database fails as a result of a corrupt data file, the database can only be restored to the last export backup available. This requires recreating the database along with the import.

Architectural Overview of RMAN Backups

By default, the management server does not backup the database automatically. If you enabled the database archive mode and RMAN backup as described in Changing the Archive Mode on page 395, the management server backs up the Oracle instance for the management server every three days and saves the backup for two weeks. When the management server first performs a scheduled backup, the backup is saved in `%ORA_HOME%\rman\backup1` on Windows and `$ORACLE_HOME/rman/backup1` on UNIX systems. The next time the management server performs a scheduled backup, it is saved in `%ORA_HOME%\rman\backup2` on Windows and `$ORACLE_HOME/rman/backup2` on UNIX systems. The management server saves the backup in alternating directories (backup1 and backup2), so you always have a copy of the last backup and the previous backup. To learn how to change the frequency of the scheduled backups, see Scheduling RMAN Hot Backups on page 386.

You can back up the database at any time by clicking the **Backup Now** button on the Database tab. When you back up the database using this technique, the backup is saved only in `%ORA_HOME%\rman\current` on Windows and `$ORACLE_HOME/rman/current` on UNIX systems. To recover the database, contact customer support. See the table at the end of this section.

Assume you recently installed the management server and did not do any backups. You scheduled the backups to take place every three days. You performed a backup, and it is stored in the backup1 folder. The next scheduled backup occurs on day 4, and it is saved in the backup2 folder. If your database fails, you can restore the database from day 1 or day 4. If you have a scheduled backup on day 7, it is saved to the backup1 folder. This backup replaces the backup from day 1. The available backups are now from day 4 and 7. If you click the **Backup Now** button on day 8, the backup is saved in the CURRENT folder because the backup is recording the current state of the database. If your database fails, you can restore the database from day 4, 7, or 8, as the following table shows.

Backup Example

Day	Backup Type	Backup 1	Backup 2	Current	Available Backup
Day 1	Scheduled	Day 1 backup	----------	----------	Day 1 Backup
Day 4	Scheduled	----------	Day 4 backup	----------	Backups from Days 1 and 4
Day 7	Scheduled	Day 7 backup	----------	----------	Backups from Days 4 and 7
Day 8	Backup Now	----------	----------	Day 8	Backups from Days 4, 6, and 8
Day 10	Scheduled	----------	Day 10 backup	----------	Backups from Days 7, 8, and 10

Keep in mind the following:

- Only one user at a time can back up the database.
- The management server archives files for the backup in a separate directory. Do not modify the files in this directory (`\oracle\oradata\APPIQ\archive` on Windows and `$ORACLE_HOME/oradata/APPIQ/archive` on UNIX systems).
- For average database activity, the management server requires at least 100 GB of disk space for archive files. If there is higher database activity than average, more disk space could be required.

Data Saved During a Backup

The following are saved during the backup:

- **Management server RMAN backup files**: Contain information about the elements monitored by your management server.
- **Oracle Network Configuration Files**: tnsnames.ora and listener.ora.
- **CIM Repository**
- **Property files, such as appiq_jboss.properties**
- **rmanbackup.log file**
- **spfileappiq.ora file**

Backing Up the Database Manually

To back up the database manually:

1. Before starting, enable the database archive mode and RMAN backup as described in Changing the Archive Mode on page 395.
2. In the management server, click **Configuration** > **Product Health**.
3. Select **RMAN Backup** in the Product Health tree.
4. Click **Backup Now**. The database is backed up.

Performing an RMAN Hot Backup

You can perform an RMAN hot backup instantly. The backup is referred to as being “hot” because the management server is still running. When you perform an RMAN hot backup, the following files are backed up:

- **Database files**
- **Control files**
- **Redo files**
- **Archive files**
- **Oracle network configuration files (tnsnames.ora, listener.ora)**

- **CIM repository**
- **File System Viewer**

The buttons on the RMAN backup page appear disabled when the database archive mode is disabled. See Changing the Archive Mode on page 395 for more information about changing the archive mode.

To perform an RMAN hot backup:

1. Verify that you enabled the database archive mode and RMAN backup as described in Changing the Archive Mode on page 395.
2. Click **Configuration** > **Product Health**.
3. Select **RMAN Backup** in the Product Health tree.
4. Click **Backup Now**. The management server performs an RMAN backup.

When the backup is complete:

- The Setup tab is refreshed with the status of the manual hot backup.
- The Results tab is updated with the status of the RMAN hot backup and displays the status of the previous RMAN hot backups (manual and scheduled).

Scheduling RMAN Hot Backups

The management server enables you to schedule an RMAN hot backup of the database, as described in the following steps. The backup is referred to as “hot” because the management server is still running.

Before you can schedule the RMAN backup, you must enable the database archive mode and RMAN backup as described in Changing the Archive Mode on page 395. The buttons on the RMAN backup page appear disabled if the database archive mode is disabled.

To learn more about backing up the database, see About the Database Admin Utility on page 388.

To schedule an RMAN hot backup:

1. Verify that you enabled database archive mode and RMAN backup as described in Changing the Archive Mode on page 395.
2. Click **Configuration** > **Product Health**.
3. Select **RMAN Backup** in the Product Health tree.
4. Click the calendar icon  .
5. In the Time box:

a. Enter the time in 24-hour format.

b. Click the date on which you want to run the next backup of the database. Today's date is highlighted in pink.

c. Click **Set**. The date and time appear in the Next Scheduled Run box in the yyyy–mm–dd format.

   If you change the date in the box to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2010-11-31, the software assumes the date is 2010-12-01.

6. In the Repeat Interval box, enter an interval. Select one of the following units from the menu:
   - **Second(s)**
   - **Minute(s)**
   - **Hour(s)**
   - **Day(s)**
   - **Week(s)**

   The minimum interval you can schedule is one day. Whether you select Hour(s), Minute(s), or Second(s), you must enter an interval that equals more than a day. For example, if you select Hour(s), you must enter 24 or more. If you select Minute(s), you must enter 1440 or more.

7. Click the **Enable** option.

8. Click **Save Schedule**.

   You can always disable the schedule by clearing the Enable option.

   When the scheduled RMAN hot backup is complete, the Results tab is updated with the status of the backup. The status of the previous RMAN hot backups (manual and scheduled) is displayed in the Results tab.

Viewing Results from RMAN Backup

The results of an RMAN backup can be determined by checking the Results tab for the date, time, and status of the backup.

RMAN backups cannot be used from earlier releases.

To view the results of RMAN backup:

1. Click **Configuration** > **Product Health**.
2. Select **RMAN Backup** in the Product Health tree.
3. Click the **Results** tab in the RMAN Backup window. The following information is displayed:

- Date/Time of the backup
- Status of the backup
- Backup folder

About the Database Admin Utility

The Database Admin Utility allows you to manage your database, from restoring it from a cold backup to resetting the temp tablespace. The tool provides flexible importing and exporting features, which lets you save time.

Note: You can also perform the Database Admin operations by running commands at the command prompt. For more information, see Performing Database Admin Operations using Command Line Interface on page 402

To learn more about the Database Admin Utility, see the following topics:

Accessing the Database Admin Utility

To access the Database Admin Utility:

1. Stop the AppStorManager service.
   - **Windows**:
     i. Go to the **Administrative Tools** > **Services** window.
     ii. Right-click **AppStorManager**.
     iii. Select **Stop** from the menu.
   - **Linux**:
     i. Open a command prompt window.
     ii. To stop the management server, enter the following:

     ```
     /etc/init.d/appstormanager stop
     ```

     iii. To see the status of the management server, enter the following:

     ```
     /etc/init.d/appstormanager status
     ```

2. Access the database utility by doing the following on the management server:
   - **Linux**:
     i. Set the display if you are accessing the Database Admin Utility remotely.

     To set Perl in your path, enter the following command at the command prompt:

     ```
     # eval `/opt/<SE Install Dir.>/install/uservars.sh`
     ```

     In this instance, /opt/<SE Install Dir.> is the directory containing the software. It is defined by `$APPIQ_DIST`.

     ii. Go to the `$APPIQ_DIST/Tools/dbAdmin` directory and enter the following at the command prompt:

     ```
     perl dbAdmin.pl
     ```

   - **Windows**:

     Go to the `%MGR_DIST%\Tools\dbAdmin` directory and double-click **dbAdmin.bat**.

Refreshing the Database Admin Utility

The Database Admin Utility requires that JBoss and CIMOM not be in use when it is operating. If either of these components is in use, the Database Admin Utility will not work. If you are shown an error message when you start the utility, stop the AppStorManager service and click the **Refresh** button.

Checking the Database Status

To find the database status:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on previous page.
2. Click **Check Database Status** in the left pane.
3. Click **Check Database Status** in the right pane.

   The database status is shown in the lower pane.
4. To clear the pane, click **Clear All**.

The Check database status option is available even when the management server is running.

Changing System Account Passwords

Change the passwords to the following accounts to prevent unauthorized access.

Use only the Database Admin Utility to change the passwords. If a password for one of the following accounts is changed through a method other than the Database Admin Utility, the Database Admin Utility might not work correctly. See Changing the Password of System Accounts on page 310 for information about changing database passwords.

- **RMAN_USER** – Used for RMAN backup and restore. This user has sys privilege. Default password: backup
- **DB_SYSTEM_USER** – Used for all the database activity, including establishing a connection to the management server database. Default password: password

You must change the passwords of the RMAN_USER and DB_SYSTEM_USER accounts using the Database Admin Utility, so the management server is aware of the changes. Do not change the password for any of these accounts using Oracle. Make sure you keep the new passwords in a safe location.

The management server requires the password to have the following characteristics:

- A minimum of three characters
- Start with a letter
- Contain only letters, numbers, and underscores (_)
- Cannot start or end with an underscore (_)

To change the password of a system account:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Change Passwords** in the left pane.
3. Select an account name from the **User Name** box.
4. Type the current password in the **Old Password** box.
5. Type the new password in the **New Password** box.
6. Retype the password in the **Confirm Password** box.
7. Click **Change**. The Database Admin Utility changes the password for the specified account.

Exporting the Database

Use the Database Admin Utility to save the database in a format that can be imported.

If you are interested in backing up the database for disaster recovery, see Restoring a Cold Backup on page 394.

To export the database:

1. Stop the AppStorManager service.
   - **Windows**:
     i. Go to the **Administrative Tools** > **Services** window.
     ii. Right-click **AppStorManager**.
     iii. Select **Stop** from the menu.
   - **Linux**:
     i. Open a command prompt window.
     ii. To stop the management server, enter the following:

     ```
     /etc/init.d/appstormanager stop
     ```

     iii. To see the status of the management server, enter the following:

     ```
     /etc/init.d/appstormanager status
     ```

2. Access the database utility by doing the following on the management server:
   - **Linux**:
     i. Set the display if you are accessing the Database Admin Utility remotely.

     To set Perl in your path, enter the following command at the command prompt:

     ```
     # eval `/opt/<SE Install Dir.>/install/uservars.sh`
     ```

In this instance, /opt/<SE Install Dir.> is the directory containing the software. It is defined by `$APPIQ_DIST`.

ii. Go to the `$APPIQ_DIST/Tools/dbAdmin` directory and enter the following at the command prompt:

```
perl dbAdmin.pl
```

- **Windows**:

  Go to the `%MGR_DIST%\Tools\dbAdmin` directory and double-click **dbAdmin.bat**.

3. Click **Export Database** in the left pane.
4. Click **Browse** to select a file path, enter a file name in the **File name** box, and click **Open**.

   Select a directory outside of the directory tree of the management server. If you remove the management server, you will not lose the saved database.

   The file name with its path is displayed in the Database Admin Utility. The .zip file extension is automatically added to the file name.
5. Select **Exclude Report Cache** if you do not want the report cache to be included with the database you are exporting. When you import this database, the report cache will be empty until it is refreshed (**Configuration** > **Reports** > **Report Cache**). This option can save time with exporting the database if your database includes a large amount of report data.
6. Click **Export Database**.

   The Database Admin Utility creates a zip file containing the database.

Importing the Database

You can revert to an earlier configuration by uploading a file (*.zip) containing the database information.

The software stores a snapshot of the data in its database. This file is a snapshot of the network at a certain time and, therefore, might not contain your most current network configuration. To view an up-to-date network configuration and the latest information about the elements, perform discovery and Get Details.

To import a database:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Import Database** in the left pane.
3. Click **Browse**, select a database file with a .zip extension to import, and click **Open**. The file name is displayed in the Database Admin Utility.
4. Select **Populate Report Cache** to refresh the reports cache during the import. Keep in mind that the amount of time to import the database might increase if you select this option when the database has a large amount of data for reports. You can instead refresh the report cache from the management server (**Configuration** > **Reports** > **Report Cache**).

5. (*For non-production systems only*): Select **Include Product Health Data** to include management data for the management server host. This option, which is for support purposes only, provides data on the health of the product so you can determine if there is an issue with the management server itself.

   Do not use this feature on production systems. Selecting this option affects product health reporting.

6. Click the **Import Database** button.

Re-initializing the Database

Caution: Re-initializing the database removes everything from the database. This is not recommended unless you are sure about what you are doing. It is strongly suggested you export the database before you re-initialize it. See Exporting the Database on page 391 for more information on how to save the database.

Keep in mind the following:

- When you re-initialize the database, all users are logged out of the management server.
- Ignore the warning messages in the command prompt window that pop up when the Database Admin Utility runs. See Warning Messages During Reinitializing the Database on page 399 for more information.

To re-initialize the database:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Re-initialize Database** in the left pane.
3. Click the **Re-initialize Database** button.

You are notified when the re-initialization is complete.

Resetting the Temp and Undo Tablespace

The temporary and undo tablespace might grow large due to high database activity. You should regularly reset the temp and undo tablespace to its initial value.

To reset the tablespace:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Reset Temp Tablespace** in the left pane.
3. Click **Reset Temp Tablespace** in the right pane.

Defragmenting the Database

The indexes and tables in the management server database can become fragmented over time. This can eventually negatively affect performance.

To defragment the database:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Tune CMS Database** in the left pane.
3. Click **Tune CMS Database** in the right pane.

   The following message appears when the tuning is complete:

   ```
   ******* Tune CMS Database completed on <Date> <Time> *******
   ```

Restarting the Database

You might sometimes need to restart the AppIQ instance of the database. Use this feature in the Database Admin Utility when the database is down or when you need to shut down and restart the database.

To restart the database:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Restart Database Server** in the left pane.
3. Click **Restart Database Server** in the right pane.

Clearing Archives

Archive files can require large amounts of disk space and should be deleted periodically. Use the Clear Archive option to manage the disk space used by the Oracle database for the management server.

To clear archives:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Clear Archives** in the left pane. The archives are deleted.

Restoring a Cold Backup

If you performed an RMAN cold backup, follow the steps in this section to restore the RMAN cold backup. The backup is referred to as “cold” because the management server is not running while the backup is occurring. For information about changing the archive mode, see Changing the Archive Mode on the facing page.

To restore a cold backup:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Restore Cold Backup** in the left pane.
3. Click **Browse**.

4. In the File Name box, provide the directory path containing the cold backup (which might automatically be populated in some Web browsers), but do not provide a file name. In this release the cold backup is saved in the COLDBACKUP directory in the path specified by the person who did the cold backup. The Database Admin Utility automatically notices the backup files in the directory provided.
5. Click **Run Cold Backup**.

Changing the Archive Mode

By default, the management server database runs in no-archive mode, which requires you to backup the database manually using a cold RMAN backup. A cold RMAN backup is an RMAN backup without the management server running. If you want to take an RMAN hot backup of the database, change the database to archive mode. An RMAN hot backup is done with the database running while the backup is occurring.

If you decide to leave the database in no-archive mode, see No-Archive Mode on page 383 for additional important information.

To change to the archive mode:

1. Access the Database Admin Utility, as described in Accessing the Database Admin Utility on page 389.
2. Click **Change Archive Mode** in the left pane. A text message indicating the current status of the archive mode is displayed.
3. Do one of the following:
   - The **Enable RMAN Backup** button is displayed if the archive mode is currently disabled. Select this option if you plan to run automated backups while the management server is running. If the database is already in archive mode, you can use this option to clean up archive files. See About the Database Admin Utility on page 388 for more information about the automatic backups.

     Changing the database to archive mode resets the logs, sets the archive parameter in the database parameter file, and enables the RMAN backup scheduler. After switching to archive mode, take a current RMAN backup (**Configuration** > **Product Health** > **Disk Space** > **RMAN Backup** > **Backup Now**) as described in Performing an RMAN Hot Backup on page 385.
   - The **Disable RMAN backup** button is displayed if the archive mode is currently enabled. Select this option if and only if you always shut down the management server prior to a backup. With the database running in no-archive mode, you can only cold back up your database. Changing the database to no-archive mode resets the logs SCN (System Change Number), sets the archiver parameter in the database parameter file, and disables the RMAN backup scheduler. See Restoring a Cold Backup on previous page for more information. Be sure to export the database after switching to no-archive mode. See Exporting the Database on page 391.
4. Click **Change Settings**.

Restore an RMAN Backup

The Database Admin Utility enables you to restore the management server database from a previously scheduled (hot) RMAN backup, which is stored in the following directories by default: backup1, backup2, and current. You might not have a previously scheduled (hot) RMAN backup if you selected **Disable Database Archive Mode and RMAN Backup**.

To restore the database:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Restore RMAN Backup** in the left pane.
3. Select one of the following directories to restore:
   - **Current** – The last restore from when the **Database Server Backup** button on the Database tab was clicked.
   - **Backup1***
   - **Backup2***

   *Information from the automatic backup (alternating weeks). See About the Database Admin Utility on page 388 for more information about the automatic backup.
4. Click **Restore RMAN Backup**. Database Admin Utility restores the selected database.

Running a Cold Backup

If you are running the database in no-archival mode, you should perform a cold backup frequently. The backup is referred to as being “cold” because the management server is not running during the backup.

The cold backup does not run if archive mode, which runs an RMAN backup periodically, is enabled. See About the Database Admin Utility on page 388 and Changing the Archive Mode on previous page for more information.

The following data is saved during a cold RMAN backup:

- **Management server RMAN backup files**: Contain information about the elements your management server monitors.
- **Oracle Network Configuration Files**: tnsnames.ora and listener.ora
- **CIM Repository**

To run a cold backup:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Run Cold Backup** in the left pane.
3. Click **Browse** to select a file path.

4. In the File Name box, provide a directory path (might automatically be populated in some Web browsers), but do not provide a file name.

   The management server saves the backup in a directory called COLDBACKUP in the path you specified. Any pre-existing content in this directory, such as previous cold backups, is removed.

5. Click **Run Cold Backup**.

Changing the Archive Destination

The default archive directory is `\oracle\oradata\APPIQ\archive` on Windows and `$ORACLE_HOME/oradata/APPIQ/archive` on UNIX systems. Over time, your database will grow. If you run out of space, you can add a new volume and change the archive destination to a new volume.

To change the archive destination:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Set Archive Destination** in the left pane.
3. Click **Browse** to select a file path.
4. Click **Set Archive Destination**.

Downloading Log Files

Use the Download Logs option to view log files generated by the management server. Download Logs saves all the log files in a zip file, which is stamped with the date and time (24-hour clock).

Some log files are generated only when you run certain features. For example, the reports.log file is only generated when you run reports.

The Download Logs option is available when the management server is running.

To download all the log files for the management server:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. In the Directory box at the top of the screen, enter the path to an existing directory or browse to a directory using the **Browse** button.
3. Click **Download Logs**. The log files are saved to a zip file which is copied to the directory you specified.

Viewing the Database Admin Utility Log File

To view the Database Admin Utility log file in a separate window:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.

2. Click the **View Log** button at the bottom of the Database Admin Utility window. The logs are displayed in a separate window.

Resetting/Clearing the Database Admin Log File

When you click the **View Log** button at the bottom of the Database Admin window, the Database Admin Utility log file is displayed.

To reset/clear or refresh the Database Admin log information:

1. To refresh the log information, click the **Refresh** button.
2. To clear this window, click the **Reset Log** button at the bottom of the Database Admin Utility window and click **Yes** to confirm that you want to refresh the log information.

Recreating the Admin User account for the Management server

If the admin user account is deleted and none of the other users are able to login into Management server, you can recreate the default admin account using this option. A new admin user account is created with default password, which is **password**.

Once you recreate the account, you should change the admin password to prevent unauthorized access. The admin user account provides full access to the management server. Changing the Password for a User Account on page 298

To recreate the admin user:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389
2. Click **Recreate Product Admin User** in the left pane.
3. Click the **Recreate Product Admin User** button. A new admin user account is created.

Resetting the Admin Password for the Management Server

If you no longer know the password for the admin user account and you need the admin user account to access the management server, you can reset its password to the default, which is **password**.

Once you reset the password, you should change the admin password to prevent unauthorized access. The admin user account provides full access to the management server. Changing the Password for a User Account on page 298

To reset the admin password:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Reset Product Admin Password** in the left pane.

3. Click the **Reset Product Admin Password** button. The admin password is reset to **password**.

Warning Messages During Reinitializing the Database

When you use the Database Admin Utility to re-initialize the database, warning messages similar to the following appear in the command prompt window that pops up when the utility runs. You can ignore these messages.

```
Connected.

Creating FSRM DATA tablespace

Creating FSRM INDX tablespace

Connected.

Warning: View created with compilation errors.
```

Generating a Support Database

The Database Admin Utility allows you to generate a support database that can be used during support calls. This support database is only useful to customer support representatives for the management server.

Never import a support database to the management server. The Support Database is only intended for use by Customer Support.

To generate a support database for customer support use:

1. Access the Database Admin Utility as described in Accessing the Database Admin Utility on page 389.
2. Click **Generate Support Database** in the left pane.
3. Click **Browse**, select a path, and enter a file name in the File name box. Select a directory other than the directory where the management server resides. This way if you remove the management server, you will not lose the saved database.
4. Click the **Open** button. The file name and path is displayed in the Database Admin Utility. The .zip file extension is automatically added to the file name.
5. Select **Exclude Report Cache** if you do not want the report cache to be included with the support database. When a customer support representative imports the support database, the report cache will be empty until the report cache is refreshed. Generating the support database is much faster if you omit the report data.
6. (*Optional*) Select the **Include File Server Data** check box if you want to capture the File Server data. File Server SRM data files are excluded by default in the support zip file because this data might include confidential directory and file names.
7. Click **Generate Support Database**. The support database is created based on your selections. See About Importing a Customer Support Database on next page for more information.

About Importing a Customer Support Database

Do not import a support database. This feature is for only Customer Support.

The Import Support Database functionality is for customer support use only. Importing a support database to your management server, results in the following:

- Only the admin user will be able to log into the management server, no other user will be able to log in to the system. To enable other users to log in, the admin user would need to re-enter the passwords for all other users.
- The passwords for all discovered elements in the database—hosts, switches, and storage systems—would become invalid and discovery would fail for all elements.

Use **Export Database** to generate a valid database backup for use on your management server. See Exporting the Database on page 391.

Troubleshooting Listener and Database Connection Problems

If there is a problem connecting to the Oracle database from the management server or you see a JBoss connection problem in appstorm.<timestamp>.log, first verify that the listener and database are running.

Checking Listener Status

If you are not able to start the listener contact technical support with the error message and network files (tnsnames.ora,listener.ora) from the following directory:

Windows:

```
%ORA_HOME%\network\admin\listener.ora
```

UNIX:

```
$ORACLE_HOME/network/admin/listener.ora
```

To verify that the listener and database are running:

Windows:

From the command line, enter:

```
lsnrctl status
```

This shows you the status of listener. If the listener is not running, enter:

```
lsnrctl start
```

From the Database Admin Utility. See Checking the Database Status on page 390.

UNIX:

From the command line, log on as an Oracle user and enter:

```
su - oracle
```

```
lsnrctl status
```

This shows the status of listener. If the listener is not running, enter:

```
lsnrctl start
```

From the Database Admin Utility. See Checking the Database Status on page 390.

Checking Database Status

If you receive an error message while performing these steps, contact technical support.

Windows:

From the command line, enter the following commands:

```
Sqlplus /nolog

Sql>connect sys/change_on_install@appiq as sysdba

Sql> startup force;
```

From the Database Admin Utility. See Checking the Database Status on page 390.

UNIX:

From the command line, log on as an Oracle user and enter

```
su – oracle
```

Enter the following commands:

```
Sqlplus /nolog

Sql>connect sys/change_on_install@appiq as sysdba

Sql> startup force;
```

From the Database Admin Utility. See Checking the Database Status on page 390.

Generating a Support Database for Customer Support

When contacting Customer Support, you might be asked to generate support database. See Generating a Support Database on page 399.

"AppIQ_Oracle_Database.dmp: Permission denied" Error

If you receive an "AppIQ_Oracle_Database.dmp: Permission denied" error, or if the import aborts soon after the "extracting zip file" message:

1. Stop the Database Admin Utility.
2. Delete the AppIQ_Oracle_Database.dmp and Sim_Oracle_Database.dmp files from the $MGR_DIST/install/database/Oracle directory.

Known Issues with the Database Admin Utility

This section provides information about open issues with the Database Admin Utility (DB Admin Tool).

- Database Consistency Tool Limited to the Command line. The user interface for the Database Consistency Tool has been disabled in this release. The tool is accessible only from the command line.
- NullPointerException Opening an Empty Folder and Clicking in the File Name Field. While using the Database Admin Utility, you might notice a "NullPointerException at javax.swing.JComponent.repaint" message in the command window when you use the enter key to open an empty folder and click in the "File name" field. The NullPointerException can be ignored, because there is no effect on the operation of the Database Admin Utility.
- Do Not Use NAS Volume or Network Share DBAdmin Archive Destination Path . Using a NAS Volume or Network Share for the DBAdmin Archive Destination Path is not supported. If you set the archive path to such a location, the Oracle database will shut down when you start the archive operation, and you will see a series of messages starting with "ORA-16032: parameter LOG_ARCHIVE_DEST_1 destination string cannot be translated" in the DBAdmin log.
- Extra Files might be Removed after Importing a Database . Files of .dmp, .ora, and .txt in the install/database/Oracle directory on the management server might be removed after you use DBAdmin to import a database.
- Extra Files might be Removed after Importing a Database . Files of .dmp, .ora, and .txt in the install/database/Oracle directory on the management server might be removed after you use DBAdmin to import a database.

  DBAdmin Reset Admin Password Feature Should Not be Used in AD/LDAP Mode. The Reset Admin Password feature within DBAdmin does not apply when you are running the management server with AD/LDAP authentication.

Performing Database Admin Operations using Command Line Interface

The Database Admin operations can be performed using the command line interface on the management server. To perform these operations, you must be logged on to the computer where the management server is running.

The Database Admin operations can be performed by running the perl commands from the command prompt.

Following Database Admin operations can be performed using the Command Line interface:

- Checking the Database Status on page 404
- Changing System Account Passwords on page 404
- Exporting the Database on page 404

Pre-requisites for Performing Database Admin Operations

Before performing any Database Admin operations, perform the steps mentioned as follows:

On Windows

1. Stop the **AppStorManager** service. To stop, perform the following:
   a. Go to the **Administrative Tools > Services window**.
   b. Right-click **AppStorManager**.
   c. Select **Stop** from the menu.
2. Open a command prompt window.
3. Change your working directory to `<SE Install Dir.>\install\database\Oracle` directory on the management server.

On Linux

1. Change your working directory to `<SE Install Dir.>/install/database/Oracle` on the management server.
2. Stop the **AppStorManager** service. To stop, perform the following:
   a. To stop the management server, type the following:
      `/etc/init.d/appstormanager stop`
   b. To see the status of the management server, type the following:
      `/etc/init.d/appstormanager status`
3. Type the following to set the Perl path:
   `eval `$APPIQ_DIST/install/uservars.sh``

In this instance, `$APPIQ_DIST` defines the path `/opt/HP/<SE Install Dir.>` that is, the directory containing the software.

Checking the Database Status

To check the database status:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on previous page.
2. Run the following command from the command prompt:
   `perl getDBState.pl`

Changing System Account Passwords

Change the passwords to the following accounts to prevent unauthorized access.

- **RMAN_USER** – Used for RMAN backup and restore. This user has sys privilege. Default password: backup
- **DB_SYSTEM_USER** – Used for all the database activity, including establishing a connection to the management server database. Default password: password

Do not change the password for any of these accounts using Oracle. Make sure you keep the new passwords in a safe location.

The management server requires the password to have the following characteristics:

- A minimum of three characters
- Start with a letter
- Contain only letters, numbers, and underscores (_)
- Cannot start or end with an underscore (_)

To change the system account passwords:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on previous page.
2. Run the following command:
   `perl changePassword.pl <Username> <Old Password> <New Password>`

   Parameters

   `Username` – Username of the account.

   `Old Password` – Old password for the account.

   `New Password` – New password for the account.

Exporting the Database

To export the database:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl expappiq.pl <Filename> [<Exclude Report Cache>] [<Obsolete Parameter Ignored>] [<Include File Server Data>] [<Mode>]`

   Parameters

   `Filename` – Specifies the logical disk path and the file name with a **.zip** extension, where the exported database is to be stored.

   `Exclude Report Cache {YES|NO}` – Specifies whether to exclude the report cache data in the database being exported. Default value is YES. Excluding report cache can save time with exporting the database if your database includes larger amount of report data.

   `Obsolete Parameter Ignored` – Specify IGNORE as the value for this parameter.

   `Include File Server Data {YES|NO}` – Specifies whether to include the file server data in the support database being generated. Default value is NO.

   `Mode` – Specifies the database mode. Specify EXPORT as the value to export the production database.

Note: To change the default value of any parameter, you must type the value for each and every parameter in the displayed sequence.

Importing the Database

You can revert to an earlier configuration by uploading a file (*.zip) containing the database information.

The software stores a snapshot of the data in its database. This file is a snapshot of the network at a certain time and, therefore, might not contain your most current network configuration. To view an up-to-date network configuration and the latest information about the elements, perform discovery and Get Details.

To import the database:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl impappiq.pl <Filename> [<Placeholder Ignored Password>] [<Populate Report Cache>] [<Placeholder Ignored schema name to Import>] [<Placeholder Ignored sim compatibility check>] [<Import Type>] [<Modecheck>] [<Revert to Brocade API>] [<Include Management DB>]`

   Parameters

   `Filename` – Specifies the logical disk path and the file name with a **.zip** extension, from which the database should be imported.

`Placeholder Ignored Password` – Specify IGNORE as the value for this parameter.

`Populate Report Cache {YES|NO}` – Specifies whether the report cache should be populated during the database import. Default value is NO. Populating the report cache can increase the time to import the customer support database substantially, if the database contains large amount of data for reports.

`Placeholder Ignored schema name to Import` – Specify IGNORE as the value for this parameter.

`Placeholder Ignored sim compatibility check` – Specify IGNORE as the value for this parameter.

`Import Type` – Specifies the database import type. Specify PRODUCTION as the value to import the production database.

`Modecheck {YES|NO}` – Specifies whether to allow a compatibility check on the production or support database being imported into the management server. The management server performs the compatibility check by comparing the Command Line operation being performed against the value specified in the Import type option.

For example, if a Import database operation is being performed, the management server checks whether the database being imported is production or a support database. If the value specified is Support, it shows an error.

`Revert back to Brocade API` – Specify NO as the value for this parameter.

`Include Management DB {YES|NO}` – This parameter is for non-production systems only. It specifies whether to include the HP Storage Essentials server product health data in the imported database for management server host or not. Default value for this parameter is NO.

Do not set this parameter to value YES, if you are importing the database on a production system. It affects the Product Health reporting on management server.

Note: To change the default value of any parameter, you must type the value for each and every parameter in the displayed sequence.

Re-initializing the Database

Caution: Re-initializing the database removes everything from the database. This is not recommended unless you are sure about what you are doing. It is strongly suggested you export the database before you re-initialize it. See Exporting the Database on page 404 for more information on how to save the database.

Keep in mind the following:

- When you re-initialize the database, all users are logged out of the management server.
- Ignore the warning messages in the command prompt window that pop up when the Database Admin Utility runs. See Warning Messages During Reinitializing the Database on page 399 for more information.

To re-initialize the database:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl reinit.pl`

   You are notified when the re-initialization is complete.

Resetting Temp and Undo Tablespace

The temporary and undo tablespace might grow large due to high database activity. You should regularly reset the temp and undo tablespace to its initial value.

To reset the temporary and undo tablespace:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl setTempTbs.pl`

Defragmenting the Database

The indexes and tables in the management server database can become fragmented over time. This can eventually negatively affect performance.

To defragment the database:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl datapump.pl`

   Parameters

   `schema name` – Specifies the name of the database to be defragmented. APPIQ is the default value.

Restarting the database

You might sometimes need to restart the AppIQ instance of the database. Use this feature in the Database Admin Utility when the database is down or when you need to shut down and restart the database.

To restart the database:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl reStartDB.pl`

Clearing the Archives

Archive files can require large amounts of disk space and should be deleted periodically. Use the Clear Archive option to manage the disk space used by the Oracle database for the management server.

To clear the archives:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:

   ```
   perl clearArchives.pl
   ```

Restoring a Cold Backup

To restore a cold backup:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perlfullRMANRestore.pl <Backup_dir>`

   Parameters

   `Backup_dir` – Specifies the logical disk path of the directory from where the backup database must be restored.

Running a Cold Backup

To run a cold backup:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl coldRMANBackup.pl <Backup_dir>`

   Parameters

   `Backup_dir` – Specifies the logical disk path of the directory, where the backup database must be stored.

   A new directory named COLDBACKUP is created inside the mentioned directory and the database backup is saved inside it. If the COLDBACKUP directory already exists in the mentioned directory, any pre-existing content in that directory, such as previous cold backups, is removed and the new database backup is saved inside it.

Changing the Archive Mode

To change the archive mode:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run one of the following commands at the command prompt, based on the current archive mode:
   - To enable the archive mode, run the following command:
     `perl setArchiveMode.pl`

     Run this command if you plan to run automated backups while the management server is running. Changing the database to archive mode resets the logs, sets the archive parameter in the database parameter file, and enables the RMAN backup scheduler. After switching to archive mode, take a current RMAN backup.
   - To disable the archive mode, run the following command:
     `perl setNOArchiveMode.pl`

     Run this command if and only if you always shut down the management server prior to a backup. With the database running in no-archive mode, you can only take cold backup of your database. Changing the database to no-archive mode resets the logs SCN (System Change Number), sets the `archiver` parameter in the database parameter file, and disables the RMAN backup scheduler. Be sure to export the database after switching to no-archive mode. See Exporting the Database on page 404

Restoring an RMAN Backup

To restore an RMAN backup:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl restoreRMAN.pl <Dir> <Time_of_Backup>`

   Parameters

   `Dir` – Specifies the logical disk path of the backup directory from where the RMAN database must be restored.

   `Time_of_Backup` – Specifies the date and time in `yyyymmddhh:min` format, until which you want to restore the RMAN database. Here, enter `hh:mm` in 24 hours format.

Example

```
perl restoreRMAN.pl C:\Backup 2011052219.00
```

In the above example, the management server database backed up until seven in the evening on May 22, 2011 is restored from the Backup directory on the C drive.

Changing the Archive Destination

The default archive directory is `\oracle\oradata\APPIQ\archive` on Windows and `$ORACLE_HOME/oradata/APPIQ/archive` on UNIX systems. Over time, your database will grow. If you run out of space, you can add a new volume and change the archive destination to a new volume.

To change the archive destination:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command:
   `perl setArchiveDest.pl <Archive_dir>`

   Parameter

   `Archive_dir` – Specifies the logical disk path of the archive directory where the archived database must be stored.

Generating a Support Database

The Database Admin Utility allows you to generate a support database that can be used during support calls. This support database is only useful to customer support representatives for the management server.

Never import a support database to the management server. The Support Database is only intended for use by Customer Support.

To generate a support database:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl expappiq.pl <Filename> <Exclude Report Cache> <Obsolete Parameter Ignored> <Include File Server Data> <Mode>`

   Parameters

   `Filename` – Specifies the logical disk path and the file name with a **.zip** extension, where the support database is to be stored after it is generated.

   `Exclude Report Cache {YES|NO}` – Specifies whether to exclude the report cache data from the support database being generated. Default value is YES.

   When a customer support representative imports the support database, the report cache will be empty until the report cache is refreshed. Excluding report cache can save time with generating the database if your database includes larger amount of report data.

   `Obsolete Parameter Ignored` – Specify IGNORE as the value for this parameter.

   `Include File Server Data {YES|NO}` – Specifies whether to include the file server data in the support database being generated.

   `Mode` – Specifies the database mode. Specify SUPPORT as the value to import the customer support database.

Note: To change the default value of any parameter, you must type the value for each and every parameter in the displayed sequence.

Example

```
perl expappiq.pl C:/SupportDB/Database.zip YES IGNORE YES SUPPORT
```

Running the above command generates a support database file with file server data, and exports it to SupportDB directory on the local C drive.

Importing a Customer Support Database

To import a customer support database:

1. Perform the steps as listed in the Pre-requisites for Performing Database Admin Operations on page 403.
2. Run the following command at the command prompt:
   `perl impappiq.pl <Filename> <Placeholder Ignored Password> <Populate Report Cache> <Placeholder Ignored Schema name to Import> <Placeholder Ignored sim compatibility check> <Import Type> <Modecheck> <Revert to Brocade API> <Include Management DB>`

   Parameters

   `Filename` – Specifies the logical disk path and the file name with a **.zip** extension, from which the customer support database must be imported.

   `Placeholder Ignored Password` – Specify IGNORE as the value for this parameter.

   `Populate Report Cache {YES|NO}` – Specifies whether the report cache should be populated during the database import. Default value is NO. Populating the report cache can increase the time to import the customer support database substantially, if the database contains large amount of data for reports.

   `Placeholder Ignored Schema name to Import` – Specify IGNORE as the value for this parameter.

   `Placeholder Ignored sim compatibility check` – Specify IGNORE as the value for this parameter.

   `Import Type` – Specifies the database import type. Specify SUPPORT as the value to import the customer support database.

   `Modecheck {YES|NO}` – Specifies whether to allow a compatibility check on the customer support database being imported into the management server. The management server performs the compatibility check by comparing the CLI operation being performed against the value specified in the Import type option.

   For example, if Import Customer Support database operation is being performed, the management server checks whether the database being imported is a support database or not. If the value specified is Support, it continues the operation and imports the customer support database, else displays an error.

   `Revert back to Brocade API` – Specify NO as the value for this parameter.

   `Include Management DB {YES|NO}` – This parameter is for non-production systems only. It specifies whether to include the HP Storage Essentials server product health data in the imported customer support database for management server host or not. Default value for this parameter is NO.

Do not set this parameter to value YES if you are importing the customer support database on a production system. It affects the Product Health reporting on management server.

Note: To change the default value of any parameter, you must type the value for each and every parameter in the displayed sequence.

Example

```
perl impappiq.pl C:/SupportDB/Database.zip IGNORE YES IGNORE IGNORE
SUPPORT YES NO YES
```

Running above command imports the support database file with file server data from SupportDB directory on local C drive.

10 Viewing Element Topology and Properties

This section contains the following topics:

About System Manager

System Manager is the gateway to many features that enable you view details about the discovered elements. It provides a topology that enables you to view how the devices in your network are connected. For example, direct-attached storage connections are displayed by a dotted line.

Another example is the display of Inter-Switch Link (ISL) trunking for supported Brocade switches. The topology screens, as well as other related displays, show the ISLs between switches and indicate the total number of ISLs and how many of them are trunked (for supported switches). For example, 6(3 trunked) means 6 is the total number and 3 is how many of them are trunked. ISL trunking information for supported switches is also provided by switch port Properties, switch port Detail table on the Navigation page, and by Reporter in various predefined reports.

To view direct-attached storage, you must enable the  button (see Toolbar in System Manager on page 417).

Use the utilities provided in the toolbar to modify the topology. For example, you can filter out fabrics and change the placement of elements in the topology through drag and drop functionality (see Toolbar in System Manager on page 417).

The following tabs, located in the middle pane, provide additional information:

- **List** – Provides information about the elements by SAN, fabric, and domain (see List Tab on page 420).
- **Access** – Provides information about zone entries, persistent bindings, and storage system LUN masking. You can also manage zone, zone aliases, and zone sets from this tab (see Access Tab on page 440).
- **Path** – Provides information about an element's path (see About the Path Tab on page 446).

When you right-click an element in the topology or in the List, Access, or Path tab, a menu appears. This menu provides additional functionality, depending on the type of element clicked, such as telnet or the creation of zone sets (see About the Right-Click Menu Options on page 447).

If a switch has more than one connection to a host or storage system, the number of connections is displayed on the line connecting the elements. If there is only one connection, no number is displayed, because the line indicates that a connection exists. For HP blade servers, loop-based connections are represented by a solid line with the word LOOP, and NPIV-based connections are represented by a solid line with the word NPIV.

Keep in mind the following:

- If your Java plug-in control panel cache is set to 50 MB, it is recommended you increase this setting to 100 MB or more. Increasing this setting improves the reloading performance of System Manager.
- ISL connections are not shown as connected ports between two McDATA switches that are not both managed by EFC Manager in a fabric.
- The user interface might load slowly while the topology is being recalculated.
- NAS stitching to Windows XP hosts does not appear in the topology.
- The Access tab might show array LUN access for volumes and hosts that are not part of your current organization. The Navigation page for a host security group displays hosts and

volumes that belong to the host security group even if those hosts and volumes are not in your current organization.

- On the Switch Navigation Ports page, a Brocade switch shows an L-Port as an FL-Port.

By double-clicking an element in the topology, you have access to the following features:

You can also access much of this information by using Element Manager. For details, see the entry for Element Manager in Management Server Components on page 47.

- **Navigation** – The Navigation tab provides information about an element and how it relates to other elements in its path (see About the Navigation Tab on page 489).
- **Properties** – The Properties tab provides a detailed status of the element (see Viewing Element Properties on page 498).
- **Events** – The Events tab enables you to view events for an element (see About Event Manager on page 529 and About the Events Tab on page 511).
- **Topology** – The Topology tab provides a graphical representation of an element's path. It displays additional information not found in System Manager, such as adapters, slots, and Fibre Channel ports (see Viewing Element Topology on page 501).
- **Asset Manager** – The Asset Manager tab enables you to keep track of asset attributes, such as contact information for the element's owner (see Asset Attributes of an Element on page 512).
- **Collectors** – The Collectors tab enables you to start a collector for a report and view the collector's corresponding reports once the information has been gathered (see About the Collectors Tab on page 514).
- **Monitoring** – The Monitoring tab enables you to access performance information about your storage elements (see About the Monitoring Tab on page 515).
- **Provisioning** – The Provisioning tab enables you to manage zones, zone sets, and zone aliases, in addition to pools, volumes, LUNs, and LUN mappings (see Provisioning Tab on page 511 and About the Provisioning Manager on page 675).
- **Policies** – The Policies tab enables you to create utilization polices, which can send an e-mail, generate an event, or run a custom script when a set threshold for an element is triggered (see About Policy Manager on page 579).
- **Presented Storage** – The Presented Storage tab displays the storage presented to the port WWNs of a host, even if the host cannot access the storage as a result of an incorrect zoning configuration (see About the Presented Storage Tab on page 515).

Grey Screen When Attempting to Access System Manager

Errors can occur if the client computer you use to access the management server has software that blocks JavaScript or pop-ups. You might be shown a grey screen when attempting to access System Manager. Other errors include not being able to get past the login screen, view topology, or perform many other functions. Set your blocking software appropriately to allow the user interface to function properly.

Accessing System Manager

To access System Manager, click **System Manager** () in the left pane.

Keep in mind the following:

- If you are unable to access System Manager, make sure your Web browser is set to allow JavaScript and cookies.
- Java 2 Runtime Environment is required to access several features in the management server, such as System Manager. If you are accessing the management server and you do not have the Java 2 Runtime environment, you are asked to install it if your Web browser is on Windows. If your Web browser is on a Linux or Solaris system, you must manually install the Java plug-in. See Installing the Java Plug-in on page 63 for more information.
- (Windows only clients) If you do not have the Java plug-in already installed and you are running Firefox, you must use Microsoft Internet Explorer to install the plug-in. After you install the plug-in, you can use Firefox to run the plug-in.
- When you are asked if you want to trust the signed applet for the software, click **Always**. The Always option prevents this message from being displayed every time you access System Manager.

About the User Interface for System Manager

This section contains the following topics:

About the User Interface

System Manager displays an easy-to-use interface, which provides the following:

- **Toolbar** – Buttons and menus to help you modify the topology in System Manager.
- **Tabs** – Information about individual elements. The following tabs are provided:
  - **List** – Information about the elements by type and by SAN, fabric, and domain. See List Tab on page 420.

- **Access** – Access to tools that enable you to provision and view information about provisioning. See Access Tab on page 440. Provisioning Manager might not be available in your version of the product. To determine if you can access the Provisioning Manager feature, access the List of Features, which is accessible from the Documentation Center (**Help** > **Documentation Center**).
- **Path** – Information about an element's path. See About the Path Tab on page 446.

- **Right-click menu** – Features you can use to manage that element. See About the Right-Click Menu Options on page 447.

Toolbar in System Manager

The System Manager toolbar provides the following menu options.

Toolbar Menu Options

Button	Description
	Prints the topology. See Printing the Topology on page 508.
	Exports the topology to an XML file that can be viewed in Microsoft Visio. See Exporting the Topology to Microsoft Visio on page 460.
	Magnifies the view.
	Decreases the magnification.
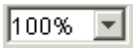	Sets the magnification to a percentage of the default magnification.
	Opens a smaller pane, which provides a global view of the topology. This enables to you position the main view to a certain section of the topology. See Using the Global View on page 458.
	Fits the topology to the window, so you can see the entire topology.
	Moves an element in the topology. See Arranging Elements in the Topology on page 456.
	Moves multiple elements at once. This button is not accessible from the Topology tab. See Arranging Elements in the Topology on page 456.
	Moves the entire topology at once. Click the **Pan** () button, and then click any place in the topology. Drag the mouse to a new location.

Button	Description
	Updates the layout of the topology and removes the last saved layout from the database. Elements that were manually moved might revert to their original position. This button is not accessible from the Topology tab.
	Saves the current topology, so that when you return to System Manager, the saved layout is restored. This option can be especially useful if you moved elements in the topology and you want to keep their current location. This button is not accessible from the Topology tab. When you click the button, you are asked if you want the layout to apply to all users. • **Yes** – All users who log in to the management server see the topology you created. Only users with system configuration capability can save their layout for all other users • **No** – Other users cannot view the topology you saved. The saved topology appears the next time you log into the management server.
	Opens a new window containing the topology. Enables you to view different domains of the topology at once. This button is not accessible from the Topology tab. See About the New Window Option on page 507.
	Displays only selected SANs in the topology. This button might not be accessible from the Topology tab. See Filtering SANs on page 468.
	Change Observer button – Monitors changes in the database status on the server. When changes are detected, the button turns amber. Click the amber button and a pop-up window displays the elements that changed on the server. When no changes are detected, the button is grayed out.
	Reloads the Change Observer button to display the latest changes to elements on the server.

Button	Description
	Finds an element or fabric in the topology. You can enter part of the information, and the management server highlights the elements and fabrics that match. To specify the scope of your search, select All, Element Name, Fabric Name, or Element WWN from the drop-down list next to the search box. After you populate the search box, click  or press **Enter**. To expand the Search box, close the left pane. See Opening and Closing the Left Pane on page 51.
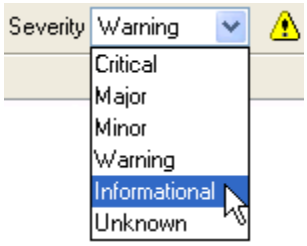	Displays the event severity icons for the elements displayed in the topology. See Viewing Event Status in the Topology on page 468. This feature is disabled for Performance Manager and Capacity Manager.
	Calculates the topology paths. Clicking this button also enables you to view direct attached storage in System Manager. Direct attached storage is indicated by dotted lines. If any of the paths are not fully calculated, a pop-up window displays a list of all the hosts with partially calculated paths.
	Displays the backup topology. The backup topology is also displayed in Backup Manager. For more information about the backup protection provided in this product, see About Backup Manager on page 729.

Icons Displayed in the Topology

The System Manager topology is represented by the following icons.

Icons Displayed in the Topology

Icon	Description
	Indicates an application.

Icon	Description
	Indicates a host. This particular icon is for a host running Microsoft Windows. If the host has a question mark and the word "inferred" after its name, the host was discovered through agentless discovery, as described in Agentless Discovery on page 253. If the host has a question mark next to it and its name contains an underscore followed by several numbers, the host is considered a generic host since HP Storage Essentials could not obtain additional information about the host in Discovery step 3.
	Indicates a Solaris container.
	Indicates a virtual machine.
	Indicates a storage system or subsystem.
	Indicates a switch. This particular icon is for a Brocade switch.
	Indicates a tape library.

List Tab

The List tab provides information about elements by type, cluster, SAN, fabric, and domain.

To find the fabrics in a domain, expand the domain node. The child nodes of the domain node are the SAN nodes. SAN nodes are containers for one or more physical or virtual fabrics. When there is an ISL between multiple switches, these switches are considered to be in a single SAN.

The "unknown" Fabric lists elements that have a Fibre Channel port connected to an undiscovered Fabric or that have a Fibre Channel port that remains unconnected.

When you right-click an element in the List tab, a menu is displayed. The options displayed depend on the type of element you clicked. For an explanation of the options in the menu, see About the Right-Click Menu Options on page 447.

Fabric Virtualization with Cisco and Brocade Switches

Fabric virtualization works differently for Cisco switches and Brocade switches.

Regardless of the type of switch you are using, when a SAN node is selected, the entire SAN and all of the fabrics and their member elements are highlighted in the topology view:

Highlighting a SAN's Members

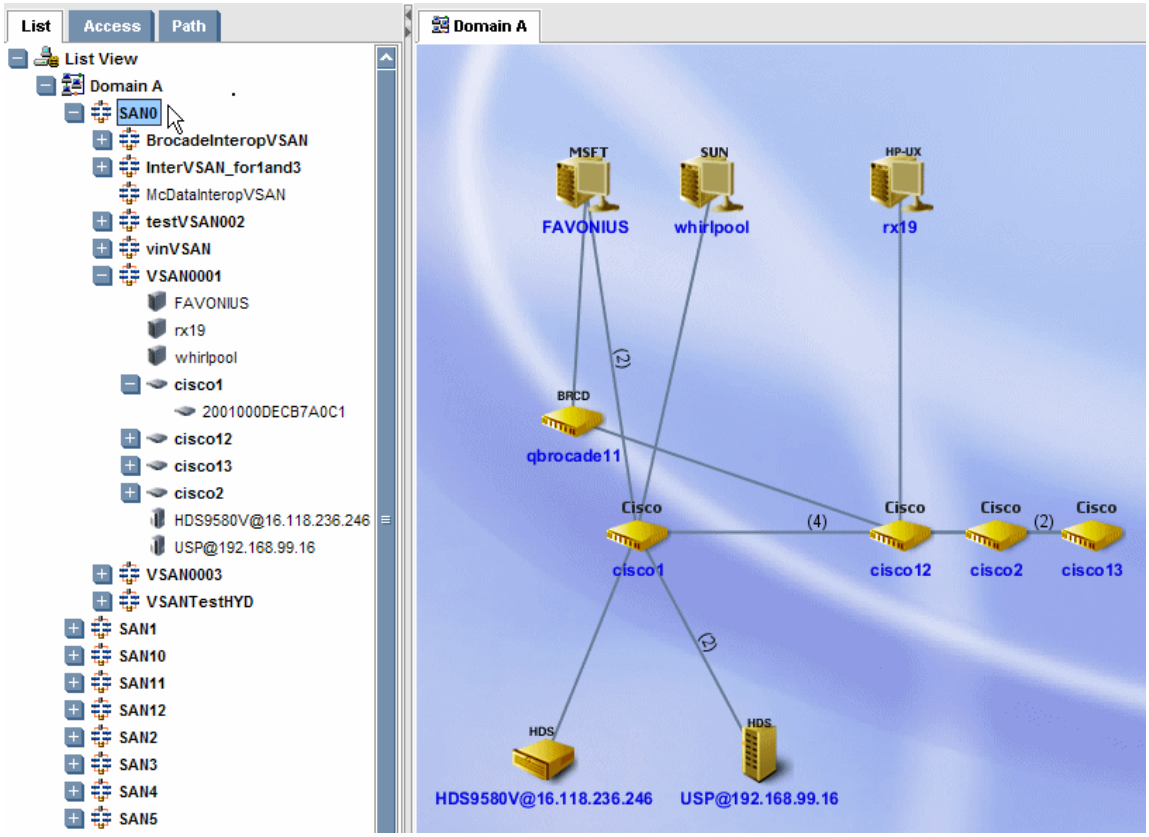

In the List tree, virtual switches are listed under physical switches. In the example, cisco1 is the name of a physical switch, and 2001000DECB7A0C1 is the name of a virtual switch.

The following sections describe how the various types of virtual fabrics are represented in HP Storage Essentials.

Cisco VSANS

- San nodes represent the SAN of a single switch, or multiple switches connected via ISL.
- Fabric nodes represent the VSANs created by a Cisco switch administrator.

Cisco's VSAN consists of multiple VSANs in a single SAN. This is also referred to as partitioning of the SAN. A single switch might have one or more VSANs. On the List tab, each VSAN node is listed as a child node of the SAN node. Elements belonging to a VSAN are listed as child nodes of the VSAN. When you select a VSAN on the List tab, the VSAN and its member elements are highlighted in the topology view:

Highlighting a Cisco VSAN's Members

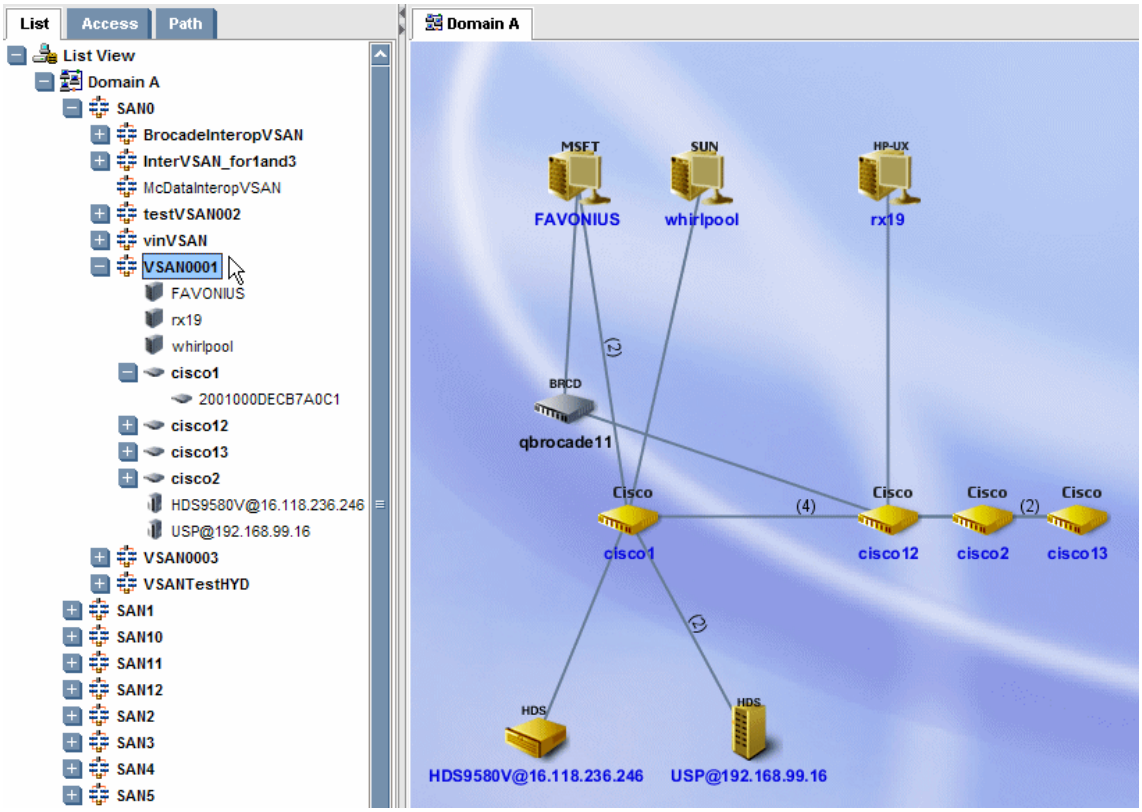

Unlike Brocade, Cisco does not partition the physical switch into virtual switches with individual switch WWNs. For this reason, all Cisco VSANs have different WWNs that are derived from their parent SAN. Each VSAN has an individual name provided by the switch administrator.

To view a the virtual switches associated with a physical switch, expand the node for the physical switch.

Brocade Virtual Fabrics

- San nodes represent the SAN of a single switch, or multiple switches connected via ISL.
- Fabric nodes represent the virtual fabrics created by a Brocade switch administrator.
- HP Storage Essentials displays the logical switch name for Brocade switches in virtual fabrics.

In Brocade's virtual fabric model, a physical switch is virtualized into multiple virtual switches, each with its unique switch WWN. You can tell the difference between virtual switches and physical switches displayed in the topology by placing your mouse over the switch. A popup window indicates if a switch is virtual, as shown in the following figure.

Determining if a Brocade Switch is a Virtual Switch

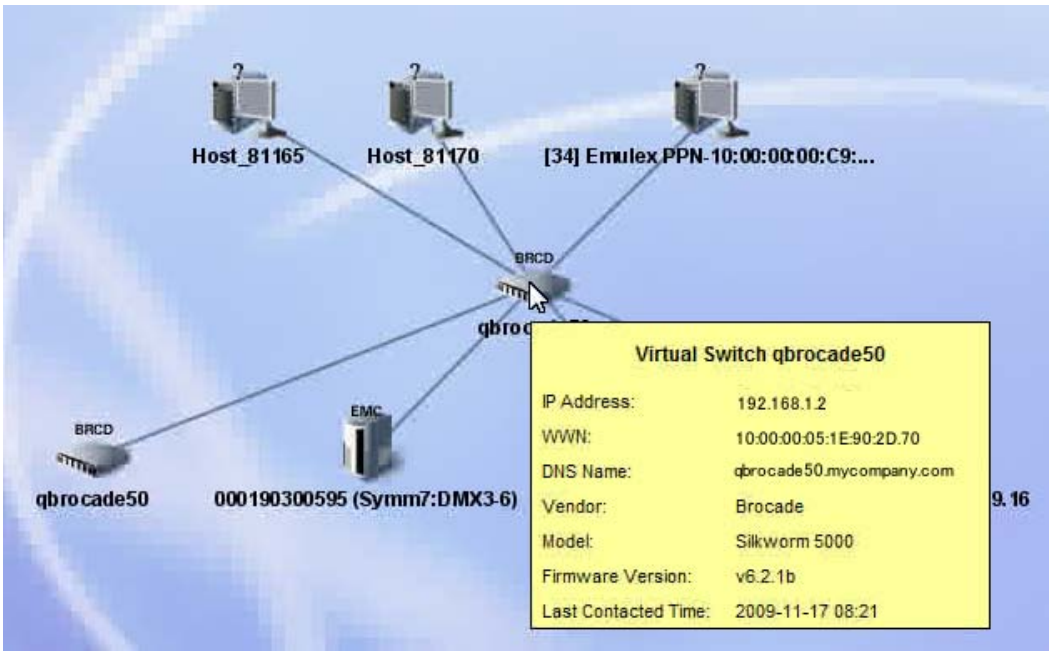

The fabric associated with a virtual switch is called a virtual fabric. Each virtual fabric is also associated with a unique principal switch WWN. When a virtual fabric node is selected on the List tab, the elements included in that virtual fabric are highlighted:

Highlighting a Brocade Virtual Fabric's Members

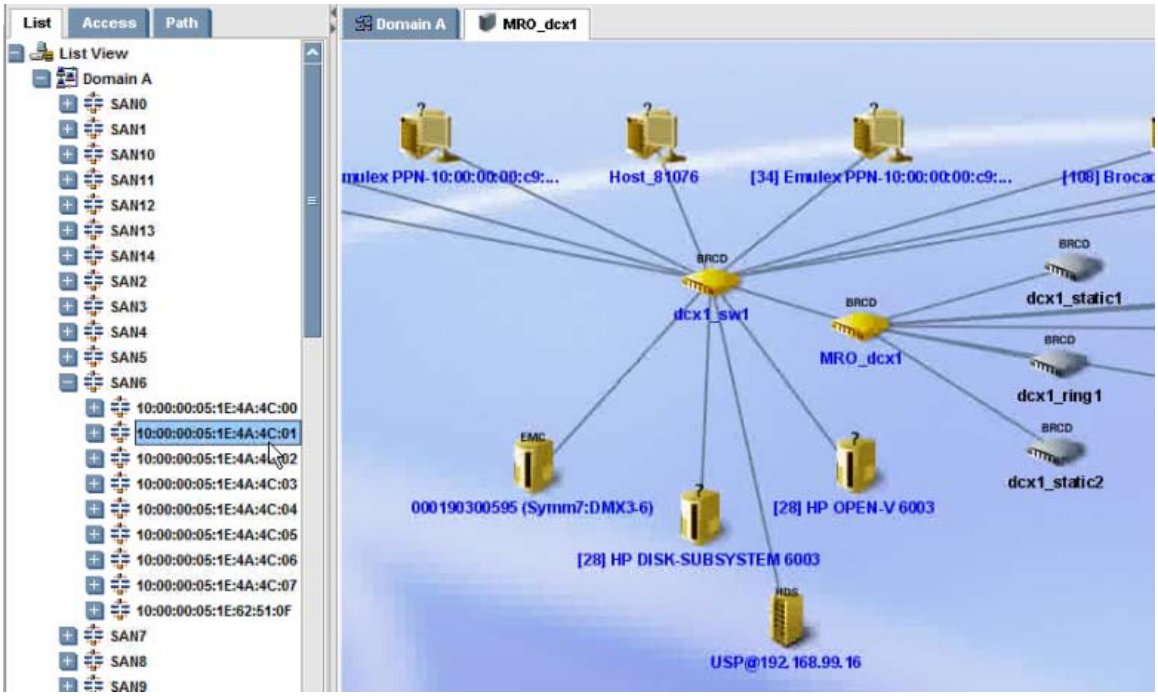

Brocade LSAN

- San nodes represent the aggregation of two or more physically distinct Brocade fabrics
- Fabric nodes represent the contained SANs in an aggregated LSAN.

A Brocade LSAN is the aggregation of two or more SANs into one single SAN. Links between these fabric switches are shown as an IFL link. LSANs are viewed from the Access tab rather than the List tab:

Highlighting a Brocade LSAN Fabric's Members

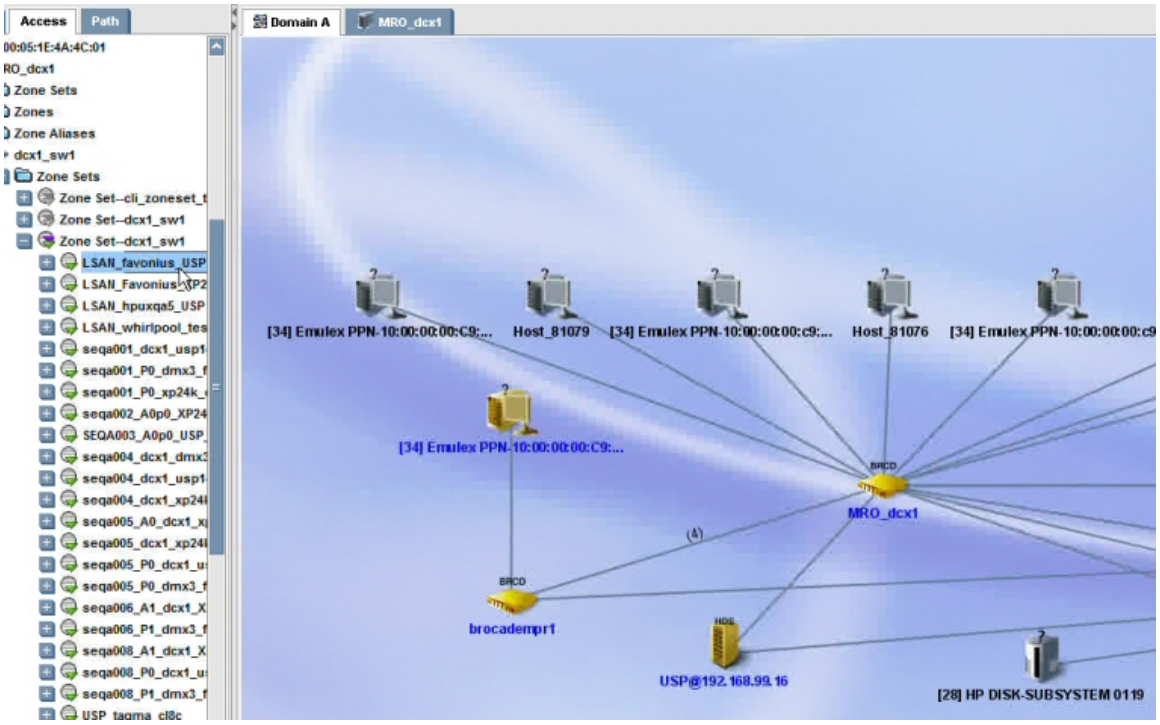

LSANs can aggregate fabrics from different vendors without losing the vendor specific functionality. For example, a Brocade switch can be connected to a McData switch via IFL. The management server displays these switches in a single fabric.

Viewing Clustered Elements

You can view elements that are part of a host cluster or application cluster. To find the elements in a cluster, expand the Host Clusters or Application Clusters node. When you click a cluster name in the tree, its members are highlighted in the topology.

Viewing Virtual Hosts

Virtual servers and their associated virtual machines are listed under the node for the fabric to which they belong. They are also listed under the Virtual Servers node under the All Elements node. When you click a virtual server name in the tree, it is highlighted in the topology. By default, the virtual servers are displayed in a minimized state, so the virtual machines are not automatically displayed. To expand the virtual server to show the associated virtual machines, click the (**+**) button.

Viewing External Storage for Virtualized Storage Arrays

You can view external storage for virtualized storage arrays in the System Manager using either the Path tree or the Topology map.

To use the Path tree:

1. With the List tab open, expand either the **SAN** or **All Elements / Storage Systems** node in the Path tree and locate the storage system for which you want to view external storage elements.

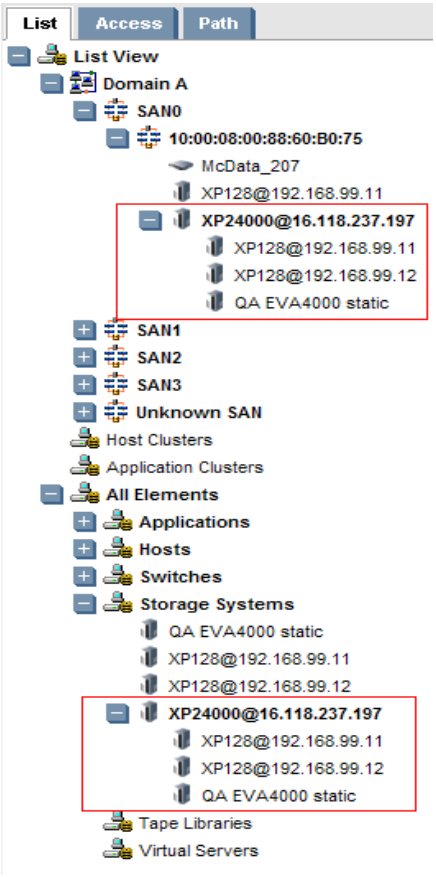

2. Right-click the storage system name and select **Show External Storage**. The Show External Storage menu option is displayed only when a storage system has external storage. If this option does not appear in the drop-down menu, the storage system has no external storage to view.

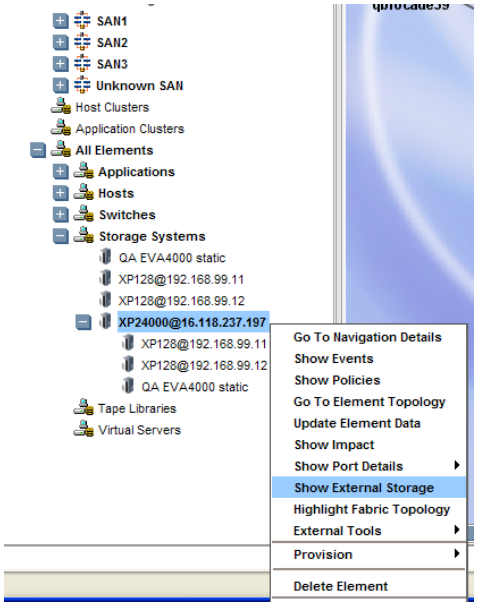

A new Topology window pops up showing the external storage elements in both the Path tree and Topology map.

To use the Topology map:

1. In the System Manager Topology map, locate the storage system icon for which you want to view external storage.
2. Right-click the icon and select **Show External Storage**. If you do not see this option, it means there are no external storage elements discovered for the system.

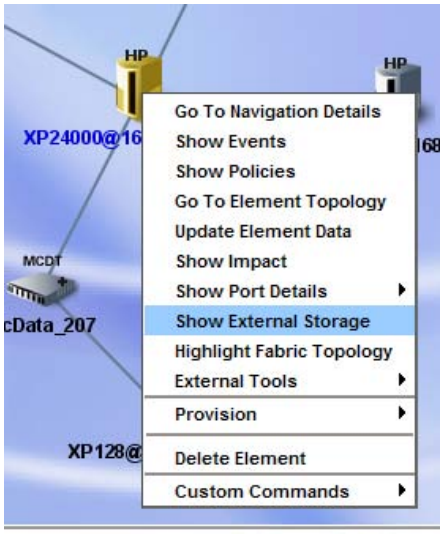

The external storage elements are displayed in a separate Topology tab.

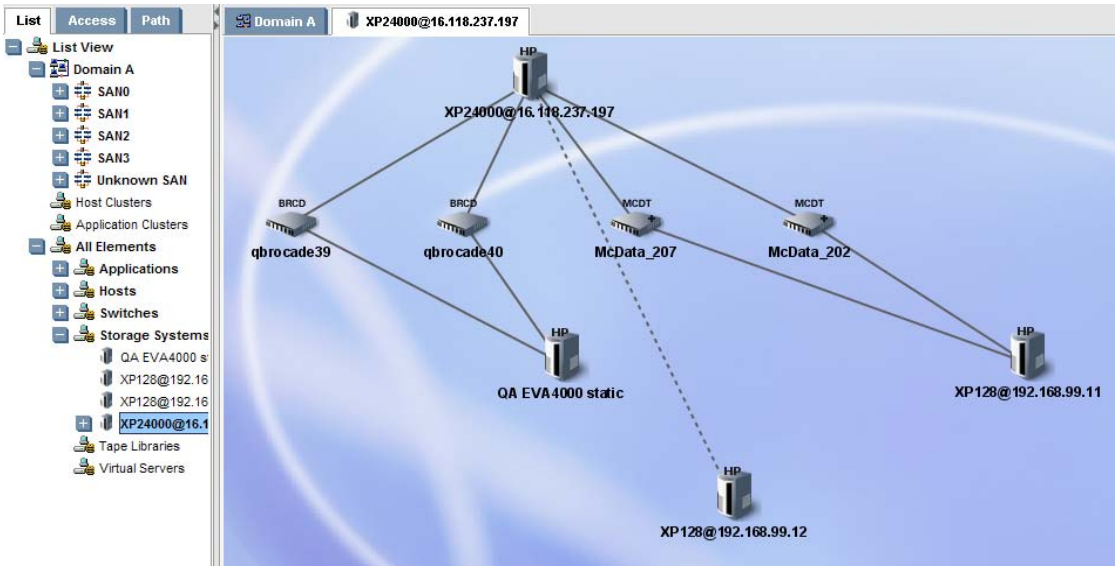

Viewing Back End Topology

To view the topology of a back-end storage system:

1. Using either the Path tree or Topology map, select a storage system element that has back-end storage.
2. Right-click the storage system element (for example, the XP2400 system) and select **Go To Element Topology**.

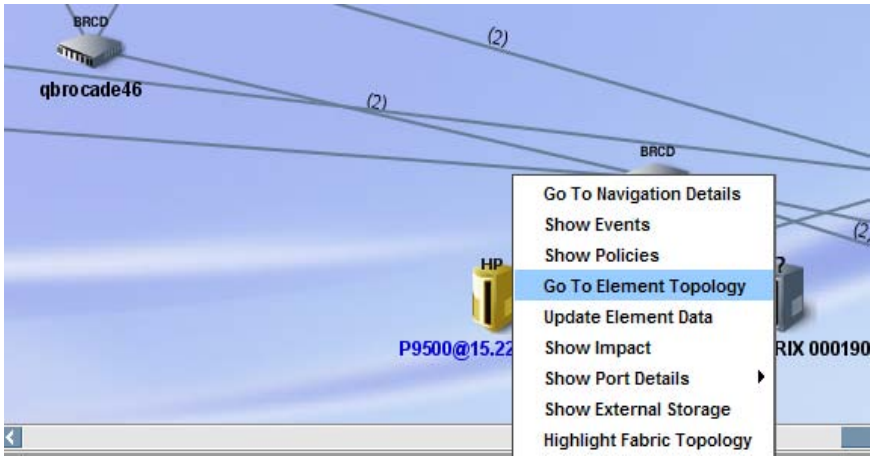

The Element Topology window appears with the storage system you selected highlighted. The name of each back-end storage element is underlined.

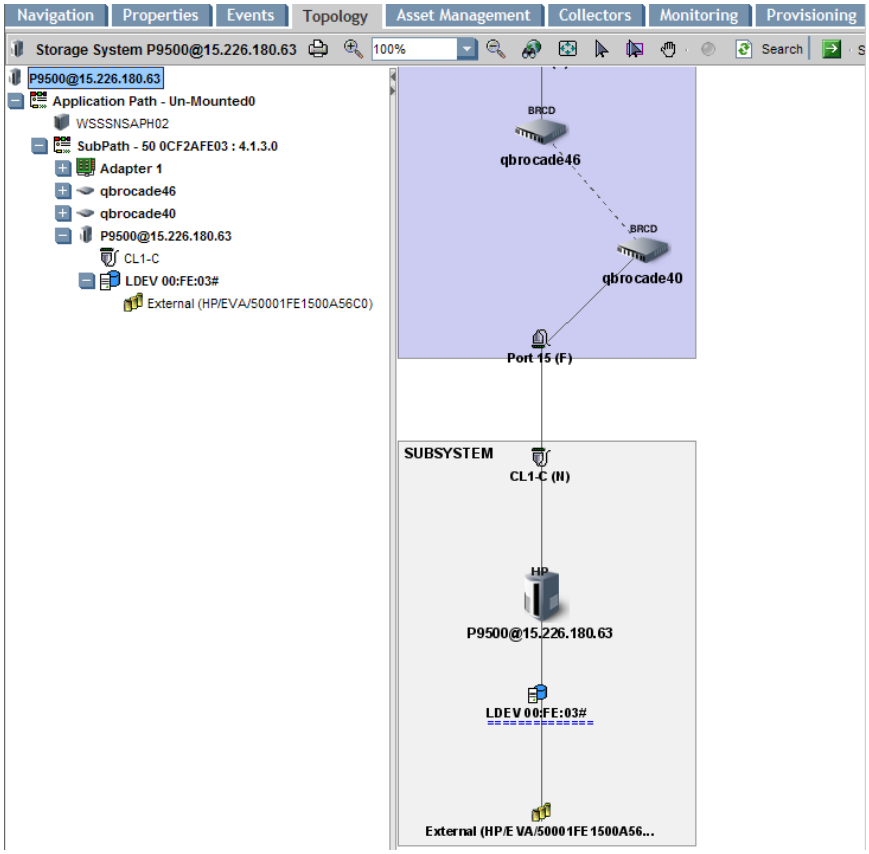

3. To view more details, right-click a back-end element and select **Show Back End Topology**.

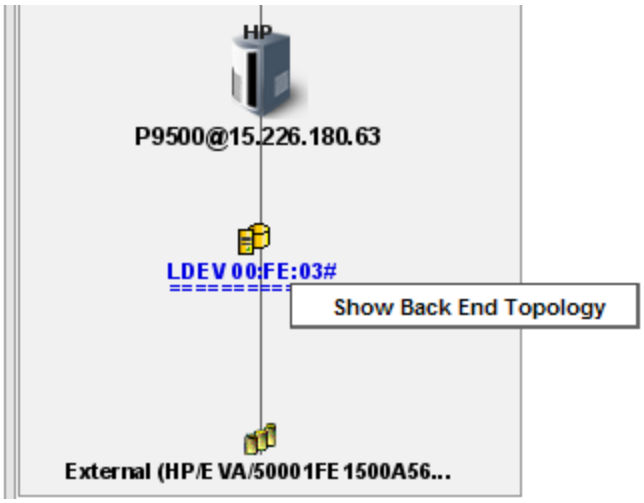

The Back End Topology window for the selected element appears as follows.

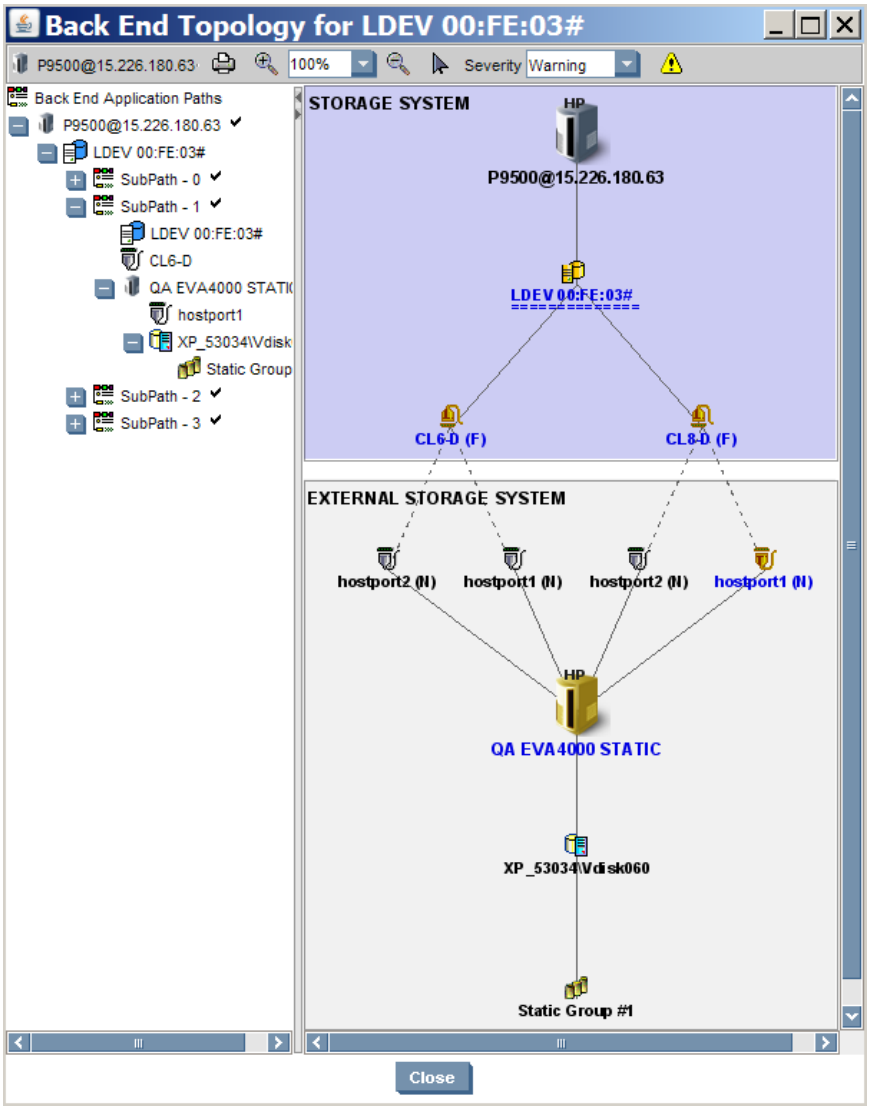

Hiding and Showing Back-End Topology Paths

You can show or hide back-end topology paths by selecting and then right-clicking a path or subpath. Choose **Display Path on Topology** to show the path's topology elements in the right pane. Choose **Hide Path on Topology** to hide it. This enables you to create views of only those paths in which you are interested.

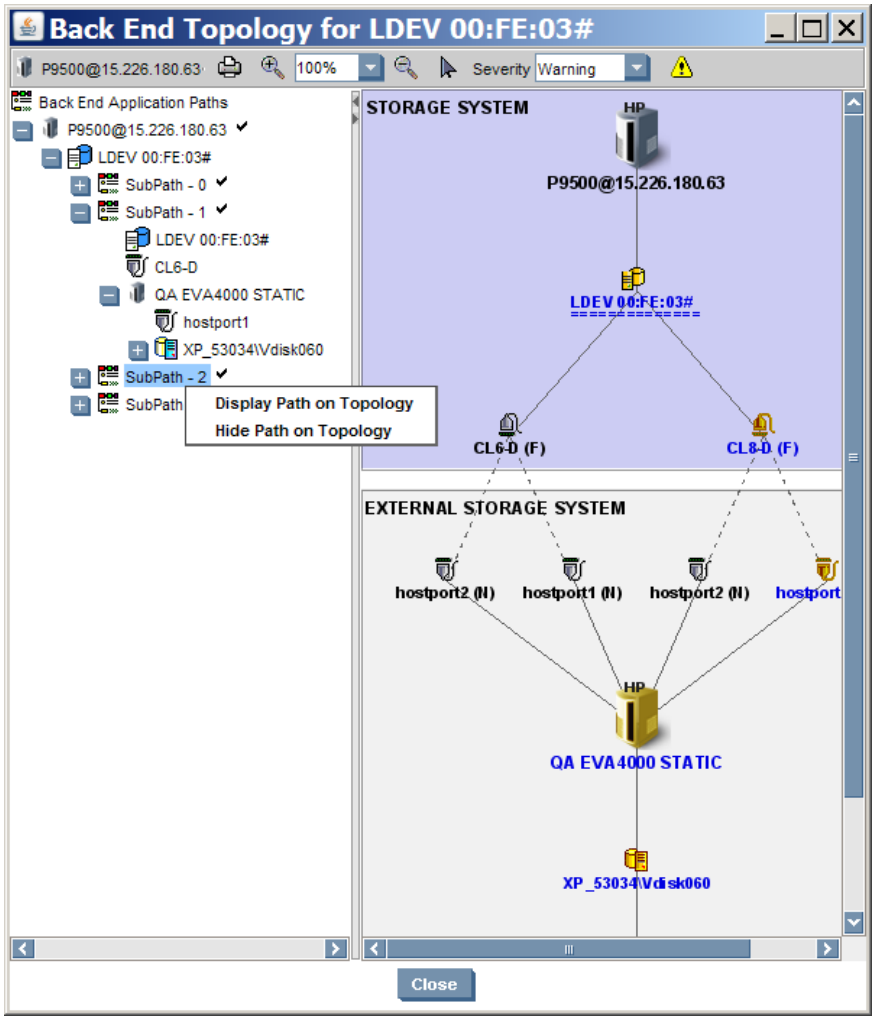

In the left pane, a check mark indicates that the path is displayed as shown in the figure. No checkmark indicates that the path is hidden in the topology view.

Viewing Front End and Back End Dependent Storage Information

In a virtualized storage environment, a front-end storage array (acting as a *storage virtualizer*) serves as the access point for several storage arrays behind it (called the *backend*). Virtualized storage enables administrators to efficiently manage storage volumes and data access for improved performance and cost reductions. HP Storage Essentials collects information about front-end and back-end storage system dependencies and makes it available for viewing on the Navigation page.

Using this information, an administrator can analyze the effectiveness of virtualized storage configurations. It will allow for detailed investigation of data bottlenecks and single points of failure, that can then be remediated for improved throughput and access.

Use either the Navigation Path tree or Topology map to select a storage system for which to view dependencies (See Viewing External Storage for Virtualized Storage Arrays on page 425). Right-click a storage system element that has virtualized storage and select **Go To Navigation Details**. The Navigation page opens.

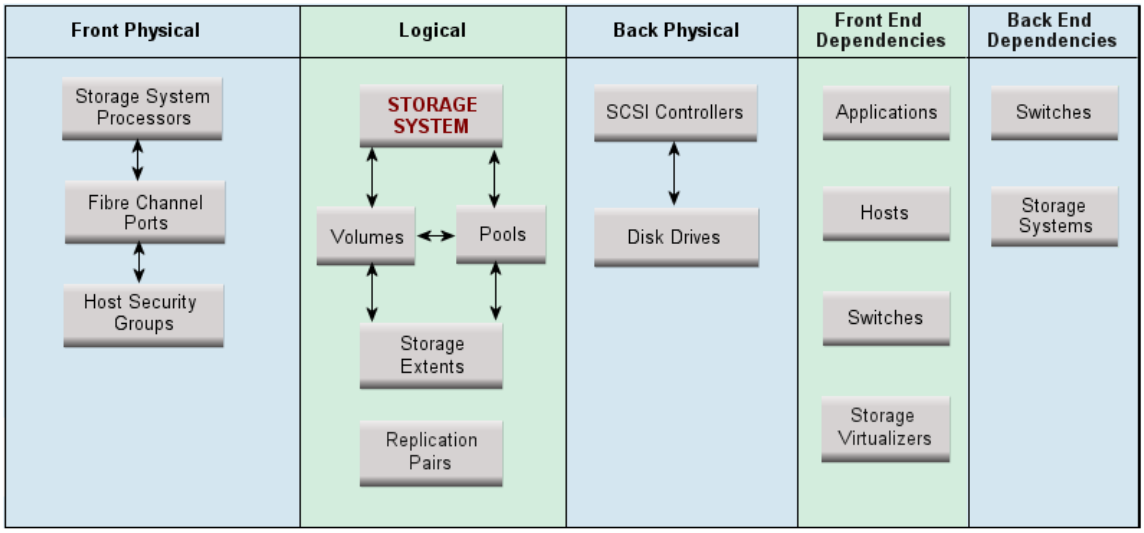

About Front End Dependencies

Front End Dependencies is equivalent to the Dependencies view for a non-virtualized storage system. A storage system acting as a front-end storage virtualizer appears like a regular storage system to its dependents. Navigation results are, therefore, displayed from this perspective. Front-end dependencies information appears exactly as if you were viewing these dependencies on a regular (non-virtualized) storage system.

To view details for front-end dependent storage, click one of the following buttons in the Front End Dependencies column:

Data Button	Description	More Information
Application	Lists front-end dependent applications and shows properties	Finding Front-End Dependent Application Information on the facing page
Hosts	Lists front-end dependent hosts and shows properties	Finding Front-End Dependent Host Information on the facing page
Switches	Lists front-end dependent switches and shows properties	Finding Front-End Dependent Switch Information on page 435
Storage Virtualizers	Lists front-end dependent storage systems which are acting as storage virtualizers and shows properties	Finding Front-End Dependent Storage Virtualizer Information on page 436

About Back End Dependencies

To view back-end details for a dependent storage system or switch, click one of the following buttons in the Back End Dependencies column:

Data Button	Description	More Information
Switches	Lists back-end dependent switches and shows properties	Finding Back-End Dependent Switch Information on page 437
Storage Systems	Lists back-end dependent storage systems and shows properties	Finding Back-End Dependent Storage System Information on page 438

Finding Front-End Dependent Application Information

You can find information about virtualized storage system with front-end dependent applications on the Navigation page (see Viewing Front End and Back End Dependent Storage Information on page 431).

To view dependent application information:

1. In the Path tree or Topology map, right-click a storage system that has external storage elements.
2. Select **Go to Navigation Details**. The Navigation page displays with the name of the storage system at the top.
3. In the Front End Dependencies column, click **Applications**. The following is an example of an Applications information table:

Application	Host	Mount Point	HBA Port	Storage System Port	Storage Volume	LUN	Composition
WSSSNSAPH02	WSSSNSAPH02	J:	Adapter 1 Port 0	CL2-B	LDEV 00:FE:B1#	13	External
WSSSNSAPH02	WSSSNSAPH02	J:	Adapter 0 Port 0	CL1-B	LDEV 00:FE:B1#	9	External

Finding Front-End Dependent Host Information

Host information for virtualized storage systems can be found on the Navigation page (see Viewing Front End and Back End Dependent Storage Information on page 431).

To view dependent host information:

1. In the Path tree or Topology map, right-click a storage system that has front-end storage elements.
2. Select **Go to Navigation Details**. The Navigation page displays the name of the storage system at the top.
3. In the Front End Dependencies column, click **Hosts**. The following information is available

for host dependencies:

Attribute	Definition
Host Name	Unique name that identifies the host machine on this storage network
Type	Indicates whether it is a Host, Virtual Server or Virtual Machine
Virtual Server	If Type is Virtual Machine, this column shows the name of the virtual server on which it is hosted. Data displays in this column only if there are virtual machines discovered.
Mount Point	NTFS file system object which provides the entry point (references root directory of the mounted volume) to other volumes; one volume can host multiple volume mount points.
HBA Port	Host bus adapter port numbers
Storage System Port	Storage system port numbers
Storage Volume	Unique name of the volume. Note that names of external volumes end with a pound (#) sign. This end character is automatically generated for all external volume names.
LUN	Logical unit number of the volume
Composition	Indicates whether storage is External, Internal or Hybrid. You can sort on this column (click column header) to group together those hosts which use external storage volumes.

The following figure shows a sample Host information table:

Dependent Hosts

Hosts associated through the Path

Host Name	Type	Virtual Server	Mount Point	HBA Port	Storage System Port	Storage Volume	LUN	Composition
WEXNMSAPH01	Host			Adapter 1 Port 0	CL2-B	LDEV 00:FE:72#	4 (0x4)	External
WEXNMSAPH01	Host			Adapter 1 Port 0	CL2-B	LUSE 00:FE:34#	1 (0x1)	External
WEXNMSAPH01	Host		O:	Adapter 1 Port 0	CL2-B	LUSE 00:0A:0A	0 (0x0)	Internal
WEXNMSAPH01	Host			Adapter 1 Port 0	CL2-B	LDEV 00:FE:3F#	3 (0x3)	External
WEXNMSAPH01	Host			Adapter 0 Port 0	CL1-B	LDEV 00:FE:72#	4 (0x4)	External
WEXNMSAPH01	Host			Adapter 0 Port 0	CL1-B	LDEV 00:FE:3F#	3 (0x3)	External
WEXNMSAPH01	Host			Adapter 0 Port 0	CL1-B	LDEV 00:00:30	0 (0x0)	Internal
WEXNMSAPH01	Host			Adapter 0 Port 0	CL1-B	LUSE 00:FE:34#	1 (0x1)	External
WEXNMSAPH01	Host		I:	Adapter 0 Port 0	CL1-B	LDEV 00:FE:BB#	5 (0x5)	External
WEXNMSAPH01	Host			Adapter 1 Port 0	CL2-B	LUSE 00:FE:38#	2 (0x2)	External
WEXNMSAPH01	Host			Adapter 0 Port 0	CL1-B	LUSE 00:FE:38#	2 (0x2)	External
WSSSNSAPH02	Host		O:	Adapter 1 Port 0	CL2-B	LDEV 00:FB:0E#	16 (0x10)	External
WSSSNSAPH02	Host		O:	Adapter 0 Port 0	CL1-B	LDEV 00:FB:0E#	14 (0xe)	External
WSSSNSAPH02	Host		W:	Adapter 0 Port 0	CL1-B	LUSE 00:FE:06#	23 (0x17)	External
WSSSNSAPH02	Host		W:	Adapter 1 Port 0	CL2-B	LUSE 00:FE:06#	24 (0x18)	External
WSSSNSAPH02	Host		Y:	Adapter 0 Port 0	CL1-B	LDEV 00:FE:FC X	21 (0x15)	Hybrid
WSSSNSAPH02	Host		Y:	Adapter 1 Port 0	CL2-B	LDEV 00:FE:FC X	22 (0x16)	Hybrid
WSSSNSAPH02	Host		X:	Adapter 0 Port 0	CL1-B	LDEV 00:FE:FB X	20 (0x14)	Hybrid
WSSSNSAPH02	Host		X:	Adapter 1 Port 0	CL2-B	LDEV 00:FE:FB X	21 (0x15)	Hybrid

Finding Front-End Dependent Switch Information

Switch information for virtualized storage systems with front-end storage is available on the Navigation page. Only switches directly connected to the front-end storage system are shown. For information about displaying the Navigation page, see Viewing Front End and Back End Dependent Storage Information on page 431.

To view dependent switch information:

1. In the Path tree or Topology map, right-click a storage system that has front-end storage elements.
2. Select **Go to Navigation Details**. The Navigation page displays with the name of the storage system at the top.
3. In the Front End Dependencies column, click **Switches**.

The following information is displayed for front-end switch dependencies:

Attribute	Definition
Switch	Unique name of a switch device on the network. The name typically includes the name for device type, such as Brocade or Cisco, for easy recognition.
Switch Port Connected	Name of port to which the switch device connects.
Storage System Port	Name of port that connects this device to the storage system.

A Switch information table looks like the following.

Switch	Switch Port Connected	Storage System Port
qbrocade39	Port 6	CL1-B
qbrocade40	Port 6	CL2-B

Finding Front-End Dependent Storage Virtualizer Information

The Navigation page provides a detailed view of storage virtualizer front-end dependent information. See Viewing Front End and Back End Dependent Storage Information on page 431.

To view front-end storage virtualizer dependencies:

1. In the Path tree or Topology map, right-click a storage system that has front-end storage elements.
2. Select **Go to Navigation Details**. The Navigation page displays with the name of the storage system at the top.
3. Click **Storage Systems** in the Front End Dependencies column. This button is available only if the storage system has front-end storage elements discovered.

The following information is displayed for front-end storage virtualizer dependencies:

Attribute	Definition
Frontend Volume	Unique volume name which may include the logical device (LDEV) *number*that makes up the volume
Initiator Port	Local port on the storage system virtualizer
Switches	Switches connected on path of storage virtualizer
Target Port	Port to which LDEV is mapped on the storage system
Frontend Port	Port through which storage volumes are presented on the storage system frontend
Frontend Storage System	Name uniquely identifies a front-end storage system

Filtering Information to Display

You can filter information by selecting the criteria for displaying available data.

To filter information:

1. Click the **+** sign that precedes **Filter** in the left corner above the properties table. This displays the Filter fields.

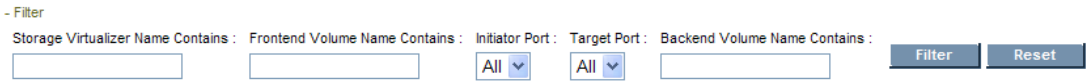

2. Select the criteria you want to use to display information for the front-end storage virtualizer. You can use the following filters:
   - Storage Virtualizer Name Contains – Type a character string on which to search.
   - Frontend Volume Name Contains – Type a character string on which to search; for example, type LUSE to find mapped volumes which include "LUSE" in their volume name.
   - Initiator Port – From the drop-down menu, select an initiator port.
   - Target Port – From the drop-down menu, select a target port to show all volumes which map to the port.
   - Backend Volume Name Contains – Type a character string on which to search available back-end volume names.
3. Click **Reset** if you want to clear the settings and re-specify the filter conditions.
4. When you are ready, click **Filter** to display the filtered information on the Navigation page.

Finding Back-End Dependent Switch Information

Back-end switch information for virtualized storage systems is available on the Navigation page. Only switches directly connected to the back-end storage system are shown.

To view dependent switch information:

1. In the Path tree or Topology map, right-click a storage system that has back-end storage elements.
2. Select **Go to Navigation Details**. The Navigation page displays with the name of the storage system at the top.
3. In the Back End Dependencies column, click **Storage Systems**.

The following information is displayed for back-end switch dependencies:

Attribute	Description
Switch	Unique name of a switch device on the network. The name typically includes the name for device type, such as Brocade or Cisco, for easy recognition.
Switch Port Connected	Name of port to which the switch device connects.
Storage System Port	Name of port that connects this device to the storage system.

A Dependent Back End Switches information table looks like the following:

Switch	Switch Port Connected	Storage System Port
McData_202	Port: 2	CL8-D
McData_207	Port: 2	CL4-D
qbrocade39	Port 12	CL6-F
qbrocade40	Port 12	CL2-F

Finding Back-End Dependent Storage System Information

HP Storage Essentials collects back-end dependent storage system information and makes it available from the Navigation page.

To view back-end storage system dependencies:

1. In the Path tree or Topology map, right-click a storage system that has discovered back-end storage elements.
2. Select **Go to Navigation Details**. The Navigation page displays with the name of the storage system at the top.
3. Click the **Storage Systems** button in the Back End Dependencies column. This button is available only if the storage system has back-end storage elements discovered.

The following information is displayed for back-end storage system dependencies:

Attribute	Definition
Backend Volume	Unique volume name for the SCSI LUN exposed on the fabric by a Storage Controller (typically a RAID array) to the SAN Volume Controller
Initiator Port	Local port on the storage system virtualizer
Switches	Switches on fabric of storage system backend
Target Port	Port on back-end array to which the volume is mapped
Backend Port	Port through which storage volumes are presented from back-end storage systems
Backend Storage System	Name uniquely identifies a back-end storage system

The following figure shows a sample Back End Dependent Storage System information table:

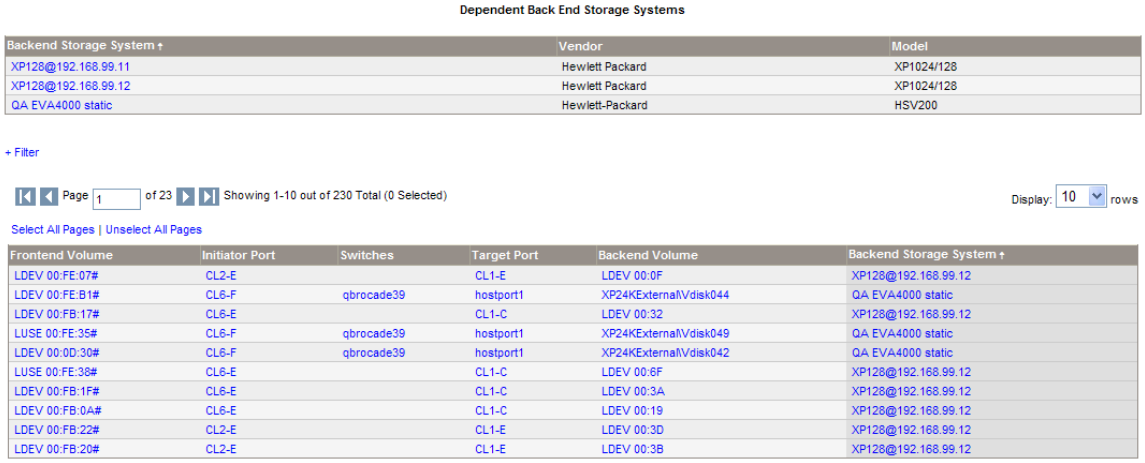

Filtering Information to Display

You can filter information by selecting the criteria you want to use to display available data.

To filter information:

1. Click the **+** sign that precedes **Filter** in the left corner above the properties table. This displays the Filter fields.

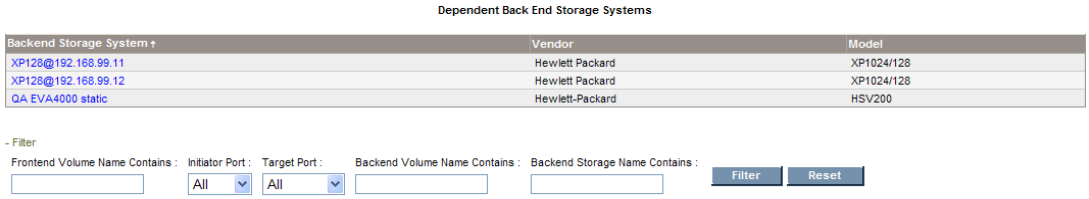

2. Select the criteria you want to use to display information for the back-end storage system. You can use the following filters:
   - Frontend Volume Name Contains – Type character string on which to search. For example, type LUSE to find mapped volumes which include "LUSE" in their volume name.
   - Initiator Port – From the drop-down menu, select an initiator port.
   - Target Port – From the drop-down menu, select a target port to show all volumes which map to the port.
   - Backend Volume Name Contains – Type a character string on which to search available backend volume names.
   - Backend Storage Name Contains – Type a character string on which to search available backend storage system names.
3. Click **Reset** if you want to clear the settings and re-specify the filter conditions.
4. When you are ready, click **Filter** to display information filtered on your selected criteria.

Viewing Elements by Type

You can view elements by type under the All Elements node on the List tab. This is especially helpful in determining how many elements you have of a specified type, such as the number of storage systems.

When you expand the tree of the element type node, all elements of that type are listed. If you select the element type node, all elements of that type are selected in the topology. For example, assume you want to determine the number of applications that the management server monitors. When you expand the tree of the Applications node, the applications are listed. When you select the **Applications** node, the applications are highlighted in the topology, as shown in the following figure.

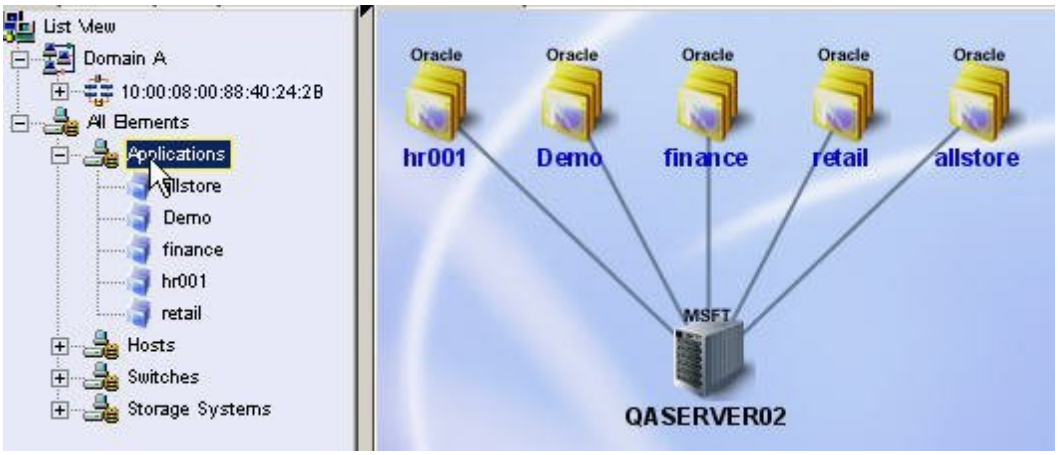

If you select an element in the left pane, the element is highlighted in the topology. You also have access to additional functionality by right-clicking the element. For more information, see About the Right-Click Menu Options on page 447.

Access Tab

The Access tab provides information about the following:

- Zone entries
- Host bindings
- Storage system LUN masking

You can also manage zones, zone aliases, and zone sets from this tab by right-clicking an element. See About the Right-Click Menu Options on page 447.

Obtaining Information about Zone Sets

To view the zone sets in a domain, expand the tree for the domain, zone entry, SAN, fabric, switch, and zone set. Select the zone set node to see the members of the zone set highlighted in the right pane, as shown in the following figure:

Members of a Zone Set

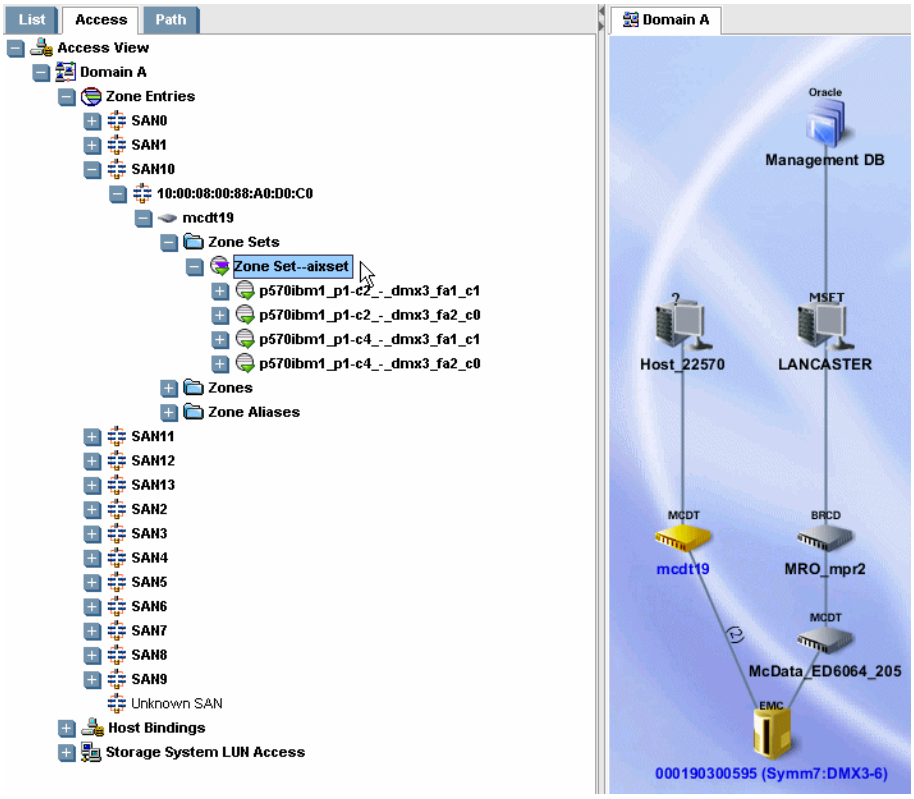

To view members of a zone, do one or more of the following:

- Expand the node of the zone in the tree. The software displays the zone members underneath the node of the zone.
- Click the node of the zone in the tree. The software highlights the zone members in the right pane.

Displaying a Zone Member and Its Switch

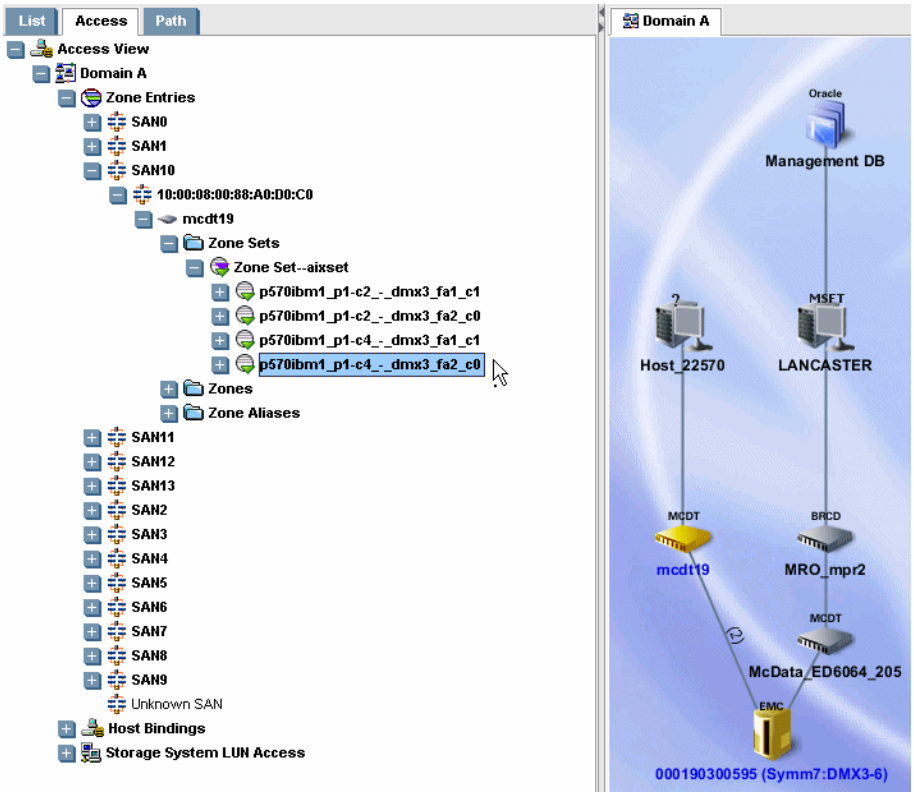

A Brocade LSAN can belong to multiple fabrics. In this case, the LSAN node displays multiple times in the access tree. When you select an LSAN node in the tree, its zone members in every fabric are highlighted in the topology

To view information about a zone member's port, expand the zone member node, as shown in the following figure. When you select the zone member node in the tree, it appears highlighted in the right pane.

Obtaining Information About a Zone Member's Port

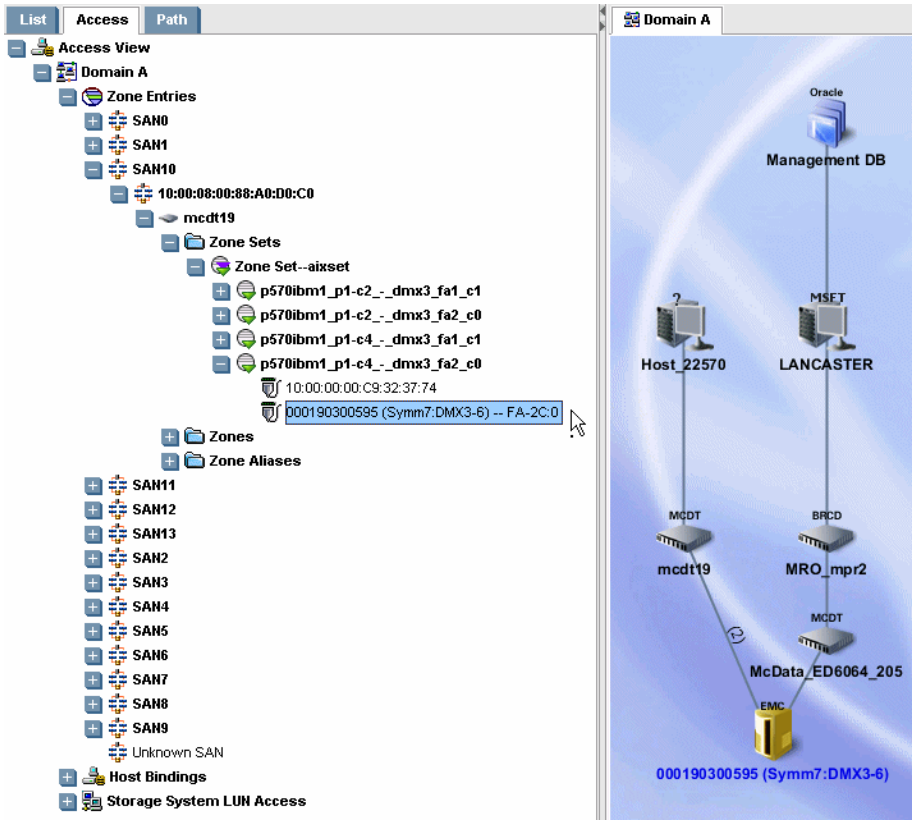

Obtaining Information about Host Bindings

To view the elements that have host bindings, click the **Host Bindings** node in the tree. Elements with host bindings are highlighted in the right pane, as the following figure shows.

Highlighting Elements with Host Bindings

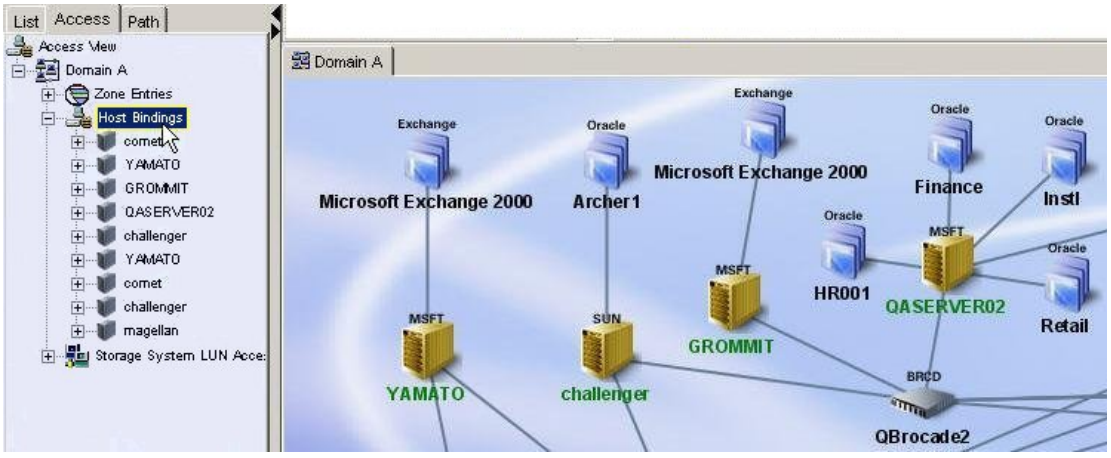

To view a host binding, click the **HBA** node in the tree. The HBA node is under the element node.

When you click the HBA node, the host and the element to which it has the binding are highlighted. A green line between the two elements indicates they have a binding, as shown in the following figure:

Displaying Host Bindings

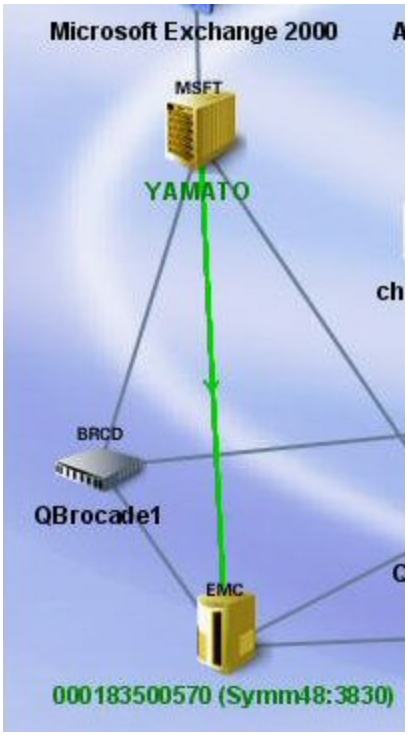

To view information about the ports on an HBA card, expand the HBA node in the tree, as shown in the following figure:

HBA Port Properties

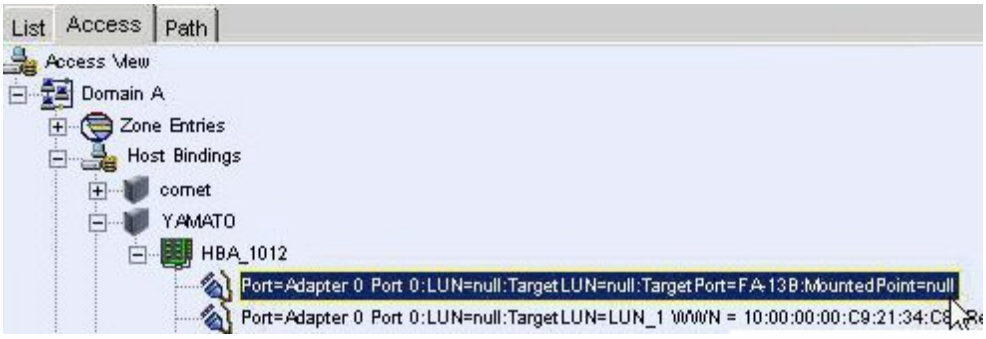

Obtaining Information about Storage System LUN Masking

To obtain information about a storage system's LUN masking, first expand the fabric node, and then expand the storage system node, and then click the **Fibre Channel** port. The values of the WWNs are displayed under the node, and the storage system is highlighted in the right pane, as shown in the following figure.

WWN Properties

The software displays properties of the WWN. If the LUN has a LUN masking, expand the LUN node to obtain information about the LUN masking, as shown in the following figure.

WWN Properties

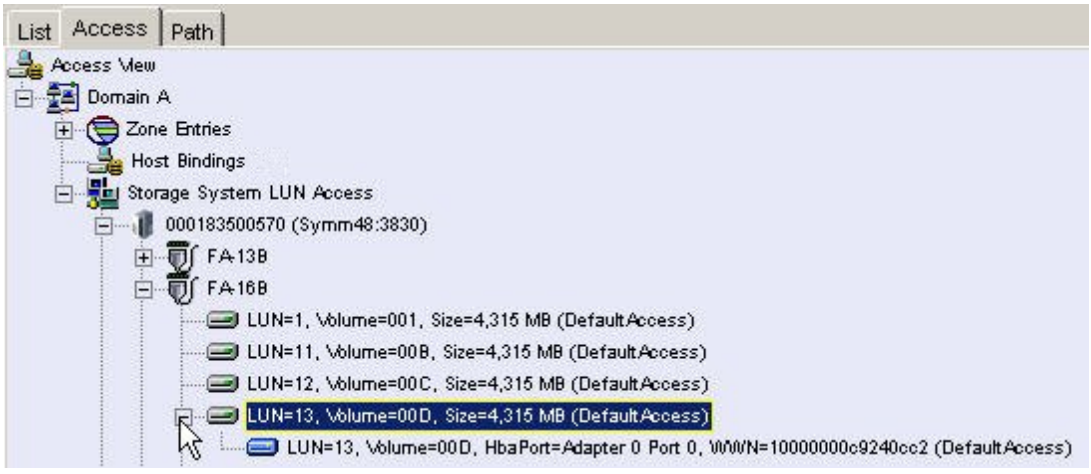

To view a LUN masking, expand a LUN node.

About the Path Tab

The Path tab provides information about an element's path. By clicking a host's node, you can determine the host's path in the application.

When you expand a domain node, if any of the paths for hosts are not fully calculated, a pop-up dialog box displays a list of all the hosts with partially calculated paths. In addition, the current state of the path calculation is appended to the node name.

When you click a host node in the tree, the elements in the host's path appear highlighted in the right pane, as shown in the following figure.

Obtaining Path Information

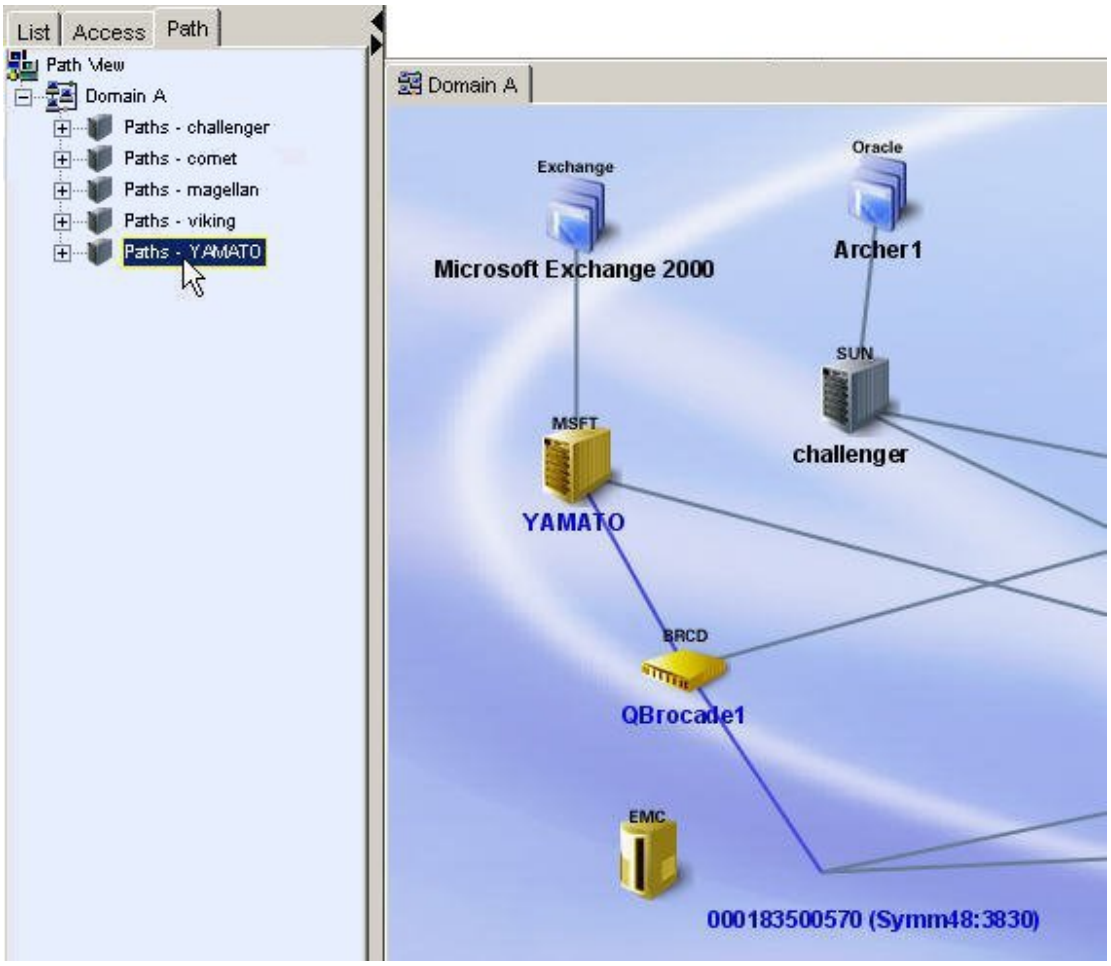

Note: If you are using the Load-on-Demand feature, you must load the path first by expanding the tree node. If a path tree node is not loaded, there are no elements under the path node and no highlighting occurs. For more information about the Load-on-Demand feature, see System Manager, Capacity Manager and Performance Manager Preferences on page 301.

You can also determine the elements in a hosts path by expanding the Application Path and Path nodes under the host node in the tree, as shown in the following figure.

Path Information Visible in the Tree

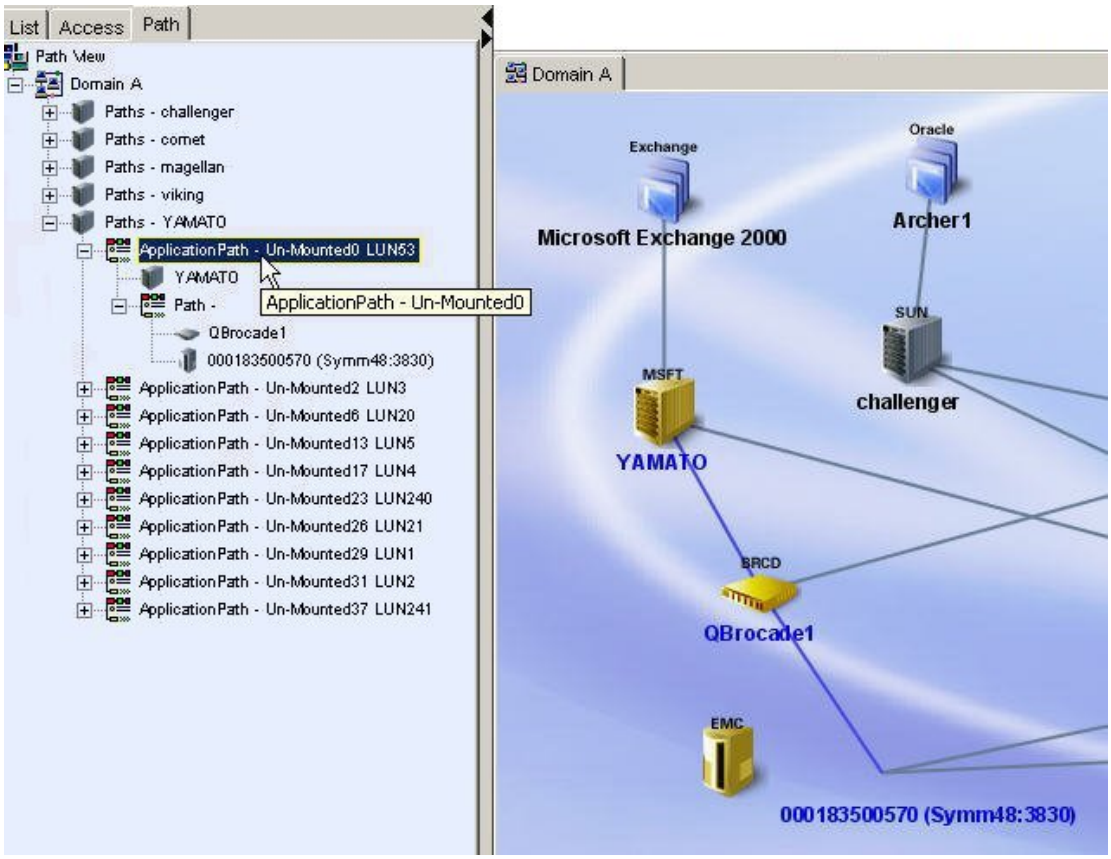

When you right-click an element in the List tab, a menu is displayed. The options displayed depend on the type of element you clicked. For an explanation of the options in the menu, see About the Right-Click Menu Options below.

About the Right-Click Menu Options

When you right-click an element in the topology pane or in the List, Access, or Path tab, you see a menu. The options displayed in the menu depend on the type of element clicked.

See the following table for an explanation of the options displayed for elements right-clicked in the topology or in the List or Path tab. For the options on the Access tab, see the table About the Right-Click Menu Options above.

Right-click menu options are not available to undiscovered fabrics.

Menu Options Accessible from the Topology*

Menu Option	Description
Go to Navigation Details	This menu option redirects you to the Navigation page. If the element is labeled discovered, you are shown the Properties page. Elements are labeled discovered when the management server has discovered the element but cannot obtain additional information about it. See About the Navigation Tab on page 489. **Note**: You can access this option for fabrics by right-clicking the fabric on the list tab.
Go to Element Topology	Displays a graphical representation of the path of an element. This also includes multipathing. See Viewing Element Topology on page 501.
Show Events	Displays the events for an element. See About the Events Tab on page 511.
Show Policies	Displays the Policy tab for the element. You can then view, add, modify, and delete polices assigned to the element. See About the Policies Tab on page 515.
Update Element Data	The management server gathers new and changed details from the element and then redraws the topology with the updated information. The Update Element Data functionality does not detect element components that were removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you right-click the storage system and select **Update Element Data**, the LUNs still appear in the user interface. You must perform Get Details for the deleted LUNs to be removed from the user interface. See Get Details on page 151. Update Element Data does not work for virtual machines or virtual servers.
Show Impact	Highlights the elements that are impacted. See Showing the Impact of an Element on page 464.
Show Cluster Impact	Highlights the elements that are impacted. See Showing the Impact of an Element on page 464.
Show Virtual Switches	Opens a new tab that displays the relationship between the physical and virtual switches and the connected devices. Only applies to Cisco and Brocade switches.
Show Port Details (for switches)	Opens a new tab that displays all of the ports on the switch.

Menu Option	Description
Show Port Details (for virtualized storage arrays)	Provides a sub-menu from which you can select Front End Port Details or Back End Port Details. Selecting Front End Port Details opens a new tab that displays connections between the selected storage array and its front-end hosts. Selecting Back End Port details opens a new tab that displays port connections between the selected virtualized storage array and external storage arrays. A dotted line indicates a directly attached storage array.
Show Front End Port Details (for non-virtualized storage arrays)	Opens a new tab that displays connections between the selected storage array and its front-end hosts.
Show Back End Port Details	Opens a new tab that displays port connections between the virtualized storage array and external storage arrays. A dotted line indicates that the external storage is directly attached to the virtualized storage array.
Show External Storage	Opens a new tab that displays the back-end topology from the selected virtualized storage array to the linked external storage arrays. A dotted line indicates direct attached external storage.
Highlight Fabric Topology	Highlights all the other elements in the fabric to which the selected element belongs. If you right-click a switch that belongs to multiple fabrics, you can select which fabric to highlight.
Build/Edit Cluster	Enables you to manually build or edit host and application clusters. See Host and Application Clustering on page 267.
Show Virtual Machines	Adds a Show Virtual Machines tab to the topology. This tab displays the topology for an ESX Server and its virtual machines.

Menu Option	Description
External Tools	Provides several ways to access an element: • **Telnet** – Enables you to access a host or a switch through the telnet utility. The Telnet feature is accessible only to Web browsers on Microsoft Windows operating systems. • **Browse** – Enables you to access the main Web page for a host or a switch. • **Set Up External Tools** – Enables you to add URLs for accessing the management tools for the storage system. In some instances, the management tool for the storage system is accessible from this menu. For example, HiCommand for HDS storage systems and Command View for HP XP storage systems are accessible from the External Tools menu. See Using External Tools on page 488.
Discovered Elements	Enables you to group unnamed generic hosts and storage devices. See Hiding and Showing Generic Hosts on page 475 and Hiding and Showing Generic Storage Devices on page 477.
Add Virtual Application	Enables you to add a virtual application so you can monitor it. A virtual application is a placeholder you create for an application. For example, you could create a virtual application for an application that was created just for your company. See Adding a Virtual Application on page 455.
Set Business Cost (Applications only)	Enables you to assign a business cost to an application. See Assigning a Business Cost to an Application on page 466.
Delete Element	Removes an element and its discovery instance from the system. It also removes other elements discovered through the removed element.
Group together with other elements	Enables you to group “Discovered” hosts and storage systems. See one of the following topics: • Grouping Discovered Hosts on page 470 • Grouping Discovered Storage Systems on page 472

Menu Option	Description
Ungroup into multiple elements	Enables you to ungroup "Discovered" hosts and storage systems. See one of the following topics: • Ungrouping Discovered Hosts on page 471 • Ungrouping Discovered Storage Systems on page 473
Recalculate Topology	Enables you to know about topology changes. The management server contacts the elements on the topology list (**Discovery** > **Topology**) to determine topology changes. The management server uses this updated information to redraw the topology.
Change Fabric Name (Only Available from the List Tab)	Enables you to change the name of the fabric. See Changing the Fabric Name on page 474.
Change SAN Name (Only Available from the List Tab)	Enables you to change the name of the SAN. See Changing the SAN Name on page 474 for more information.
Go to Backup Manager	Opens Backup Manager in a new window. Using Backup Manager on page 729

*Additional menu items can appear for types of automators and advisors, such as Reachable Storage. The descriptions of the right-click menu options for a fabric appear in the following section, since these menu options are only accessible from the List tab.

Right-clicking a fabric in the List tab reveals the following options:

- **Change Fabric Name** – Enter the new fabric name in the Change Fabric Name window, and then click **OK**.
- **Go to Navigation Details** – Displays the navigation tab for the fabric. See About the Navigation Tab on page 489 for more information.
- **Go to Properties** – Displays the properties of the fabric. See Viewing Element Properties on page 498 for more information. It also provides access to the Events and Provisioning Manager tabs. The Events tab displays events occurring within the fabric. The Provisioning

Manager tab enables you to set up and manage zone provisioning. Provisioning Manager might not be available in your kit. To determine if you can access Provisioning Manager, access the List of Features, which is accessible from the Documentation Center.

- **Delete This Fabric** – Deletes a fabric. When you are asked if you want to delete the fabric, click **Yes** if you do not mind waiting for the management server to recalculate the topology. If the elements in the deleted fabric do not belong to another fabric, they are moved to the "unknown" node on the List tab.

Right-clicking a SAN in the List tab reveals the following options:

- **Go to Properties** – Displays the properties of the fabric. See Viewing Element Properties on page 498 for more information.
- **Delete This SAN** – Deletes a SAN. When you are asked if you want to delete the SAN, click **Yes** if you do not mind waiting for the management server to recalculate the topology. If the elements in the deleted SAN do not belong to another SAN, they are moved to the "unknown" node on the List tab.

Provisioning features are available from the right-click menu in the Access tab. When you right-click a zone, zone alias, or zone set in the Access tab, a menu is displayed. The provisioning options in this menu depend on the type of element clicked.

If you do not see the provisioning features from the right-click menu, your license is not allowing you to access provisioning.

The following table describes the menu options. To learn more about each task and its required steps, see Provisioning Manager on page 675.

Task	Right-Click the Following to Perform Task
Activate a zone set	A zone set under the Zone Entries node 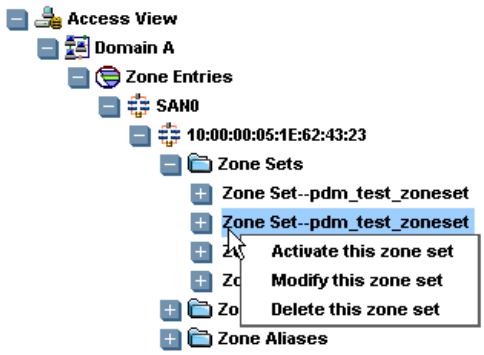
Modify a zone set	
Delete a zone set*	

Task	Right-Click the Following to Perform Task
Create a zone	The Zones folder under a fabric node 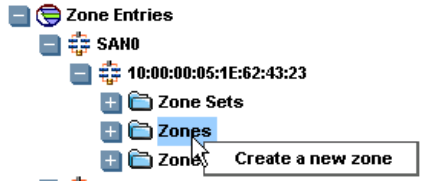
Create an LSAN zone	A SAN or fabric node 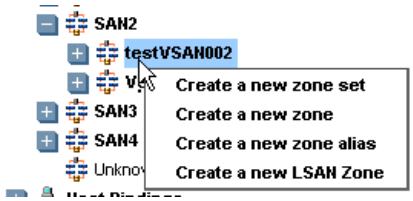
Delete a zone Modify a zone Show zone details	A zone under the Zone node 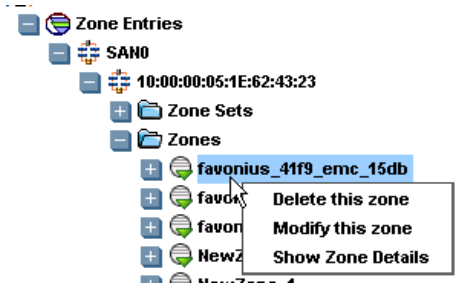
Create a zone alias	The Zone Alias folder under a fabric node 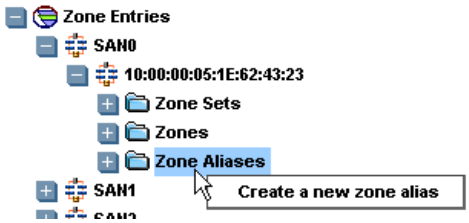
Delete a zone alias Modify a zone alias Show zone alias details	A zone alias under the Zone Aliases node.

Task	Right-Click the Following to Perform Task
Provision a storage pool	A storage system under the Storage System LUN Access node 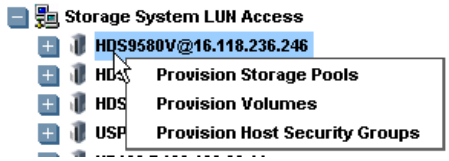
Provision a volume	
Provision a host security group	

*The delete zone set option is disabled for active zone sets.

Viewing Storage Elements

System Manager has a wide range of features to help you in viewing your storage elements in the topology. For example, you can filter by SANs, arrange elements in the topology, and search for elements in the topology. You can even view the status of elements in the topology and find the impact of removing an element.

See the following topics for more information:

- Adding a Virtual Application on the facing page
- Adding Information for Discovered Hosts on the facing page
- Arranging Elements in the Topology on page 456
- Closing Topology Windows on page 458
- Using the Global View on page 458
- Printing the Topology on page 459
- Exporting the Topology to Microsoft Visio on page 460
- Updating Element Data on page 463
- Viewing Ports on page 464
- Showing the Impact of an Element on page 464
- Assigning a Business Cost to an Application on page 466
- Expanding the Topology Pane on page 467
- Filtering SANs on page 468
- Viewing Event Status in the Topology on page 468
- Custom Name for a Switch Truncated in the Topology on page 470
- Managing Groups on page 470
- Managing SANS and Fabrics on page 473
- Hiding and Showing Generic Hosts on page 475

Adding a Virtual Application

The management server enables you to monitor applications not listed in the support matrix. For example, assume your company has created an internal application, and you want to be able to use the management server to monitor that application. You can create a virtual application for that product. A virtual application is a placeholder you create for an application.

Only a user belonging to a role that has System Configuration selected on the Edit Role page (such as the Domain Administrator role) is allowed to create a virtual application.

It is not possible to add a virtual application to a virtual machine.

Once you create the virtual application, it will appear as connected to a host in your topology.

1. Select a host.
2. Right-click, and select **Add Virtual Application**.
3. Enter the following information for the virtual application:
   - **Name**
   - **Product**
   - **Description**
   - **Vendor**
   - **Version**
4. Select the storage volume for the application.

   You can view the properties of a volume by clicking its link.
5. Click **OK**.

The virtual application appears connected to the selected host.

Adding Information for Discovered Hosts

The software labels a host as discovered when it cannot obtain additional information about a host it has discovered. To learn why the software was unable to obtain information about the element, see Troubleshooting Discovery and Get Details on page 860.

When a host has been discovered but no more additional information can be obtained from it because it does not have a CIM extension installed, it is referred to as a generic host.

If you have more than one discovered host, it can be difficult to differentiate them. To make them easier to identify, you might want to add information about the host, such as the following:

- Custom Name - Once you provide a custom name to a generic host, you will be able to access the Asset Attributes of an Element on page 512 and About the Presented Storage Tab on page 515 tabs for that host.
- IP Address
- DNS Name

- Operating System
- Version of the operating system

Do not add information for generic elements during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details by looking at the label near the status button.

1. To access System Manager, click **System Manager** () in the left pane.
2. Double-click a "Discovered" host in the right pane. The Properties tab is displayed.
3. In the custom name box, enter a name for the element.

   Keep in mind the following:

   - The name must contain 1 to 64 characters.
   - The following characters and symbols are accepted: letters, numerals (0 to 9), ~, @, *, _, -, +, ., < >, (), [ ], { }, |.
   - The name is case sensitive. For example, "Element1" and "element1" are different elements.
4. In the IP Address box, enter an IP address for the element.
5. In the DNS Name box, enter a DNS name for the element.
6. In the Version box, enter the version of the operating system.
7. In the Operating System box, enter one of the following operating systems:
   - **AIX** – corresponds to IBM AIX
   - **HP-UX** – corresponds to all versions of HP-UX
   - **IRIX** – corresponds to SGI IRIX
   - **Linux**
   - **Windows** – corresponds to Microsoft Windows
   - **Solaris** – corresponds to Sun Solaris
   - NonStop
8. Click **Save**.

When you access System Manager, the information you entered appears in the topology.

Arranging Elements in the Topology

To improve usability, you can arrange the topology to suit your environment. For example, if you plan to filter by various SANs, you could arrange the topology so that elements are arranged by SAN. When you filter the SANs, large gaps will not appear in the topology. You can arrange elements individually or in groups.

The topology displays direct attached connections as a dotted line from the host to the storage system. To view direct attach storage, you must enable the button. For more information, see Toolbar in System Manager on page 417.

To arrange elements individually:

1. Click the button.
2. Click the element you want to move and drag it to a new location.
3. Repeat the previous step for each element you want to move.

   The management server provides buttons to help you with viewing and arranging the topology. To learn more about those buttons, see Toolbar in System Manager on page 417.
4. Once you finish arranging the topology, click the button to save it.

To learn more about filtering SANs, see Filtering SANs on page 468.

To arrange elements in a group:

1. Click the button.
2. Holding down the mouse button, move the cursor diagonally across the elements you want to move.

   A square encloses the elements, as shown in the following figure. To redo the square, click outside of the square and retry.

   Enclosing the Elements

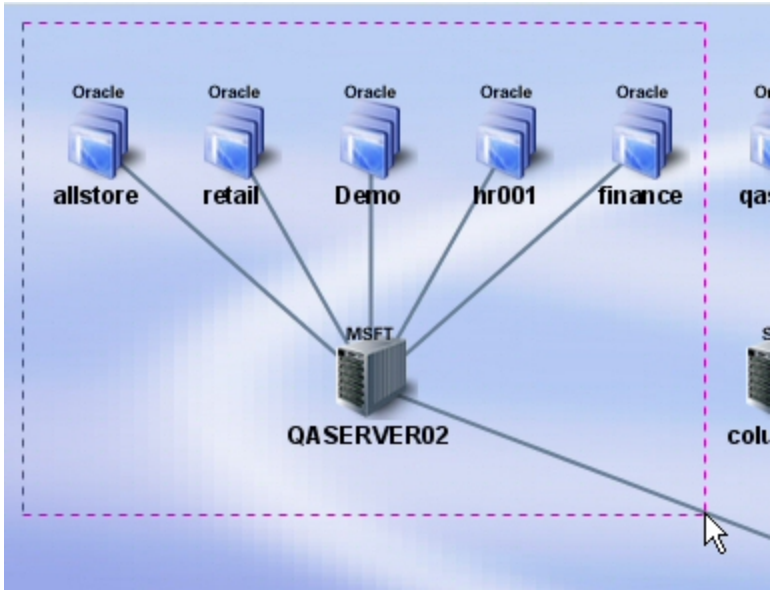

3. To move the elements within the square, click within the square. Holding down the mouse button, drag the elements to the new location.

Dragging Multiple Elements to Their New Location

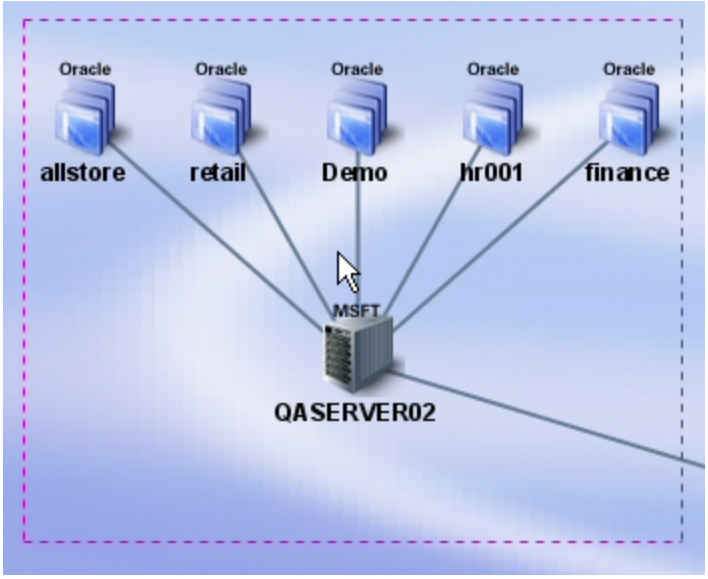

Closing Topology Windows

Whenever you select a new topology view, the software creates a pane for that view.

To lessen the number of panes open:

1. Right-click the tab of one of the views.
2. Select one of the following from the menu:
   a. **Close** – Closes the current topology pane in System Manager.
   b. **Close All** – Closes all of the topology panes in System Manager.
   c. **Close All But Current** – Closes all of the topology panes in System Manager, except the current one.

Using the Global View

If you have a large storage network, navigating it can be daunting, especially if it cannot fit in the pane. The global view ( ) provides a high-level view of the network. With this view, you can move the viewing area to a certain section of the network.

To move the viewing area:

1. Click the button at the top of the screen.

   A smaller pane displaying high-level view of the topology appears.
2. Move the brackets so they enclose the area of the network you want to view in the main pane.

Printing the Topology

To print the topology displayed in System Manager:

1. To access System Manager, click System Manager () in the left pane.
2. If the topology spans more than one screen, arrange the elements so they are closer together, preferably on one screen. This prevents the printout from appearing too stretched. To move an element, click the  button and then click the element you want to move. Drag the element to its new location.
3. Click the  button. The Paper tab shows the page setup.

   If you want the default settings, click **Default**. You can modify the following settings:

   Before you change the margins, decide on a unit of measurement.

   - **Paper format** – Select the paper size from the menu.
   - **Unit** – Select cm (centimeters) or inch for the margins.
   - **Paper width** – To modify the width of the paper, select the Custom option in the Paper format menu.
   - **Paper height –** To modify the measurement in this box, select the **Custom** option in the Paper format menu.
   - **Top margin** – Enter a measurement.
   - **Bottom margin** – Enter a measurement.
   - **Left margin** – Enter a measurement.
   - **Right margin** – Enter a measurement.
   - **Orientation** – Click an orientation for the printout.

   A preview of the printout is displayed in the right pane.
4. When you are done, click **Apply**.
5. To see how the printout will appear on the page, click the **View Selection** tab. If you want the default settings, click **Default**. You can modify the following settings:

   - **Start x** – Determines the horizontal placement of the printout on the page, with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. You can also enter negative numbers. Anything more than zero expands the printout to another page.
   - **Start y** – Determines the vertical placement of the printout on the page, with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom. You can also enter negative numbers.

- **Width** – Determines the width of the printout.
- **Height** – Determines the height of the printout.

To remove extra space around the topology, click **Trimmed**.

A preview of the printout is displayed in the right pane.

6. When you are done, click **Apply**.
7. The Pages tab shows how many pages the printout will use. If you want the default settings, click **Default**. You can modify the following settings:

   Before you change the margins, decide on a unit of measurement.

   - **Unit** – Select cm (centimeters) or inch for the margins.
   - **Position/Size** – Enables you to change the position and size of the printout so that it spans several pages:
     - **Start x** – Same as in step 6.
     - **Start y** – Same as in step 6.
     - **Width** – Determines the width of the printout. If the width entered does not fit on the page, the printout wraps around to another page.
     - **Height** – Determines the height of the printout. If the height entered does not fit on the page, the printout wraps around to another page.
   - **Resolution (pixel/unit)** – Enables you to change the resolution so that the printout spans several pages.
   - **Page** – Enables you to expand the printout so it prints on several pages without modifying the graphic.

   A preview of the printout is displayed in the right pane.
8. When you are done, click **Apply**.
9. To preview your pages, click the **Preview** tab. Click the page you want to preview. The page appears in the right pane.
10. When you are ready to print, click **Print**.
11. Click **Close**.

To return to all of the original settings, click the **Default** button next to the **Print** button.

Exporting the Topology to Microsoft Visio

To export the topology to an XML file that can be viewed in Microsoft Visio:

1. To access System Manager, click **System Manager** () in the left pane.
2. Click **Export to Visio**.

Before exporting to Visio, make sure the fabric names are unique in the topology.

3. Name the file, and select the directory where you want the file to be saved.
4. Click **Save**.

Backup topology can be exported only through Backup Manager, not through System Manager.

Viewing the Topology in Microsoft Visio

Note: Use Visio 2007 to import large topology XML files. Importing large topology XML files (for example file size greater than 13 MB) to Visio 2003 fails due to its limitation of zooming in the diagram drawn.

Storage Planner (Export to Visio) does not export details for all configurations shown by the management server:

- Blades are displayed as regular hosts
- VMware configurations are not supported with Storage Planner
- Virtual fabrics are not supported with Storage Planner
- Configurations with Solaris Containers are not supported with Storage Planner
- Properties are not shown for Clusters within Visio

After you export the topology to an XML file, follow these steps to view it in Microsoft Visio:

1. Install the necessary components, as described in Installing Storage Planner below.
2. Configure Visio, as described in Configuring Visio to View Exported Topology on next page.
3. Start Visio, and select **Storage Planner** > **Import XML File**.
   Note: Organizational Chart Wizard add-on must be enabled before importing the XML file. To enable, Configuring Visio to View Exported Topology on next page.
4. Browse to the XML file that you would like to view, and click **Open**. The selected topology is displayed in Visio.
5. Right-click any element, and select **Properties**. The Custom Properties window opens and displays additional information about each element.
6. Each fabric is displayed in a separate layer. Select **View** > **Layer Properties** to display the Layer Properties window. This allows you to customize your view of the various fabrics. For example, you can change the color or visibility settings for each fabric.

Installing Storage Planner

Microsoft Visio must be installed before installing Storage Planner.
Also install the recommended XMLLite patch for your version of Windows operating system.

To install Storage Planner:

1. From the UtilitiesCD/Windows directory on the StorageEssentialsDVD, run StoragePlanner.exe. The Welcome to the Storage Planner Setup Wizard window opens.
2. Click **Next**. The Select Destination Location window opens.
3. Click **Next**. The Select Components window opens.
4. Click **Next**.
5. Click **Install**. The Storage Planner component is installed.
6. If you did not install Microsoft XML 6.0 Parser, select the check box.
7. Click **Finish**.

   If you are installing Microsoft XML 6.0 Parser, the MSXML 6.0 Parser Setup wizard is displayed. If MSXML 6.0 parser was installed previously, you will not see the following steps.
8. Select the **I Accept** option button, and click **Next**.
9. Enter your name and company information. Click **Next**.
10. Click **Install**. Microsoft XML 6.0 Parser is installed.
11. Click **Finish**.

Configuring Visio to View Exported Topology

Before you can view the exported topology, follow these steps to correctly configure Visio 2003 and Visio 2007:

1. Open Microsoft Visio.
2. Select **Tools** > **Options**.

   If you are configuring Visio 2007, go to Step 7.
3. Click the **Security** tab.
4. Click **Macro Security**.
5. Click the **Medium** radio button.
6. Click **OK**. You are returned to the previous window.
7. Click the **Advanced** tab.
8. Click **File Paths**.
9. Click the button to the right of the Add-ons box. Browse to your Visio installation directory, and select the following path:

   ```
   <Visio_installation_directory>\1033\Solutions\StoragePlanner
   ```

   In this instance, <Visio_installation_directory> is the installation directory for Visio; for example:

   ```
   For Visio 2003 C:\Program Files\Microsoft Office\Visio11
   ```

10. Click the button to the right of the Start-up box. Browse to your Visio installation directory, and select the following path:

    `<Visio_installation_directory>\1033\Solutions`

    In this instance, <Visio_installation_directory> is the installation directory for Visio

11. Click **OK**. You are returned to the previous window.
12. Click **OK**.
13. Restart Visio, and select **Enable Macros** when prompted.

If the Storage Planner menu item does not appear when you restart Visio, go to **Tools** > **Add-Ons** > **Storage Planner** and select **Storage Planner**.

You must enable the Organizational Chart Wizard add-on every time before importing a XML file in Visio 2007. By default, it is disabled. To enable the Organizational Chart Wizard add-on in Visio 2007, follow these steps:

1. Select **Tools** > **Add-Ons**> **Run Add-On**.
2. Select **Organizational Chart Wizard** from the list of displayed add-ons.
3. Click **OK**.
4. A pop-up window appears on the screen. Click **Cancel**to continue.

Updating Element Data

System Manager enables you to update data about elements directly from this screen. When you update element data, the management server updates infrastructure details from the element and then redraws the topology with the updated information.

Keep in mind the following:

- Updating element data does not detect element components that were removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you right-click the storage system and select **Update Element Data**, the LUNs still appear in the user interface. For the deleted LUNs to be removed from the user interface, you must perform Get Details. See Get Details on page 151.
- Update Element Data does not correctly update switch ISL and topology changes. To obtain switch ISL and topology changes, you must perform Get Details.
- Update Element Data does not work for virtual machines or virtual servers.

To update an element:

1. To access System Manager, click **System Manager** () in the left pane.
2. Right-click the element you want to update.
3. Select **Update Element Data** from the menu. The software begins to update its database with the updated infrastructure details from the element.

During this process, the status button appears red, and “Getting Details” appears next to it. When the process is complete, the status button returns to green.

Viewing Ports

When you are looking at an element on the network, such as a switch, it can be difficult to determine how the ports are used. System Manager provides a view that enables you to determine the use of each port for all elements in the network.

To view the ports:

1. Do one of the following:

   Access System Manager – by clicking System Manager () in the left pane.

   Or

   Access the Topology page (for an element) in one of the following ways:

2. Double-click an element in System Manager, and then click the **Topology** tab.

   Or

   Right-click an element, and then select **Show Element Topology** from the menu.

3. Right-click an element in the topology.
4. From the menu, select **Show Port Details**.

Showing the Impact of an Element

You can display an element's impact. Assume you want to replace a switch. You can use this feature to determine which elements in the network would be impacted by taking the switch off the network.

To find an element's impact:

1. Right-click the element from which you want to obtain impact information.
2. Select **Show Impact** (or **Show Cluster Impact** if you right-clicked a cluster) from the menu.

The other elements in the path of the element you right-clicked are highlighted in yellow. Assume you right-clicked the Oracle instance klu2e (shown in the following figure) and selected **Show Impact** from the menu. The elements on which klu2e is dependent are highlighted in yellow. This means that if any of these highlighted elements are removed from the network, klu2e might have difficulty functioning.

Showing the Impact of an Element

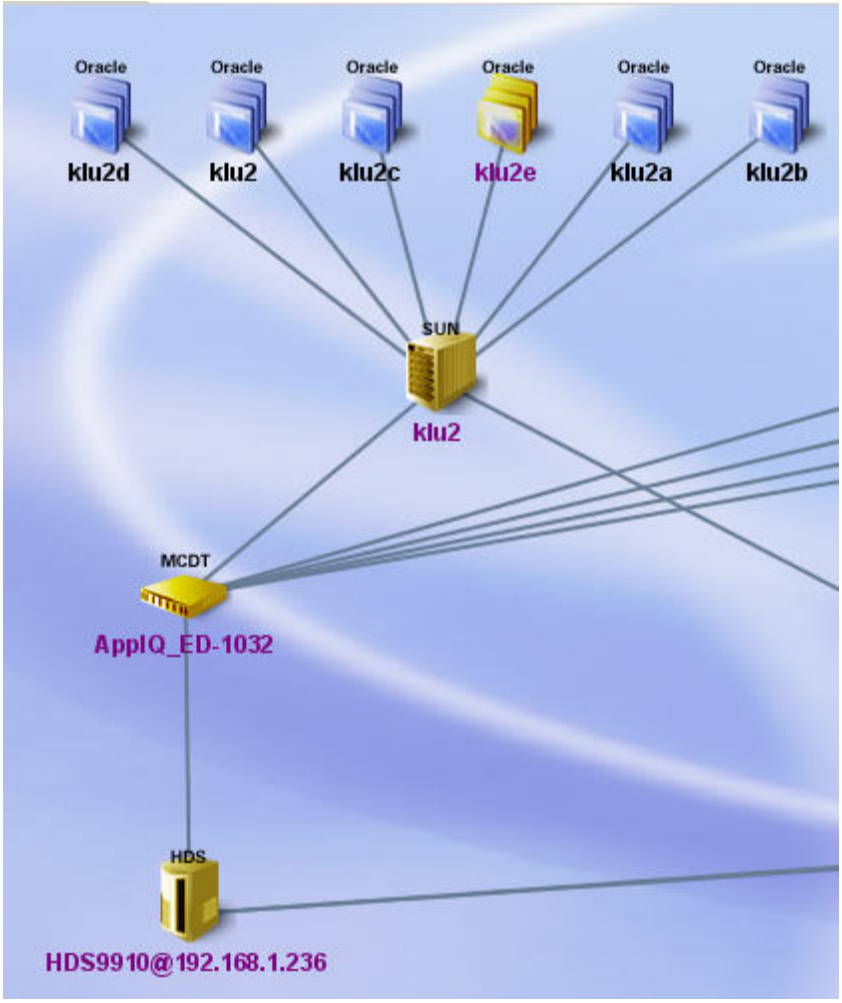

The Show Impact feature not only displays the elements on which an element is dependent, it also displays the other elements dependent on it. Assume you right-clicked a switch and selected Show Impact from the menu. Each highlighted element would include its dependent elements, such as the hosts, applications, and storage systems connected to it. These elements might have difficulty communicating with one another if the switch were removed.

Likewise, if you decided to show the impact of a host, each highlighted element would not only include its dependent elements, such as its applications, but also the elements on which it is dependent, such as switches.

Use the following considerations to help you determine whether the highlighted elements are dependent or required.

If you select Show Impact for...	The software highlights...
An Application (virtual or real)	Elements required by the application, such as its host and a switch.

If you select Show Impact for...	The software highlights...
A Host	Elements that are dependent on the host, such as its applications. Elements that are required by the host, such as switches.
A Switch	Elements that are dependent on the switch, such as hosts and storage systems
A Storage System	Elements that are dependent on the storage system, such as hosts.

Assigning a Business Cost to an Application

The management server enables you to assign a business cost to an application, including virtual applications. This information is used in Event Manager for ranking events from elements. Event Manager determines the rank of an event by taking into account the business cost of the application and the severity of the event. You can sort events by rank in Event Manager by clicking the Rank column.

For example, assume you assigned a business cost of $40 to one application and a business cost of $30 to another application. If an event with the same severity level occurs from both applications, the $40 application has a higher rank because it has a higher business cost.

The events from the elements in the path of the applications also inherit a business cost from the applications that use it. For example, assume a host has an application assigned a cost of $30. The host would have a business cost of $30. If the host has two applications (both valued at $30), the host would be valued at $60, because the two applications are using it. Likewise, the switch connected to the host would also have a value of $60, because the two $30 applications use it. If a switch has a $40 application on one host and two $30 applications on another host, the switch has a value of $100.

The cost of the storage system is determined by the applications in its path. Two storage systems connected to a switch can have different business costs, based on the applications in their path. For example, a storage system has a value of $60 if two $30 applications are in its path, as shown in the following figure.

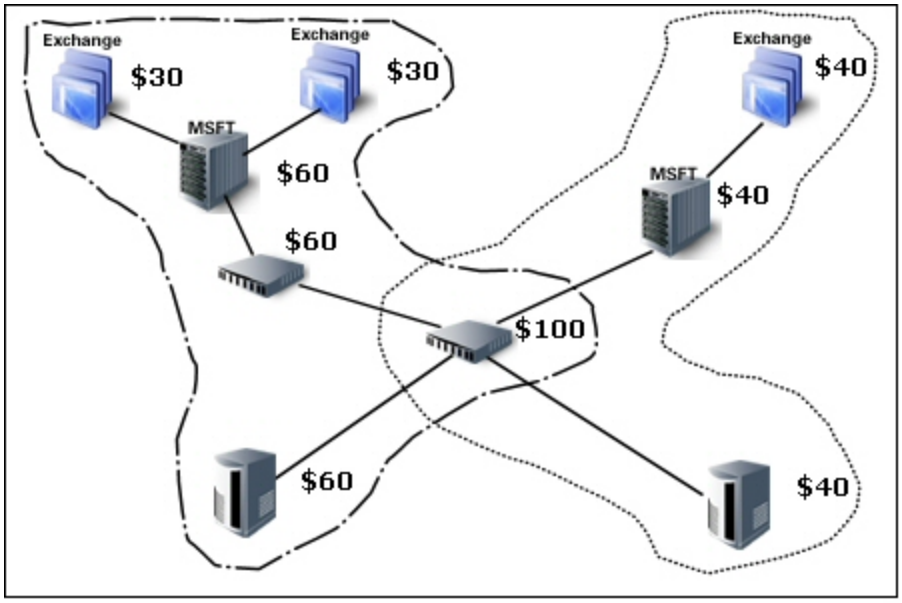

To assign a business cost to an application:

1. Do one of the following:
    a. Right-click an application in System Manager, and then select **Set Business Cost** from the menu.

    Or

    b. Double-click an application in System Manager. Click the **Properties** tab, and then click the **Change** button next to the Business Cost box.

    Or

    c. Click an application in Application Viewer. Click the **Properties** tab. and then click the **Change** button next to the Business Cost box.
2. In the Business Cost box, enter an amount; for example, 35.25.
3. Click **OK**.

Expanding the Topology Pane

To increase screen space for viewing the topology, hide the List, Access, and Path tabs by clicking the arrow pointing left on the border between the pane containing the tabs and the main pane.

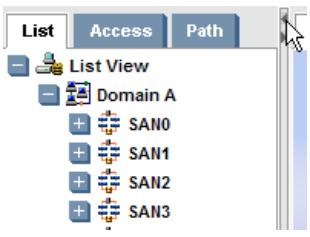

To obtain more screen space, close the left pane, as described in Opening and Closing the Left Pane on page 51.

To display these tabs, click the arrow pointing right on the border for the left pane.

Filtering SANs

To view a specified SAN in the topology

1. To access System Manager, click **System Manager** ( ) in the left pane.
2. Click the button near the top of the screen.
3. Deselect the SANs you do not want to view.

Filtering SANs

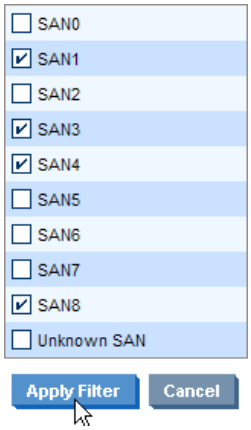

4. Click **Apply Filter**. System Manager displays the selected SANs.

Viewing Event Status in the Topology

You can obtain a status of the events occurring on the elements displayed in System Manager by clicking the button located on the toolbar. Elements with events that occurred within the last 5 minutes display next to them an icon that indicates the severity of the event. The following table shows the icons and describes their meaning.

The Event Status button ( ) is disabled in Capacity Manager and Performance Manager.

Severity Levels

Icon	Severity Level	Description
	The event has a critical impact.	Denotes elements that have a critical severity level. The elements might also have events of lower severity levels. Example: A Brocade switch has a failed firmware download. The failure reason code for each respective switch is displayed.

Icon	Severity Level	Description
	The event has a major impact.	Denotes elements that have a major severity level. The elements might also have events of lower severity levels. Example: one or more physical fabric objects (device port, switch, or fabric) have disappeared.
	The event has a minor impact.	Denotes elements that have a minor severity level. The elements might also have events of lower severity levels. Example: A physical fabric object (switch port or fabric) has changed state.
	The event is providing a warning.	Denotes elements that have a warning severity level. The elements might also have events of lower severity levels. Example: One or more new physical fabric objects (device port, switch, or fabric) have appeared.
	The event is providing information.	Denotes elements that have an informational severity level. The elements might also have events of lower severity levels. Example: A progress report event for a firmware download operation is currently in progress.
	The severity of the event is not known.	Denotes elements that have an unknown severity level. It displays icons for the following severity levels: • Critical • Major • Minor • Information

The icon corresponding to the highest severity is shown. For example, the management server displays an icon for a critical event next to an element that might also have minor events.

Since the severity level for an element is set by the manufacturer, the meanings of the severity levels vary. It is best to view the description of the event.

Use the Severity menu to filter which type of events you want to view. It displays the severity icons with the selected severity level or higher. For example, you can be notified of only critical and major events by selecting **Major** from the Severity menu and clicking the  button.

If you select a severity, click the  button, and then leave System Manager, System Manager remembers your selection.

Custom Name for a Switch Truncated in the Topology

If a custom name for a switch is long, its name might appear truncated in the topology. The full name appears once the cursor is positioned over the switch with the custom name.

Managing Groups

This section contains the following topics:

About Groups

System Manager enables you to group together hosts and storage systems that were labeled Discovered, so the management server sees them as one element in the topology. The management server labels an element as “Discovered” when it has discovered the element, but it cannot obtain additional information about it. Grouped elements preserve space in the topology, since only one element is displayed to represent the group. It also provides a way to keep track of all your “Discovered” hosts and storage systems.

Grouping Discovered Hosts

If you have several discovered hosts, you might want to group them together, so the management server sees them as one element in the topology. Grouped elements preserve space in the topology, since only one Discovered host is displayed to represent the group. Grouping also provides a way to keep track of your “Discovered” hosts.

Keep in mind the following:

- A user's role must include an access level of Element Control or Full Control for hosts. See Editing Roles on page 303 for more information about the access level of a role.
- Grouped elements are still seen as Discovered, so the management server is unable to monitor or manage them.
- Do not create groups during Get Topology or Get Details. You can determine if the management server is getting either the topology or all element details, because the status button appears red during both operations.
- You can determine if a host is generic by double-clicking the host in System Manager and then clicking the Properties tab. If a host is generic, it is listed as Generic Host for its description.

To group Discovered hosts:

1. Access System Manager, as described in Accessing System Manager on page 416.
2. Right-click a Discovered host, and select **Group together with other hosts** from the menu.
3. In the Custom Name box, enter a custom name for the group.
4. In the IP Address box, enter an IP address for the group.
5. In the DNS Name box, enter the DNS name for the group.
6. In the Version box, enter a version number for the group.
7. In the Operating System box, enter the operating system for the hosts in the group.
8. Select the hosts you want to be a part of the group, and click the button with the greater than sign (>). The hosts are added to the group. You can sort the hosts by:
   - **Name**: Click the **Hosts** column heading.
   - **Port**: Click the **Ports** column heading.
   - **Connected Switches**: Click the **Connected Switches** column heading.

   An arrow appears next to the column heading that sorts the hosts. For example, if the hosts are being sorted by name, an arrow appears next to the Hosts column heading. If the arrow next to the column heading is pointing up, the hosts are sorted in ascending numerical and alphabetical order. If the arrow is pointing down, the hosts are sorted in descending numerical and alphabetical order. Click the column heading to change the direction of the arrow.
9. To remove hosts from the group, click the button with the less than sign (<).
10. Click **OK**.

The management server no longer displays the grouped elements in the topology individually. A host icon with the group name on the bottom represents the group. The group cannot be monitored or managed.

Ungrouping Discovered Hosts

If one of the hosts in a group is going to change, you might want to ungroup Discovered hosts. An example of such a change would be when a host will be taken off line.

Keep in mind the following:

- Do not ungroup elements during Get Topology or Get Details. You can determine if the management server is getting either the topology or all element details, because the status button appears red during both operations.
- A user's role must include an access level of Element Control or Full Control for hosts. For more information about the access level of a role, see Editing Roles on page 303.

To ungroup multiple elements:

1. Access System Manager, as described in Accessing System Manager on page 416.
2. Right-click the host icon for a group, and select **Ungroup into multiple hosts** from the menu.
3. When you are asked if you want to ungroup the elements, click **OK**. The elements are ungrouped.

Grouping Discovered Storage Systems

If you have several Discovered storage systems, you might want to group them together so the management server sees them as one element in the topology. The management server labels a storage system as Discovered when it has found the storage system, but it cannot obtain additional information about it. Grouping elements preserves space in the topology, since only one Discovered element is displayed to represent the group. It also provides a way to keep track of all your Discovered storage systems.

Keep the following in mind:

- A user's role must include an access level of Element Control or Full Control for storage systems. See Editing Roles on page 303 for more information about the access level of a role.
- Grouped elements are still seen as "Discovered", so the management server is unable to monitor or manage them.
- Do not group storage systems during Get Topology or Get Details. You can determine if the management server is getting either the topology or all element details by looking at label near the status button.
- To determine if a storage system is generic, double-click the storage system in System Manager, click the Properties tab, and look at the storage system's description.

To group Discovered storage systems:

1. Access System Manager, as described in Accessing System Manager on page 416.
2. Right-click a Discovered storage system, and select **Group together with other** from the menu.
3. In the Custom Name box, enter a custom name for the group.
4. In the Vendor box, enter the vendor names for the storage systems in the group.
5. Select **Tape Library** if you want the entire group to be considered a tape library.

   This tape library will be considered as discovered. It will not be managed or monitored by the management server.
6. Select the storage systems you want to be part of the group and click the button with the greater than sign (>). The storage systems are added to the group. You can sort the storage systems by:
   - **Name**: Click the **Storage Systems** column heading.
   - **Port**: Click the **Ports** column heading.

- **Connected Switches**: Click the **Connected Switches** column heading.

An arrow appears next to the column heading that sorts the storage systems. For example, if the storage systems are being sorted by name, an arrow appears next to the Storage Systems column heading. If the arrow next to the column heading is pointing up, the storage systems are sorted in ascending numerical and alphabetical order. If the arrow is pointing down, the storage systems are sorted in descending numerical and alphabetical order. Click the column heading to change the direction of the arrow.

7. To remove storage systems from the group, click the button with the less than sign (<).
8. Click **OK**.

The management server no longer displays the grouped elements in the topology individually. A storage system icon with the group name on the bottom represents the group. The group cannot be monitored or managed.

Ungrouping Discovered Storage Systems

If one of the storage systems in a group is going to change, you might want to ungroup Discovered hosts. An example of such a change would be when a storage system will be taken off line.

Keep in mind the following:

- A user's role must include an access level of Element Control or Full Control for storage systems. For more information about the access level of a role, see Editing Roles on page 303.
- Do not ungroup elements during Get Topology or Get Details. You can determine if the management server is getting the topology or all element details by looking at label near the status button.

To ungroup multiple elements:

1. Access System Manager, as described in Accessing System Manager on page 416.
2. Right-click a storage system icon for a group, and select **Ungroup into multiple storage** from the menu.
3. When you are asked if you want to ungroup the elements, click **OK**. The elements are ungrouped.

Managing SANS and Fabrics

This section contains the following topics:

Changing the Fabric Name

To change a fabric name:

1. Access System Manager.
2. Click the **List** tab.
3. Right-click a fabric name.
4. Select **Change Fabric Name** from the menu.
5. In the Enter Fabric Name box, enter a new fabric name.
6. Click **OK**.

Changing the SAN Name

To change a SAN name:

1. Access System Manager.
2. Click the **List** tab.
3. Right-click a SAN name.
4. Select **Change SAN Name** from the menu.
5. In the Enter SAN Name box, enter a new SAN name.
6. Click **OK**.

Deleting Fabrics

When you delete a fabric, the elements in the fabric are not removed. After you delete the fabric, the management server recalculates the entire topology. The recalculation might take some time, especially if you have a large topology.

To delete a fabric:

1. Access System Manager.
2. Click the **List** tab.
3. Right-click a fabric name.
4. Select the **Delete This Fabric** option from the menu.
5. When asked if you want to delete the fabric, click **Yes**. The management server recalculates the topology. If the elements in the deleted fabric do not belong to another fabric, they are moved to the “unknown” node on the List tab.

Deleting SANs

When you delete a SAN, the elements in the SAN are not removed. After you delete the SAN, the management server recalculates the entire topology. The recalculation might take some time, especially if you have a large topology.

To delete a SAN:

1. Access System Manager.
2. Click the **List** tab.
3. Right-click a SAN name.
4. Select the **Delete This SAN** option from the menu.
5. When asked if you want to delete the SAN, click **Yes**. The management server recalculates the topology. If the elements in the deleted SAN do not belong to another SAN, they are moved to the "unknown" node on the List tab.

Hiding and Showing Generic Hosts

You can reduce the amount of time it takes to arrange your topology, by using the Hide unnamed generic hosts feature to hide hosts that a switch has detected. An element is considered to be generic if the management server can detect the element but it cannot obtain additional information about the element during Getting the Topology or Get Details.

When you use the show/hide feature, your changes persist to the next time you log into the management server. If you log in as another user, you will not see your changes. This feature allows each user to arrange the topology as he or she wishes.

The show/hide feature pertains only to unnamed generic hosts. If you name a generic host, you cannot use this feature to hide the named host. The Hide generic element feature also does not work for grouped named generic hosts and missing elements. To learn how to give a custom name to an unnamed generic host, see Assigning a Custom Name on page 500. To learn more about groups, see About Groups on page 470.

Hiding Generic Hosts

HP Storage Essentials provides two variations of the hide generic hosts feature:

- **Hiding Generic Hosts for One Switch:** This feature hides unnamed generic hosts detected by a switch. The management server detects an element by looking at the ports on a switch. If it cannot find additional information about the element, it marks it as generic by displaying a question mark above its icon.

  For example, assume you have a switch with 10 discovered elements. The management server detected these elements by looking at the ports on the switch and determined the type of element connected. Discovered elements appear with a question mark above their icon in the topology. The question mark indicates that the management server has detected the element, but it cannot obtain additional information.

To learn how to use this feature, see Hiding Generic Hosts for a Switch below.

- **Hiding Generic Hosts for All Switches:** This feature hides unnamed generic hosts within a domain. To learn how to use this feature, see Hiding Generic Hosts for All Switches below.

Hiding Generic Hosts for a Switch

Simplify your topology by hiding unnamed generic hosts connected to a switch. If you have an unnamed generic host connected to more than one switch and you want to hide the generic element, you must repeat the following steps for each switch connected to the generic host. You can hide all unnamed generic hosts at once by using the Hide Generic Hosts for All Switches feature. See Hiding Generic Hosts for All Switches below for more information.

To hide generic storage devices connected to a switch:

1. Right-click the storage device.
2. Select **Discovered Element** > **Hide Generic Storage Device for the Switch** from the menu.

   A "+" icon is added to the icon of the switch you right-clicked to indicate that it has hidden generic hosts.

Hiding Generic Hosts for All Switches

To hide all unnamed generic hosts and unnamed generic Cisco switches:

1. Right-click a switch.
2. Select **Discovered Element** > **Hide Generic Hosts for All Switches** from the menu. All unnamed generic hosts are hidden. A “+” icon is added to the icon of all switches that have generic hosts that were hidden.

Showing Generic Hosts

HP Storage Essentials provides two variations of the expand generic host feature:

- **Showing generic hosts for a switch:** This feature displays previously hidden generic hosts. To learn how to use this feature, see Expanding Generic Storage Devices for a Switch on page 478
- **Showing generic storage device for all switches:** This feature displays previously unnamed generic storage devices detected by a switch. To learn how to use this feature, see Expanding Generic Hosts for All Switches on the facing page.

Expanding Generic Hosts for a Switch

To display hidden generic hosts connected to a switch:

1. Right-click a switch with a "+" icon. The "+" icon indicates that the switch has hidden generic storage devices and/or hidden generic hosts.

2. Select **Discovered Element** > **Expand Generic Hosts for the Switch** from the menu. The hidden elements for the switch appear in the upper-right corner of the topology.

Expanding Generic Hosts for All Switches

To display hidden generic hosts for a domain:

1. Right-click a switch with a "+" icon. This "+" icon indicates the switch hidden generic storage devices and/or hidden generic hosts.
2. Select **Discovered Element** > **Expand Generic Hosts for All Switches** from the menu. The hidden elements for the domain appear in upper right corner of the topology.

Hiding and Showing Generic Storage Devices

You can reduce the amount of time it takes to arrange your topology, by using the Hide unnamed generic storage devices feature to hide storage devices that a switch has detected. An element is considered to be generic if the management server can detect the element but it cannot obtain additional information about the element during Getting the Topology or Get Details.

When you use the show/hide feature, your changes persist to the next time you log into the management server. Your changes will also persist in a saved topology view. If you log in as another user, you will not see your changes. This feature allows each user to arrange the topology as he or she wishes.

The feature described in this section pertains only to unnamed generic storage devices. If you name a generic storage device, you cannot use this feature to hide the named storage device. The Hide generic element feature also does not work for grouped unnamed generic storage devices and missing elements. To learn how to give a custom name to an unnamed generic storage device, see Assigning a Custom Name on page 500. To learn more about groups, see About Groups on page 470.

Hiding Generic Storage Devices

HP Storage Essentials provides two variations of the hide generic storage device feature:

- **Hiding generic storage devices for a switch:** This feature hides unnamed generic storage devices detected by a switch. The management server detects an element by looking at the ports on a switch. If it cannot find additional information about the element, it marks it as generic by displaying a question mark above its icon.

  For example, assume you have a switch with 10 discovered elements. The management server detected these elements by looking at the ports on the switch and determined the type of element connected. Discovered elements appear with a question mark above their icon in the topology. The question mark indicates that the management server has detected the element, but it cannot obtain additional information.

- **Hiding generic storage devices for all switches:** This feature hides unnamed generic hosts within a domain. To learn how to use this feature, see Hiding Generic Hosts for All Switches.

Hiding Generic Storage Devices for a Switch

Simplify your topology by hiding unnamed generic storage devices connected to a switch. If you

have an unnamed generic storage device connected to more than one switch and you want to hide the generic element, you must repeat the following steps for each switch connected to the generic storage device. You can hide all unnamed generic storage devices at once by using the **Hide Generic Storage Devices for All Switches** feature. See Hiding Generic Hosts for All Switches for more information.

To hide generic storage devices connected to a switch:

1. Right-click the storage device.
2. Select **Discovered Element** > **Hide Generic Storage Device for the Switch** from the menu.

   A "+" icon is added to the icon of the switch you right-clicked to indicate that it has hidden generic hosts.

Hiding Generic Storage Devices for All Switches

To hide all unnamed generic storage devices:

1. Right-click a switch.
2. Select **Discovered Element** > **Hide Generic Hosts for All Switches** from the menu. All unnamed generic hosts are hidden. A “+” icon is added to the icon of all switches that have generic hosts that were hidden.

Expanding Generic Storage Devices

HP Storage Essentials provides two variations of the expand generic storage device feature:

- **Showing generic storage devices for a switch:** This feature displays previously hidden generic storage devices. To learn how to use this feature, see Expanding Generic Storage Devices for a Switch below.
- **Showing generic storage device for all switches:** This feature displays previously unnamed generic storage devices detected by a switch. To learn how to use this feature, see Expanding Generic Storage Devices for All Switches below.

Expanding Generic Storage Devices for a Switch

To display hidden generic storage devices connected to a switch:

1. Right-click a switch with a "+" icon. The icon indicates that the switch has hidden generic storage devices or hidden generic hosts.
2. Select **Discovered Element** > **Expand Generic Storage Devices for the Switch** from the menu. The hidden elements for the switch appear in the upper-right corner of the topology.

Expanding Generic Storage Devices for All Switches

To display hidden generic storage devices for a domain:

1. Right-click a switch with a "+" icon. The icon indicates that the switch has hidden generic storage devices and/or hidden generic hosts.

2. Select **Discovered Element** > **Expand Generic Storage Devices for All Switches** from the menu. The hidden elements for the domain appear in upper right corner of the topology.

Setting Up Custom Commands

This section contains the following topics:

About Custom Commands

Custom commands enable you to run a command that you create on the management server. The command could point to an executable or a script that does not use the graphical user interface. For example, if you create a script that backs up a storage system, you can run that script from System Manager.

You can also use environment variables in your scripts. For example, you can use the variables to obtain information about a host, such as its total physical memory and the number of processors.

Important Considerations

Keep in mind the following:

- Run scripts at your own risk. The management server enables you to run any script, including those that can disable the management server.
- The custom command always runs on the management server unless you are running the telnet utility. You can obtain information about an element on which you right-click by using the software's environment variables. See Software Environment Variables for Scripting on page 483.
- Custom commands only supports executables and scripts that do not use the graphical user interface.
- *Management servers on Windows only*: If you leave the remote console (cmd /k) open after running a script, users can traverse the directory structure of the management server.
- To run a Perl script as a custom command on a UNIX system, prefix the script with the Perl executable; for example, `perl myscript.pl`. In this instance, myscript.pl is the script you want to run.

A best practice is to prefix the script with the path to Perl and the Perl executable; for example, `perl/bin/perl myscript.pl`. In this instance, perl/bin/ is the directory containing the Perl executable, perl is the executable, and myscript.pl is the script you want to run.

To run a Perl script as a custom command on Microsoft Windows, prefix the script name with the complete path to Perl. The management server already has a directory containing the Perl executable inside the folder `%JBOSS4_DIST%\server\appiq\remoteScripts\perl\bin`.

Prefix the script name as follows:

```
.\perl\bin\perl myscript.pl
```

In this instance, .\perl\bin\ is the directory containing the Perl executable in the RemoteScripts directory, perl is the executable, and myscript.pl is the script you want to run.

- The Print Buffer and Standard Error (STDERR) must be flushed each time you run any Perl script through custom commands. Flushing helps Perl script to behave consistently with custom command applet.

  For flushing, add the following lines at the start of the code of every custom command:

  ```
  $|=1; #unbuffer stdout
  select STDERR; $|=1; select STDOUT; #unbuffer stderr
  ```

Adding a Custom Command

Before adding a custom command, make sure you are aware of the considerations listed in Important Considerations on previous page.

To add a custom command:

1. Right-click an element in System Manager.
2. Select **Custom Commands** > **Set Up Custom Commands** from the menu.
3. (*Optional*): If you plan to use a command to activate a file, such as a script, the file must be uploaded to the management server, as follows:
   a. In the Custom Command Setup window, click **Browse** to find the file containing the custom command.
   b. Click **Open**, and then click **Upload to server**.

      The file is saved on the management server.
4. Click the **Add Command** button in the upper-right corner of the window.
5. In the Add Custom Command window, enter a name for the command in the Name box, for example, backup command.
6. In the Description box, enter a description; for example, "This command activates a script that backs up an element."

7. In the Command Line box, enter a command. This could be a command required to start a script; for example:

   ```
   myscript.bat
   ```

   The remote console automatically becomes inactive once the command finishes.

   Windows only: If you want the remote console to stay open, prefix the command with the following:

   ```
   cmd /k
   ```

   Here is an example:

   ```
   cmd /k dir
   ```

   The file is appended to the command line.

   (Optional): If you plan to use a file in the command, select the file from the **Files** menu, and then click **Append To Command Line**.

   If the file is missing, repeat step 3.

8. Select one of the following options to determine the elements for which you want the command to be visible. For example, if you select the **All Elements** option, the command is visible in the menu when you right-click any element.

   - **Name of the Element** – Select the name of the element if you want the command to be visible in the menu only when you right-click this element.
   - **All Elements** – Select this option if you want the command to be visible in the menu when you right-click any element.
   - **Selected element types and filter criteria** – Select this option if you want to narrow the filtering criteria for an element type. For example, you could specify that the command is only in the menu when a Brocade switch is right-clicked. The options are as follows:
     - **Applications** – If you want the command to be visible in the menu when a particular application is right-clicked, enter the name of the product in the Product Name box. To make sure you enter the correct product name, enter the product name displayed in the Product Name box on the Properties tab, accessible by double-clicking the application in System Manager and then clicking the **Properties** tab.
     - **Hosts** – If you want the command to be visible in the menu when a particular host is right-clicked, enter the name of the operating system in the OS Name box. To make sure you enter the correct operating system, enter the operating system displayed in the Target Operating System box on the Properties tab, accessible by double-clicking the host in System Manager and then clicking the **Properties** tab.
     - **Switches** – If you want the command to be visible in the menu when a switch from a particular vendor is right-clicked, enter the name of the vendor in the Vendor Name box. To make sure you enter the correct vendor name, enter the vendor name

displayed on the Properties tab, accessible by double-clicking the switch in System Manager and then clicking the **Properties** tab.

- **Storage Systems** – If you want the command to be visible in the menu when a storage system from a particular vendor is right-clicked, enter the name of the vendor in the Vendor Name box. To make sure you enter the correct vendor name, enter the vendor name displayed on the Properties tab, accessible by double-clicking the storage system in System Manager and then clicking the Properties tab.

9. Click **OK**.

To run a command:

1. Right-click an element.
2. Select **Custom Commands** from the menu.
3. Select the command from the Custom Commands menu. A remote console displays the result of the command.

You can stop a command by clicking the **Stop** button in the remote console. Once a command is executed, the console becomes inactive. The software assumes you are in the `%MGR_DIST%\JBossandJetty\server\appiq\remotescripts` directory on the management server when the script is executed.

Editing a Custom Command

To edit a custom command:

1. Right-click an element in System Manager.
2. Select **Custom Commands** > **Set Up Custom Commands** from the menu.
3. Click the  button corresponding to the custom command you want to edit.
4. Make the appropriate changes in the Edit Custom Command window.
5. Click **OK**. The custom command is modified.

Deleting a Custom Command

To delete a custom command:

1. Right-click an element in System Manager.
2. Select **Custom Commands** > **Set Up Custom Commands** from the menu.
3. Click the  button corresponding to the custom command you want to delete. The custom command is deleted.

Software Environment Variables for Scripting

The software provides environment variables for you to put in your scripts. For example, assume you have a script that backs up a host. You could use variables to obtain information about the host.

The software gathers information about the element you right-click. For example, if you use the variable APPIQ_ELEMENT_ID, the management server obtains the element ID of the element you right-click.

The following variables can be used to gather information for all elements. If an application resides on the host, these variables provide information about the application. For variables that return information about the host, see Variables for Applications Only on page 485.

Variables for All Elements

Variable	Value
APPIQ_ELEMENT_ID	Identifier of an element
APPIQ_ELEMENT_NAME	Name of the element
APPIQ_ELEMENT_STATUS	The following statuses are available: • Managed • Generic • Missing • Virtual Application • Asset
APPIQ_ELEMENT_DESCRIPTION	Description for the element
>APPIQ_ELEMENT_MANAGEMENT_IP_ADDRESS	IP address of the first access point used to discover the element
APPIQ_ELEMENT_VENDOR	Vendor for the element
APPIQ_ELEMENT_TYPE_NAME	Type of element; for example, switch, application, or host

The following variables can be used to gather information for storage systems, switches, and hosts only. If an application resides on the host, the variables in this table provide information about the application. For variables that return information about the host, see Variables for Applications Only on page 485.

Variables for Storage Systems, Switches, and Hosts Only

Variable	Value
APPIQ_ELEMENT_IP_ADDRESS	IP address of the element
APPIQ_ELEMENT_DNS_NAME	DNS name of the element
APPIQ_ELEMENT_MODEL	Model of the element
APPIQ_ELEMENT_VERSION	Version of the element

The following variables can be used to gather information for switches only.

Variables for Switches Only

Variable	Value
APPIQ_ELEMENT_SWITCH_ID	Identifier for the switch
APPIQ_ELEMENT_IP_GATEWAY	IP gateway of the switch
APPIQ_ELEMENT_IP_NETWORK_MASK	IP network mask for the switch
APPIQ_ELEMENT_SWITCH_STATUS	Status of the switch
APPIQ_ELEMENT_DOMAIN_ID	Domain identifier of the switch

The following variables can be used to gather information for hosts only. If an application resides on the host, the variables in this table provide information about the application. For variables that return information about the host, see Variables for Applications Only on the facing page.

Variables for Hosts Only

Variable	Value
APPIQ_ELEMENT_OPERATING_SYSTEM	Operating system of the host
APPIQ_ELEMENT_NUMBER_OF_ PROCESSORS	Number of processors used by the host
APPIQ_ELEMENT_TOTAL_PHYSICAL_ MEMORY	Total physical memory of the host
APPIQ_ELEMENT_DOMAIN	Domain of the host

The following variables can be used to gather information for applications. Use the variables with the “APPIQ_HOST” prefix when you are using variables from the first table to gather information about the application. For example, if you are running a script containing APPIQ_ELEMENT_STATUS on a host, it would obtain information about the status of the application. You would need to run APPIQ_HOST_STATUS to obtain information about the status of the host on which the application resides.

Variables for Applications Only

Variable	Value
APPIQ_ELEMENT_PRODUCT_NAME	Name of the application
APPIQ_HOST_NAME	Name of the host on which the application resides
APPIQ_HOST_ID	Identifier of a host on which the application resides
APPIQ_HOST_STATUS	Status of the host on which the application resides
APPIQ_HOST_DESCRIPTION	Description of the host on which the application resides
APPIQ_HOST_VENDOR	Vendor of the host on which the application resides
APPIQ_HOST_TYPE_NAME	Type name of the host on which the application resides
APPIQ_HOST_IP_ADDRESS	IP address of the host on which the application resides
APPIQ_HOST_DNS_NAME	DNS name of the host on which the application resides
APPIQ_HOST_MODEL	Model of the host on which the application resides
APPIQ_HOST_VERSION	Version of the host on which the application resides
APPIQ_HOST_OPERATING_SYSTEM	Operating system of the host on which the application resides
APPIQ_HOST_NUMBER_OF_PROCESSORS	Number of processors on the host on which the application resides
APPIQ_HOST_TOTAL_PHYSICAL_MEMORY	Total physical memory of the host on which the application resides
APPIQ_HOST_DOMAIN	Domain of the host on which the application resides

Using the Remote Console

This section contains the following topics:

About the Remote Console

Whenever you run a custom command on the management server, the remote console appears. The remote console displays the result of a custom command. For example, you can use the remote console to start a remote command prompt on the management server.

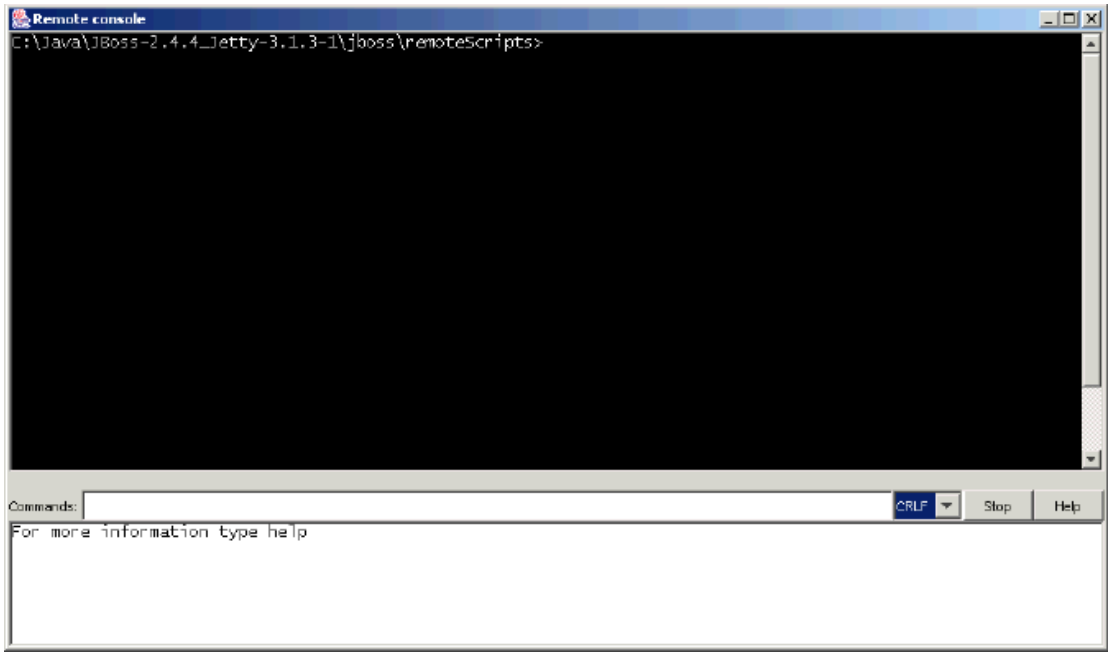

Keeping the Remote Console Active

Note: This section is for management servers running on Microsoft Windows only.

The remote console become inactive when the custom command finishes its execution. To use the menus and buttons in the remote console, the remote console must be kept active.

If you leave the remote console (cmd /k) open after running a script, users can traverse the directory structure of the management server.

To keep the remote console window active, follow these steps to create a remote command prompt:

1. Right-click an element in System Manager.
2. Select **Custom Commands** > **Set Up Custom Commands** from the menu.
3. Click the **Add Command** button in the upper-right corner of the window.
4. In the Add Custom Command window, enter a name for the command in the Name box, for example, prompt.
5. In the Description box, enter a description; for example, “Accesses the remote console.”
6. In the Command Line box, enter the following command, which will run on the management server:

   ```
   cmd /k
   ```

7. Select the **All elements** option.
8. Click **OK**.

To run the remote command prompt:

1. Right-click an element from which you want to obtain information.
2. Select **Custom Commands** and select the command from the menu. The software displays the remote console on the management server.

To enter a command in the remote console:

1. Enter the command in the Commands box.
2. Press **Enter**.

To stop a command, click the **Stop** button in the remote console.

You can quickly access information about the element you right-clicked by typing the following at the command prompt:

```
set appiq
```

The software ships with a utility called plink. To view the commands for plink, enter **plink** in the Commands box and press **Enter**.

Buttons on the Remote Console

The remote console provides the Stop and Help buttons, as described in Buttons on the Remote Console on next page.

Buttons on the Remote Console

Button	Description
Stop	Stops a command. Once a command executes, the console becomes inactive.
Help	Provides the following information about the remote console: • **Clearing the remote console** – Enter `CLS` in the Commands box of the remote console. • **Copying text to the Commands box** – Place the cursor at the end of the line in the window below the Commands box, and then press ENTER. The command is copied to the Commands box. If you are viewing the remote console on Microsoft Windows, you can copy text by using CTRL + C and CTRL + P to paste it.

Menu Options

The remote console provides the following menu options:

Option	Description
CRLF	(Default setting) Provides a carriage return and a line feed. Do not use this option when using Telnet to access another computer. You must select the CR option after you enter a user name. To enter a password, if you leave the setting at CRLF, the software enters a carriage return and a line feed when you click **OK**. No value is entered for the password.
CR	Provides a carriage return.
LF	Provides a line feed.

Copying Text from the Remote Console

To copy text from the remote console:

1. Select the text in the remote console.
2. Right-click the top frame in the remote console.
3. Select **Copy** from the menu. The text is stored in the buffer of your computer to be pasted elsewhere.

Using External Tools

This section contains the following topics:

External Tools Feature

The management server ships with an external tools feature that enables you to:

- **Browse the element** – Access a host or a switch through its main Web page. The software assumes the host or switch has a Web page at http://hsIPaddress, where hsIPaddress is the IP address of the host or switch. To access the main Web page of the host or switch, right-click the element in System Manager and select **External Tools** > **Browse to 192.168.1.2**. In this instance, 192.168.1.2 is the IP address of the host or switch.
- **Telnet to the element** – Access a host or a switch through the telnet utility. Telnet must be already enabled on the element. The command uses telnet://hsIPaddress, where hsIPaddress is the IP address of the host or switch. To telnet to a host or switch, right-click the element in System Manager and select **External Tools** > **Telnet to 192.168.1.2**. In this instance, 192.168.1.2 is the IP address of the host or switch.
- **Set up external tools** – Enables you to add a URL for accessing management software, such as Hitachi HiCommand Device Manager and EMC ControlCenter Navisphere. See Setting up External Tools below for more information.
- **Access the management tool for the storage system** – In some instances, the management tool for the storage system is accessible from this menu. For example, HiCommand for HDS storage systems and Command View for HP XP storage systems are accessible from the External Tools menu.

Setting up External Tools

You can add URLs for accessing external tools used for managing an element, such as Hitachi HiCommand Device Manager and EMC ControlCenter Navisphere for sentertorage systems.

When you add a URL, it applies only to the element you originally right-clicked.

To add a URL for accessing external tools:

1. Access System Manager.
2. Right-click the element, and select **External Tools** > **Set Up External Tools**.
3. Click **Add New Management URL**.
4. In the Description box, enter the name of the product you plan to access.
5. In the URL box, enter the URL that is used to access the product.
6. Click **OK**.

When you right-click the element and select **External Tools**, the external tool is listed.

To delete the URL for an external tool, click the corresponding button in the External Tools window.

About the Navigation Tab

To access the Navigation tab, see Accessing the Navigation Tab on page 497.

The Navigation tab provides information about an element and illustrates how the element relates to other elements in its path. For example, the Navigation page displays logical and physical components, such as ports, zone sets, zones and zone aliases. It also displays the dependencies for switches, as shown in the following figure.

Obtaining Information About a Host

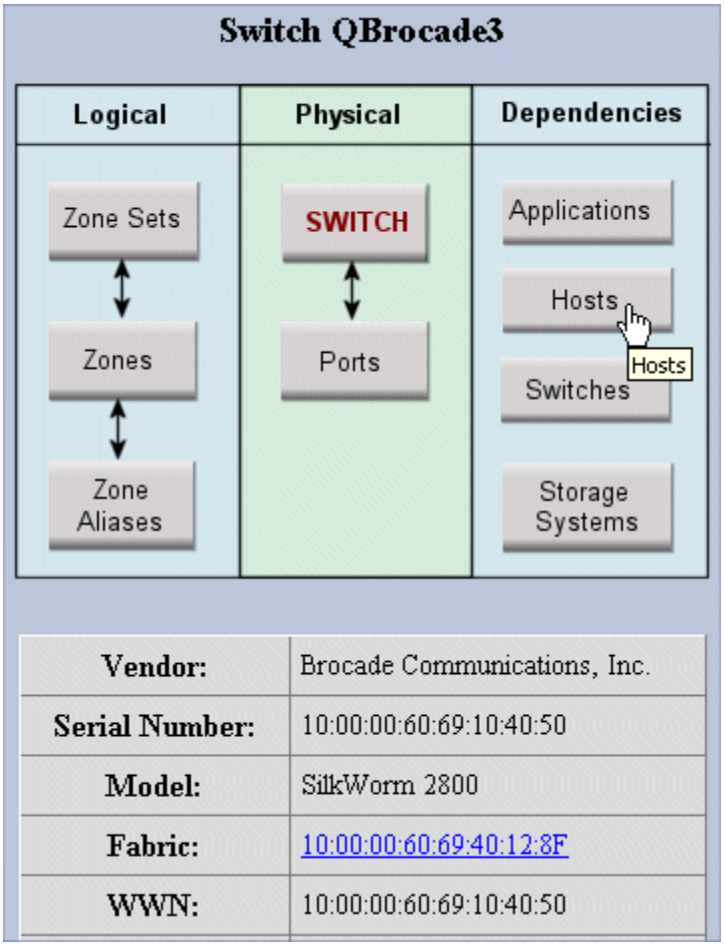

If you see a message that zone aliases are not supported on a Brocade switch, perform Get Details. The management server does not gather provisioning information from a fabric until Get Details is performed.

You can learn more about a component by clicking it in the Navigation page. For example, to learn which hosts are dependent, click **Hosts** (as shown in the previous figure), and the following information is displayed:

Details of a Host Connected to a Switch

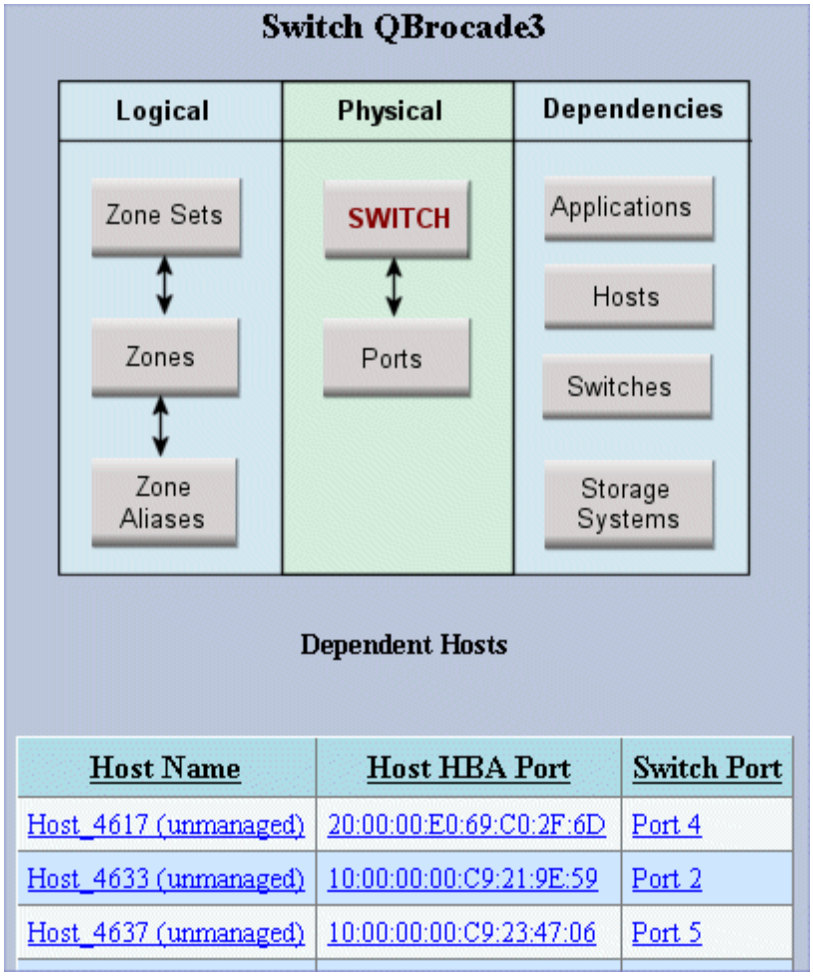

The following information is provided for each type of element:

Element	Dependencies	Front End Dependencies	Back End Dependencies	Front Physical
Applications	✓			
Fabrics	✓			
Hosts*	✓			
Switches	✓			
Storage Systems (non-virtualized)	✓			✓

Element	Dependencies	Front End Dependencies	Back End Dependencies	Front Physical
Storage Systems (virtualized)		✓	✓	

Element	Back Physical	Logical	Physical	Fabric
Applications				
Fabrics		✓		✓
Hosts*				
Switches		✓	✓	
Storage Systems (non-virtualized)	✓	✓		
Storage Systems (virtualized)				

*The management server displays **cxfs** for SGI IRIX computers if it detects CXFS on the cluster. On individual IRIX computers `cxfs` is not displayed when you enter the following at the command prompt:

```
df -k
```

Switches with Multiple Fabrics

If a physical switch has multiple fabrics, a Switch icon and a Fabric icon are displayed in the Logical column. These represent virtual switches and fabrics. When you click the **Switches** icon, the virtual switch names and the fabrics that they belong to are displayed. When you click the **Fabrics** icon, the virtual fabric names and the active zone sets that they belong to are displayed.

The navigation page for a switch with multiple fabrics looks like the following.

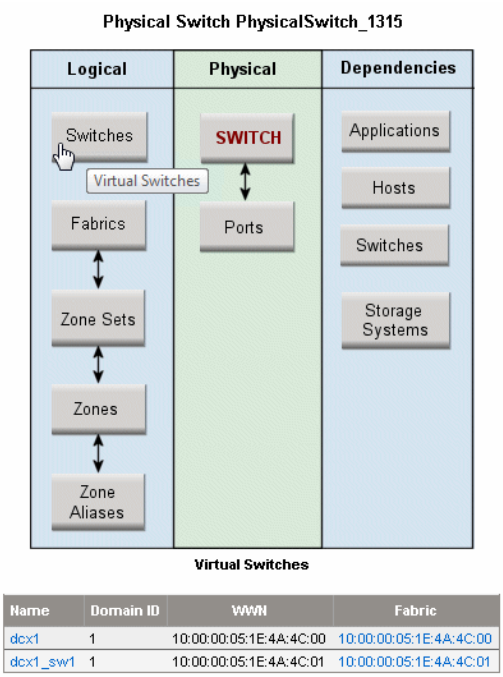

When you click **Zone Sets**, **Zones**, or **Zone Aliases** for a switch that has multiple fabrics, you can filter the list by selecting a fabric from the Fabric drop-down list:

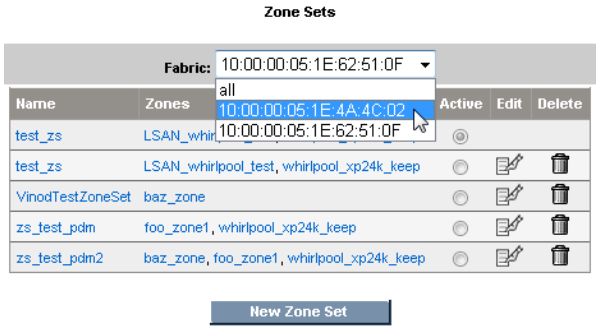

Finding the Status of a Port on a Switch

The management server can detect the status of a switch. This can be especially useful if a port is being a problem.

To find the status of a port on a switch:

1. Access System Manager as described in Accessing System Manager on page 416.
2. Double-click a switch in the right pane.
3. Click the **Ports** button in the Physical column in the Navigation tab.

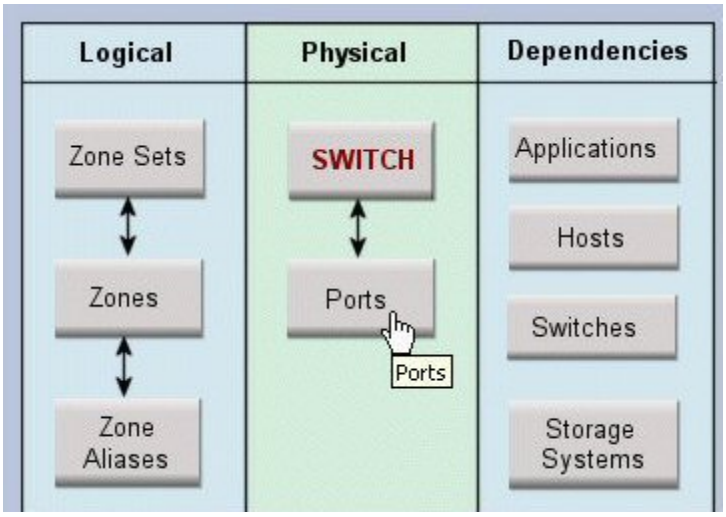

4. Under the Name column in the Ports table, click the port whose status you want to obtain.

   On the Properties page, the status of the port is displayed in the right column. The status of the port can be online, off line, or unknown.

Port Status Definitions

Status	Definition
Online	Port is physically installed with node connections
Offline	Port is physically installed, but without node connections. Brocade switches also display this status if the port is not physically installed (Gigabit Interface Converter (GIBIC) is not installed).
Not Installed (McDATA SWAPI connections only)*	Port is not physically installed (Gigabit Interface Converter (GIBIC) is not installed).
Unknown (McDATA SNMP connections only)*	Port is not installed.

*An ES4500 switch displays its status differently when a port is not installed:

- **SWAPI connection** – Unknown status
- **SNMP connection** – Offline

Finding Information on Replication Pairs

HP Storage Essentials collects information about replication pairs and makes it available from the storage device navigation screen via the Replication Pairs button, highlighted below in yellow.

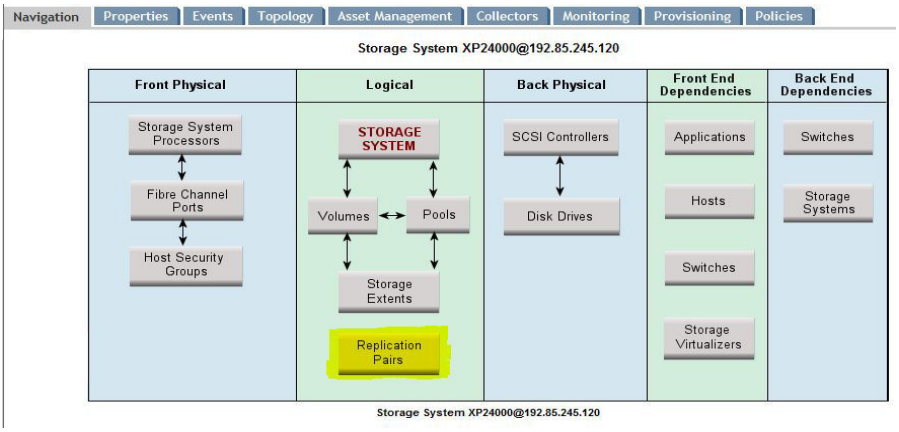

Clicking the Replication Pairs button produces the following table, which shows the replication pairs that the device is involved in.

Replication Pairs

Showing 1-10 out of 10 Total (0 Selected)

Display: 10 rows

Source	Target	Copy Type	Replica Type	When Synced	Sync State	Sync Maintained	Locality	Remote System Id	Sync State Collection Time
LDEV 00:50:03 X	LDEV 00:00:D5	Sync	Full Copy		Idle	false	Local Pair		2010-05-19 16:34
LDEV 00:60:00	LDEV 00:70:00	Async	Full Copy		Idle	false	Local Pair		2010-05-19 16:34
LDEV 00:70:01	LDEV 00:70:05	Sync	Full Copy		Idle	false	Local Pair		2010-05-19 16:34
LDEV 00:70:00	LDEV 00:70:04	Sync	Full Copy		Idle	false	Local Pair		2010-05-19 16:34
LDEV 00:70:00	LDEV 00:70:30 V	UnSyncAssoc	After Delta	2010-04-17 06:37	Idle	false	Local Pair		2010-05-19 16:34
LDEV 00:70:00	LDEV 00:70:34 V	UnSyncAssoc	After Delta	2010-04-17 06:38	Idle	false	Local Pair		2010-05-19 16:34
LDEV 00:70:00	LDEV 00:70:32 V	UnSyncAssoc	After Delta	2010-04-17 06:37	Idle	false	Local Pair		2010-05-19 16:34
LDEV 00:00:96	LDEV 00:01:1C	Sync	Full Copy	2010-02-08 13:19	Synchronized	true	Local Pair		2010-05-19 16:34
LDEV 00:70:00	LDEV 00:70:31 V	UnSyncAssoc	After Delta	2010-04-17 06:37	Idle	false	Local Pair		2010-05-19 16:34
LDEV 00:70:00	LDEV 00:70:33 V	UnSyncAssoc	After Delta	2010-04-17 06:37	Idle	false	Local Pair		2010-05-19 16:34

The following attributes are displayed for each replication pair.

Attribute	Definition
Source	The source storage volume for the replication pair
Target	The target storage volume for the replication pair
Copy Type	SMI-S term used to describe the Replication Policy. Values include: Async: Create and maintain an asynchronous copy of the source Sync: Create and maintain a synchronized copy of the source UnSyncAssoc: Create an unsynchronized copy and maintain an association to the source

Attribute	Definition
Replica Type	SMI-S term that provides information about how the Replica is being maintained. Values include: Full Copy: A full copy of the source object is (or will be) generated. Before Delta: The source object will be maintained as a delta data from the replica. After Delta: The replica will be maintained as delta data from the source object. Log: The replica object is being maintained as a log of changes to the source. Not Specified: The method of maintaining the copy is not specified.
When Synced	Date on which the replication pair was last synchronized. Not all devices report this value.
Sync State	State of the synchronized pair
Sync Maintained	Whether or not the synchronization is maintained
Locality	Whether the replication pair spans two devices and, if it does, whether the target or source is on this device
Remote system ID	If the pair spans several devices, the id of the remote device. This can be useful if HP Storage Essentials has not yet discovered the other device.
Sync state collection time	The last time the sync state field was updated

With the exception of sync state, each of these attributes is updated whenever a Get Details is triggered within HP Storage Essentials. Sync state is updated whenever a Get Details is triggered and whenever the replication collector is scheduled to run. This allows the sync state to be updated frequently without incurring the cost of collecting all data from a device.

HP Storage Essentials adds replication information for NetApp, EMC, HP XP, and EVA storage arrays (see Device-Specific Replication Information on page 163).

Storage Volumes on Virtual Machines

When you view the storage volumes associated with a virtual machine, the volumes are marked as local even though the virtual server's volume store data is located on external datastores. The virtual machines use a SCSI adapter for all of their I/O commands, and the SCSI adapter does not care where the datastore is located. The use of the SCSI adapter tells the management server that the volume is local.

Virtual Pools on Centera Devices

The following are the capacity values for virtual pools in the Navigation page.

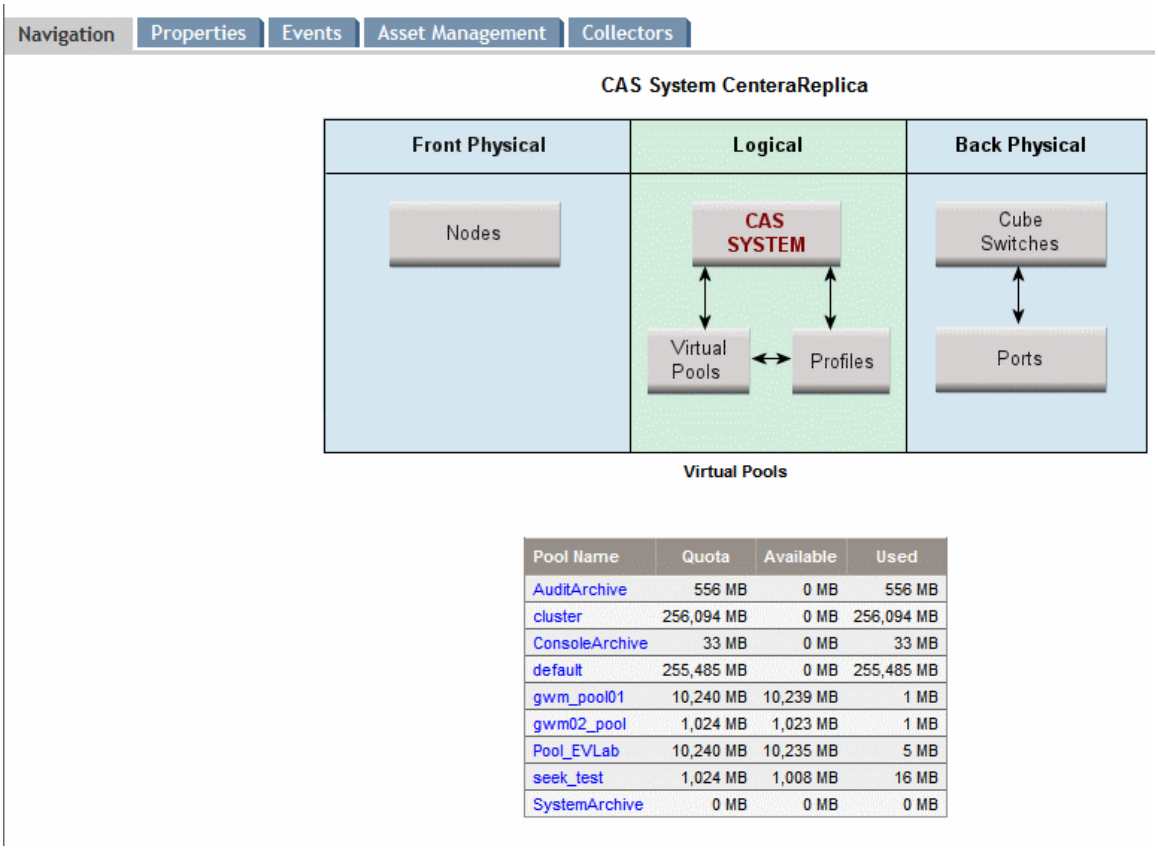

Pool Name	Quota	Available	Used
AuditArchive	556 MB	0 MB	556 MB
cluster	256,094 MB	0 MB	256,094 MB
ConsoleArchive	33 MB	0 MB	33 MB
default	255,485 MB	0 MB	255,485 MB
gwm_pool01	10,240 MB	10,239 MB	1 MB
gwm02_pool	1,024 MB	1,023 MB	1 MB
Pool_EVLab	10,240 MB	10,235 MB	5 MB
seek_test	1,024 MB	1,008 MB	16 MB
SystemArchive	0 MB	0 MB	0 MB

The Quota column shows the Total Capacity of the Centera virtual pools.The Available Capacity is calculated as the difference between the Total Capacity and the Used Capacity. However, the Total and Used values calculated here do not represent the actual capacity values returned by the Centera devices.

Accessing the Navigation Tab

To access the Navigation tab:

1. Access the management server.
2. Access the Navigation tab by doing one of the following:
   - Click an element in Application Viewer.
   - Double-click an element in Capacity Manager, Performance Manager, or System Manager.
   - Click one of the following elements in Backup Manager, and then click **Navigation** in the

lower-right corner.

- Backup Client
- Backup Library
- Host
- Master backup server
- Master backup media

3. Click the **Navigation** tab. (This is not necessary if you accessed the Navigation tab from Backup Manager.)

Viewing Element Properties

This section contains the following topics:

About the Properties Tab

The Properties tab provides detailed information about an element. Since the information obtained from each type of element varies, the Properties tab displays only information relevant to that type of element. For example, the Properties tab for fabrics lists the zones, zone sets, switches, and zone aliases, as compared to the Properties tab for a host, which lists the processors, cards, applications, and storage volumes the host uses. For supported Brocade switches, trunking ISL ports have Trunking State set to 2, and non-trunking ISL ports have Trunking State set to 1.

The Properties tab usually provides the following, although this list does vary from element to element:

- **Assign a Custom Name** – To make it easier to identify the element, assign the element a custom name. See Assigning a Custom Name on page 500. This option is not available to all elements.
- **IP Address** – Enter an IP address for the generic or inferred host.
- **DNS Name** – Enter a DNS name for the generic or inferred host.
- **Version** – Enter a version number for the generic or inferred host.
- **Operating System** – Enter the basic name for the operating system on the generic or inferred host, such as the following:
  - Windows
  - Solaris

- Linux
- AIX

If you enter a more complex name for the operating system, such as "Win2KSP2," the product displays the name of the operating system on the Properties tab, but not anywhere else in the product such as in System Manager.

- **Vendor** – Enter the vendor name.
- **View element properties** – Enables you to view the element properties for the type of device. The properties provided vary according to the element. The following information is usually provided:
  - **Record Creation** – The first time the software contacted this element.
  - **Discovery Status** – The status of the discovery of the element, for example “Contacted.”
  - **Vendor** – The name of the vendor.
  - **IP Address** – The IP address of the element.
  - **DNS Name** – The element's DNS name.
  - **Provider Name** – The name of the provider.
  - **Model** – The model of the element.
  - **Update Element Data** – To update the displayed properties, click the **Update Element Data** button at the bottom of the screen. The management server gathers new and changed details from the element and then redraws the topology with the updated information. This button is not shown for agentless hosts and custom generic hosts.

    The Update Element Data functionality does not detect element components that were removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you perform Update Element Data for the storage system, the LUNs still appear in the user interface. For the deleted LUNs to be removed from the user interface, you must perform Get Details. See Get Details on page 151.

    Update Element Data does not work for virtual machines or virtual servers.

    Red Hat Linux IA64 hosts might not return a serial number in the host properties tab or in the Asset Summary reports. This is due to an issue in the operating system on the host. If the serial number information is required, install the kernel-utils-2.4.13.99.ia64 package on the Linux host, then remove the access point for the host from the management server Discovery user interface and run discovery Step 1 and Step 3 against the host.
  - **Virtual Memory** – The virtual memory on the virtual machine.
  - **Is Virtual Server** – If it is a virtual machine, the value is false.
  - **Virtual Server** – If the Property Tab is for a virtual machine, the name of its virtual server is displayed. Click the name of the virtual server to access its property information.

Accessing the Properties Tab

To access the Properties tab:

1. Access the management server.
2. Access the Properties tab for an element by doing one of the following:
   - Click an element, except a file server, in Application Viewer.

     Or

   - Double-click an element in System Manager.
3. (*For SANs or fabrics only*) Access the Properties tab as follows:
   a. In System Manager, click the **List** tab.
   b. Right-click a SAN or fabric name in the List tab.
   c. Select **Go to Properties** from the right-click menu.
4. (*Not applicable to SANs or fabrics*) Click the **Properties** tab.

Assigning a Custom Name

To make it easier to identify a generic element instance in the system, assign it a custom name. The custom named element also appears in Chargeback Manager and can be tracked as an asset.

Since all users query the same database, this name is displayed to others using the software, so you might want to make them aware of the custom name.

When you assign a custom name to a generic host, you are granted access to the Topology and Asset Management tabs for that host. You can also view asset-based chargeback for that host.

To assign a custom name:

1. Access the custom name box by double-clicking the element in System Manager and then clicking the **Properties** tab.
2. In the custom name box, enter a name, keeping in mind the following:
   - The name can contain one to 64 characters.
   - The following characters and symbols are accepted: letters, numerals (0 to 9), ~, @, *, _, -, +, ., < >, (), [ ], { }, |.
   - The name is case-sensitive; for example, “Element1” and “element1” are different elements.
3. Click **Save**.
4. Refresh the page.

If you created a custom name for a generic host, HP Storage Essentials displays the Topology and Asset Management tabs. Generic hosts are displayed as follows:

```
custom_name(Discovered)
```

In this instance, the custom_name is the name you assigned to the generic host.

Viewing Element Topology

This section contains the following topics:

Topology Tab

The Topology tab shows the graphical representation of an element path. It displays information not found in System Manager, such as adapters, slots, and Fibre Channel ports.

For example, assume you want to view the topology of a server called QASERVER02, and it contains seven fixed local disks. If you double-click the server in System Manager and then click the **Topology** tab, you can see the path of the server. The Topology tab also displays the drives of the server's fixed local disks, as well as the adapter used to connect the server with the switch, as shown in the following figure. The following example shows that server QASERVER02 can access three storage systems: LSI, EMC, and HDS.

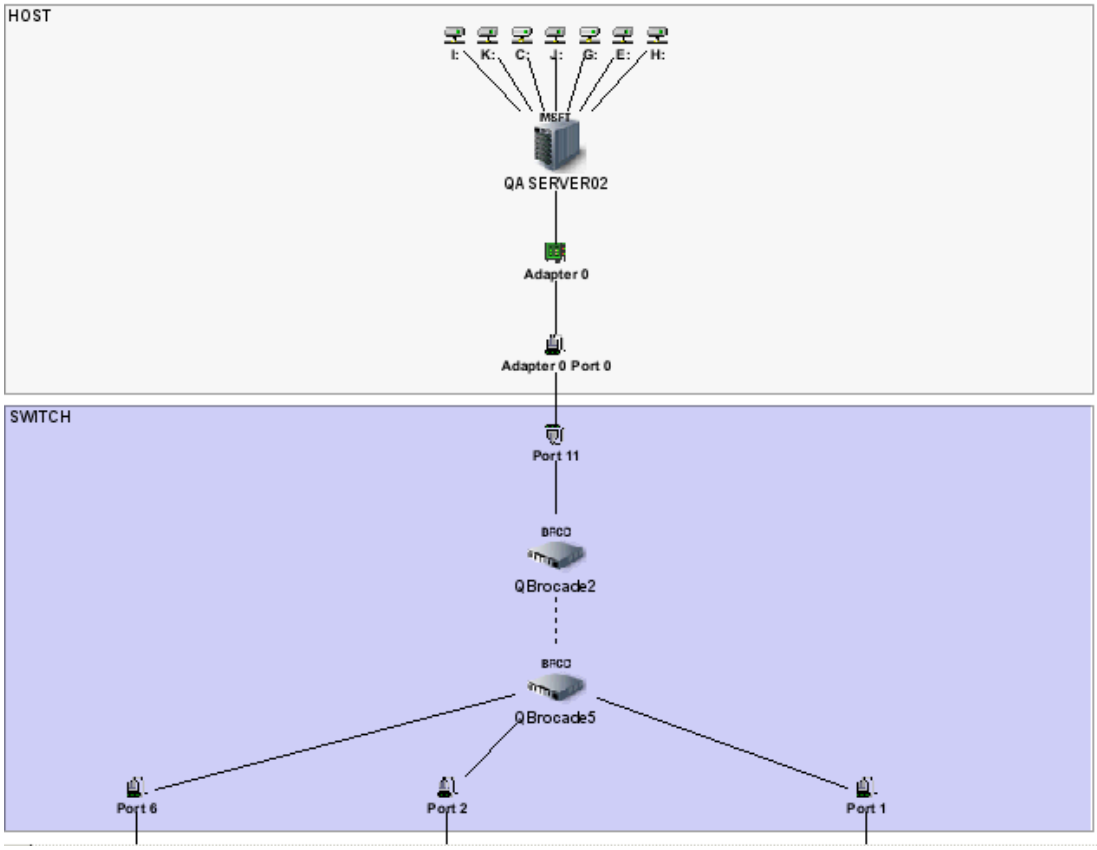

The topology extends the length of the screen.

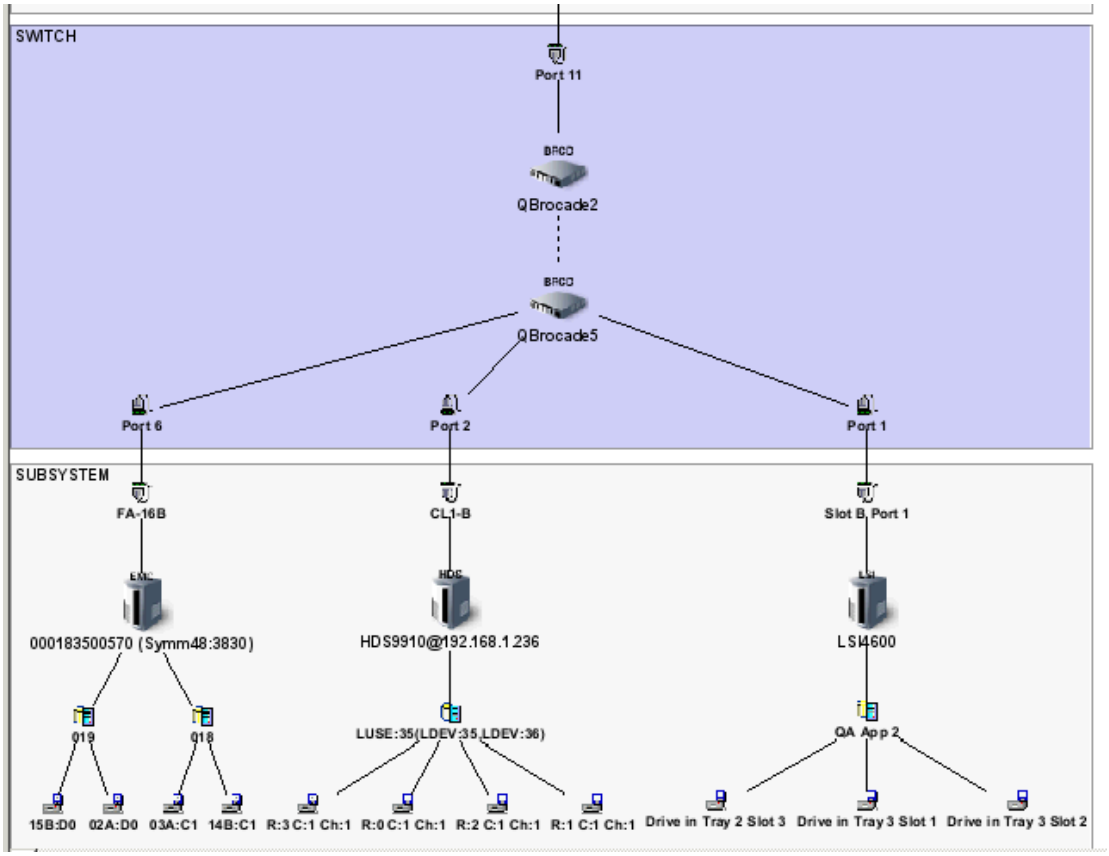

If any of the paths are not fully calculated, a pop-up dialog box displays a list of all hosts with partially calculated paths. In addition, the current state of the path calculation is appended to the node name in the left pane.

Drilling down into EVA VDisks mounted to a host will reveal that the VDisks comprise all disks on the EVA. This reflects the information provided to the management server by the EVA Provider.

Multipathing

Multipathing is the process of providing a server more than one path to a storage system, so that in case of an emergency, the server has continuous access to the storage system. Multipathing can be done many ways. One example of multipathing is providing redundant switches for a host to access a storage system. Another is providing redundant paths from the host to the switch. To determine if your multipathing software is supported, see the support matrix, which is available from the Documentation Center.

SANtricity Manager Utilities must be installed on the host running RDAC for the management server to obtain RDAC information.

HDLM on Sun Solaris requires the storage array to be included in discovery to report the correct information back to the bindings page. The management server does not recognize when EMC storage is managed by Hitachi HDLM. EMC-based HDLM devices are not shown on the management server. See Known Host Issues on page 902 for more information about HDLM.

The software supports VERITAS Volume Manager without VxDMP, but VxDMP is required to do multipathing.

Microsoft Windows 2003 hosts with Service Pack 1 and IBM TotalStorage DS6800 arrays do not stitch properly as a result of the Subsystem Device Driver (SDD) appearing on the same disk. The multipathing page returns the Windows SDD path as something similar to `600507630efe01a80000000000001104:c0t0d0p3`. This makes it difficult to match it up to your SDD path names.

The following figure shows how the software detects multipathing for a server running VERITAS Volume Manager. **MP** is displayed on the path of the redundant volumes:

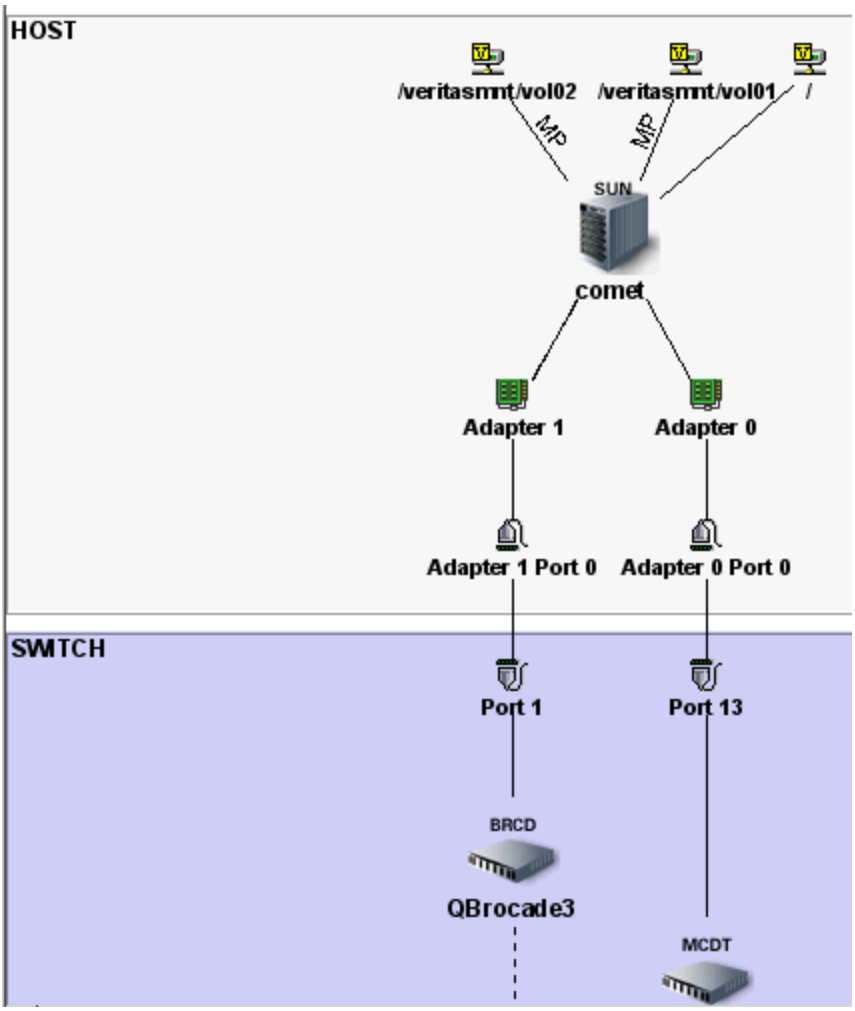

The topology extends the length of the screen. The following figure displays the second portion of the topology:

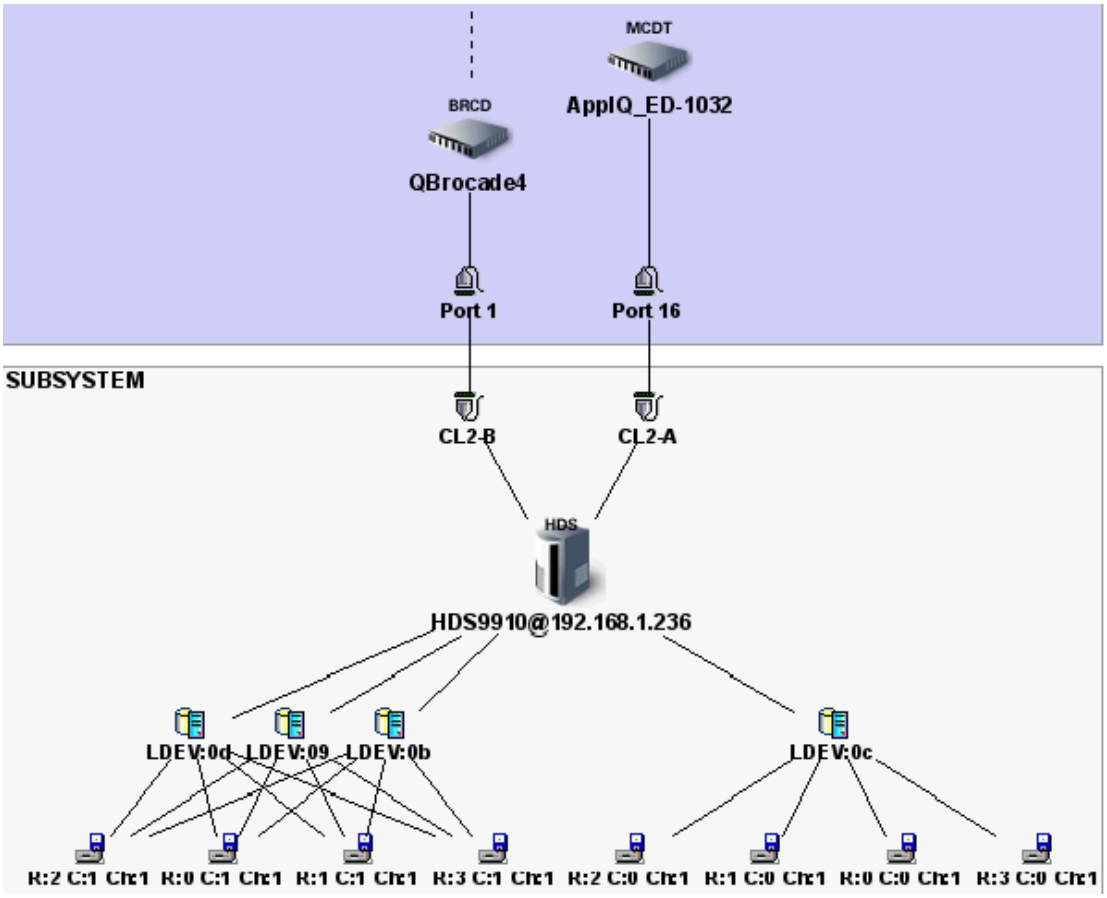

If you do not see all of the elements in the path displayed, verify that they were discovered and details were obtained from them. Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries on page 71.

For multipathing issues regarding certain devices, see Known Host Issues on page 902.

The management server displays only the active path for an RDAC host. It displays only the active path when PowerPath is running on a host connected to a CLARiiON storage system. The management server does not support RDAC configurations for monitoring disk statistics.

Virtual Machines

Virtual machines can store data files on the following types of storage:

- Data stores on the virtual server (logical disks)
- Raw disk partitions created using raw device mapping

The management server represents both types of storage in the topology.

Data Stores and Raw Disk Partitions on Virtual Machines

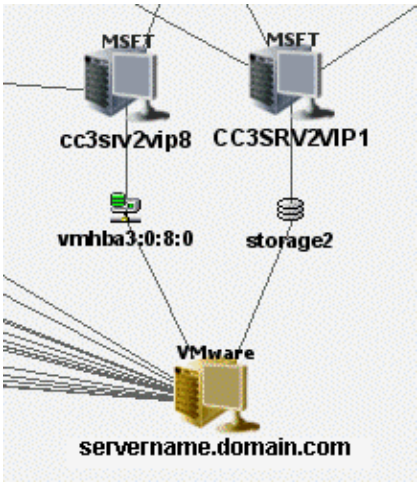

The icon represents a raw disk partition, and the icon represents a data store.

If a virtual machine does not have CIM extensions installed, the management server does not have enough information to determine the volume application path beyond the virtual machine. If the virtual machine does have CIM extensions installed, the management server can determine the entire application path.

If you use RDM to create a virtual disk, only the raw disk partition's dependency is shown. The dependency on the datastore storing the mapping file is not shown.

The element topology normally shows volume-to-spindle associations, but for virtual volumes it shows volume-to-pool associations instead.

Direct Attached Storage

To view direct attach storage, you must enable the button. See Toolbar in System Manager on page 417 for more information.

Once the button is enabled, the management server displays the link between the storage system port and the port to the host as a dotted line, as shown in the following figure:

Direct Attached Storage in the Topology

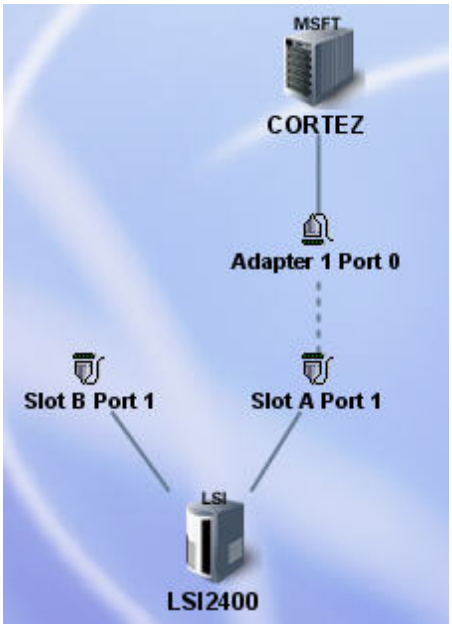

In this figure, Slot A Port 1 belongs to the storage system, and Adapter 1 Port 0 belongs to the host. The dotted line indicates that the storage system is directly attached to the host.

Filers

Element topology for a filer shows the connection from a host to the filer going through an IP cloud, which represents the IP network.

Accessing the Topology

To access the Topology tab:

1. Access the management server.
2. Access the Topology tab by doing one of the following:
   - Select an element in Application Viewer and click the **Topology** tab. The Topology tab is not available for clustered file servers.

     Or
   - Double-click an element in Capacity Manager or System Manager and then click the **Topology** tab.

     Or
   - Right-click an element in System Manager, and then select **Go To Element Topology** from the menu.

Accessing the Topology above describes the icons on the toolbar.

Icon	Description
	Prints the topology. For more information, see Installing Storage Planner on page 461.
	Magnifies the view
	Decreases the magnification
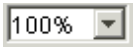	Enables you to set the magnification to a percentage of the default magnification
	Opens a smaller pane, which provides a global view of the topology. This enables you to position the main view to a certain section of the topology. See Using the Global View on page 458.
	Enables you to drag an element in the topology.
	Enables you to move the entire topology at once. Click the **Pan** button (), and then click any place in the topology. Drag the mouse to a new location.
	Enables you to find an element or fabric in the topology. You can enter part of the information, and the management server highlights the elements and fabrics that match. To specify the scope of your search, select All, Element Name, Fabric Name, or Element WWN from the drop-down list next to the search box. To expand the Search box, close the left pane. See Opening and Closing the Left Pane on page 51.
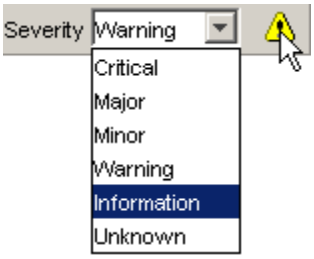	Displays the event severity icons for the elements displayed in the topology. This feature is disabled for Performance Manager and Capacity Manager. See Viewing Event Status in the Topology on page 468.

About the New Window Option

The New Window option in System Manager enables you to view several sections of the topology at once. Click the  button. A new window pops opens. Use this window to view another section of the topology.

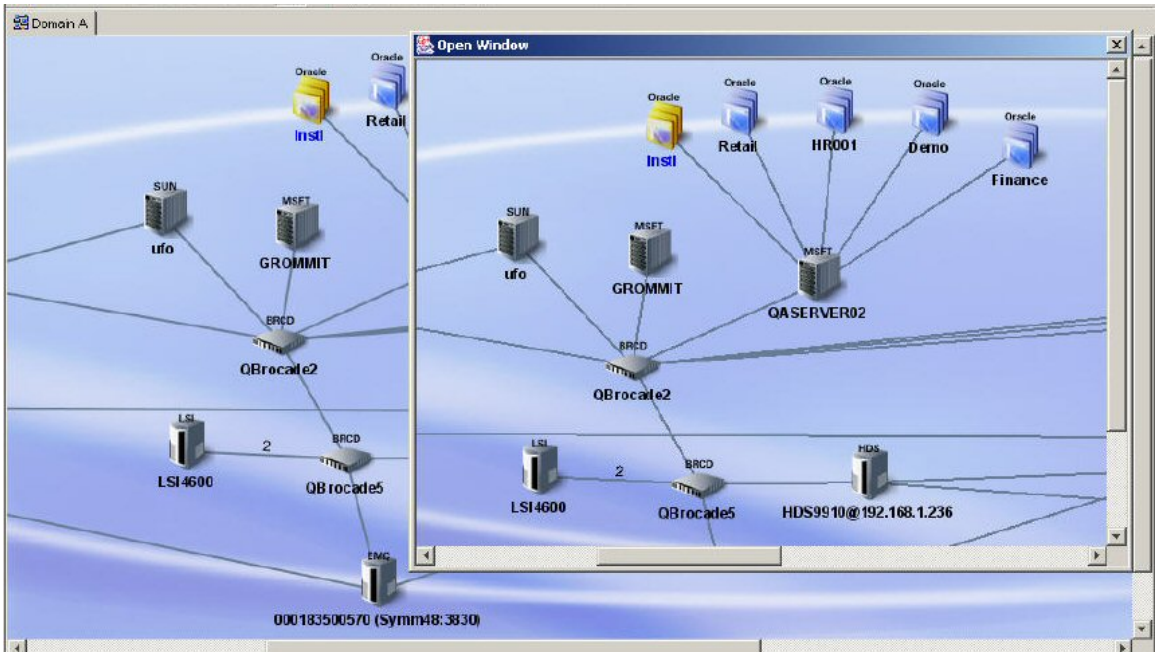

Printing the Topology

The software enables you to print the topology. This option is extremely helpful when you want to show someone the layout of the network, such as in a presentation.

To print the topology:

1. Access the management server.
2. To access the Topology tab, do one of the following:
   - Select an element in Application Viewer, and click the **Topology** tab.

     Or

   - Double-click an element in Capacity Manager or System Manager, and then click the **Topology** tab.
3. If the topology spans more than one screen, arrange the elements so they are closer together, preferably on one screen. To move an element, click the button and then the element you want to move. Drag the element to its new location. Moving elements closer together provides a more compact printout.
4. Click the button.

   The Paper tab shows the page setup. If you want the default settings, click **Default**. You can modify the following settings:

   Before you change the margins, decide on a unit of measurement.

   - **Paper format** – Select the paper size from the menu.
   - **Unit** – Select cm (centimeters) or inch for the margins.

- **Paper width** – To modify the width of the paper, select the Custom option in the Paper format menu.
- **Paper height** – To modify the measurement in this box, select the **Custom** option in the Paper format menu.
- **Top margin** – Enter a measurement.
- **Bottom margin** – Enter a measurement.
- **Left margin** – Enter a measurement.
- **Right margin** – Enter a measurement.
- **Orientation** – Click an orientation for the printout.

A preview of the printout is displayed in the right pane.

5. When done, click **Apply**.
6. To see how the printout will appear on the page, click the **View Selection** tab. If you want the default settings, click **Default**. You can modify the following settings:
   - **Start x** – Determines the horizontal placement of the printout on the page, with zero being the closest to the right margin. For example, if the value is 50 for **Start x**, the printing starts at 50 inches or centimeters (depending on what you selected) from the right margin. You can also enter negative numbers. Anything more than zero expands the printout to another page.
   - **Start y** – Determines the vertical placement of the printout on the page, with zero being the closest to the bottom margin. For example, if the value is 50 for **Start y**, the printing starts at 50 inches or centimeters (depending on what you selected) from the bottom. You can also enter negative numbers.
   - **Width** – Determines the width of the printout.
   - **Height** – Determines the height of the printout.

   To remove extra space around the topology, click **Trimmed**. A preview of the printout is displayed in the right pane.
7. When done, click **Apply**.
8. The Pages tab shows how many pages the printout will use. If you want the default settings, click **Default**. You can modify the following settings:

   Before you change the margins, decide on a unit of measurement.

   - **Unit** – Select cm (centimeters) or inch for the margins.
   - **Position/Size** – Enables you to change the position and size of the printout so that it spans several pages:
     - **Start x** – Same as in step 6.
     - **Start y** – Same as in step 6.

- **Width** – Determines the width of the printout. If the width entered does not fit on the page, the printout wraps around to another page.
- **Height** – Determines the height of the printout. If the height entered does not fit on the page, the printout wraps around to another page.

- **Resolution (pixel/unit)** – Enables you to change the resolution so that the printout spans several pages.
- **Page** – Enables you to expand the printout so it prints on several pages without modifying the graphic.

A preview of the printout is displayed in the right pane.

9. When done, click **Apply**.
10. To preview your pages, click the **Preview** tab, and then click the page you want to preview. The page appears in the right pane.
11. When you are ready to print, click **Print**.
12. Click **Close**.

To return to all of the original settings, click the **Default** button next to the **Print** button.

Creating a Virtual Application

The management server enables you to keep track of unsupported applications. For example, assume your company has created an internal application, and you want to be able to use the software to keep track of that application. You can create a virtual application for that product. A virtual application is a placeholder you create for an application.

Once you create the virtual application, it will appear connected to a host in your topology.

1. Access System Manager by clicking the **System Manager** button in the left pane.
2. Right-click the host that contains the application you want to monitor.

   If the host is not in the topology, verify that you discovered the element and obtained element details. See Get Details on page 151.
3. Select **Add Virtual Application** from the menu.
4. Enter the following information for the virtual application.
   - **Name**
   - **Product**
   - **Description**
   - **Vendor**
   - **Version**
5. Click **Next**.

6. Select a storage volume containing the application for which you are creating the virtual application.

   You can view the properties of a volume by clicking its link.

7. If applicable, choose a disk partition by clicking the **Disk Partitions** tab or the **Next** button and then selecting a disk partition.
8. Click **Finish**.

Provisioning Tab

The provisioning tab provides different functionality, depending on the type of element you double-click in System Manager or click in the Provisioning pages. You can also access the provisioning table by right-clicking a fabric, selecting the **Go to Properties** option, and clicking the **Provisioning** tab.

If you selected a switch or a fabric, you are shown zone provisioning tools that let you manage zones, zone aliases, and zone sets. These tools provide a wide range of functionality, such as the following:

- Creating a Zone Alias on page 683
- Creating a Zone in a Fabric on page 686
- Creating a Zone Set on page 688
- Activating a Zone Set on page 691

For more information about setting up zones, see SAN Zoning Overview on page 679.

If you double-click a storage system, you are shown storage provisioning tools that let you create storage pools, volumes, and host security groups. These tools provide a wide range of functionality, such as the following:

- Managing Storage Pools on page 701
- Managing Volumes on page 704
- Rules for Creating Host Security Groups on page 712
- Managing Host Security Groups on page 716

About the Events Tab

The Events tab enables you to view, clear, sort, and filter events for an element. An event can be anything that occurs on the element; for example, a device connected to a Brocade switch has gone off-line. The Events tab provides the following information about the events:

- **ID**- The identification number assigned to the event
- **Severity** – The severity level
- **Time** – The time the event was recorded

- **Summary Text** – A brief explanation of the event. When you click the summary text, the details of the event are displayed.

The Events tab enables you to use Event Manager to:

- **View Event Details** – See Viewing Event Details on page 535.
- **Clear Events** – See Clearing Events on page 538.
- **Delete Events** – See Deleting Events on page 540.
- **Sort Events** – See Sorting Events on page 541.
- **Select a Severity for Filtering** – See Setting Up a Filter on page 545 and Setting Up Advanced Filtering on page 550.

To view all events, click the **Event Manager** button in the left pane. See About Event Manager on page 529.

Asset Attributes of an Element

Depending on your license, Chargeback Manager might not be available. To determine if you have access to Chargeback Manager, see the List of Features, which is accessible from the Documentation Center (**Help** > **Documentation Center**).

Chargeback Manager provides a handy way for you to keep track of your asset information for an element. You can easily store warranty and licensing information, as well as contact information for the element. For example, assume a switch on the network is having some problems, and you want to contact the person in charge of that switch. You can use the element's asset record to find both the contact information for that switch and the location of the switch.

To access asset information for an element, do one of the following:

- Click an element (except a file server) in Application Viewer, and then click the **Asset Management** tab.
- Double-click an element in Capacity Manager, Performance Manager, or System Manager, and then click the **Asset Management** tab.
- Click a discovered host in Backup Manager, and then click **Chargeback** in the lower-left corner.
- Click an element in Chargeback.

  The Asset Management tab displays general asset information about an element. It also provides access to other screens that provide additional asset information, such as staff, geographic, licensing, and warranty information. You can access these other screens by expanding the Asset Record node and clicking one of its children, as shown in the following figure. To learn more about these other screens, see the following topics:

  - Adding Asset Information on page 819
  - Adding General Information on page 820
  - Adding Staff Information on page 821

- Adding Geographic Information on page 822
- Adding Licensing and Warranty Information on page 822
- Adding Custom Information on page 822

Viewing Asset Records

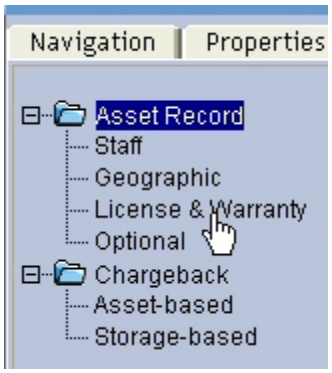

To set up chargeback, expand the Chargeback Manager node, and click **Asset-based** or **Storage-based**. To learn more about each type of Chargeback Manager, see Setting Up Asset-Based Chargeback on page 824 and Setting Up Storage-Based Chargeback on page 829.

The boxes on the Asset Management tab are as follows. When you are done with adding information on this page, click the **Save Changes** button at the bottom of the page.

The boxes that accept input cannot contain more than 250 characters.

- **Custom Name** – A name you assign to the element. See Assigning a Custom Name on page 500 for more information.
- **Date Created** – Date the element was discovered.
- **Date Last Modified** – Date the record was last modified.
- **Description** – A description of the element. This description cannot be more than 250 characters.
- **Status** – The current status of the element. If the status of the element has changed, select the new status from the Status menu.
  - **New** – This is the default category for all detected elements.
  - **Missing** – The element is no longer detectable through discovery.
  - **Repaired** – The element is being repaired. The software does not automatically select this status.
  - **In Use** – The element is in use.
- **Vendor** – The vendor for the element.
- **Model** – The model of the element.
- **Serial Number** – Serial number of the element.

- **Barcode Number** – The barcode on the device.
- **Asset Code** – The asset code assigned to the element.
- **Asset Type** – The asset type assigned to the element.
- **Asset Tag** – The asset tag assigned to the element.
- **Asset Category** – The asset category assigned to the element.
- **Geographic Location** – The location of the element; for example, Boston, Massachusetts.
- **(Storage Systems Only) Storage Tier Classification** – Click the **Set Storage Tier Cost** link to set up storage tiers.

About the Collectors Tab

The management server uses collectors to gather information. The Collectors tab provides information about the collectors for a particular element.

To start collectors and view reports for an element:

1. To access the Collectors page, do one of the following:
   - Click an element in Application Viewer, and then click the **Collectors** tab. (For file servers, click the **Scan Schedule** tab).

     Or
   - Double-click an element in Capacity Manager, Performance Manager, or System Manager, and then click the **Collectors** tab.

     Or
   - Click an element in Chargeback Manager, and then click the **Collectors** tab.

     Or
   - Click a discovered host in Backup Manager, and then click **Collectors**in the lower-left corner.
2. To change a collector's start time, modify the time and date entered in the Next Scheduled Run box. If you decide to change the start time, make sure the date is in yyyy-mm-dd format and the time in 24-hour format. There should be a space between the date and the time, as shown:

   ```
   2005-06-26 09:41
   ```

   After the collector runs, the value in this column is updated to the next time the collector runs.
3. To change how often the collector runs, enter the number of minutes in the Interval box.

   Do not make the interval too short. Running a collector too frequently uses up space on the management server and impacts its performance.
4. To enable the collector, click **Start**.
5. To stop a collector, click **Stop**.

About the Monitoring Tab

To easily access performance information about an element:

1. Double-click the element in System Manager or Application Viewer.
2. Click the **Monitoring** tab. The element appears highlighted in Performance Manager.
3. Select one of the monitoring options in the lower pane to view specific performance data about the element.

See Performance Manager on page 595 for more information about Performance Manager.

About the Policies Tab

The Policies tab enables you to view the utilization policies for an element. Utilization policies can automatically send an e-mail, generate an event, or run a custom script when an element is being overused. If the policy table is unpopulated, no policies exist for the element.

The Policies tab enables you to use Policy Manager to do the following. See About Policy Manager on page 579 for more information:

- Add Policies
- Test Policies
- Edit Policies
- Delete Policies

To access the Policies tab, do one of the following:

- Double-click an element in Capacity Manager, Performance Manager, or System Manager, and then click the **Policies** tab.
- Right-click an element in Capacity Manager, Performance Manager, or System Manager, and then select **Show Policies** from the menu.
- Click a discovered host in Backup Manager, and then click **Policies** in the lower-left corner.

To access utilization policies for other elements and to create other types of policies, click the **Policy Manager** button in the left pane.

About the Presented Storage Tab

The Presented Storage tab displays the storage presented to the port WWNs of a host, even if the host cannot access the storage as a result of an incorrect zoning configuration.

For example, assume you used HP Storage Essentials to discover a host that is connected to two storage devices and that has host security enabled. You then change the zoning so the host can detect the LUN from only one storage device, but you do not change the storage presented to this host from the array that is removed from the zone. You reboot the host and make sure that the host now sees storage from only one array. After running Get Details again to view the updated information, the Presented Storage tab still displays information about both the storage devices.

Assume you added a multipathing configuration. You will not see the change in the topology until you run an update on the Host tab. The storage calculations displayed on the Presented Storage tab may also change as well to account for new configurations. See Viewing Agentless Hosts on page 262 for more information about the Host tab and the update functionality.

An update looks at the WWNs on the storage array. If you had a WWN that belonged previously to a host's HBA port, HP Storage Essentials will report that host with an empty HBA port column. Host inference is only as good as the configuration of the zoning and host security groups.

To access the Presented Storage tab:

1. Double-click a host in System Manager that is one of the following:
   - A discovered host
   - An inferred host, which is a host without a CIM extension and which has been discovered through agentless discovery
   - A custom generic host, which is a host without a CIM extension and which has been since named
2. Click the **Presented Storage** tab.

Presented storage is only available when LUN security is enabled. If LUN security is disabled for any ports on an array, HP Storage Essentials cannot determine which hosts have storage presented to them over those ports.

The Presented Storage tab provides the following information:

- **Name of the storage system connected to the host**
- **The vendor for the storage system**
- **The model of the storage system**
- **Gigabytes presented on the storage system**
- **The total amount of storage available to the host**

  If the host has only one storage system connected to it, this number is most likely the same as the number provided in the GB Presented column.

- **Details button**

  Click the **Details** button to obtain the following information about the volumes on the storage system:

  - Volume name
  - Size of the volume
  - Name of the Ports through which this volume is accessible
  - Name of the Pool that this volume belongs to
  - Is the volume on a mainframe?
  - Number of Storage Logical Partitions (SLPR) on the volume
  - Number of Cache Logical Partitions (CLPR) on the volume
- Volume Flag provides information about a volume, such as whether a volume is a member of a replication pair, thin provisioned volume, command device or reserved for some other special purpose.

Determining If a Host Belongs to a File System

You can determine if a host is a member of a file system such as CXFS on the Navigation tab or in Capacity Manager.

To use the Navigation to determine if a host is part of a file system:

1. Access System Manager as described in Accessing System Manager on page 416.
2. Double-click the host.
3. Click the **Navigation** tab.
4. Click **Storage Volumes**. The system type, such as CXFS, is listed in the File System Type column.

The following information about the storage volume is also provided:

- Name of the storage volume
- Description of a storage volume
- Drive Type

To use Capacity Manager to determine if a host is part of a file system:

1. Access Capacity Manager as described in Accessing Capacity Manager on page 555.
2. Select the host.
3. Scroll to the bottom of the page. If a storage volume is a member of a shared file system, such as CXFS or XFS, it is listed in the Storage Volume column.

   You might need to expand the Storage Volume column if the volume names are long.

About the Data from CXFS File Systems

The management server can only monitor CXFS file systems from the host generating the input/output. For example, assume the elements in the following figure are part of a CXFS file system. When you generate input/output into the metadata server into /folder, only the metadata server is able to monitor the file system. For example, if the metadata server generates a 100-KB write, the management server displays a 0-KB write for the /folder on the metadata client.

The information in the /folder on the metadata server is actually being mirrored to the /folder on the metadata client. The management server, however, does not detect the changes being mirrored to the /folder on the metadata client.

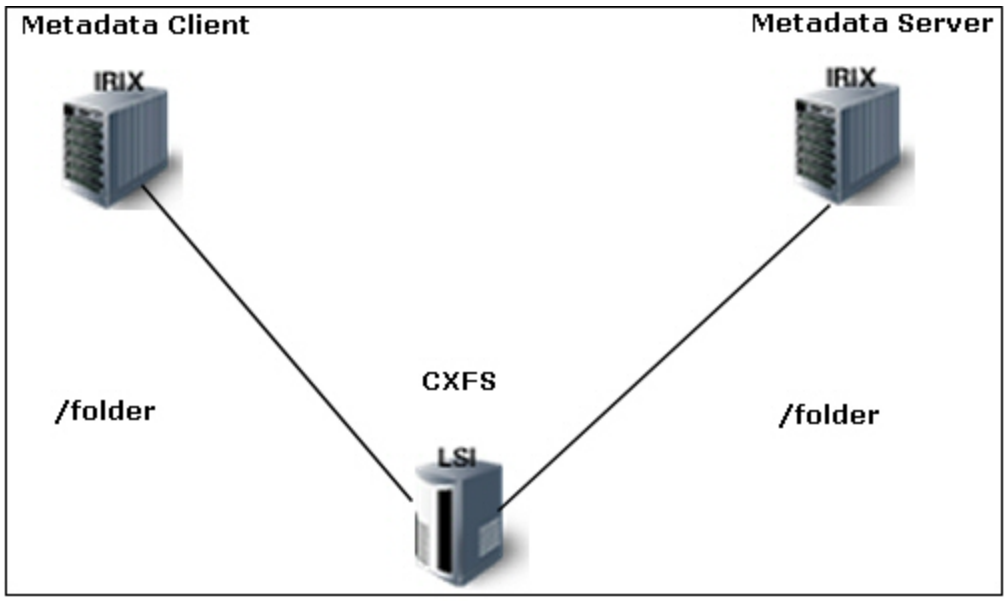

11 Using Element Manager

Use Element Manager to do the following:

- **Search tab**. Lets you quickly find elements and their details. See Searching for Elements below.
- **Groups tab**. Lets you create hierarchical groupings of elements. You can group elements for reporting, billing and inventory purposes. For example, you can group elements from multiple data centers. The subgroups could then contain elements from each data center. For example, the top group could be named data centers. The second level of groups can be named after each data center. See Creating Groups in Element Manager on page 522.

Searching for Elements

Use the search functionality in Element Manager to find information about an element quickly.

To search for an element:

1. Click **Element Manager**.
2. Click the **Search** Tab.
3. Enter text in the search box, and click **Search**.

Element Manager lists all of the elements whose name includes the string of text that you entered.

The following information is provided for each element:

- Vendor
- DNS Name
- IP Address

The following links are provided for each element:

- Properties
- Events
- Asset Management
- Collectors
- Monitoring
- Policies
- Provisioning (if applicable)
- Build Cluster (if applicable)

To view a list of dependent elements, open the collapsible Dependent Elements section by clicking **+ Dependent Elements** (from path). Click a dependent element's name to access the navigation page for that element.

Viewing the Hierarchical Grouping of Elements

To access hierarchical grouping:

1. Click **Element Manager** in the left pane .
2. Click the **Groups** tab.

The hierarchical grouping feature allows you to use folders to create hierarchical groups of elements for reporting, billing, and inventory management. The folders are created with a parent-child relationship. When a parent folder is selected, all of the elements and elements of the child folders are listed. You can nest folders up to six levels deep. See Groups and Reports below for more information about using groups within reports.

Figure 3 Example of a Folder Structure with Four Levels

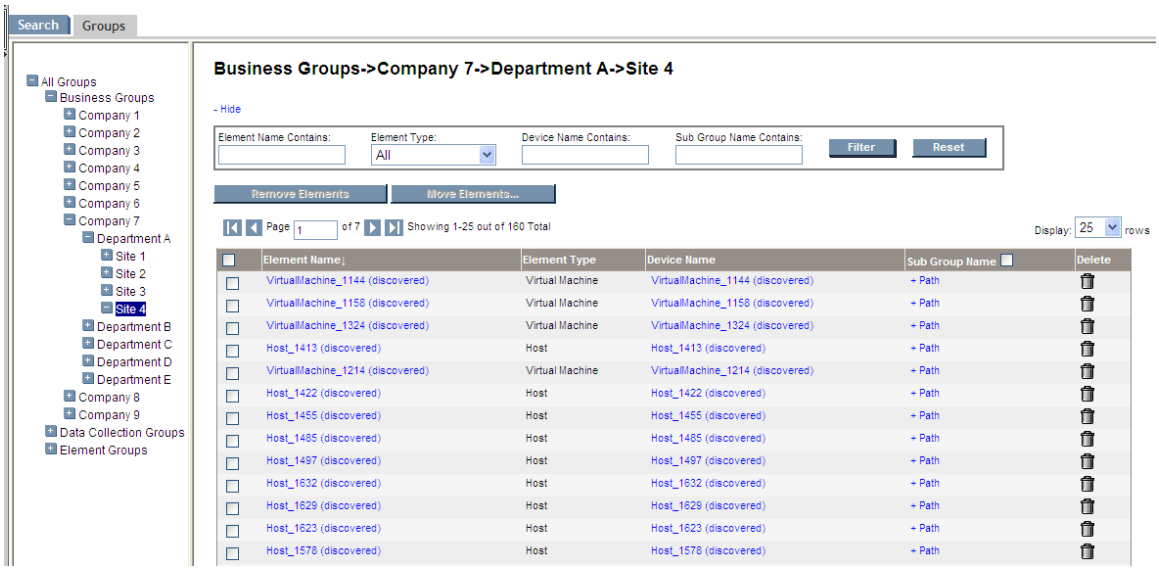

Groups and Reports

Use the information from groups to create reports in Report Optimizer.

For example, the following figure shows a report being created that will display some of the information from groups, such as group name, group level, and group path.

Figure 4 Creating a Report with Group Information

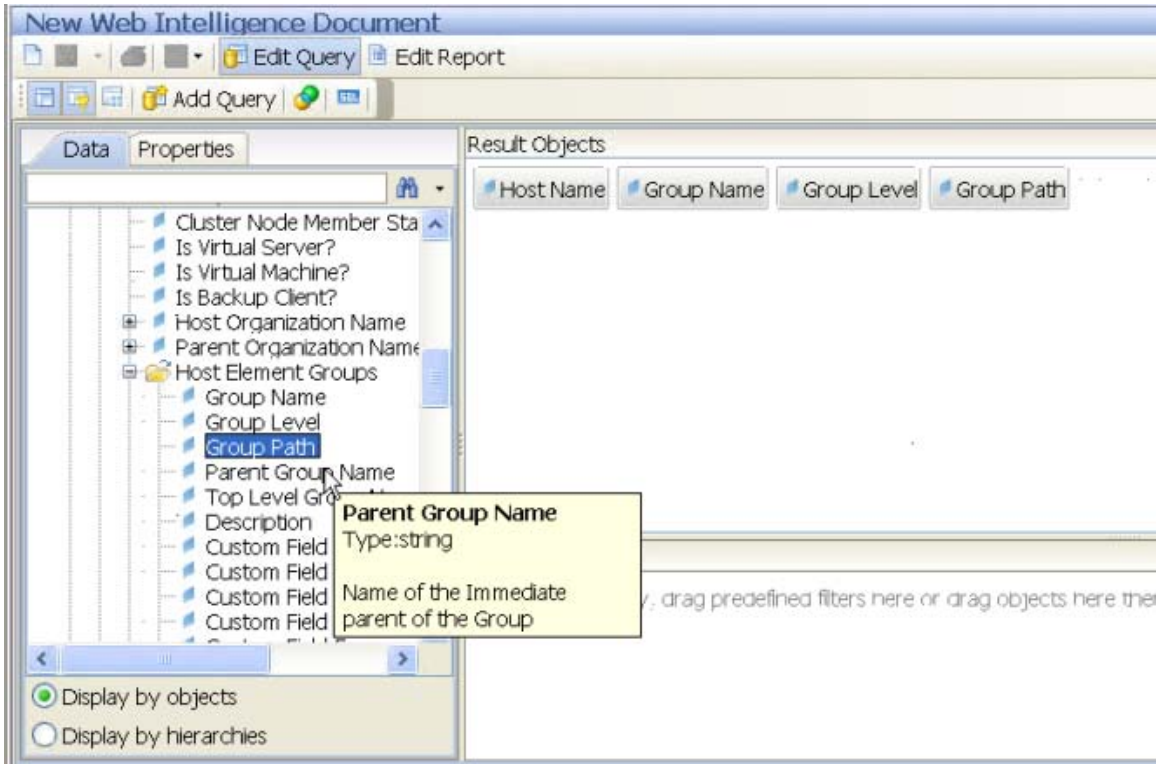

The following figure shows the output of the report that was created. Notice how information from the groups is used in the report, such as group name, group level and group path, in addition to a custom field.

Figure 5 Report Displaying Group Information

Report Title

Host Name	Group Name	Group Level	Group Path	Custom Field 1	Total Capacity in GI	Used Capacity in G	Free Capacity in GE
buggigadu.ind.hp.c	Reports	6	HP BSA->Storage E		67.64	8.1	59.54
QA229	Applications	6	HP BSA->Storage E		143.81	44.99	98.82
QA40	COST	6	HP BSA->Storage E	XYZ	44.17	8.61	35.56
QA61	COST	6	HP BSA->Storage E	XYZ	968.84	382.72	586.12
QA73	CPE	6	HP BSA->Storage E		37.27	34.83	2.44
QA86	Providers	6	HP BSA->Storage E		66.56	45.77	20.79

When you search for a report, you can search by group path, as shown in the following figure.

Figure 6 Searching for Reports by Group Path

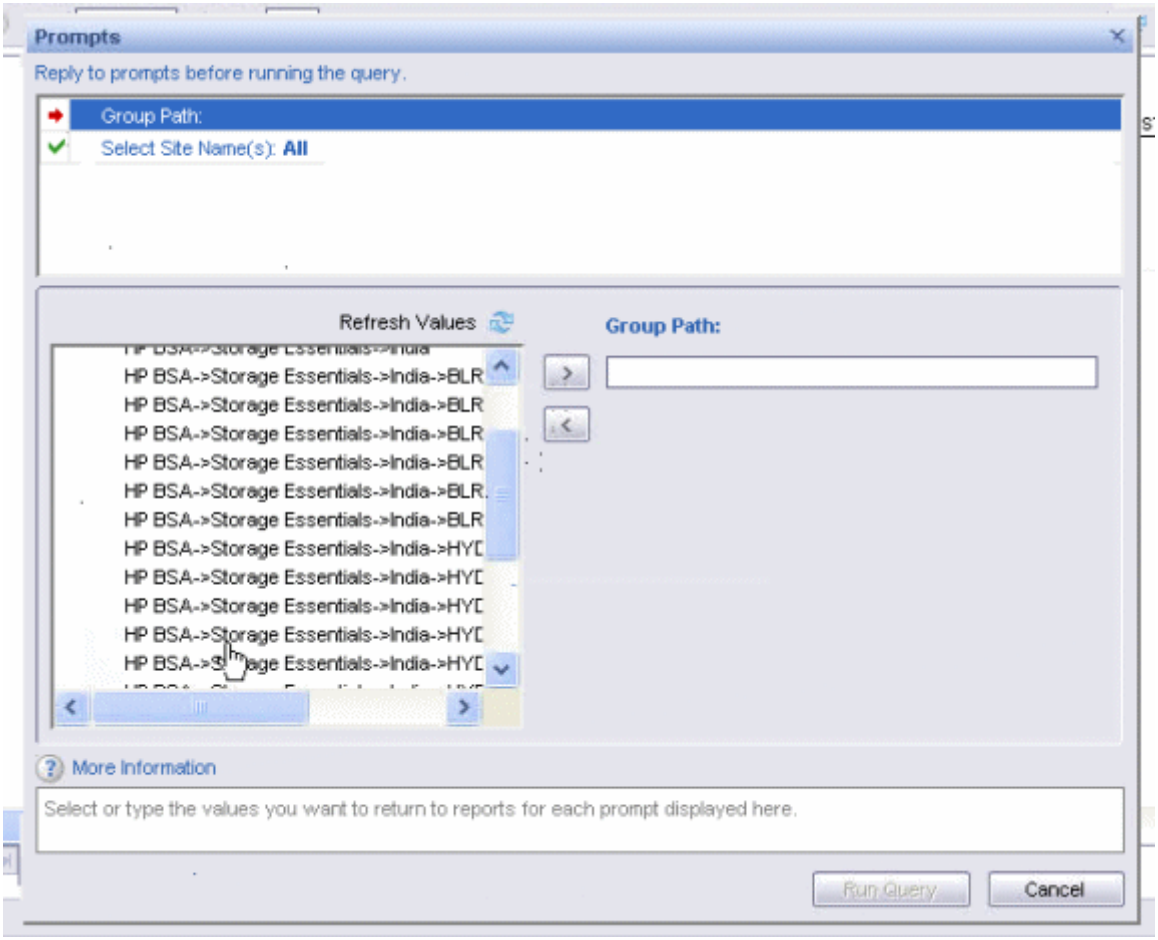

Creating Groups in Element Manager

You can group elements for reporting, billing and inventory purposes. For example, you can group elements from multiple data centers. The subgroups could then contain elements from each data center. For example, the top group could be named data centers. The second level of groups can be named after each data center.
An element can be assigned to multiple groups, and groups can be created up to six levels deep in a hierarchy.

Drop-down lists in the filter for Element Manager do not list the expected storage pools, pool types, or capabilities when creating a group.

To create a group in Element Manager:

1. Click Element Manager
2. In the left pane, right-click the parent group under which you want to create the new group.

Each group name must be unique from those at the same level; however, a group can have the same name as another group at a different level.

3. Select **Create Group** from the menu.
4. In the Properties section, enter a name for the group in the Name box.
5. *Optional*: Enter a description for the group in the Description box.
6. *Optional*: Enter information in the custom field boxes. The custom field boxes allow you to add information that can be used when generating reports. You can rename the custom fields. See Renaming the Custom Fields on page 526 for more information.
7. In the Available Elements section, select an element type from the Element Type drop-down list.
8. Depending on the element type, use filters to limit the elements that are listed as available.
9. Select the elements that you want to add to the group.
10. Click the **Add Selected Items to Group** button. The selected elements are added to the Group Members section.
11. Click **OK**. The group is created with the members you specified.

Editing a Group in Element Manager

To edit a group:

1. Click Element Manager.
2. In the left pane, right-click the group you want to edit.
3. Select **Edit Group** from the menu.
4. Make the desired changes.
5. Click **OK**. The edits to the group are saved.

You can change the name of the group, its description and custom fields. You can rename the custom fields (see Renaming the Custom Fields on page 526).

To add elements to the group:

1. Select an option from the Element Type menu.
2. Select the elements you want to add.
3. Click the **Add Selected Items To Group** button.

To remove elements from the group:

1. Select the group members you want to remove from the Group Members pane.
2. Click the **Remove Selected Items From Group** button.
   Elements can be moved or removed from the current group only if the same elements are assigned to the child groups. They cannot be deleted.

Importing Groups in Element Manager

To recreate your group hierarchy on other servers running HP Storage Essentials, import the group xml file to those other servers.

To import groups:

1. Click **Element Manager**.
2. Right-click the **All Groups** node and select **Import Groups** from the menu.
3. Browse to the location of the XML file that contains the groups that you want to import. Select the file, click **Open**, and click **OK**.

   You are shown the number of groups that were saved/imported, skipped, or ignored.

   Groups are saved/imported if the product does not find any group under the given parent in the same level with the same name.

   Groups are skipped if the product identifies a pre-existing group with the same name at the same level. Group names must be unique at the same level. For example, assume the group xml file is imported twice. The first time the product imports all the groups defined in the xml file. During the second time, the product does not import any of the groups because the groups already exist.

   Groups are ignored if the XML in the file is not correct. For example if the <NAME> tag for a group is missing in the XML file, then the product ignores that group and its children. A group might also be ignored if the group is more than six levels deep. The product supports groups at most six levels deep.
4. Click **Close**.

Exporting Groups Created in Element Manager

You can export the group hierarchy you created and imported it into another instance of the management server.

To export groups:

1. Click Element Manager.
2. Right-click the **All Groups** node and select **Export Groups**.
3. Click **Save** to save the file containing the listing of groups.

Moving Elements in Element Manager

Element Manager lets you move elements from one group to another. Elements can be moved to an existing group or to a new group.
Elements can be moved or removed from the current group only if the same elements are assigned to the child groups. They cannot be deleted.
To move elements into a different group:

1. Click **Element Manager**.
2. In the left pane, select the group that contains the elements that you want to move.
3. In the right pane, select the elements that you want to move, and click **Move Elements**.
4. Complete one of the following actions:

   To move the elements into an existing group, select the "The Group selected below" radio button. In the group table, select the radio button for the group into which you want to move the selected group.

   To move the elements into a newly created group, select the "A new group" radio button, and specify the name of the new group in the available box.
5. Click **OK**.The elements are moved into the specified group.

Moving Groups in Element Manager

To change the location of a group in the folder structure:

1. Click **Element Manager**.
2. In the left pane, select the group that you want to move.
3. Click **Move Group**.
4. Complete one of the following actions:

   To move the group into an existing group, select the "The Group selected below" radio button. In the group table, select the radio button for the group into which you want to move the selected group.

   To move the group into a newly created group, select the "A new group" radio button, and specify the name of the new group in the available box.
5. Click **OK**.The group is moved into the specified group.

Deleting Groups from Element Manager

When you delete a group, its child groups and the corresponding mappings are removed.

To delete a group from Element Manager:

1. Click Element Manager.
2. Right-click the group you want to delete and select **Delete Group**.
3. When you are asked if you want to delete the group, click **OK**.

Renaming the Custom Fields

You can rename of the custom fields displayed on the page for creating or editing a group in Element Manager. The product has a separate property for the name of each custom field at each level. For example, the following property is for naming the first custom field from the left on the Create Groups and Edit Groups windows and at the first group level:

```
group.level1.CustomField1=Custom Field1
```

In this instance:

- `level1` refers to the first group level.
- `CustomField1` refers to the first custom field from the left on the Create Groups and Edit Groups windows.
- `Custom Field1` refers to the name assigned to the field.

To rename the custom fields:

1. Determine which custom fields you want to rename and at which group levels.
2. Make your changes in a text editor so you can easily paste them into the management server interface. Open a text editor, such as Notepad and paste the following property:

   group.level1.CustomField1=Custom Field1

   Change the following:

   - The number after `level`, so it corresponds to the group level for the custom field you want to rename.
   - The number after `CustomField`, so it corresponds to the name of the custom field you want to rename.
   - Rename `Custom Field` to the new name for the custom field.

     Create a separate property for each custom field you want to rename at each group level. If you want to rename all custom fields at all levels, you would have 30 properties.

     For example, assume you want to rename the third custom field at all group levels, you would list the modified properties as follows:

     group.level1.CustomField3=New Field Name

     group.level2.CustomField3=New Field Name

     group.level3.CustomField3=New Field Name

     group.level4.CustomField3=New Field Name

group.level5.CustomField3=New Field Name

group.level6.CustomField3=New Field Name

Keep in mind the following:

- The number after the word `level` refers to the group level.
- The number after the word `Field` refers to which field will be renamed. In this case we are modifying the name of the third custom field from the left on the pages for creating and editing a group in Element Manager.
- `New Field Name` is the new name of the field.

3. Select **Configuration** > **Product Health**, and then click **Advanced** in the **Disk Space** tree.
4. Click **Show Default Properties** at the bottom of the page.
5. Paste the properties you modified into the Custom Properties box.
6. Click **Save**. The updated field name is displayed for the third custom field at all group levels, as shown in the following figure.

12 Event Management

This section contains the following topics:

About Event Manager

Event Manager enables you to view, clear, sort, and filter events from managed elements. An event can be anything that occurs on the element, such as when a device connected to a Brocade switch goes offline. Event Manager provides the following information about the events:

- **ID** – The identification number assigned to the event.
- **Severity** – Identifies the severity level.
- **Time** – The time the event was recorded.
- **Element** – The source of the event. An element can be a switch, host, application, fabric or anything else on the network.
- **Summary Text** – A brief explanation of the event. When you click the text, the details of the event are displayed.
- **Event Type** – Specifies whether the source of this event is an application, a host, etc.
- **Count** – The total count of similar events.
- **Cleared** – Specifies whether an event is cleared. See Clearing Events on page 538.
- **Delete** – Click the  icon to remove an event.

Event Manager also provides several buttons at the top of its screen:

- **Delete Selected** – Deletes all selected events.
- **Delete All** – Deletes all events lists.

- **Clear Selected** – Marks the selected events as cleared.
- **Clear All** – Marks all events as cleared.
- **Un-clear Selected** – Removes the clear status from events that are selected.
- **Un-clear All** – Removes the clear status from all events.

Event Manager Summary Metrics

The top of the Event Manager pane provides a subtotal of the severity rating of the events and the number of events from the different element types. If you do not see the Event Manager summary metrics, you might need to expand the Summary node located at the top of the page.

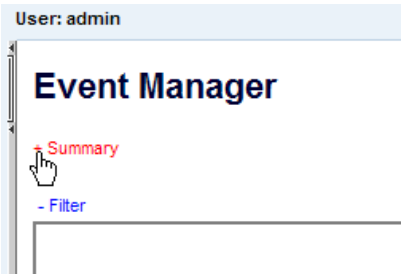

The Event Manager summary metrics are displayed after the Summary node is expanded.

You might need to click the **Refresh** button occasionally for the latest Event Manager summary metrics.

The definition for each severity level varies according to the type of element; however, there are some generic definitions:

Severity Level	Description
Unknown	The event does not fall into the other categories. Further information can not be obtained from it.
Informational	Provides informational data; for example, for a Brocade switch, a list of switches that successfully completed firmware download.
Warning	Provides warning data; for example, for a Brocade switch, one or more new physical fabric objects (device port, switch, or fabric) appeared.
Minor	Provides a message to indicate a minor problem; for example, for a Brocade switch, a physical fabric object (switch port or fabric) changed state.
Major	Provides a message to indicate a major problem; for example, for a Brocade switch, one or more physical fabric objects (device port, switch, or fabric) disappeared.

Severity Level	Description
Critical	Provides a message to indicate a major problem; for example, for a Brocade switch, a device connected to the switch went offline.

Accessing Event Manager

To access Event Manager, do one of the following:

- Click **Event Manager** () in the left pane.

  Or

- To view events from a specific element, take one of the following actions:
  - Double-click the element in Capacity Manager, Performance Manager or System Manager, and then click the **Events** tab. Only events from the element that are double-clicked are displayed.

    You cannot access the **Element Type** filter using this method.
  - Right-click the element in System Manager and select the **Show Events** option in the menu.
  - Select a discovered host in Backup Manager and click **Events** in the lower-right corner.

To change your user preferences for Event Manager, select **Configuration** > **Events**. See Controlling the Display of Cleared and Deleted Events on page 353.

Event Manager Icons

The following icons are displayed in Event Manager.

Icon	Description
	Event is marked cleared. See Clearing Events on page 538 for more information.
	The severity of the event is not known.
	The event is informational.
	The event might have some impact.
	The event has a minor impact.
	The event has a major impact.
	The event has a critical impact.

Events Supported

Event Manager does not support events from all discovered elements. The following elements are supported.

Hardware	Events Supported?	Additional Information
3PAR	N	
Brocade switches	Y	Not all events that show up in the webtool appear in Event Manager.
Cisco switches SNMP	Y	Need to configure the switch or proxy to send traps to the management server. See Configuring the Management Server to Receive SNMP Events on page 536.
Cisco switches SMI-S	Y	
McData switches SMI-S	Y	Only a subset of the available events display in event manager.
CLARiiON storage systems	Y	
LSI storage systems	Y	
HDS storage systems	Y	
HP EVA	Y	HP Insight Remote Support is required with Command View EVA 9.x and the SMI-Provider. See HP Insight Remote Support Required with CommandView EVA 9.x and the SMI-S Provider on page 534 for more information.
HP XP with Command View AE	Y	
HP XP with XP Provider	Y	

Hardware	Events Supported?	Additional Information
HP MSA 1000/1500 storage systems	N	
HP MSA P2000 G2 (2312fc/2324fc) storage systems	Y	
HP MSA P2000 G3 FC storage systems	N	
IBM DS storage systems	N	
IBM XIV	N	
Symmetrix storage systems	Y	
Xiotech storage systems	N	
HP NAS Filers	N	
NetApp Filers	Y	To receive events from NetApp filers, you must add the management server address as an SNMP trap host on the NetApp filer.
Sun NAS Devices	N	
Tape Libraries	N	
EMC Centera	Y	
EMC Celerra	Y	
VMware ESX Server	Y	To receive events related to server Link up, Link down, you must configure the management server to receive SNMP traps from the VMware server.

Viewing Events from the Management Server

By default the management server displays events from all of the elements, regardless of the user's organization. However, it does not display its own events.

To view events from the management server:

1. Select **All** or the **<Product Name of the Management Server>** option from the Show Element Type menu.
2. Click the **Customize** button next to the **Show Element Type** menu in Event Manager.
3. Select the management server and then click **OK**.
4. When asked if you want to apply your changes, click **Yes** and **Apply Filters**.

Avoiding Excessive Notification

The management server provides separate event notification for every event that is reported from the devices it is monitoring. Excessive notification could delay provisioning, as the providers are kept busy notifying the management server of the events. If you do not want the management server to be notified of every event, modify the event threshold of the devices to filter out some of the events. Refer to the documentation accompanying the device for more information about setting the threshold.

Issues with NetApp Filers

To enable the management server to receive events from a NetApp Filer, add the IP address of the management server CIMOM to the NetApp configuration. The management server CIMOM runs on the same computer running the management server by default.

HP Insight Remote Support Required with CommandView EVA 9.x and the SMI-S Provider

With CommandView EVA version 9.2 or later and the SMI-S provider, SNMP traps are no longer used to convey events directly from the CommandView EVA GUI as host alerts and notifications to the HP Storage Essentials CMS (Central Management Server). The HP Storage Essentials server ignores any unexpected SNMP events. You must instead install the latest version of HP Insight Remote Support (IRS) to obtain events from Command View EVA SMI-S formatted as alert indications.

The Web Based Enterprise Services (WEBES), which is the application that performs monitoring for EVA events, is now bundled within IRS. WEBES was previously a standalone download. The new bundling with HP Insight Remote Support has many advantages, including the ability to automatically submit actionable cases to HP including the device's entitlement. If you are using the standalone WEBES download, you are strongly advised to replace your version of WEBES with HP Insight Remote Support. Further details can found at the following website:

http://www.hp.com/go/insightremotesupport

Note: You must configure IRS/WEBES to send “type 3" traps to CommandView EVA (using the “desta snmp on” command) so that the CommandView EVA SMI-S provider will include event severities in the Alert Indications sent to the HP Storage Essentials CMS; otherwise, the events will all have an unknown or informational status.

You must make sure the Array-Based Management (ABM) controller is accurate and that the time zone setup in IRS is accurate; otherwise, there can be very long delays in receiving events from the EVA array. Obtaining SNMP Traps Using Command View EVA on page 120 for more information.

Issues with Brocade Switches

Event Manager does not display events from Brocade switches with the firmware version 3.0. This firmware version is not supported by Event Manager. You can, however, specify that Event Manager not display events from additional firmware versions.

Events from Brocade Fabric Watch are not supported when the Brocade switch is discovered through the SMI-S provider.

Issues with CLARiiON Storage Systems

When you manage a CLARiiON storage system, extraneous events appear in Event Manager such as CRU Bound, CRU Enabled, and various Success messages. These do not indicate problems and can safely be ignored.

Viewing Event Details

To access event details:

1. Access Event Manager as described in Accessing Event Manager on page 531.
2. In Event Manager click the event summary, as follows:

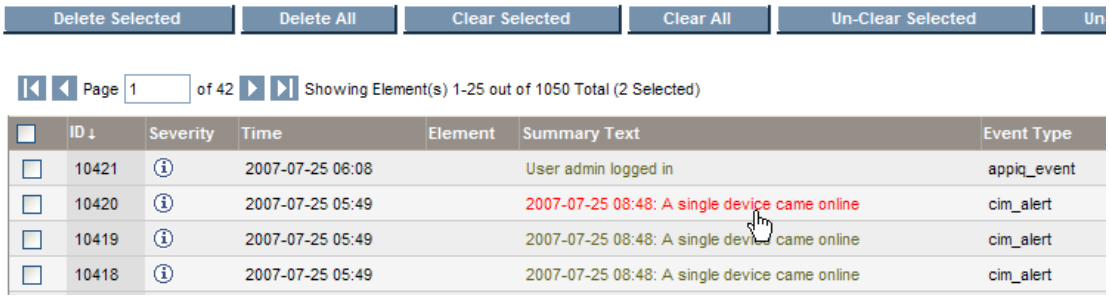

The Event Details pane provides information on one or more of the following:

- **Serial Number** – The number assigned to the event.
- **Status** – Indicates whether the event has been cleared.
- **Component** – If the event came from a component of an element, the component is listed.

- **Element** – The source of the event. An element can be a switch, host, application, fabric, or anything else on the network. If this box is blank, the event did not come from an element.
- **Severity** – The severity level, which can be one of the following:
  - Clear
  - Unknown
  - Informational
  - Warning
  - Minor
  - Major
  - Critical
- **Time Reported** – The time and date the event was reported to the management server.
- **Type** – A brief label of the event.
- **Summary Text** – An explanation of the event.
- **Probable Cause** – An explanation of a probable cause.
- **Recommended Actions** – Provides recommendations.
- **Journal Entries** – Use this box to enter additional information and then click **Add Journal Entry**. This box is limited to 4,000 characters. See Adding Journal Entries on page 541.

Events listed in Event Manager might not be attributed to the correct source until Get Details has completed.

Configuring the Management Server to Receive SNMP Events

Caution: You will not receive SNMP notifications from your EVA if you are running Command View 9.1 or later.

When you move a device discovered by SNMP to a non-default discovery list, configure the management server so that it can still receive SNMP events from the device. The new CIMOM managing the device will not be able to listen on port 162 because the default CIMOM is already listening on that port. The CIMOM will no longer receive the traps because the configured CIMOM is not hosting the device.

Configure the management server by doing one of the following:

- Set the cimom.winsnmpTrapService property to true (cimom.winsnmpTrapService=true). See Setting the cimom.winsnmpTrapService Property to True on the facing page.

Or

- Set the CIMOM to listen on a different port for each of the SNMP events. See Setting the CIMOM to Listen on a Different Port below.

You would want to use each of the methods in the following situations:

Using Windows SNMP Service (Preferred Method)	Multiple CIMOMs
Set the cimom.winsnmpTrapService property to true.	Set the CIMOM to listen on a different port for each of the SNMP events.

Setting the cimom.winsnmpTrapService Property to True

Set the cimom.winsnmpTrapService property to true (cimom.winsnmpTrapService=true) if:

- you are using the Windows SNMP trap service (Windows management server).

And

- you want the management server to still receive SNMP events from the device after you move the device to a nondefault discovery group.

The Windows SNMP trap service must be running on the management server for it to be able to receive SNMP events from devices.

To set the cimom.winsnmpTrapService to true:

1. Select **Configuration** > **Product Health**.
2. Click **Advanced** in the Disk Space tree.
3. Paste the following text into the Custom Properties box. How you paste the text depends on your Web browser.

   ```
   cimom.winsnmpTrapService=true
   ```

4. When you are done, click **Save**.

   The product notifies you if a restart of the AppStorManager service is required.

Setting the CIMOM to Listen on a Different Port

Set the CIMOM to listen on a different port if:

- you are not using the Windows SNMP trap service

And

- you want the management server to still receive SNMP traps from the device after you move the device to a nondefault discovery group.

Each discovery group has its own CIMOM. For example, assume you have two discovery groups. One named Group1 and the other named Group 2. Devices in Group1 are using a different CIMOM than the devices in Group2. You need to tell each CIMOM to listen on a different port for SNMP traps.

To tell the CIMOMs to listen on a different ports:

1. Edit the `CIMOMConfig.xml` which is located in the JBossAndJetty\server\appiq\conf directory.

   This file contains a section for each of the different CIMOMs (Default, Discovery Group 1, Discovery Group 2, etc.).

   In each section there are optional parameters designated by the tags `<option> … </option>`.

2. Add a line to each section specifying the port that will be used to listen for SNMP traps for that CIMOM. The format is as follows:

   ```
   <option>-Dcimom.snmpTrapListenerPort=162</option>
   ```

3. Repeat the previous step for each CIMOM. Each time select a different port. Make sure the port you select is available.

   As you add the ports, make a list of the ports you are adding.

4. Use your list of ports that you assigned to each CIMOM to set up the devices that are managed by that CIMOM. For example, if Switch_A is in Discovery Group 1, you must set Switch_A to the port for Discovery Group 1.

Clearing Events

When you review an event, you might want to mark it as “cleared” so you can keep track of which events you reviewed.

To clear events, do one of the following:

- To clear an event – Click the **Not cleared** text for the event. The text turns red.

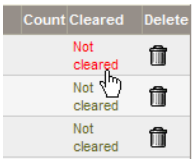

- To clear several events – Select the events and then click **Clear Selected**.
- To clear all events – Click **Clear All**.

When events are cleared, a clear icon (○) appears in the Cleared column for the event.

Configuring the Clearing of Events

Depending on the severity of an event, the management server might mark the event as clear after 60 minutes. Events designated as Major and Critical are never marked as clear. You can change the time delay in clearing an event, and you can specify that the management server never mark an event as clear.

To help filter events, you can have unimportant events marked as cleared rather than automatically deleted. Depending on how you configured the deletion of events, you can view the cleared events at a later time.

The default settings for clearing events are the following.

Severity Level	Default Time Delay to Clear the Event (Hours)
Unknown	1
Informational	1
Warning	1
Minor	1
Major	Never
Critical	Never

To change the default time delay before clearing an event:

1. Select **Configuration** > **Events** to access the Events page.
2. Do one of the following:
   - If you never want an event of a specified severity level marked as cleared, select the **Never** option next to the severity level in the Automatic Clear Delay column.

     Or

   - To change the delay time in clearing an event, select one of the following units of measurement from the combo box and enter the number in the adjacent box in the Automatic Clear Delay column:
     - Seconds
     - Minutes
     - Hours
     - Days
     - Weeks
3. Click **Save Changes**.

Configuring the Deletion of Events

The management server automatically deletes events after two weeks by default. For each severity level you can specify different time periods for deleting events. For example, you can modify the management server to delete events with the Information severity level every two days. You can also specify the management server to never delete events with the Critical severity level.

To change the default time delay to delete an event:

1. Select **Configuration** > **Events** to access the Events page.
2. Do one of the following:
   - If you never want an event of the specified severity level automatically deleted, select the **Never** option under the Automatic Delete Delay column.
   - To change the delay time for deleting an event, select one of the following units of measurement from the combo box and enter the number in the adjacent box:
     - Seconds
     - Minutes
     - Hours
     - Days
     - Weeks

   For example, to have events that are a week old deleted, select **Weeks** in the combo box in the Automatic Delete Delay column and enter **1**.
3. Click **Save Changes**.

Deleting Events

To delete an event, do one of the following:

- To delete one element at a time – Click the trash can icon next to the event.

  Or

- To delete specific elements – Select the events you want to delete, and then click the **Delete Selected** button at the top of the screen.

  Or

- To delete all elements – Click the **Delete All** button at the top of the screen.

Sorting Events

In Event Manager, you can sort events. For example, to see the most severe event on a page, click the Severity column header link. Click it again to sort events in the reverse order.

To sort events:

1. Access Event Manager as described in Accessing Event Manager on page 531.
2. In the Event Manager table, click the column title corresponding to the attribute you want to sort on:
   - **ID** – The identification number assigned to the event.
   - **Severity** – The severity level.
   - **Time** – The time and date the management server was aware of the event. The time in date are displayed in the following format: YYYY-DD-MM HH:MM.
   - **Element** – The source of the event. An element can be a switch, host, application, fabric, or anything else on the network.
   - **Summary Text** – A brief explanation of the event. When you click the text, the details of the event are displayed.
   - **Event Type** – Specifies whether the source of this event is an application, a host, etc.
   - **Count** – The total count of similar events.
   - **Cleared** – Indicates whether an event has been cleared.

Adding Journal Entries

While you are tracking an event, add journal entries to make others aware of what you are doing and to prevent others from repeating your steps. For example, if a host went down, you could use journal entries as a way to track your steps. Others would know what you did to get the host running. They could use this information to solve problems with other hosts.

To add a journal entry to an event:

1. Access Event Manager as described in Accessing Event Manager on page 531.
2. To access the Event Details page for an event, click the text for the event in the **Summary Text** column.
3. In the Journal Entries box, type the entry for the event. This box is limited to 4,000 characters.
4. Click **Add Journal Entry**. The entry is added with the user's account name and the date and time it was added.

Changing the CLARiiON Event Polling Interval

You can change how frequently the management server polls the CLARiiON storage systems by modifying the `cimom.ClariionEventPollInterval` property. You can change this interval if you are receiving too many “information” messages from the CLARiiON storage system.

Do not set a very long time interval. The management server does not become aware of events occurring on CLARiiON storage system until it polls the storage system. For example, if you set the polling interval to every two days, a serious issue could occur on the first day, but you would not know about it until the second day because you set such a long time.

To change the polling interval:

1. Select **Configuration** > **Product Health**, and click **Advanced** in the **Disk Space** tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the cimom.ClariionEventPollInterval property. How you copy the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, select the text, right-click the selected text, and then select **Copy**.
4. Repeat step 1 to return to the Advanced page.
5. Paste the copied text into the Custom Properties box. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the box and select **Paste**.
6. Make your changes in the Custom Properties box. Remove the hash (#) symbol in front of the property to make sure the property is not commented out.
7. Change the value assigned to the cimom.ClariionEventPollInterval property. Note that the value is in milliseconds. In the example below, the polling interval is set to 5 minutes.

   ```
   cimom.ClariionEventPollInterval=300000
   ```

8. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

While the AppStorManager service is stopped, users are not be able to access the management server, and the management server is unable to monitor elements.

Brocade Events

Brocade Switch Events

When a Brocade switch generates an event, it assigns a code instead of an event severity level to the event. The software assigns an event severity level to the event according to the event's code. This enables you to filter Brocade switch events by severity level in Event Manager, as described in the following table.

Events regarding firmware downloads are removed from the following table since the management server cannot be made aware of those events.

Events from changes in the port status now includes information on the connected device, such as device name, device type, and port name.

*Code	Event Severity Level	*Name	Description
0	Minor	EV_OBJ_ CHANGED	A physical fabric object (switch port or fabric) changed state.
1	Major	EV_OBJ_ DELETE	One or more physical fabric objects (device port, switch, or fabric) disappeared.
2	Warning	EV_OBJ_ CREATE	One or more new physical fabric objects (device port, switch, or fabric) appeared.
3	Critical	EV_ CONNECTED_ OBJECT_ OFFLINE	A device connected to a switch went offline.
4	Major	EV_ CONNECTED_ OBJECT_ ONLINE	A device connected to a switch has come online.
5	Info	EV_RSCN	For those RSCN events not covered by EV_OBJ_ xxx codes listed above; for example, fabric, domain, connected area state unknown, connected device state unknown.
7	Minor	EV_API_ HEART_ CONDITION	Monitoring proxy switch.
11	Major	EV_STATE_ CHANGE	State changes such as login, logout, login failed, configuration change, track on/off, port up/down, fabric segmentation, security violations, zone change.
12	Major	EV_ PLATFORM_ CHANGE	The platform database changed.

*This term does not appear in the event description, but is provided for clarity.

Supported Brocade Events

Event Manager displays the following events from Brocade switches:

- **RSCN** – Events about the state of the switch, such as its being offline.
- **Fabric Access library** – Events about proxy switch health.
- **Zoning** – Events about zoning, such as zone-related state change notification.

Cisco Switch Events

Events from changes in the port status include information on the connected device, such as device name, device type, and port name.

Event Severity Level	Name	Description
Warning	fcTrunkIfDownNotify	A trunk interface status goes down.
Information	fcTrunkIfUpNotify	A trunk interface status goes up.
Warning	vsanStatusChange	A VSAN status goes down.
Information		A VSAN status goes up.
Information	vsanPortMembershipChange	A port is added to VSAN.
Information	dmNewPrincipalSwitchNotify	A new principal switch is elected in VSAN.
Information	fcMgmtNotification	A port in a switch is online.
Warning	A port status in a switch is offline.	
Information	cieLinkUp	A communication link on a FC port comes up; for example, a connected device is switched on.
Information	cieLinkDown	A communication link on a FC port goes down.
Information	zoneActivateNotify	A zone set is activated on a VSAN.
Information		A zone set is deactivated on a VSAN.
Warning		A zone set activation fails on a VSAN.
Warning		A zone set deactivation fails on a VSAN.
Information		A zone set activation or deactivation is in progress.
Information	ciscoFeatureOpStatusChange	A switch feature is enabled or disabled.

Event Severity Level	Name	Description
Major	cefcPowerStatusChange	A switch FRU is powered off either administratively or due to insufficient system power, power translation errors, temperature problems, etc.
Information		A switch FRU is powered on.
Major		A switch FRU is in failed state.
Warning		A switch FRU is powered on, but the fan failed.
Information	linkDown	Link Down generic SNMP Trap handled to track link status of ethernet and management interfaces is down.
Information	linkUp	Link Up generic SNMP Trap handled to track link status of ethernet and management interfaces is up.

Filtering Events

Setting Up a Filter

The management server provides several types of event filters that enable you to specify which events you want Event Manager to display.

You can use all the event filters at once or just one. You can filter events by:

- **Time period**
- **Severity level**
- **Element type**
- **Summary text**
- **Element name**
- **Cleared status**
- **Specific element**

To set up a filter:

1. To access the filter for Event Manager, click the Filter heading at the top of the page.

   Figure 7 Accessing the Filter Feature

   

   The filtering feature appears.

   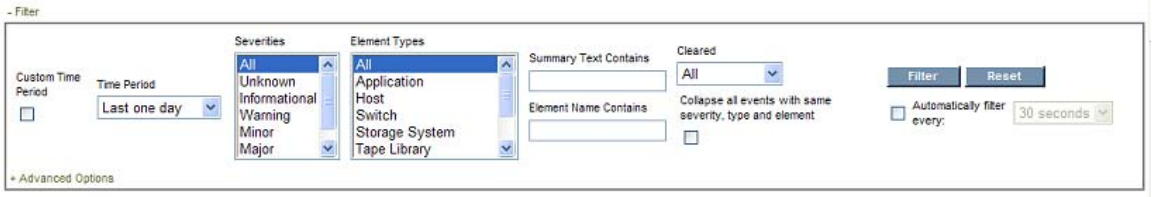

2. Select a time from the **Time Period** combo box. You can also select a customized time as described in Selecting a Custom Time Period on page 548.

   Figure 8 Selecting a Time Period

   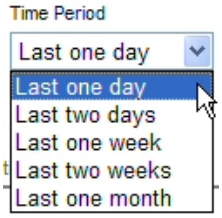

3. Select which events of a severity type to display. Use the Control and Shift keys to select multiple severities. The following options are provided:
   - **All** – Events of all severities are displayed.
   - **Unknown** – Only events of severity type unknown are displayed.
   - **Informational** – Only events of severity type informational are displayed.
   - **Warning** – Only events of severity type warning are displayed.
   - **Minor** – Only events of severity type minor are displayed.
   - **Major** – Only events of severity type major are displayed.
   - **Critical** – Only events of severity type critical are displayed.
4. Select which events of an element type to display. Use the Control and Shift keys to select

multiple severities. The following options are provided:

- **All**
- **Application**
- **Host**
- **Switch**
- **Storage System**
- **Tape Library**
- **Fabric**
- **Other**
- **HP Storage Essentials**

5. To set the filter by summary text, enter the text you want used for the filter in the **Summary Text Contains** field.

   You can use this option when you see several events that span over several elements or severity levels. For example, to determine if someone else is logging into the management server as admin, you can find how often the admin user logged into the management server over the past few days by entering the following text in the Summary Text Contains field:

   ```
   User admin logged in
   ```

6. To set the filter by element name, enter text in the **Element Name Contains** field.

   This feature can be helpful if you are interested in events from elements that have similar names. For example, say you have a naming convention for hosts, where all hosts belonging to the engineering group begin with engineering. You could enter "engineering" in the Element Name Contains field so that Event Manager would display only events from those hosts.

7. To set a filter by cleared status, select one or more of the following using the Control and Shift keys:

   - **All** – All events, regardless of their clear status
   - **All But Clear** – All events, except for those marked clear
   - **Clear** – Only events marked clear

8. To have events merged together, select **Collapse all events with same severity, type and element**.

   This feature is very useful when numerous events of the same severity, type, and element are being picked up by HP Storage Essentials. When you select this feature, HP Storage Essentials displays just unique events with a total count in front of them instead of listing all the events individually.

9. To have Event Manager filter elements automatically at a set interval, select the **Automatically filter every** option and one of the following:

- 30 seconds
- 1 minute
- 5 minutes
- 10 minutes

10. When done setting your options, click the **Filter** button.

You can update the events displayed by clicking the **Filter** button.

Selecting a Custom Time Period

You might want to select a custom time period for the filter when troubleshooting an issue.

To select a custom time period:

1. Select the **Custom Time Period** option.

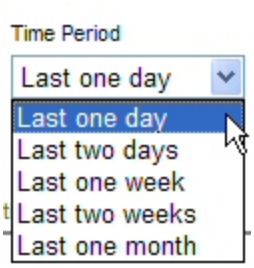

2. Event Manager only displays events that occur after the start date and time and before the end date and time:
   a. Click the Calendar icon to the right of the **Start Date** field.

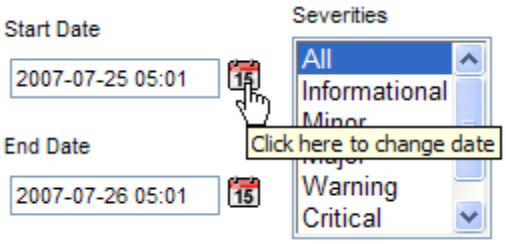

   b. In the Calendar, select the start date.

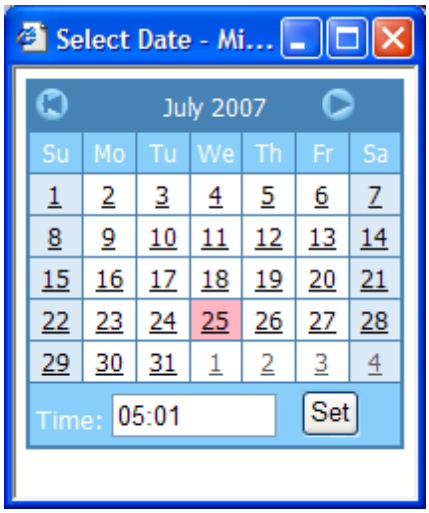

c. In the **Time** field, enter the start time. The time is based on a 24-hour clock. For example, to have Event Manager display events occurring after 8 p.m. on the specified date, enter 20:00.

d. When done setting the start time and date, click **Set**.

3. Set the end date and time for the filter.

   Event Manager only displays events that occur after the start date and time and before the end date and time.

   a. Click the Calendar icon to the right of the **End Date** field.

   b. In the Calendar, select the end date.

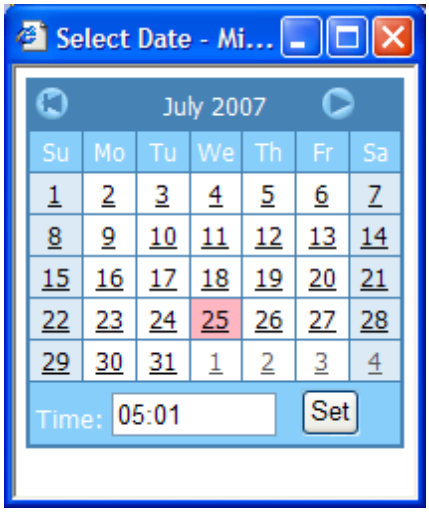

c. In the **Time** field, enter the end time. The time is based on a 24-hour clock. For example, to have Event Manager display events occurring before 8 p.m. on the specified date, enter 20:00.

d. When done setting the end time and date, click **Set**.

Resetting a Filter

To revert to the default settings for the filter:

1. To access the filter for Event Manager, click the **Filter** heading at the top of the page, as shown in the following figure.

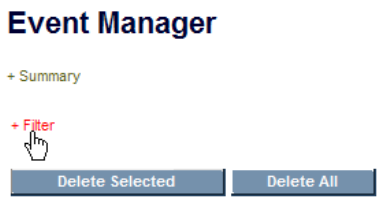

The filtering feature appears.

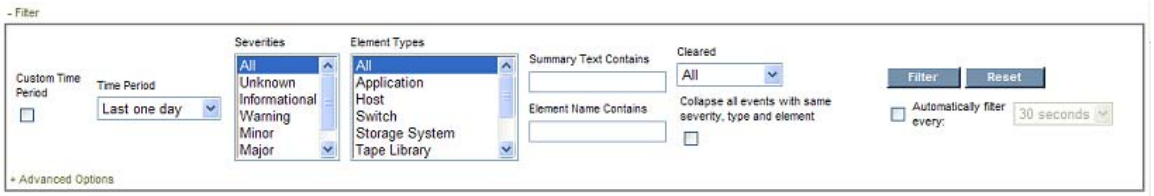

2. Click the **Reset** button. The filter is reset.

Setting Up Advanced Filtering

Event Manager has an advanced filtering feature that lets you provide detailed information for filtering.

To set up advanced filtering:

1. To access the filter for Event Manager, click the **Filter** heading at the top of the page.

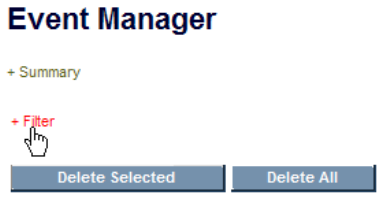

The filtering feature appears.

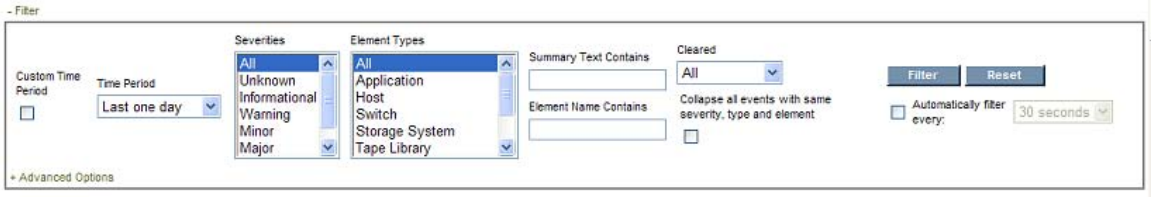

2. Expand the **Advanced Options** heading.

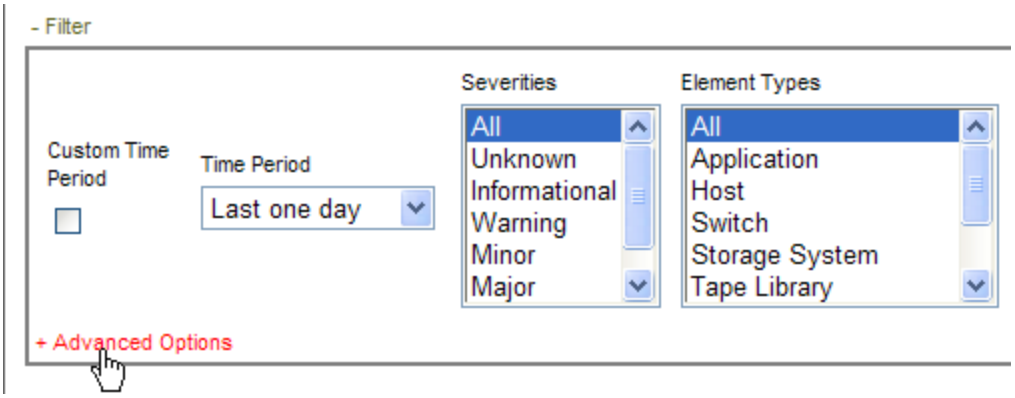

Once the Advanced Options heading has been expanded, the Element Name Contains field becomes inactive. Any values entered in the Element Name Contains field are ignored as long as the Advanced Options heading is expanded.

3. Click the **Add** button.

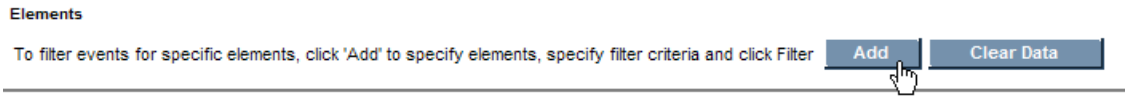

4. Select the element type you want to add and click **Next**.
5. Select one or more elements to add to the filter and click **Next**.

   Only the elements of the type you specified in the previous window are displayed. For example, if you selected hosts in the previous window, only hosts are displayed in this screen.

   If the element you are looking for does not appear on the first page of the table, use the navigation tools at the top of the table to page through the list of elements.

   The hosts you select are listed in the navigation filter. For example, the following figure shows that Host_13081 and Host_10380 were selected for advanced filtering, so they are listed under the Advanced Filtering heading.

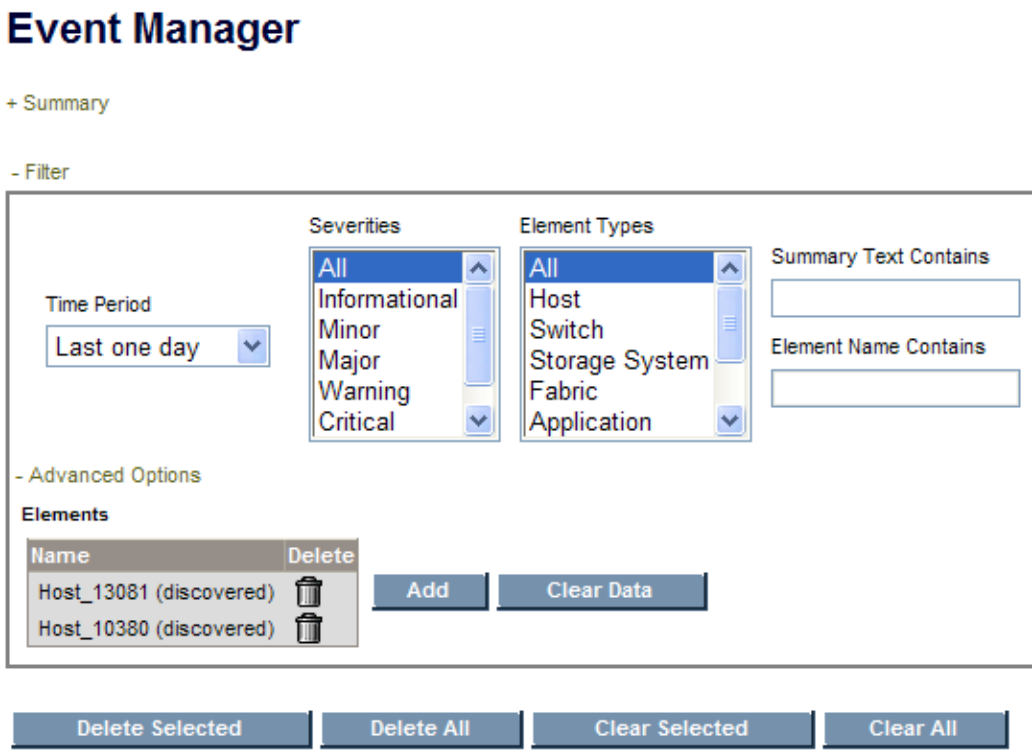

6. Verify that the **Advanced Options** heading is expanded; otherwise, advanced filtering will not work.

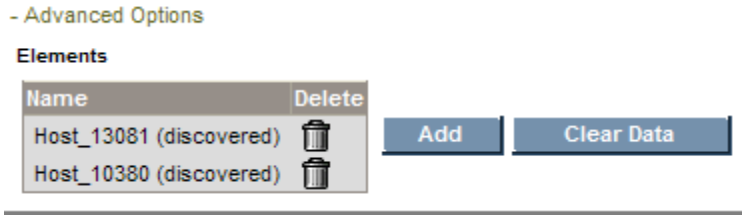

7. Click the **Filter** button. Event Manager displays the elements specified under **Advanced Options**.

Clearing Advance Filtering Options

You can clear the filtering set for advanced options, by clicking the **Clear Data** button under the Advanced Options heading.

13 Finding an Element's Storage Capacity

This section contains the following topics:

Capacity Manager

Depending on your license, Capacity Manager might not be available. See the List of Features to determine if you have access to Capacity Manager. The List of Features is accessible from the Documentation Center (**Help** > **Documentation Center**).

Capacity Manager provides a graphical representation of an element's storage capacity in the storage network with the following tabs:

- List
- Path
- Capacity Data
- Capacity Chart

Capacity Manager provides a different set of information for NetApp NAS devices. For more information, see Capacity Information for NetApp NAS Devices on page 561.

Capacity Definitions

Capacity Manager uses the following capacity definitions:

- **Partition Capacity** – The sum of the disk partitions or volume manager volume capacities used by applications and virtual machines.
- **Total Capacity** – The sum of the total storage of each logical disk + partition capacity.
- **Total Used Capacity** – The sum of used storage of logical disks of the host + partition capacity.
- **Available Capacity** – The unused capacity form the logical disks.

- **Unmounted Volume Capacity (Reachable Capacity)** – The sum of storage on unmounted paths (storage mapped from arrays to HBA ports, but not used as a partition, volume manager volume, or logical disk).

Keep in mind the following:

- The Capacity Manager displays the total capacity of an application, including the network drives. If you look at the capacity of the application in Chargeback Manager, the capacity differs. Chargeback Manager provides only network capacity with the exception of Windows 2000 hosts. For more information, see How Capacity Differs in Chargeback Manager and Capacity Manager on page 833.
- Volume names from ambiguous automounts on Solaris hosts are not displayed in Capacity Manager. See Volume Names from Ambiguous Automounts Are Not Displayed on page 854.
- Capacity Manager takes extra time to load the first time you access it after restarting the management server. The extra time is required for the management server to calculate the element capacity data. Capacity Manager loads faster during subsequent times because the element capacity data was already calculated with the exception of Windows 2000 hosts.
- After discovering new elements, the capacity data for those elements shows up as null until the data collectors run. Data collectors are set to run every 15 minutes by default.
- For virtual servers, Capacity Manager only displays datastore volumes. Raw disk partitions are not displayed.

List Tab

The List tab enables you to quickly access an element. For example, to quickly access a host in the topology, expand the List View Tree node, and then select your host in the tree by expanding the All Elements and Hosts nodes. When you select an element in the tree, it is highlighted in the topology.

Path Tab

The Path tab provides information about an element's path. By clicking a host's node, you can determine the host's path in the application. When you click a host node in the tree, the elements in the host's path appear highlighted in the right pane.

Capacity Data Tab

The Capacity Data Tab provides information about the capacity of an element. You can quickly determine the general capacity of an element by its colors.

Color Coding for Capacity Manager

Color	Description
	Description varies according to element type: • For storage systems – The space is mapped. • For all elements except storage systems – The space is used.
	Unallocated
	Unused raw (storage systems only)
	Unmapped (storage systems only)

The colors indicate that the element displayed in the following figure is about 75 percent available. The rest of it is being used.

Figure 9 Capacity of an Element

You can obtain more detailed information about an element by clicking it in the right pane or in the Capacity Manager tree, as explained in Finding the Capacity of an Element on page 557.

Capacity Chart Tab

The Capacity Chart tab enables you to create bar or line charts to view your capacity data. You can use these charts to display trending information. For more information, see Viewing Capacity Charts on page 574 and Viewing Trending Information for Storage Capacity on page 576.

Accessing Capacity Manager

To access Capacity Manager, click **Capacity Manager** ().

Toolbars in Capacity Manager

Capacity Manager provides two toolbars: one in the upper pane and another in the lower pane. The toolbar in the upper pane is the same as the one in System Manager. For information about the toolbar in the upper pane, seeToolbar in System Manager on page 417.

When the Capacity Chart tab is active, the toolbar in the lower pane provides the following information.

Capacity Manager provides additional tabs for NetApp NAS devices. The toolbars are the same.

Icon	Description
	Enables you to switch between the following: • Bar – Displays each data point as a bar. The data for the different elements is displayed side by side. • Line – Displays each data point as a dot with a line connected to the previous data points. The data for the different elements for a specific point in time is displayed in the same column.
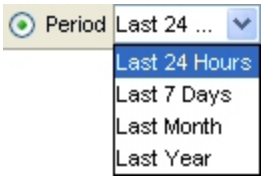	Enables you to format the graph to provide data within the time period specified. Select the option to the left of the Period combo box. Select one of the following from the menu. To update the chart, click the  button. • Last 24 Hours – Information collected in the last 24 hours is reported. This option is only available to hosts and applications. • Last 7 Days – Information collected in the last seven days is reported. • Last Month – Information collected in the last month is reported. • Last Year – Information collected in the last year is reported.
	Enables you to format the graph to provide data within the starting and ending time specified.

Icon	Description
	*Applications and hosts only*: Enables you to change the display frequency. The options are: • Hourly – The information is displayed in hourly increments. • Daily – The information is displayed in daily increments. • Weekly – The information is displayed in weekly increments. • Monthly – The information is displayed in monthly increments.
	Enables you to set trending information. When a switch or storage system is selected, the frequency box is set to hourly. See Viewing Trending Information for Storage Capacity on page 576. Keep in mind the following: • An element's performance can drastically change in the future. Keep in mind that the data trends are just assumptions and should not be treated as fact. • Trending requires at least two sets of data gathered within the frequency specified.
	Applies period, frequency, and trending information.
	Enables you to filter out the additional data series if the chart contains more than one series of data.
	Enables you to print a graph.

Finding the Capacity of an Element

This section contains the following topics:

Overview

Capacity Manager rounds the data it displays. As a result, the totals you add for a property might be different from the data displayed in the Summary column. For example, if you add the total capacity from each data pool and compare that total to the number for Total Capacity displayed under the Summary heading, you will most likely find that the results differ.

To find the capacity of an element:

1. Access Capacity Manager as described in Accessing Capacity Manager on page 555.
2. (*Optional*) To quickly view the capacity of all of the elements in a fabric or application path, click the fabric or application path displayed in the tree for Capacity Manager, as shown here:

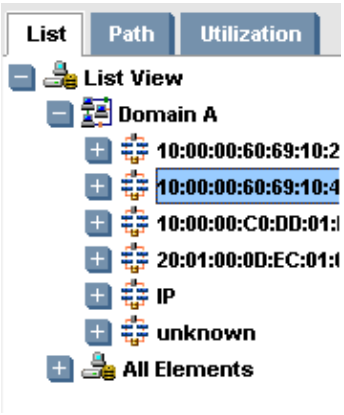

   For an explanation of the colors displayed, see Capacity Manager on page 553.
3. Do one of the following:
   a. Click an element in the right pane.

      Or

   b. Click an element in the tree for Capacity Manager.

   A pane appears along the bottom of the page displaying the capacity information.

Capacity Information for Applications

Capacity Manager displays the following information under the Capacity Summary and Capacity columns:

- **Total Capacity**
- **Total Used**
- **Available**
- **Percentage Used** – The percentage used compared to the total capacity of the storage groups or database files.

The following additional information is displayed for each storage group (Microsoft Exchange) or database file (Oracle):

- **Total Capacity**
- **Total Used**
- **Available**
- **Percentage Used** – The percentage used compared to the total capacity of the storage group or database file.

Figure 10 Capacity Information for Microsoft Exchange 2003

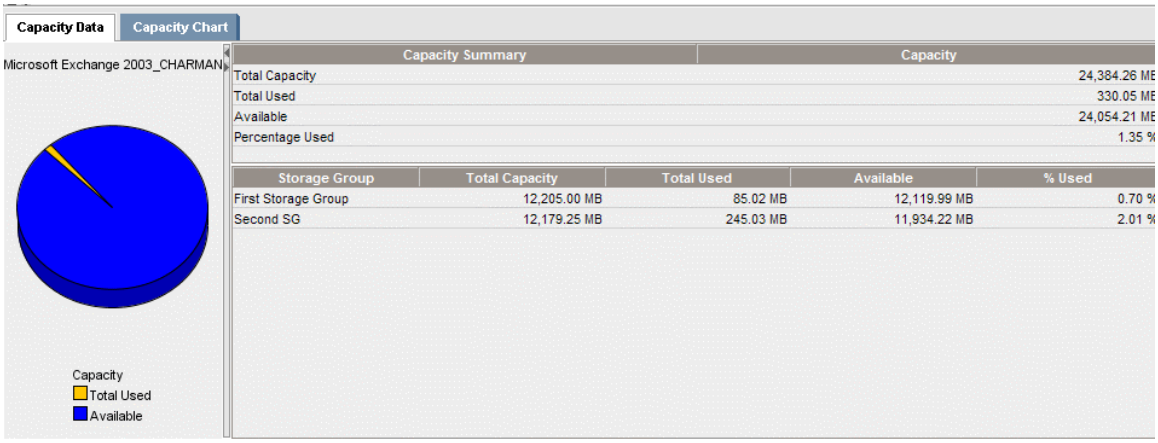

For Microsoft Exchange 2007, the Capacity Manager displays only the Total Used capacity for each storage group.

Figure 11 Capacity Information for Microsoft Exchange 2007

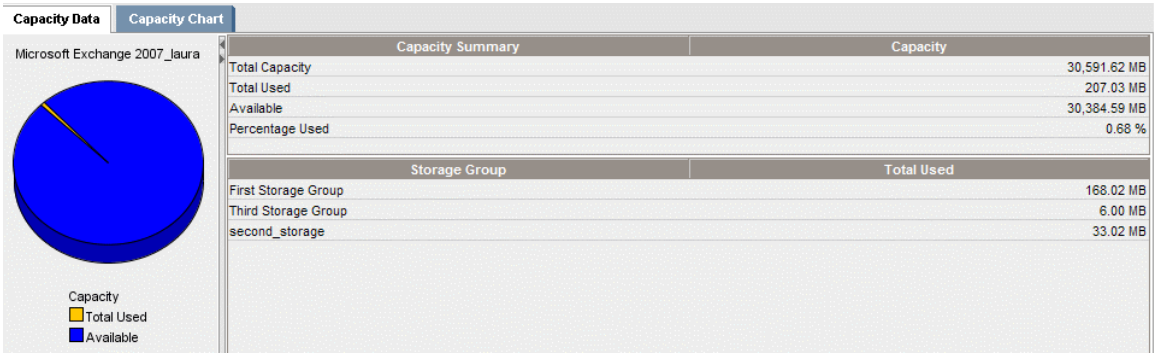

Capacity Information for Hosts

Capacity Manager displays the following information under the Capacity Summary and Capacity columns:

- **Total Capacity**
- **Total Used**
- **Available**
- **Unmounted Volume** – The amount in gigabytes of unmounted storage.

  This box automatically displays 0 GB, if you discovered the host but not the storage system connected to it. This might occur if you did not enter the IP address of the storage system when performing discovery, or your license does not allow you to discover a particular storage system. The List of Features is accessible from the Documentation Center (**Help** > **Documentation Center**).
- **Percentage Used** – The percentage used compared to the total capacity of the storage volumes.

The following additional information is displayed for each volume:

- **Total Capacity**
- **Total Used**
- **Available**
- **Percentage Used** – The percentage used compared to the total capacity of the storage volume.

By default, raw partitions are not used in capacity calculations on the host. When you assign the raw partition to an application, the raw partition is seen as being in use, and it will then be used in calculating capacity. When you create a virtual application, it is possible to select disk partitions such as /dev/sda1 in addition to the mounted file systems comprising that disk partition. If you select both when creating the virtual application, the virtual application's usage information will reflect both the mounted file system and the disk partition, in effect double-counting. Use only the mounted file systems when creating or editing a virtual application.

Capacity Information for NetApp NAS Devices

Capacity Manager displays the following information under the Capacity Summary and Capacity columns on the Volume Data tab:

- Total Aggregate Available
- Total Aggregate Used
- Total Volume Maximum
- Total Volume Used
- Percentage Total Volume Used

Capacity Manager displays the following information about each volume:

- Volume Name
- Aggregate Available
- Aggregate Used
- Volume Maximum
- Volume Used
- Percentage Volume Used

Capacity Manager displays the following information under the Capacity Summary and Capacity columns on the Aggregate Data tab:

Aggregate tabs only apply to NetApps hosts.

- Aggregate Name
- Total Allocated Capacity
- Total Available Capacity
- Total Used Capacity
- Total Reserved Capacity
- Percentage Used

Capacity Manager displays the following information about each aggregate:

- Allocated Capacity
- Available Capacity
- Used Capacity
- Reserved Capacity
- Percentage Used

Capacity Manager displays the following information under the Capacity Summary and Capacity columns on the Quota Data tab:

- Total Disk Limit
- Total Disk Used
- Percentage Disk Used
- Total File Limit
- Total File Used
- Percentage File Used

Capacity Manager displays the following information about each disk or file:

- Quota Name
- Quota Type – There are two types of quotas: disk (space) and file (count).
- Quota Limit – The amount of disk space or the number of files reserved for the target.
- Threshold – The amount of disk space that would have to be exceeded before a message is logged.
- Quota Soft Limit – The amount of disk space or the number of files that would have to be exceeded before a message is logged and an SNMP trap is generated.
- Quota Used
- Percentage Used

Capacity Manager displays the following options under the Capacity Summary and Capacity columns on the Snapshot Data tab.

Snapshot tabs only apply to NetApps hosts.

- Total Volume Space – Sum of space on all volumes.
- Total Space Reserved – Sum of the space reserved for snapshots across all volumes.
- Total Space Used – Sum of space used by snapshots on all volumes.

  Total Space Used is the sum of the largest total space used on each volume.
- Cumulative% of Total Vol – Percentage of space used for snapshots from the total volume space across all volumes.
- Cumulative% of Used Vol – Percentage of space used for snapshots from the total used space across all volumes.

Capacity Manager displays the following information about each volume snapshot:

- Snapshot Name
- % Reserved
- % of Total Vol

- % of Used Vol
- Cumulative% of Total Vol
- Cumulative% of Used Vol
- Total Space Used

Capacity Information for Celerra

Capacity Manager displays the Capacity Summary for file systems and the following detailed capacity information for each file system.

- **Total Capacity**
- **Total Used**
- **Available**
- **Unmounted Volume**(Ignore this information as it always displays a zero value)
- **Percentage Used**

Capacity Information for Centera

Capacity Managerdisplays the following information under the Capacity summary and Capacity columns on Centera Cluster Data tab:

- **Total Centera Cluster Capacity** - The total capacity of the cluster.
- **Total Centera Cluster Capacity Used** - The total used capacity of the cluster, or the capacity that is not available to store data. This includes the capacity reserved for system resources, but not assigned for storage or off-line capacity. The capacity actually used to store data and associated audit and metadata.
- **Total Centera Cluster Capacity Free** - The capacity that is free and available for storing data, or for self-healing operations in case of a disk or node failure, or for database growth of failover.

Capacity Manager displays the following information about each virtual pool:

- **Pool Name**
- **Total Capacity** - Quota of the virtual pool
- **Used Capacity**
- **Free Capacity**
- **Percentage Free**

For some pools, the values for Total Capacity and Free Capacity are shown as -1. These values are replaced by that of Used Capacity.

Capacity Manager displays the following information under the Capacity summary and Capacity Columns on the Centera Node Data tab:

- **Total Centera Cluster Capacity** - The sum of the total capacities of all Centera nodes.
- **Total Centera Cluster Capacity Used** - The sum of the total capacities of all Centera nodes.
- **Total Centera Cluster Capacity Free** - The sum of the free capacities of all Centera nodes.

Capacity Management Information for HP P4000 Devices

Using Capacity Manager, you can access the statistical data collected from an HP P4000 device. HP Storage Essentials collects volume and pool data that is summarized at the end of each Get Details collection.

There are two main view of capacity information: Post-RAID physical and Post-RAID virtual.

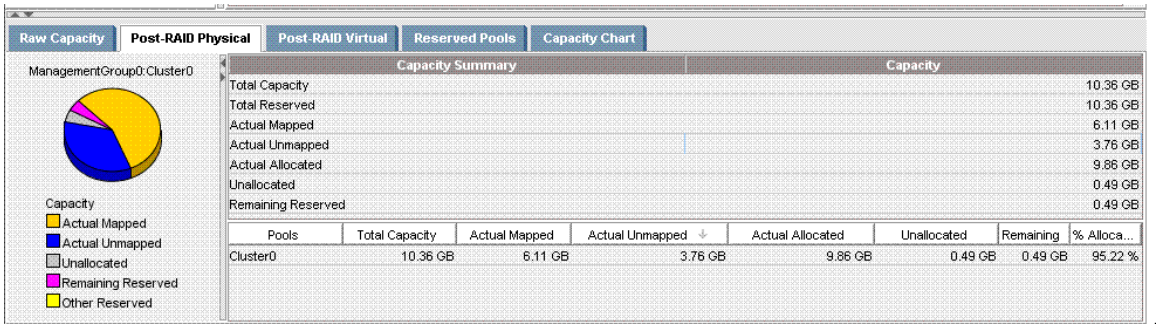

Cluster0 represents the pool of storage available on the cluster that all volumes and snapshots are created from. Mapped/Unmapped values refer to whether or not the volume was assigned to a server. Unallocated shows the remaining physical space on the cluster that is available for provisioning.

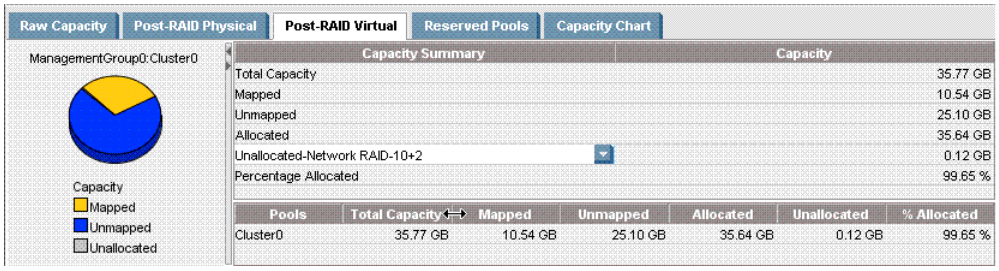

Post-RAID virtual provides a view of the storage from a presented storage or carved volume size. Combining this view with the previous view shows that this cluster is over-allocated.

The amount of over-allocation (25.29 GB) is called out in the Reserved Pools tab.

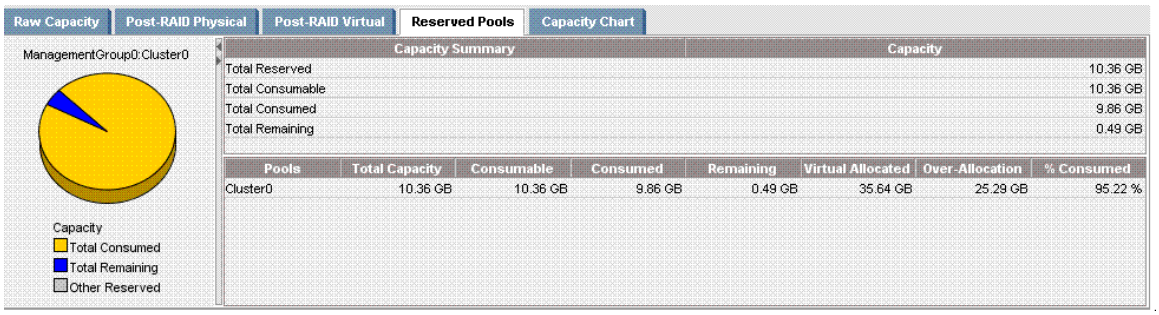

You can obtain historical trending information by viewing the Capacity chart:

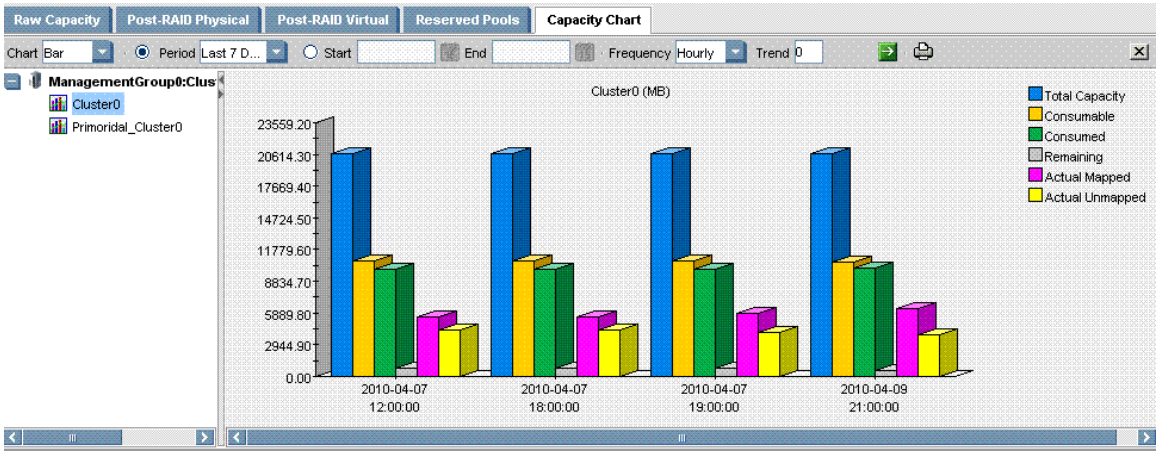

This information can be useful in tracking the remaining capacity of the cluster over a period of time.

Capacity Information for HP X9000 Network Storage Devices

Using Capacity Manager, you can access information on the capacity utilization for an HP X9000 device.

Capacity Manager displays the following information under the Capacity Summary column for the device:

- **Total Capacity** - The sum of the sizes of the physical volume.
- **Total Used Capacity** - The sum of the sizes of the physical volume that are a part of any volume group.
- **Total Available Capacity** - The sum of the sizes of the physical volumes that are not a part of any volume group.

Capacity Manager displays the following information for file systems:

- **Total Capacity**
- **Total Used Capacity**
- **Free Capacity**
- **Available Capacity**

Capacity Manager displays the following information for a file server node:

- **Total Capacity** - Sum of the sizes of the disk drives owned by the node.
- **Used Capacity** - Sum of the sizes of the disk drives owned by the node which are part of any volume group.
- **Available Capacity** - Sum of the sizes of the disk drives owned by the node which are not a part of any volume group.

Capacity Information for Storage Systems

Capacity Manager provides additional capacity information for storage systems that support thin provisioning. Capacity information for thin provisioning is provided in additional tabs that are displayed only for storage systems that support thin provisioning.

You can view statistics for only internal or external storage by selecting Internal or External from the Display drop-down menu. This information is only shown if the storage system supports both internal and external storage.

Storage system supports thin provisioning?	Where to find capacity information
Yes	About Thin Provisioning for Storage Systems below
No	Capacity Information for StorageSystems that Do Not Support Thin Provisioning on page 569

About Thin Provisioning for Storage Systems

For storage systems that support thin provisioning, post-RAID information is separated into three tabs: Post-RAID Logical, Post-RAID Allocated, and Post-RAID Usage.

Post-RAID Logical Tab

The Post-RAID Logical tab displays the capacity as seen by the host. The following information is provided on the Post-RAID Logical tab for all storage pools.Shows the logical side of the volumes which have been split into the following:

- Mapped – Sum of the volumes visible to hosts. For a volume to be mapped, it must be have a logical mapping to at least one host initiator.
- Unmapped – Sum of the volumes not visible to hosts. An unmapped volume is storage committed as a single volume but not visible or potentially visible to any initiator.
- Allocated – Sum of the logical side of the volumes allocated from the storage pools. Allocated is the sum of Mapped and Unmapped.

The following information is provided for each storage pool:

- Mapped – Sum of the volumes visible to hosts. For a volume to be mapped it must be have a logical mapping to at least one host initiator.
- Unmapped – Sum of the volumes not visible to hosts. An unmapped volume is storage committed as a single volume but not visible or potentially visible to any initiator.
- Allocated – Sum of the logical side of the volumes allocated from the storage pool. Allocated is the sum of Mapped and Unmapped.

Thin Provisioning for an SVC Array Displayed on the Post-RAID Logical Tab

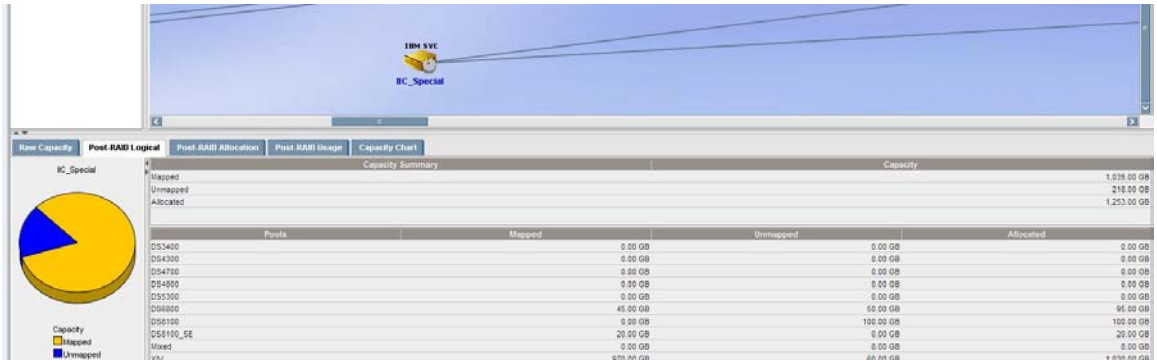

Post-RAID Allocation Tab

The following information is provided on the Post-RAID Allocation tab for all the storage pools on the array:

- Total Capacity – Sum of the physical capacity of all the configured storage pools in the array.
- Unallocated – Sum of the capacities in all storage pools available for creating volumes.

  The drop-down box provides the following information:

  - Network RAID-0
  - Network RAID-10
  - Network RAID-10+1
  - Network RAID-10+2
  - Network RAID-5
  - Network RAID-6

  At this time, RAID-5 and RAID-6 information is only applicable for P4000 v8.5 and later.
- Actual Allocated – Sum of what is physically allocated in all storage pools. Anything allocated cannot be used for creating volumes.
- Percentage Allocated – The percentage of Actual Allocated to Total Capacity of the storage pools.
- *Total Reserved Capacity – Sum of the capacity of the reserved pools, such as thinly provisioned and snapshot pools on the XP array.
- *Total Reserved Allocated – Sum of what is allocated for the reserved pools.

*These values are only displayed if the storage system contains pools that are marked as "reserved," such as thinly provisioned (THP Pools) and snapshot pools.

The following information is provided for each pool on the array:

- Total Capacity – Sum of the capacity on the pool, which includes physically allocated and unallocated capacities. For arrays that support virtualization, Total Capacity includes virtualized external storage.

- Unallocated – Sum of the capacities in the storage pool available for creating volumes.
- Actual Allocated – Sum of the physical capacity that has been allocated to the use of storage volumes.

  For pools dedicated to thin provisioning (THP pools), Actual Allocated represents the total space available for consumption by thinly provisioned volumes.

  For SVC arrays, Actual Allocated displays the same value as shown as real capacity in the native tool.
- % Allocated - The percentage of Actual Allocated to Total Capacity of the storage pool.
- Virtual Allocated - Sum of what is virtually allocated for the pool.
- Over Allocation - Difference between virtual and physical allocation. Physical allocation is the sum of Actual Allocated and Unallocated. This shows the amount of which a storage pool is over or under subscribed.
- % Over Allocated - Percentage of over allocation. If the percentage is over 100 percent, the storage is over allocated. If the percentage is under 100 percent, the storage is under allocated.

Thin Provisioning for an SVC Array Displayed on the Post-RAID Allocation Tab

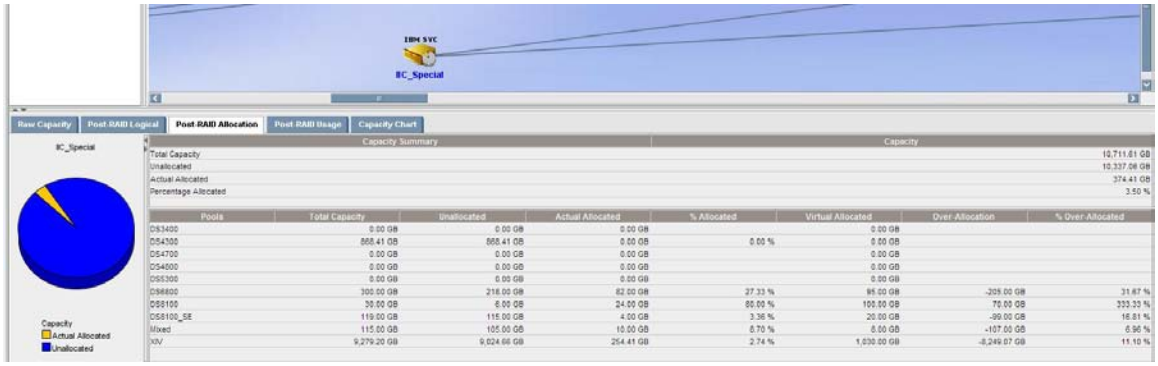

Post-RAID Usage Tab

The Post-RAID Usage tab shows the usage of what is physically allocated. The Post-RAID Usage tab provides the following information for all storage pools:

- Actual Allocated - Sum of the physical capacity that has been allocated to the use of storage volumes.
- Actual Used - Sum of the capacity actually used by the volumes.
- Actual Unused - Actual capacity not used.
- Actual Percentage Used - Percentage of what is used out of Actual Allocated.
- *Total Reserved Allocated - Sum of the reserved pools that are allocated.
- *Total Reserved Used - Sum of the reserved pools that are used.

- *Total Reserved Unused - Sum of the reserved pools that are unused.
- *Percentage Reserved Used - Percentage of the reserved pools that are used.

*These values are only displayed if the storage system contains pools that are marked as "reserved," such as thinly provisioned (THP Pools) and snapshot pools.

The following information is provided for each storage pool:

- Actual Allocated - Sum of the physical capacity that has been allocated to the use of the storage volume.
- Actual Used - Sum of the capacity actually used by the volume.
- Actual Unused - Actual capacity not used.
- % Used - Percentage of what is used out of Actual Allocated.

Thin Provisioning for an SVC Array Displayed on the Post-RAID Usage Tab

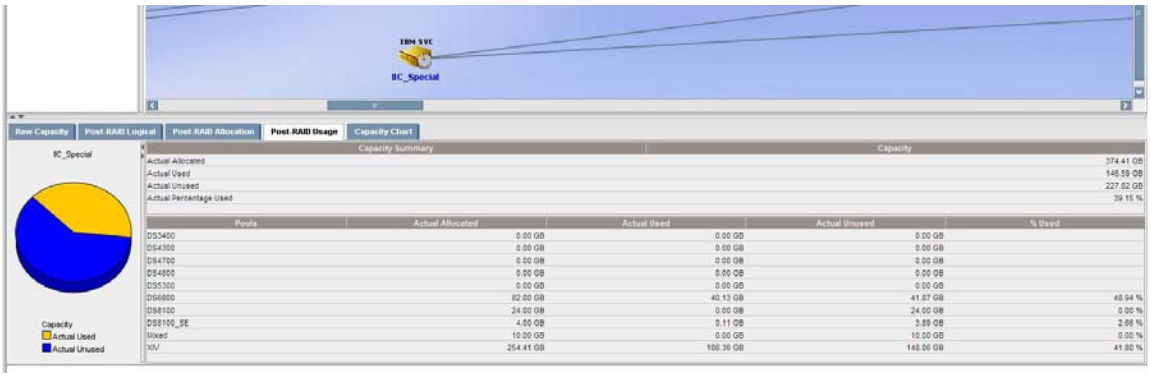

Capacity Information for StorageSystems that Do Not Support Thin Provisioning

If a storage system does not support thin provisioning, Capacity Manager displays the following tabs for the storage system:

- Raw Capacity – See Viewing the Raw Capacity of a Storage System below.
- Post-RAID – See Viewing Post-RAID Information on next page.

Viewing the Raw Capacity of a Storage System

The Raw Capacity tab displays the raw capacity in gigabytes for the ports.

- Total Raw – The sum of used raw and unused raw capacity.
- Used Raw – Raw disk capacity consumed by RAID groups or other such disk groups on the array. Disks configured for use in provisioning volumes, regardless of whether volumes were allocated from those disk groups. For example, on CLARiiON storage systems these are disks that are in RAID groups.
- Unused Raw – Raw disk capacity not currently configured into any kind of RAID or disk group. Disks that were not configured for use in provisioning volumes. For example, on the EVA or

MSA 1000/1500, P2000 G2 (2312fc/2324fc), and P2000 G3 FC, these are not part of any volume group. For enterprise arrays that are preconfigured at installation, Unused Raw is typically zero.

- Percentage Used – The percentage of raw disk capacity used.

The raw capacity values come directly from the SMI instrumentation of storage arrays, where raw capacity is modeled as primordial storage pools.

The Raw Capacity tab provides information about the raw capacity used on a storage system. A graphic displays which percentage of the storage system has used raw capacity and unused raw capacity.

Current Primordial Pool Capacity Calculations for EMC Symmetrix Arrays

The current displayed primordial pool capacity does not include any configured save devices. Save devices are not host mappable and, therefore, are not included in the primordial pool capacity calculations, which include only user visible devices.

When calculating the total capacities, include the capacity for the configured save devices. Save devices can be configured in several RAID formats, such as Mirror 2, RAID 5 (3D+1).

The appropriate conversion factor should be applied to the total capacity as recommended by EMC. For example, if you are using RAID 5 (3D+1) with no mirroring, you need 25 percent more space to account for overhead. And if you are using Mirror 2, you must multiply the space required by two. Without the save devices, if configured, the total pool capacities do not add to primordial pool capacity plus the configured save devices capacity.

Calculating Capacity for CX Arrays

Capacity for CX arrays is calculated as follows:

- Total Capacity – Sum of Pool Total Capacity rounded to two decimal points.
- Used Capacity – Total Capacity-Available Capacity rounded to two decimal points. Pools defined as Hot Spare were altered to reflect 100 percent used as they cannot be used for user space.
- Available Capacity – Total Pool Available Capacity.

Different terminology is used in the CLARiiON Free Capacity report:

- In HP Storage Essentials Total Capacity=Logical GB
- In HP Storage Essentials Available = Free GB

Viewing Post-RAID Information

Post-RAID information for storage systems that do not support thin provisioning can be viewed from the Post-RAID tab under the Capacity Summary and Capacity columns:

- Total Capacity – Sum of the storage pools in the array, which includes allocated and unallocated capacities. For arrays that support virtualization, Total Capacity includes virtualized external storage.

- Total Internal – Sum of the storage pools that are internal to the array. Does not include virtualized external storage. This statistic only displays for arrays that support virtualization.
- Mapped – Sum of the mapped volume capacities allocated from storage pools. For a volume to be mapped it must be have a logical mapping to at least one host initiator.
- Unmapped – Sum of the unmapped volume capacities allocated from storage pools. An unmapped volume is storage committed as a single volume but not visible or potentially visible to any initiator.
- Total Allocated – Sum of the capacities of the volumes allocated from storage pools.
- Unallocated – Sum of the capacities in storage pools available for creating volumes.

  Primordial pools are not available for creating volumes, so they do not contribute to this total.
- Percentage Allocated – Percentage of the total capacity that has been allocated.

  For arrays that permit RAID choice when creating volumes (for example the EVA), the concept of flexible RAID Pools is introduced. In this case, the amount of unallocated space depends on the RAID level chosen for new volumes. The values for Unallocated and Total Capacity displayed can be modified dynamically by choosing a RAID level for unallocated space.

Post-RAID Tab

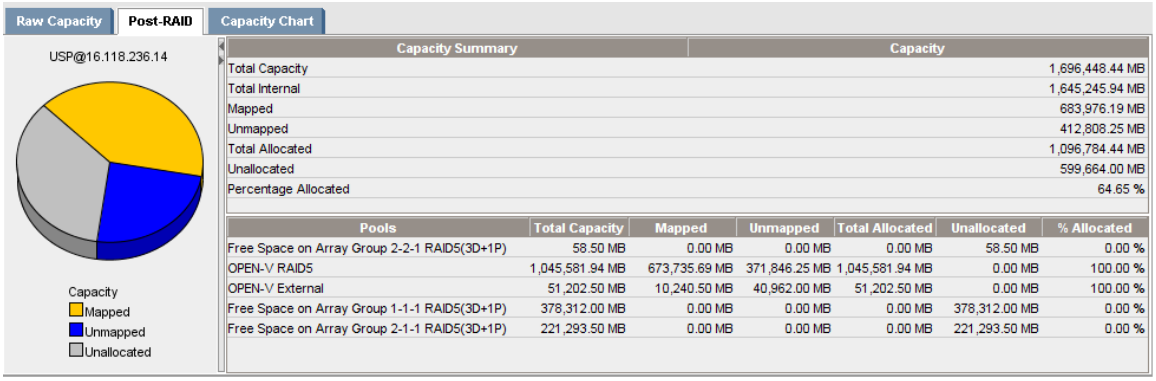

Capacity Summary	Capacity
Total Capacity	1,696,448.44 MB
Total Internal	1,645,245.94 MB
Mapped	683,976.19 MB
Unmapped	412,808.25 MB
Total Allocated	1,096,784.44 MB
Unallocated	599,664.00 MB
Percentage Allocated	64.65 %

Pools	Total Capacity	Mapped	Unmapped	Total Allocated	Unallocated	% Allocated
Free Space on Array Group 2-2-1 RAID5(3D+1P)	58.50 MB	0.00 MB	0.00 MB	0.00 MB	58.50 MB	0.00 %
OPEN-V RAID5	1,045,581.94 MB	673,735.69 MB	371,846.25 MB	1,045,581.94 MB	0.00 MB	100.00 %
OPEN-V External	51,202.50 MB	10,240.50 MB	40,962.00 MB	51,202.50 MB	0.00 MB	100.00 %
Free Space on Array Group 1-1-1 RAID5(3D+1P)	378,312.00 MB	0.00 MB	0.00 MB	0.00 MB	378,312.00 MB	0.00 %
Free Space on Array Group 2-1-1 RAID5(3D+1P)	221,293.50 MB	0.00 MB	0.00 MB	0.00 MB	221,293.50 MB	0.00 %

The properties are calculated for EVA array groups as follows:

Property Displayed for an Array Group (CIM_StoragePool)	How It Is Calculated	Explanation
Total Capacity	sum of volume.size + pool.totalRemainingSpace	The sum of the sizes of all LDEVs in the array group plus the total free space
Total Internal	(sum of volume.size – sum of external volume.size) + pool.totalRemainingSpace	The sum of the sizes of all internal LDEVs

Property Displayed for an Array Group (CIM_StoragePool)	How It Is Calculated	Explanation
Mapped	sum of volume.size for each volume that has a LUN	LDEVs that have LUNs
Unmapped	sum of volume.size for each volume that doesn't have a LUN	LDEVs that do not have LUNs but are not on the management server's “free” list
Total Allocated	sum of volume.size for all volumes	All LDEVs that are not on the management server's “free” list
Unallocated	pool.totalRemainingSpace	LDEVs on the management server's “free” list plus the total free space

Capacity Information for Solaris Virtual Machines

There are two broad configurations for Solaris virtual machine file systems. In the first type of configuration, the file systems on the virtual machines are exclusive to the virtual machines. In this scenario, Capacity Manager includes the capacity numbers for each of the file systems when calculating the capacity summary. Capacity Manager sums the total, used, and free capacity for each file system and displays the numbers.

In the following example, the exclusive file systems are /, /fs1, /fs2, and /localfs. In this case, the capacity summary shows the following:

- Total Capacity = 7878.65 + 1939.61 + 1939.61 + 1939.61 = 13697.47 GB
- Total Used = 5076.39 + 300.44 + 300.44 + 300.44 = 5977.70 GB
- Available = 2723.47 + 1580.98 + 1580.98 + 1580.98 = 7466.42 GB

Figure 12 Solaris Virtual Machines with Only Exclusive File Systems

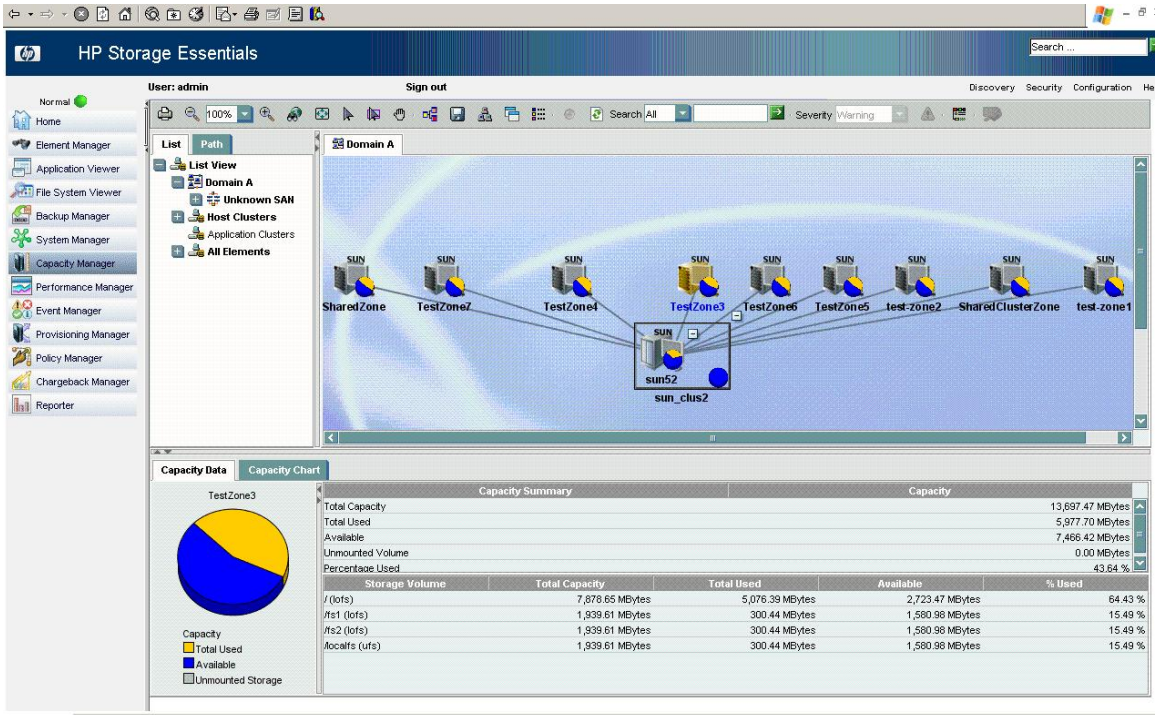

In the second type of configuration, some file systems are exclusive to the virtual machines and some are created on subdirectories of the global zone and, therefore, share storage with the parent file system of these subdirectories. The capacity numbers for each exclusive file system are considered when calculating the capacity summary numbers.

In the following example, the file systems /, /fs1, and /fs2 are exclusive to the virtual machine. The capacity numbers are considered for each when calculating the capacity summary. For all of the virtual machine file systems that point to the same file system on the virtual server, the capacity numbers are taken into account once for each parent file system on the virtual server.

In the following image, the virtual machine file systems /sharedfs1, /sharedfs2, and /sharedfs2 are created from the virtual server directories `/zone-dir/TestZone2/SharedStorage/fs1`, `/zone-dir/TestZone2/SharedStorage/fs2`, and `/zone-dir/TestZone2/SharedStorage/fs3`. These subdirectories are created from the same virtual server file system `/zone-dir/TestZone2/SharedStorage`. Therefore, the capacity values are included just once for the parent file system during the calculation.

The capacity summary shows the following:

- Total Capacity = 7878.65+1939.61+1939.61+ 64,990.91 = 76748.77 GB
- Total Used = 5084.35 + 300.44 + 300.44 + 31,325.79 = 37011.02 GB
- Available = 2715.51 + 1580.98 + 1580.98 + 33015.21 = 38892.68 GB

Figure 13 Solaris Virtual Machines with Shared and Exclusive File Systems

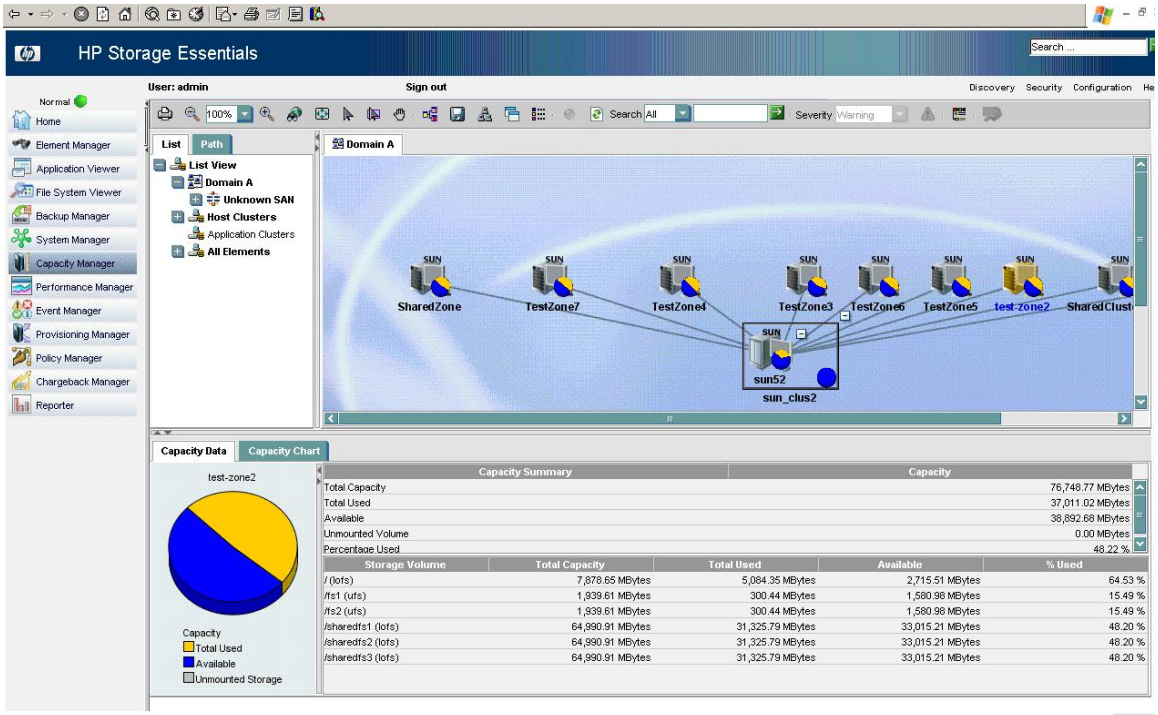

The capacity numbers displayed for each file system in Capacity Manageris the same as shown by the “df -k” command on the virtual machines.

Printing Elements in Capacity Manager

The software enables you to print the topology in Capacity Manager. For example, you can provide a printout to upper management that shows not only the topology of the network, but also the capacity of each element.

To print the elements in Capacity Manager:

1. Access Capacity Manager as described in Accessing Capacity Manager on page 555.
2. If the topology spans more than one screen, arrange the elements so they are closer together, preferably on one screen. To move an element, click the button and then the element you want to move. Drag the element to its new location. Moving elements closer together prevents the printout from appearing too stretched.
3. Click the button. For more information, see Configuring Print Settings on page 336.

Viewing Capacity Charts

Capacity Manager provides a graphical representation of the capacity history of an element, such as port summary information for switches.

The following types of capacity charts are available:

- Volume
- Aggregate
- Quota
- Snapshot

You can manipulate the charts so that they show a different reporting period and frequency. For example, you can show the capacity of a host over the past 24 hours with an hourly monitoring frequency.

Keep in mind the following:

- Verify that the performance collector for the element is enabled (**Configuration** > **Performance** > **Data Collection**). See Managing Performance Collectors on page 372 for more information about enabling performance collectors.
- An hourly roll-up of capacity chart data occurs daily at 2 a.m. Capacity chart data is not available until after Get Details is performed and the hourly roll-up occurs.
- Use the Report Cache Refresh feature to update the data supplied for generating capacity charts. The capacity charts use the same database materialized views as the reports. See Refreshing the Report Cache on page 365 for information.
- Switches and storage systems display data from the last time Get Details was performed.
- If you see the message "There is not enough data to produce a chart [chart_title] at this time," lessen the frequency option, so that the amount of time listed in the **Frequency** menu passes before you view the chart again. For example, if the Frequency menu displays hourly, you must wait an hour for data to appear in the chart.

To find the performance of an element:

1. Access Capacity Manager as described in Accessing Capacity Manager on page 555.
2. Click the element whose capacity history you want to see.
3. In the bottom split pane, click the **Chart** tab.
4. In the lower-middle split pane, click a report title. The chart for the monitoring option appears in the lower-right pane.

   To change the orientation of the chart, hold down the mouse button when you click the chart, and continue holding it while you move the mouse.
5. To change the reporting period, do one of the following and then click the  button to update the chart.
   - Display data within a time period – You can format the graph to provide data within the time period specified. Select the option to the left of the Period combo box. Select one of the following from the menu:
     - Last 24 hours – This option is not available to switches and storage systems.
     - Last 7 Days

- Last Month
- Last Year

- Display data within a starting and ending time – You can format the graph to provide data within the starting and ending time specified:

  i. Select the option to the left of the Start box.

  ii. Click the icon.

  iii. In the Time box, enter the time you want the graph to start, using the 24-hour format.

  iv. Select a date and click **Set**.

  v. Repeat the steps for the End box.

  If you change the date in the box to a date that does not exist in a month, the software automatically calculates the date to the first day of the next month. For example, if you enter 2010-11-31, the software assumes the date is 2010-12-01.

6. *Applications and Hosts only*: To change the display frequency, select one of the following from the Frequency menu:
   - Hourly
   - Daily
   - Weekly
   - Monthly
7. If the chart contains more than one series of data, you can filter out the additional data series by clicking the button.
8. To add trending information, enter an integer in the Trend box.

   The integer corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you enter "5" in the Trend box, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the Frequency menu.

   See Viewing Trending Information for Storage Capacity below.
9. Click the button to update the chart.
10. To print the chart, click the button displayed in the same pane as the chart.

Viewing Trending Information for Storage Capacity

The management server can display trending information in its reports. For example, you can configure Capacity Manager to display trending information for the next week. This information can give you an indication of an element's future capacity need based on its past capacity utilization.

Keep in mind the following:

- An element's capacity can drastically change in the future. Keep in mind that the data trends are just assumptions and should not be treated as fact.
- Trending requires at least two sets of data gathered within the frequency specified.
- Trending information is particularly useful if you need to plan LUN provisioning for virtual machines. For example, you could analyze the remote volume usage over a period of time to determine if an application running on a virtual machine needs to be moved to a different volume.

To view trending information:

1. Access Capacity Manager as described in Accessing Capacity Manager on page 555.
2. Click the element whose performance history you want to see.
3. In the bottom pane, click the **Chart** tab.
4. In the lower-middle split pane, click a monitoring option.
5. In the pane displaying the chart, enter a number in the Trend box.

   The number corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you enter "5" in the Trend box, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the Frequency menu.
6. Click the  button. The trending information is displayed.

If there is not enough data to display, Capacity Manager does not display the chart. For example, if you selected the weekly option from Frequency menu, and you have only two days of data, a chart is not displayed, regardless of the value in the Trend box. Capacity Manager does not display a chart if there is not enough data, and the trending number is ignored.

Different Results for the df -k Command and Capacity Manager

When you run the df -k command on UNIX for a storage system, you might notice that the total capacity displayed does not match the total capacity in Capacity Manager. This difference occurs because Capacity Manager calculates the total capacity differently than the df -k command. The df -k command calculates the total capacity as follows:

```
used capacity + available capacity + reserved capacity
```

Capacity Manager calculates the capacity as follows:

```
used capacity + available capacity
```

The difference between the two calculations is the capacity reserved for superuser. If a file system has a reserved capacity, the total capacity from the df -k command and Capacity Manager will differ.

For example, assume you run the df -k command for the file system `/dev/dsk/c0t0d0s0`. After you run the df -k command, the total capacity is reported as 6688076 KB. When you look at Capacity Manager, the total capacity is reported as 6621196 KB. Capacity Manager actually displays results in gigabytes, but for this example, it is easier to have the totals using the same units.

The totals differ for the following reason.

When you run the df -k command, the computer uses the "used capacity + available capacity + reserved capacity" equation with the following result:

```
1904031 KB + 4717165 KB + 66880 KB = 6688076 KB
```

In this instance:

- 1904031 KB is the used capacity.
- 4717165 KB is the available capacity.
- 66880 KB is the capacity reserved for the superuser. The percentage of the reserved capacity can be set when using the newfs -m command.

Capacity Manager, however, calculates total capacity using the "used capacity + available capacity" equation and displays the following result:

```
1904031 KB + 4717165 KB = 6621196 KB
```

In this instance:

- 1904031 KB is the used capacity.
- 4717165 KB is the available capacity.

Capacity Manager does not include the reserved capacity in its calculations; therefore, the difference between the two calculations is the capacity reserved for the superuser, or 66880 KB.

14 Managing Policies

Depending on your license, Policy Manager might not be available. See the List of Features to determine if you have access to Policy Manager. The List of Features is accessible from the Documentation Center (**Help** > **Documentation Center**).

This section contains the following topics:

About Policy Manager

Policy Manager can automatically send an e-mail, generate an event or run a custom script when an element is being overused or when one of the following events occurs:

- A new element is discovered.
- Successful provisioning occurred.
- An event occurs on one or more specified elements.

Caution: Policies for storage systems only work if you have the performance license pack installed.

If you have EVA, XP, and HDS arrays:

1. Import the Array Performance license for the EVA, XP, and HDS arrays.
2. Go to the Performance Licensing tab in **Security** > **License**.
3. Apply the license for the array performance pack.

For more information, see "License Setup for Array Performance Pack" in the *Installation Guide*.

Collectors gather information for reports, monitoring and File System Viewer. Policies are not triggered when a collector is running. If you configured a policy to be triggered when a host has a critical event and the collector is running when a critical event occurs, the policy is not triggered.

You can create the following types of policies.

Utilization policies monitor the utilization of an element. Options provided depend on the type of element.

For example, you can configure Policy Manager so you receive an e-mail message when the amount of free space on a server decreases to a specified level.

Performance policies monitor the performance of the following.

- One or more EVA storage systems.
- One or more EVA storage volumes.
- One or more EVA storage pools.
- One or more XP or HDS storage systems.
- One or more XP or HDS storage volumes.
- One or more XP or HDS array groups.

Backup policies monitor the backups on your network. For example, polices can be created to notify you if a backup failed on a backup client.

Infrastructure policies monitor the following:

- Discovery of a new element
- Successful provisioning
- Occurrence of an event on one or more specified elements

Accessing Policy Manager

Although there is only one way to access policies for discovery, provisioning, and events; there are multiple ways to access policies. This flexibility in accessing utilization policies lets you easily create and manage policies without interrupting your work flow.

Policy Manager provides the following options for accessing policies:

- To access policies from System Manager – Double-click an element in System Manager and click the **Policies** tab. This method displays the utilization policies for just the element that was double-clicked.
- To access policies for file server users – Click **File System Viewer** () in the left pane. Select **Enterprise** from the Select Context menu and **Users** from the Select View menu. Click a **User ID** in the User ID column and the **Policies** tab.
- To access policies for backup elements – Select one of the following elements in Backup Manager, and click **Policies** in the lower-right corner:
  - Backup Clients
  - Backup Library

- Master backup server
- Master backup media

To access all policies in Policy Manager:

1. Click **Policy Manager** ().
2. Click a policy in the Policy Manager tree.

Creating Policies

This section contains the following topics:

It is not possible to create policies for managed tape libraries.

Actions Available When a Policy Condition Is Fulfilled

When you create or modify a policy, you must select an action to occur when the policy condition is fulfilled. More than one action can be assigned to a policy. The following actions are available:

- **Send E-mail** – Policy Manager sends an e-mail when the condition is fulfilled. Enter a comma-separated list of e-mail addresses and click **OK**. To update an email address in the Action box, delete the outdated email address from the Action box and add the updated one.
- **Generate Event** – Policy Manager generates an event of the specified event type, and the event appears in Event Manager. After you select a severity level, click **OK**. For a list of the severity levels you can select from, see Severity Levels on next page.
- **Execute a Custom Command** – Policy Manager executes a custom command on the management server when the condition is fulfilled. Enter a command that will execute the script in the box, and click **OK**.

  The software assumes you are in the `%JBOSS4_DIST%\server\appiq\remotescripts` directory on the management server when the script is executed. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION, where POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides a description of the policy. See Software Environment Variables for Scripting on page 483 for more information.

Prefix the command with “start” if the custom command triggers a user interface component, for example, Microsoft Internet Explorer or a command prompt window. For example, to have the custom command to open a command prompt window and list the contents in the directory, prefix the command as follows:

```
start dir
```

Severity Levels

When you create or modify a policy, you can select the severity level for which you want Policy Manager to generate an event. The following severity levels are available:

The severity level for an element is set by the manufacturer; therefore, the meanings of the severity levels vary. It is best to view the description of the event.

- **Unknown** – The severity level is not known.
- **Informational** – An example of an informational event is a progress report event for firmware download operation currently in progress.
- **Warning** – An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
- **Minor** – An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** – An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** – An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

Creating a Performance, Backup or Utilization Policy

You can create a performance, utilization or backup policy that generates an event, sends an e-mail, or runs a custom command when a discovered element is being overused. For example, you can configure Policy Manager so you receive an e-mail message when the amount of free space on a server decreases to a specified level.

Caution: Policies for storage systems, including performance policies, only work if you have the performance license pack installed.

To enable performance policies for the EVA, XP, and HDS arrays:

1. Import the Array Performance license for the EVA, XP, and HDS arrays.
2. Go to the Performance Licensing tab in **Security** > **License**.
3. Apply the license for the array performance pack.

For more information, see License Setup for Array Performance Pack on page 328.

Keep in mind the following:

- If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See Enabling Email Notification on page 335.
- Policies that are triggered for virtual applications are also triggered for file server users.
- For the following XP array policies, the default CLPR used for obtaining performance information is CLPR 0. The default MP for the XP controller is MP 0.
  - Read Hits
  - Percent Cache Usage
  - Sidefile Usage
  - Write Pending Data
  - Cache Usage
  - Percent Side File Usage
  - Percent Write Pending Data

To create a utilization or Backup policy:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. In the left pane, select an element or element type to which you want the policy to apply.
3. In the right pane, click **Add**.
4. Select a policy, and click **Next**. Refer to the user interface for details about the policy.

   This step is not applicable for backup libraries.
5. In the Name box on the Policy Properties tab, enter a name for the policy or keep the default.

   To distinguish the policies from one another, include the condition of the policy in its name, for example, Array Groups Average Write Size for greater than 35 bytes. You might also want to include the name or type of devices the policy is monitoring.
6. In the Description box, modify the description for the policy or keep the default.
7. In the Re-arm Period box, specify the amount of time (in minutes) after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the policy will execute.

   For example, assume you set the collectors to run every 2 minutes. When the collectors run, the condition for the policy is met and an event is generated. You find that Event Manager is being flooded by events. You want to be notified when the condition for the policy is met, but you want to be notified every hour instead of every 2 minutes. You would enter 60 in the Re-arm Period box. The policy would then run every hour.

   Specify shorter periods for important applications.
8. If you are setting a policy for a backup client, go to Step 11.
9. In the Condition menu, specify a comparison operator.

10. (Available for only backup libraries). To the left of the comparison operator, select the media pool you want to monitor.
11. To the right of the Comparison Operator menu, take the following action, depending on the box displayed:
    - Enter a percentage.
    - Enter an amount in gigabytes (GB).
    - Type the number of available media that will trigger an alert. For example, if you want to be alerted when the number of available media for a storage pool is less than two, you would set the conditional to less than (<). You would then enter 2 in the Media box.
12. For trending policies, enter the number of days in the Historic period box (7 to 180 days) and Projection period box (min=1; max=180). The Historic period is the number of days in the past that meet the condition.
13. For trending policies, enter the number of days in the future that would meet the condition in the Projection period box (1 day to 360 days).
14. Select an action to occur when the policy condition is fulfilled. For more information, see Actions Available When a Policy Condition Is Fulfilled on page 581.
15. Click **Finish**.
16. To test a policy, click the **Test** button in the Utilization Policy table. The management server fires a test for all utilization policies associated with that element.

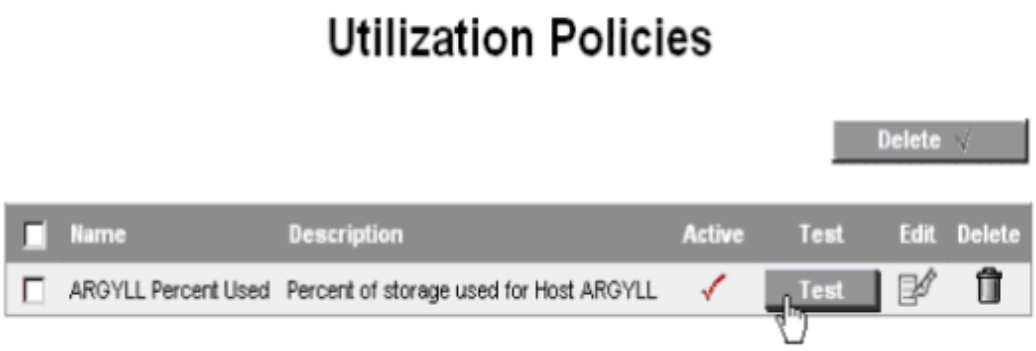

Creating Policies for Discovery

You can create an infrastructure policy that generates an event, sends an e-mail, or runs a custom command when an element is discovered.

If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See Enabling Email Notification on page 335.

To create a policy for discovery:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. In the Policy Manager tree, expand the Infrastructure Policies node, and then select **New Element Discovery**.
3. Click the **Add** button in the right pane.
4. In the Name box, enter a name for the policy.
5. In the Description box, enter a description for the policy.
6. Select one or more element types.

   When a condition is fulfilled on a selected element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to act when the event is cleared. If you do not select this, the policy is triggered when the event is received.
8. Select one of the following options from the Summary Text list to establish how Policy Manager will scan the summary text and respond:
   - **Is anything** – Regardless of the contents of the event's summary text, Policy Manager sends an e-mail, generates an event, or runs a custom command. No additional text is required.
   - **Contains** – If the event's summary text contains the text you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired text in the box that appears to the right of the Summary Text menu.
   - **Matches Regular Expression** – If the event's summary text matches the expression you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired expression in the box that appears to the right of the Summary Text menu.
9. Select an action to occur when the policy condition is fulfilled. For more information, see Actions Available When a Policy Condition Is Fulfilled on page 581.
10. Click **OK**.

Creating Policies for Provisioning

You can create an infrastructure policy that generates an event, sends an e-mail, or runs a custom command when successful provisioning occurred.

If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See Enabling Email Notification on page 335.

To create a policy for provisioning:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies, and then click **Provisioning**.

3. Click **Add**.
4. In the Name box, enter a name for the policy.
5. In the Description box, enter a description for the policy.
6. Select one or more element types.

   When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to act when the event is cleared. If you do not select this, the policy is triggered when the event is received.
8. Select an action to occur when the policy condition is fulfilled. For more information, see Actions Available When a Policy Condition Is Fulfilled on page 581.
9. Click **OK**.

Creating Policies for Events

You can create a policy that generates an event, sends an e-mail, or runs a custom command when a specific type of event occurs on one or more specified elements.

For example, you can create a policy that sends an e-mail when a new element generates a critical event.

If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See Enabling Email Notification on page 335.

To create a policy for events:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. In the Policy Manager tree in the middle pane, expand the Infrastructure Policies node and click **Events**.
3. Click **Add**.
4. In the Name box, enter a name for the policy.
5. In the Description box, enter a description for the policy.
6. In the Re-arm Period box, specify the amount of time (in minutes) after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the same actions can execute. Specify shorter periods for important applications.
7. Select one or more element types.

   When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
8. Select **Fire when event is cleared** if you want the policy to act when the event is cleared. If you do not select this, the policy is triggered when the event is received.
9. In the Severity menu, specify a comparison operator.

10. To the right of the Severity menu, select a severity level.

    The severity level for an element is set by the manufacturer; therefore, the meanings of the severity levels vary. It is best to view the description of the event.

    - **Unknown** – The severity level is not known.
    - **Informational** – An example of an informational event is a progress report event for firmware download operation currently in progress.
    - **Warning** – An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have appeared.
    - **Minor** – An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
    - **Major** – An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
    - **Critical** – An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

11. Select one of the following options from the Summary Text list to establish how Policy Manager will scan the summary text and respond:

    - **Is anything** – Regardless of the contents of the event's summary text, Policy Manager sends an e-mail, generates an event, or runs a custom command.
    - **Contains** – If the event's summary text contains the text you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired text in the box that appears to the right of the Summary Text menu.
    - **Matches Regular Expression** – If the event's summary text matches the expression you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired expression in the box that appears to the right of the Summary Text menu.

12. Select an action to occur when the policy condition is fulfilled. For more information, see Actions Available When a Policy Condition Is Fulfilled on page 581.

13. Click **OK**.

Testing a Utilization Policy

After you create or modify a utilization policy, test it to verify that it provides the results you are anticipating. To test a policy, click the **Test** button in the Utilization Policy table. The management server fires a test for all utilization policies associated with that element.

To run the Test functionality again, set the Re-arm period to zero before clicking the Test button a second time.

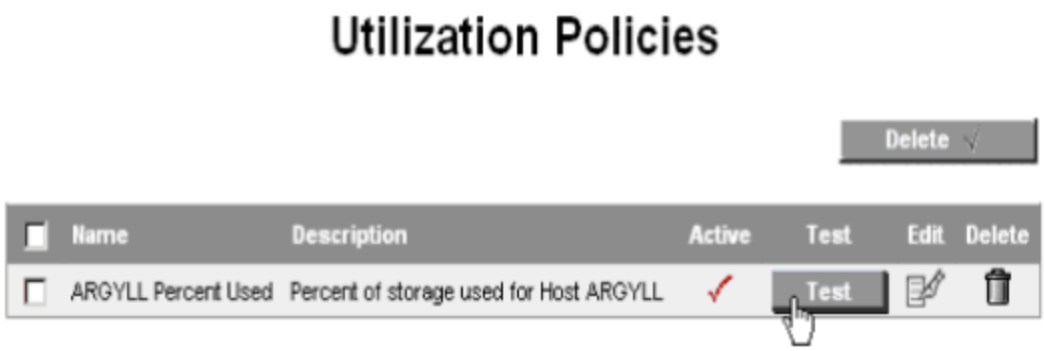

Modifying Policies

This section contains the following topics:

- Modifying Performance, Utilization, and Backup Policies below
- Modifying Discovery Policies on the facing page
- Modifying Provisioning Policies on the facing page
- Modifying Policies for Events on page 590

Modifying Performance, Utilization, and Backup Policies

To modify a performance, utilization, or backup policy:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. Click the **Edit** ( ) button corresponding to the policy you want to modify.
3. In the Name box, change the name for the policy.
4. In the Description box, change the description for the policy.
5. Select or deselect the **Active** option to activate or de-activate the policy.
6. In the Re-arm Period box, specify the amount of time (in minutes) after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the same actions will execute.

   Specify shorter periods for important applications.
7. *Skip this step for backup clients: In the Condition menu, change the conditions of the policy.*
8. For trending policies, change the number of days in the Historic period box (min=7; max=180) or Projection period box (min=1; max=180).
9. Select an action to occur when the policy condition is fulfilled. For more information, see Actions Available When a Policy Condition Is Fulfilled on page 581.
10. Click **OK**.

11. To test a policy, click the **Test** button in the Utilization Policy table. The management server fires a test for all utilization policies associated with that element.

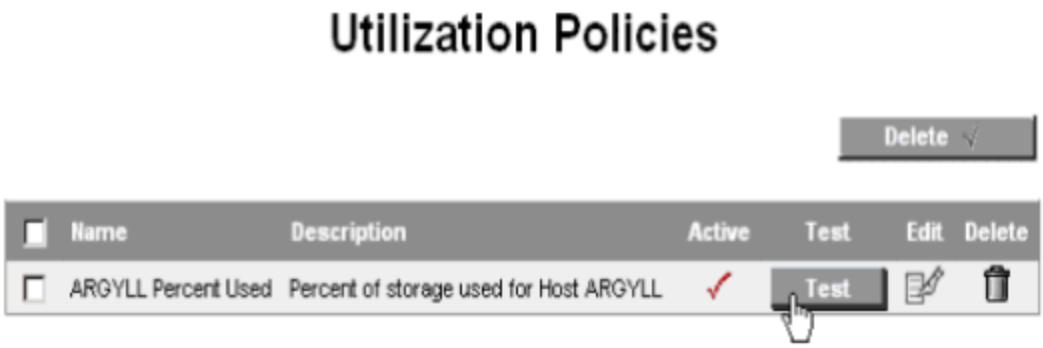

Modifying Discovery Policies

To modify a policy for discovery:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies, and then click **New Element Discovery**.
3. Click the **Edit** ( ) button corresponding to the policy you want to modify.
4. In the Name box, change the name for the policy.
5. In the Description box, change the description of the policy.
6. Select or deselect one or more element types.

   When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to act when the event is cleared. If you do not select this, the policy is triggered when the event is received.
8. Select an action to occur when the policy condition is fulfilled. For more information, see Actions Available When a Policy Condition Is Fulfilled on page 581.
9. Click **OK**.

Modifying Provisioning Policies

You can create an infrastructure policy that generates an event, sends an e-mail, or runs a custom command when successful provisioning occurred.

If you plan to use e-mail notification with your policy, first assign an SMTP server from which the management server can send its e-mail notifications. See Enabling Email Notification on page 335.

To modify a policy for provisioning:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. In the Policy Manager tree in the middle pane, expand the node, Infrastructure Policies, and click **Provisioning**.
3. Click the **Edit** () button corresponding to the policy you want to modify.
4. In the Name box, change the name for the policy.
5. In the Description box, change the description for the policy.
6. Select one or more element types.

   When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
7. Select **Fire when event is cleared** if you want the policy to act when the event is cleared. If you do not select "Fire when event is cleared," the policy is triggered when the event is received.
8. Select an action to occur when the policy condition is fulfilled. For more information, see Actions Available When a Policy Condition Is Fulfilled on page 581.
9. Click **OK**.

Modifying Policies for Events

To modify a policy for events:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. In the Policy Manager tree in the middle pane, expand the Infrastructure Policies node and click **Events**.
3. Click the **Edit** () button corresponding to the policy you want to modify.
4. In the Name box, change the name for the policy.
5. In the Description box, enter the description for the policy.
6. In the Re-arm Period box, specify the amount of time (in minutes) after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the same actions canexecute.

   Specify shorter periods for important applications.
7. Select or deselect one or more element types.

   When a condition is fulfilled on a select element, Policy Manager generates an event, sends an e-mail, or runs a custom command.
8. Select **Fire when event is cleared** if you want the policy to act when the event is cleared. If you do not select this, the policy is triggered when the event is received.
9. In the Severity menu, specify a comparison operator.

10. To the right of the Severity menu, you can change the severity level. For a list of the available severity levels, see Severity Levels on page 582.
11. To change how Policy Manager scans the summary text, select one of the following from the Summary Text menu. Policy Manager scans the summary text and responds according to one of the following actions selected:
    - **Is anything** – Regardless of the contents of the event's summary text, Policy Manager sends an e-mail, generates an event, or runs a custom command.
    - **Contains** – If the event's summary text contains the text you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired text in the box that appears to the right of the Summary Text menu.
    - **Matches Regular Expression** – If the event's summary text matches the expression you specify here, Policy Manager sends an e-mail, generates an event, or runs a custom command. Enter the desired expression in the box that appears to the right of the Summary Text menu.
12. Select an action to occur when the policy condition is fulfilled. For more information, see Actions Available When a Policy Condition Is Fulfilled on page 581.
13. Click **OK**.

Viewing Policies

To view policies, access Policy Manager as described in Accessing Policy Manager on page 580.

The policies are listed in a table whose headings have the following meaning:

Policy Table Description

Heading	Description
Name	The name of the policy
Description	A description of the policy
Active	A check mark in this column means that the policy is active.
Edit	Click the **Edit** ( ) button to edit a policy.
Delete	Click the **Delete** ( ) button to remove a policy.

When you click the NAS Hosts entry in the Policy Manager tree, the product shows the Policies for All Hosts page.

Deactivating a Policy

Policies are activated when they are created. You can deactivate a policy, but still keep it stored in the management server. For example, assume you created a policy that sends an e-mail whenever an event of type Minor is generated for a server. You could deactivate this policy before you upgrade the server.

To deactivate a policy:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. Click the **Edit** () button corresponding to the policy you want to modify.
3. Deselect the **Active** option.
4. Click **Finish**. The policy is deactivated.

Deleting Policies

To delete a policy:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. In the Policy Manager tree in the middle pane, click the element or infrastructure to view its policies.
3. Perform one of the following actions:
   - **Delete a policy** – Click the **Delete** () button of the policy you want to delete.

     Or
   - **Delete several policies at once** – Select the check boxes next to the policies you want to delete. To select all of the policies, select the check box next to the **Name** heading, and click the **Delete** () button. The policies are deleted.

Providing E-mail Notification for a Policy

You can configure Policy Manager to provide e-mail notification when a resource is being overused or when any of the following events occurs:

- New element is discovered.
- Provisioning is successful.
- Event occurs on one or more specified elements.

Keep in mind the following:

- First assign an SMTP server from which the management server can send its e-mail notifications. See Enabling Email Notification on page 335 for more information.

- The following instructions assume you created a policy. If you did not create a policy, see the following topics:
  - Creating a Performance, Backup or Utilization Policy on page 582
  - Creating Policies for Discovery on page 584
  - Creating Policies for Provisioning on page 585
  - Creating Policies for Events on page 586

To set up e-mail notification for a policy:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. Click the **Edit** ( ) button corresponding to the policy you want to modify.
3. Click **Send E-mail**.

   Policy Manager sends an e-mail when the condition is fulfilled. The software verifies that the address is entered in the correct format.
4. Enter a comma-separated list of e-mail addresses, and then click **OK**.
5. Click **Finish**.

Providing Event Generation for a Policy

You can configure Policy Manager to generate an event when an element is being overused or when any of the following occurs:

- New element is discovered.
- Provisioning is successful.
- Event occurs on one or more specified elements.

The following instructions assume you created a policy. If you did not create a policy, see the following topics:

- Creating a Performance, Backup or Utilization Policy on page 582
- Creating Policies for Discovery on page 584
- Creating Policies for Provisioning on page 585
- Creating Policies for Events on page 586

To set up event generation for a policy:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. Click the **Edit** ( ) button corresponding to the policy you want to modify.
3. Click **Generate Event**.

Select the severity level for which you want Policy Manager to generate an event. For a list of the severity levels you can select from, see Severity Levels on page 582.

4. After you select an event level, click **OK**.
5. Click **Finish**.

Providing a Custom Command for a Policy

You can configure Policy Manager to run a custom command on the management server when an element is being overused or when any of the following occurs:

- New element is discovered
- Provisioning is successful
- Event occurs on one or more specified elements

The following instructions assume you created a policy. If you did not create a policy, see the following topics:

- Creating a Performance, Backup or Utilization Policy on page 582
- Creating Policies for Discovery on page 584
- Creating Policies for Discovery on page 584
- Creating Policies for Events on page 586

To set up a custom script for a policy:

1. Access Policy Manager as described in Accessing Policy Manager on page 580.
2. Click the **Edit** () button corresponding to the policy you want to modify.
3. Click **Execute Custom Command**.
4. Enter a command that will run the script in the box, and click **OK**.

   The software assumes you are in the `%JBOSS4_DIST%\server\appiq\remotescripts` directory on the management server when the script is run. You can use environment variables in your script, such as POLICY_NAME and POLICY_DESCRIPTION, where POLICY_NAME provides the policy name and POLICY_DESCRIPTION provides a description of the policy. See Providing a Custom Command for a Policy above for more information.

   Prefix the command with “start” if the custom command triggers a user interface component; for example, Microsoft Internet Explorer or a command prompt window. For example, if you want the custom command to open a command prompt window and list the contents in the directory, you would prefix the command as follows:

   ```
   start dir
   ```

5. Click **Finish**.

Policy Manager runs a remote script on the management server when the condition is fulfilled.

15 Performance Manager

Performance Manager provides a graphical representation of the performance history of a managed element.

Depending on your license, Performance Manager might not be available. See the List of Features to determine if you have access to Performance Manager. The list is accessible from the Documentation Center (**Help** > **Documentation Center**).

When licensed, the optional EVA Performance Pack, XP and HDS Array Performance Pack, and EMC Symmetrix Performance Pack provide additional Performance Manager data collection capabilities. For information about using optional performance packs with the Performance Manager, see the HP Storage Essentials *Storage Performance Management Guide*.

This section consists of the following topics:

For more information about the Performance Manager and storage performance monitoring, see the HP Storage Essentials *Storage Performance Management Guide*.

Performance Manager and Array Performance Packs

Performance Manager provides a graphical representation of the performance history of a managed element, such as bytes transmitted per second for a switch. From this performance information, you can also generate charts and customize reports.

You can manipulate the charts, so they show a different reporting period and frequency. For example, you could show the performance of an element over the past 24 hours with an hourly monitoring frequency.

The licensed optional EVA Array Performance Pack, the optional XP and HDS Array Performance Pack, and the optional EMC Symmetrix Performance Pack provide additional performance monitoring and reporting capabilities to Performance Manager. These Performance Packs enhance both the Performance Manager and the Reporter capabilities.

The following sections provide more information about the EVA Performance Pack and the XP and HDS Array Performance Pack.

For more information, see the HP Storage Essentials Performance Management Guide. It describes how to manage the storage performance of your environment using the additional performance features of HP Performance Packs, HP Storage Essentials Performance Manager, and Report Optimizer.

EVA Array Performance Pack

The EVA Array Performance Pack provides performance data and metrics for EVA arrays.

The following sections describe important requirements and considerations for licensing and using the EVA Array Performance Pack:

- EVA Licensing Requirements and Setup below
- EVA Array-Related Software Requirements on the facing page
- EVAPerf Data Collector Requirements on the facing page
- Collecting EVA Performance Metrics on the facing page
- Specifying EVA Data Collectors on page 598
- EVA Array Discovery on page 600
- EVA Metrics-Collection on page 601
- EVAPerf Considerations on page 601
- Managing Late Data or Errors from EVA Arrays on page 602

EVA Licensing Requirements and Setup

Your EVA Array Performance Pack license determines how many EVA arrays you can select for performance monitoring. Additional licenses can be purchased if you need to monitor more arrays than your current license allows.

As part of the license setup, you must specify which EVA arrays are to be monitored. You can make changes to your selection at a later time to accommodate new monitoring and reporting needs. License setup is discussed in Managing Licenses on page 319.

Setup for the EVA Array Performance Pack also requires you to specify the data collectors used to implement your specific reporting needs. This is discussed in Collecting EVA Performance Metrics on the facing page.

EVA Array-Related Software Requirements

The following list is a set of current best practices for EVA Performance data collection:

- If you installed the Command View EVA 9.1 or later SMI-S provider, use:
  - EVA firmware XCS 6200 or later
  - CV EVA 9.1 or later
- Command View EVA versions 9.1 and later cannot obtain cumulative metrics from the VCS 3.110 firmware. Therefore, it has no metrics to provide to HP Storage Essentials. If you cannot upgrade your VCS firmware, make sure you are running Command View 8.0.2 or 9.0.1. Do not upgrade to Command View EVA 9.1 or later. If you upgrade to Command View EVA 9.1 or later, you will no longer see your EVA metrics in HP Storage Essentials.
- A maximum of 8 EVAs per CV EVA server are supported.
- Performance data for EVAs running VCS 3.x is only point in time. Each data point for performance metrics corresponding to EVAs running VCS 3.x will only represent a one second sample period.
- VCS 4.004 is not supported for EVA5000 disk arrays.
- Management server performance provider does not support EVA VCS code earlier than 4.x.

The maximum number of vdisks supported per EVA array is 512.

You must discover the arrays for the Array Performance Pack to work.

EVAPerf Data Collector Requirements

The management server requires the EVAPerf Data Collector (EVAPerf) service to be running on each Command View (CV) EVA server being monitored. This service is installed by default with the CV EVA suite, but the service is set to **enabled** and **manual startup**, which means that the service is not automatically started. Therefore, you might want to change the service startup parameters to **automatic** and **restart on failure**.

Although EVAPerf version 6.0.2 or newer is supported, it is strongly recommended that you run EVAPerf version 8.0.2, or later, to take advantage of service crash fixes, stability improvements, and miscellaneous fixes.

RPC over Port 860 must be enabled in any firewalls between the management server and the CV EVA server(s).

Collecting EVA Performance Metrics

HP recommends not enabling the collectors for EVA physical disk drive metrics. Due to the large amount of data associated with disk drive metrics, collection of these metrics can affect the responsiveness of CV EVA. Collection of EVA disk drive metrics should only be done when instructed to do so by HP service personnel or when troubleshooting a specific issue that requires examining performance data for one or more disk drives.

Collecting individual volume statistics and collecting aggregated volume statistics each cause the metrics for every volume to be collected. Each of these collections might take a couple of minutes on a heavily loaded array. Also, each of these collectors queries the array for the same volume information separately. Therefore, running both aggregated and non-aggregated collectors puts additional load on the array. For this reason, you should only run them both if you need both types of information (that is, aggregated and non-aggregated). If you do need to run both, you should schedule each set of collectors with at least a 10-minute polling interval so that one can complete before the next one starts.

Collection should only be enabled for metrics that are actually used by the customer. Enabling all collectors, or more than are necessary, can result in decreased Command View (CV) responsiveness and intermittent performance data collection failures.

Specifying EVA Data Collectors

After applying the EVA Array Performance Pack license to the target arrays, you must select which corresponding array-specific “performance data collectors” to enable in the Data Collectors for Performance page. Each EVA Storage System-specific row in the list of performance data collectors represents a distinct set of performance metrics for the corresponding EVA array. When you start a collector, you will be prompted for a sample interval and start time for the collection of the corresponding metrics. This can be edited after a collector is started by selecting the icon for the collector. Be aware that the type, quantity, and frequency of data collectors started can affect system performance. For example, EVA physical disk metrics collection is I/O intensive, so it should be avoided unless needed.

To choose collectors:

1. On the management software home screen, click **Configuration**.
2. Click **Performance**. The Data Collection screen displays.
3. Select the **Data Collection** tab.
4. Select the desired collectors from those listed in your display.
5. Click **Start Selected** for multiple collectors or click **Action** for single collectors.

   You might prefer to start and run the collectors only for specific needs, rather than running all of them all the time; otherwise it might affect overall system response.
6. To edit the start time and interval for multiple collectors, click **Spread Start Time** to stagger the starting times for multiple collectors. This minimizes the impact on system performance.
7. Review the collected data in Performance Manager. (You can view this data in Performance Manager only if you are licensed appropriately.) You can expand the device sub-elements (controllers, volume group, etc.) in the navigation tree, as shown in the following example.

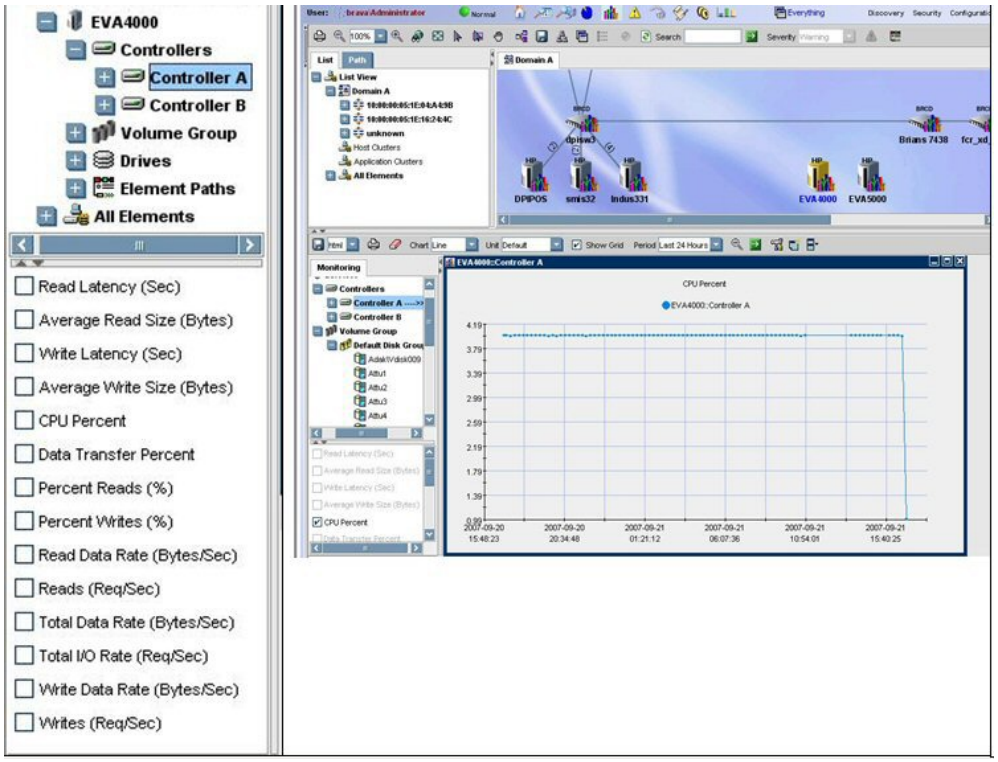

The list of sub-elements and metrics varies according to the device type. The selections determine the performance data retrieved and available for analysis.

The following example shows information about data rates associated with the highlighted EVA array host port.

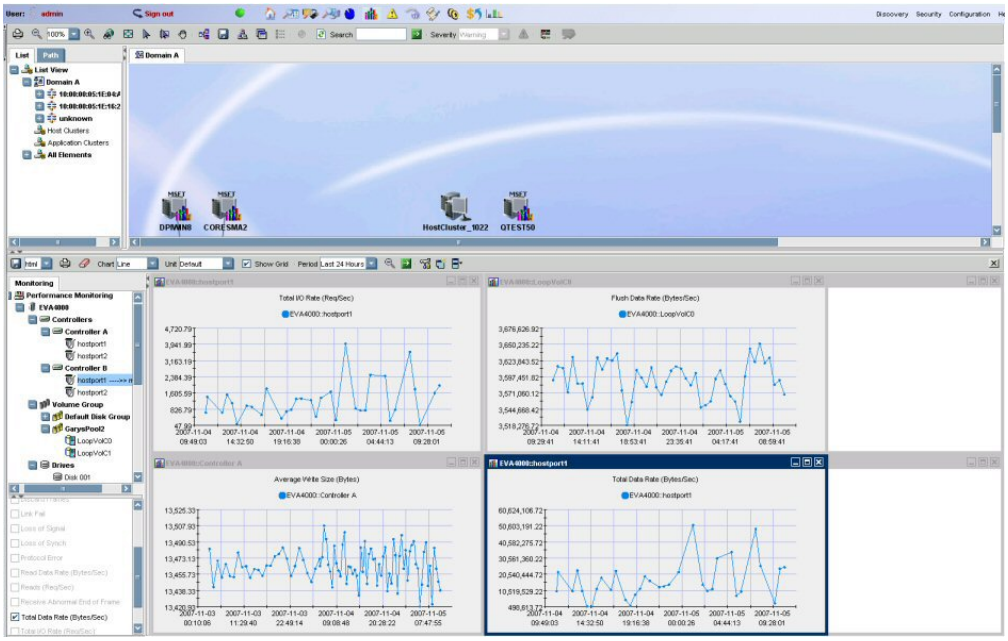

EVA Array Discovery

The following sections describe important considerations in the EVA array discovery process to support the EVA Array Performance Pack.

Discovery Considerations Using Command View EVA

Array Performance Pack for EVA requires Get Details of EVA Arrays using the supported Command View (CV) EVA versions. By default, the HP Storage Essentials management server utilizes the built-in provider for CV versions prior to 9.1, or the SMI-S provider shipped with the CV EVA Software Suite for CV versions 9.1 or newer.

If both the managed and the unmanaged (standby) arrays are discovered using the required version of CV EVA with the new built-in provider, or external SMI-S provider for CV 9.1 or later, both the managed and unmanaged arrays must remain in the same discovery group. They can be moved to a different discovery group, but both must remain together in the same discovery group.

If you performed EVA array discovery previously using the SMI-S provider (as with earlier CV EVA versions), the Array Performance Pack requires Get Details of EVA Arrays using the supported CV EVA version, which utilizes the built-in provider. (By default, Management Server version 6.x discovers EVAs using the built-in provider.)

You should have run Get Details on an EVA array before the Performance Monitoring tree will be populated with the corresponding EVA Elements used for corresponding metric selection.

Any configuration changes to the EVA array (for example, addition or deletion of mapped volumes) require a Get Details in order for the changes to be reflected in the Performance Monitoring tree, and in the actual performance data collected. Also, changes in the number of elements (for example, disks or volumes) can impact the aggregated performance metric data.

If you make changes to your selection of EVA arrays later, you must ensure the selected new EVA arrays have been discovered using the required CV EVA version. For more information regarding discovery procedures, refer to the management server installation guide.

CV EVA supports a standby configuration whereby an array can be managed by multiple CV EVA stations simultaneously, with only one of them actively managing a specific array at a time. Discovery will include both the EVA array being actively managed and the array not being actively managed (Standby CV station). In such cases, there are limitations to the information that can be gathered, based upon whether both the managed and the unmanaged (standby) arrays are discovered by the management server.

If only the unmanaged array is discovered, only the top level array information can be collected.

EVA Metrics-Collection

EVA performance data cannot be collected for EVA arrays that have controller passwords set via the Operator Control Panel (OCP).

Only EVA volumes that are mapped to a host have associated performance data.

The management server will not be able to gather host port metrics if the EVA controller has two switches attached to it with the same DomainID and the same switch port number is used in each switch.

EVAPerf Considerations

When running the EVAPerf Data Collector service, consider the following:

- EVAPerf Data Collector service can sometimes crash, especially in larger environments. This is less likely with EVAPerf 7.0.1 or later. Restarting the service will enable performance data collection to succeed. Consider changing the service startup parameters to **automatic** and **restart on failure**.
- If the total number of vdisks configured on the EVA exceeds 512 using EVAPerf versions prior to v7.0.1, some mapped vdisks might not have any metrics associated with them.
- Occasionally, reporting might not occur for a physical disk because the EVAPerf Data Collector returns an unknown identifier for it.
- Vdisk and physical disk metric data viewed in the EVAPerf CLI on the CV EVA server might appear inconsistent with the management server's data. This is due to an EVAPerf CLI situation in which only point-in-time samples are returned for sample intervals greater than 30 seconds (this is for all VCS versions).
- Socket connections between the management server and the EVAPerf Data Collector service on port 860 are not immediately closed after stopping the EVA performance collectors in the management server, or when un-licensing EVA arrays for performance data collection in the management server. However, the connection will eventually time out and be closed after the default idle time for the connection expires. In these situations, although the connections are still open, no I/O occurs over the connection.

- The minimum collection interval that can be set for the EVA performance data collectors is 1 minute. However collecting at frequencies less than 5 minutes, especially if all collectors are enabled, might result in decreased stability and failure to collect metric data. The collection interval for real-time metrics is 20 seconds and cannot be changed. Frequent collection of large amounts of performance data, combined with regular I/O, might impact CV EVA responsiveness. However, this does not impact the data I/O of the EVA array.
- Only Vdisks that are mapped to a host have performance metrics associated with them. Only grouped physical disks have performance metrics associated with them.
- The default idle time-out is 2 hours and is not configurable.

When using the internal SMI-S provider, if no performance data is collected for a period of 2 hours, the connection to the EVAPerf is closed due to an idle time-out. This will cause the EVA performance provider to be reset for the affected EVA array resulting in a failure to display a data point at the next scheduled collection interval. The failure to display the data point is because 2 consecutive data points are required after a reset in order to derive the rate data between 2 samples. Therefore, if EVA array performance collectors are scheduled at a frequency greater than 2 hours and the start times for that EVA array collectors are not staggered (that is, resulting in 2 hours periods with no collections for that EVA array), then no data will ever be displayed for that EVA array.

Managing Late Data or Errors from EVA Arrays

If you are performing real-time data collection, and the element is not returning the information in time, you are shown a message in red like the following:

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

The software cannot obtain the information in a timely manner because of one or more of the following:

- The element might be inherently slow.
- The element might be busy with other tasks.
- You are trying to collect too much information at once from the element.
- The specific element is already being monitored in real time on another chart.

You could look in the appstorm and cimom logs around the time the error occurred to see if there are any error messages that explain why the management server failed to collect the data. There is substantial error handling and logging for failed EVA performance data collection.

If you think you are trying to collect too much information from the element, you could narrow down the collection. For example, if you are trying to collect monitoring information for three disk drives on a server, you could try collecting information for one disk drive.

Performance Manager will continue to attempt to retrieve data from the element until the chart is closed.

XP and HDS Array Performance Pack

The XP and HDS Array Performance Pack provides performance data and metrics obtained from HP and XP arrays. The following sections provide important requirements and considerations for licensing and using the XP and HDS Array Performance Pack:

XP and HDS Licensing Requirements and Setup

Your XP and HDS Array Performance Pack license determines how many XP and/or HDS arrays you can select for performance monitoring. Additional licenses can be purchased if you need to monitor more XP and HDS arrays than your current license allows.

As part of the license setup, you must specify which XP or HDS arrays are to be monitored. You can make changes to your selection at a later time to accommodate new monitoring and reporting needs. License setup is discussed in Managing Licenses on page 319.

Setup for the XP and HDS Array Performance Pack also requires you to specify the data collectors used to implement your specific reporting needs. This is discussed in the Specifying XP and HDS Data Collectors on page 606.

XP and HDS Array Performance Pack Architecture and Data Collection

The XP and HDS Array performance pack collection paths are very similar. The main difference is that the management server uses the Command View Advanced Edition to collect data from the XP array. Likewise, the management server uses Hitachi Device Manager (HDvM) and the built-in HDS provider to collect data for the HDS array.

The management server uses a proxy host that is connected to the array to gather performance data. You cannot use the same proxy host for both XP and HDS arrays. The proxy host can be used either for multiple XP or HDS arrays, but not for both types of arrays.This host proxy can run Windows, Linux or HP-UX operating system, and it must have a CIM extension installed on it. The RAID Manager Library is installed on the host proxy to get the performance data from the array using a special command LUN.

Figure 14 XP Array Performance Pack Collection Paths

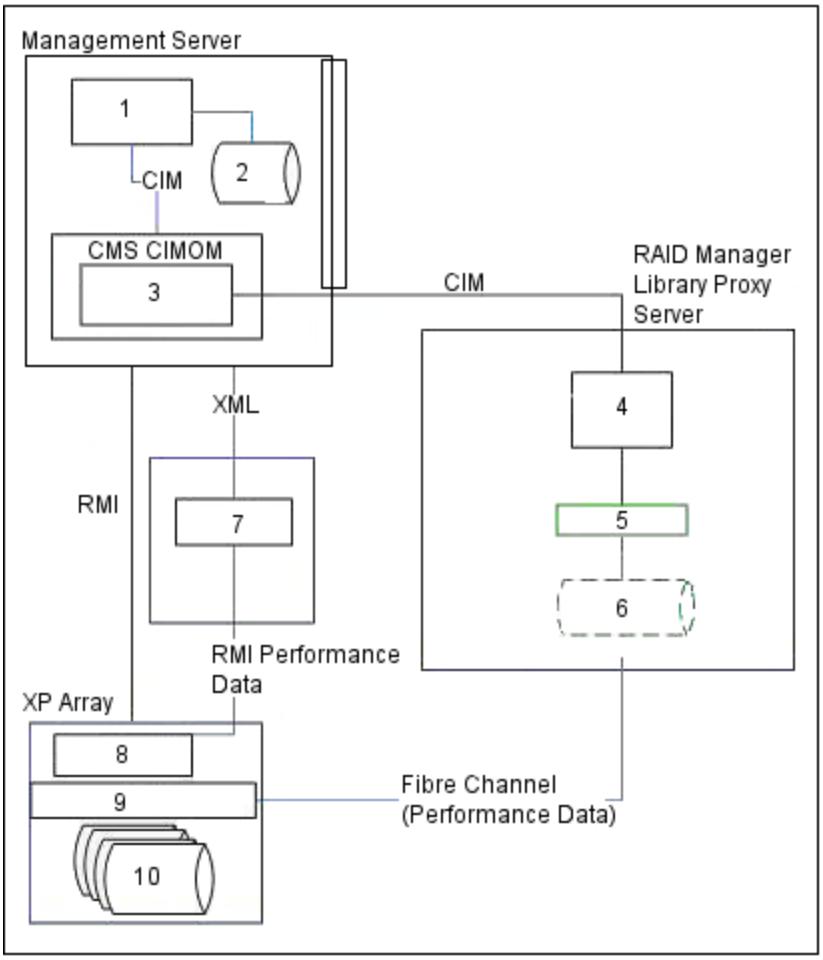

1	Management Server
2	HP Storage Essentials Oracle Database
3	XP Provider
4	HP Storage Essentials Host Agent/CIM Extension
5	RAID Manager Library
6	Command LUN
7	Command View Advanced Edition
8	Service Processor (SVP)
9	Controllers
10	Data

Figure 15 HDS Array Performance Pack Collection Paths

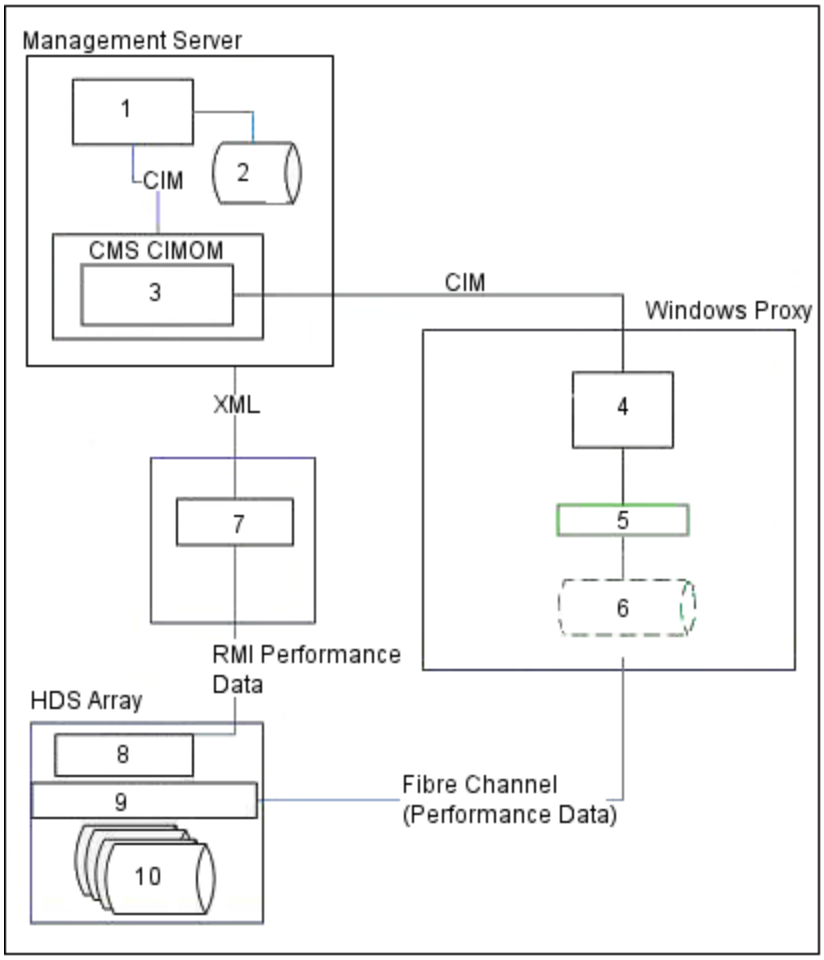

1	Management Server
2	HP Storage Essentials Oracle Database
3	HDS Provider
4	HP Storage Essentials Host Agent/CIM Extension
5	RAID Manager Library
6	Command LUN
7	Hitachi Device Manager (HDvM)
8	Service Processor (SVP)

9	Controllers
10	Data

Refer to the HP Storage Essentials Installation Guide for information about setting up the host proxy.

The array group, containing the command device LUN that was chosen to gather statistics on the host proxy, will have increased input/output traffic because this command device LUN generates additional input/output traffic as a result of its gathering statistics. Statistical data is returned to the management server where it is stored in the database.

XP and HDS Array-Related Software Requirements

The XP and HDS Array Performance Pack requires the RAID Manager Library on the host proxy, as well as CV Advanced Edition or the built-in provider.

Host Proxy requirements:

- CIM extension – The host proxy for the XP array must have a CIM extension installed and running on the proxy. The Operating System for the host proxy server is limited – Windows, Linux, HP-UX, others. See the support matrix for your edition for the full list.
- RAID Manager Lib XP (RMLIBXP) is not included in the CIM extension installation package: it must be installed separately on the host proxy server. For more information about locating and installing RAID Manager Lib XP, refer to your XP Product documentation.

Refer to the support matrix for your edition for a more detailed listing of required software and their versions, including the required version software of your XP Array.

Specifying XP and HDS Data Collectors

The new collectors provided by the XP and HDS Performance Pack include:

- Array Group Collector
- Backend Controller Collector (only if built-in provider is used)
- Frontend Controller Collector
- Storage System Collector
- Volume Collector
- For Port Collection:
  - Storage System Port Observer (for pre-RAID500 arrays, XP128K, XP512K, XP1024K, requires switch to be discovered)

    Or

  - Host Fibre Port Collector (for all other arrays)

To choose collectors:

1. On the management software home screen, click **Configuration**.
2. Click **Performance**. The Data Collection screen appears.
3. Select the **Data Collection** tab.
4. Select the desired collectors from those listed in your display.
5. Click **Start Selected** for multiple collectors, or click **Action** for single collectors.

   You could also start and run the collectors only for specific needs, rather than running all of them all the time; otherwise it might affect overall system response.
6. To edit the start time and interval for multiple collectors, click **Spread Start Time** to stagger the starting times for multiple collectors. This minimizes the impact on system performance. The most frequent data collection interval is 3 minutes.
7. Review the collected data in Performance Manager. (You can view this data in Performance Manager only if you are licensed appropriately.) See Viewing XP and HDS Array Data on next page for more information.

Configure and enable the collectors for the array(s) to be monitored. Pay particular attention to the date/time specified for the first data collection. By default this will be up to 1 hour from current time. To start the data collectors more quickly, set the start date/time to a few minutes in the future rather than the default hour.

For more information on Configuring and Enabling performance collectors, see Managing Performance Collection on page 372.

XP and HDS Metrics Collection Considerations

Performance Manager requires two points to plot the first data point. Therefore, depending on the collector setup, interval, and other factors, it can take a while for the data to begin to display, in some cases a couple of hours.

The minimum collection interval that can be set for the XP and HDS performance data collectors is 1 minute. However, collecting at frequencies less than 5 minutes will increase the amount of data in the database. HP recommends using an interval between 15 minutes and 1 hour. Optionally, set the collection interval lower while analyzing a problem, and then restore the collector interval to the default. The interval for real time is always 20 seconds.

Non-performance data can be collected from the array using either the CV SMI-S provider or the built-in XP or HDS provider; however, you are required to use the built-in provider to collect back-end controller statistics. You will only see front-end controllers in the Performance Manager display when you connect to the array in Step 1 Discovery using CV SMI-S provider. Use the built-in provider in Step 1 Discovery to collect Back-end Controller statistics.

When collecting performance statistics for XP and HDS logical array groups, be aware that a parity group is divided into several logical RAID groups. Thus, for example, parity group X-X can contain logical array groups X-X-1, X-X-2, and so on. The collecting of performance statistics for logical array groups is only done at the parity group level. Since the Performance Manager for XP and HDS array statistics does not show the parity group, all the statistics are gathered at the first logical RAID group only (in our example, X-X-1), and no statistics are gathered at X-X-2 and the other subsequent logical RAID groups.

Refer to the available XP Array white papers for more information and recommendations for your XP array, especially regarding data collection. You can find white papers at:

http://h18006.www1.hp.com/storage/softwhitepapers.html

Viewing XP and HDS Array Data

You can review the collected data in Performance Manager. You can expand the device sub-elements (controllers, array group, etc.) in the navigation tree.

The list of sub-elements and metrics varies according to the device type. The selections determine the performance data retrieved and available for analysis.

You can select multiple metrics or historical data when holding down the Ctrl button. The following example shows information about data rates associated with the highlighted XP array host port. (The IP addresses have been hidden.)

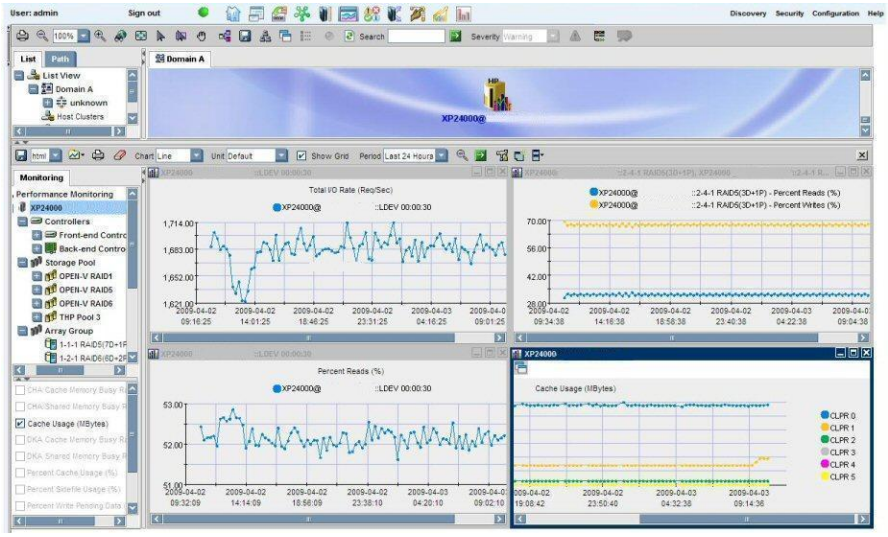

You can also right-click certain objects in the Performance tree. Pop-up menus like Properties take you to the Properties page, based on the object selected.

Managing Late Data or Errors from XP and HDS Arrays

If you are performing real-time data collection, and the element is not returning the information in time, you will see a message like the following displayed in red:

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

The message means the software cannot obtain the information in a timely manner because one or more of the following occurred:

- The built-in provider went down.
- Someone logged into the built-in provider for the array on the Service Processor (SVP) and changed the “View” mode to “Modify”.

When this happens, in most cases, statistics will continue to be collected because they are collected using the host proxy, not through the provider.

However, if Step 3 Discovery is attempted when the provider is down, statistics will stop being collected. The error that would appear in the cimom.log is the following:

```
HDS_MISSING_CONNECTION=Connection information is missing for Storage
Array with Serial Number {0}. Try rediscovering the access points in
Discovery Setup.
```

To resolve the error, restart the provider and then redo Step 1 Discovery for the array.

You also might want to look in the appstorm and cimom logs around the time the error occurred to see if there are any error messages that might explain why the management server failed to collect the data. The product provides substantial error handling and logging for failed XP and HDS performance data collections.

XP Collected Performance Statistics

The following tables show the performance statistics that are collected by the XP and HDS Array Performance Pack.

Volume and Array Group Metrics

Statistic	Units	Notes
Total IO Rate	req/sec	(difference of RandomReadIOs + difference of RandomWriteIOs + difference of SequentialReadIOs + difference of SequentialWriteIOs) / difference of StatisticTime
Total Data Rate	bytes/sec	(difference of RandomReadData + difference of RandomWriteData + difference of SequentialReadData + difference of SequentialWriteData) / difference of StatisticTime
Read Rate	req/sec	(difference of RandomReadIOs + difference of SequentialReadIOs) / difference of StatisticTime
Read Cache Hit Rate	req/sec	(difference of RandomReadHitIOs + difference of SequentialReadHitIOs) / difference of StatisticTime
Read Data Rate – Random	bytes/sec	difference of RandomReadData / difference of StatisticTime

Statistic	Units	Notes
Read Data Rate – Sequential	bytes/sec	difference of SequentialReadData/ difference of StatisticTime
Read Data Rate	bytes/sec	(difference of RandomReadData + difference of SequentialReadData) / difference of StatisticTime
Write Rate	req/sec	(difference of RandomWriteIOs + difference of SequentialWriteIOs) / difference of StatisticTime
Write Data Rate – Random	bytes/sec	difference of RandomWriteData / difference of StatisticTime
Write Data Rate – Sequential	bytes/sec	difference of SequentialWriteData/ difference of StatisticTime
Write Data Rate	bytes/sec	(difference of RandomWriteData + difference of SequentialWriteData) / difference of StatisticTime
Read Response Time	sec	difference of ReadResponseTimeCounter / (difference of RandomReadIOs + difference of SequentialReadIOs)
Write Response Time	sec	difference of WriteResponseTimeCounter / (difference of RandomWriteIOs + difference of SequentialWriteIOs)
Percent Reads	%	100*(difference of RandomReadIOs + difference of SequentialReadIOs)/ (difference of RandomReadIOs + difference of RandomWriteIOs + difference of SequentialReadIOs + difference of SequentialWriteIOs)
Percent Writes	%	100*(difference of RandomWriteIOs + difference of SequentialWriteIOs) / (difference of RandomReadIOs + difference of RandomWriteIOs + difference of SequentialReadIOs + difference of SequentialWriteIOs)
Read Cache Hit %	%	100*(difference of RandomReadHitIOs + difference of SequentialReadHitIOs) / (difference of RandomReadIOs + difference of SequentialReadIOs)

Statistic	Units	Notes
Read Cache Hit % – Sequential	%	100*difference of SequentialReadHitIOs / difference of SequentialReadIOs
Read Cache Hit % – Random	%	100*difference of RandomReadHitIOs / difference of RandomReadIOs
Average Read Size	bytes	(difference of RandomReadData + difference of SequentialReadData) / (difference of RandomReadIOs + difference of SequentialReadIOs)
Average Write Size	bytes	(difference of RandomWriteData + difference of SequentialWriteData) / (difference of RandomWriteIOs + difference of SequentialWriteIOs)
Utilization	%	100*difference of ActiveTime fields / difference of StatisticTime The Utilization statistic is calculated differently than the RAID Group Utilization Percent statistic in HP StorageWorks XP Performance Advisor software.

Controller Metrics

Statistic	Units	Notes
Processor Utilization	%	difference of BusyTimeCountern / difference of ElapsedTimeCounter. Only a single Processor Utilization is presented as a controller metric. When the metric is requested, the values from all 8 database columns are returned concurrently (and displayed concurrently on the same chart).

Host FC Port Metrics

Statistic	Units	Notes
Total IO Rate	IO/sec	difference of TotalIOs / difference of statisticTime
Total Data Rate	bytes/sec	difference of TotalIDataTransferred / difference of statisticTime

Cache Metrics

Statistic	Units	Notes
Total Read Hit Rate	IO/sec	difference of ReadHitOs/difference of StatisticTime
Percent Write Pending Data	%	WritePendingData/CacheSize
Percent Cache Usage	%	CacheUsage/CacheSize
Percent Sidefile Usage	%	SidefileUsage/CacheSize

General Considerations for Performance Manager

Check the support matrix for version requirements and the Release Notes for late-breaking documentation about Performance Manager.

Consider the following when running Performance Manager:

- If you see the message "`There is not enough data to produce a chart [chart_title] at this time`," lessen the frequency option or select **All** in the **Edit Chart Property** dialog box.
- Verify the performance collector for that element is enabled (**Configuration** > **Performance**). See Managing Performance Collection on page 372 for more information about enabling performance collectors.
- Direct attached storage ports are not displayed in the storage tree in the bottom left pane.
- Drives such as Veritas Volume Manager Volumes and Multipathing Volumes that do not support statistics collection are incorrectly shown in the Performance Manager user interface.
- Performance Manager is not available to file servers.
- To learn more about the buttons in the toolbar, see Toolbars in Performance Manager on page 614.
- If one or more scheduled data points on a graph seem to be missing, an error could ahve occurred while collecting data at that time. Check the CIMOM log for errors for the time frame covering the sample interval of the missing data points.

All collectors are stopped during Get Details. This means that during Get Details, data for Performance Manager is not updated. Historical collectors, such as those available from the Configuration tab, are restarted when they are stopped during Get Details. Charts that were active in Performance Manager when Get Details was started are not restarted.

Accessing Performance Manager

To access Performance Manager, select **Performance Manager** ().

Creating Performance Charts

To create a performance chart for an element:

1. Access Performance Manager as described in Accessing Performance Manager on previous page.
2. Select the element you want to monitor.
3. Under the **Monitoring** tab in the lower-left pane, select the element again. In some instances, you might need to select an element component, such as a port on a switch.
   Note: To display a statistic for multiple elements in the chart, hold down the SHIFT or CTRL key and select elements in the tree.
4. In the pane under the tree, select a monitoring option.
5. Use the Chart and Unit combo box to modify the chart. When done, click the  button in the lower pane.

To monitor more than one element in a chart, see Comparing the Performance of Different Elements on page 618

When bar charts throughout the product are shown for multiple objects, the first and last bar or bars in the chart might not show. The producers of the charting package used within the management server user interface have been notified of this issue.

Saving Chart Settings

To save chart settings:

1. Create performance charts with the settings you want to save. All of the charts currently open will be saved in the chart settings group.
2. Click the  button in the lower pane, and select Save Chart Settings from the drop-down menu. The Save Chart Settings window displays.
3. Enter the name of the new chart group in the Chart Group box, or select an existing chart group to overwrite that group.
4. Click **Save**. The following settings are saved:
   - Elements
   - Observations
   - Chart type
   - Unit
   - Grid
   - Period (custom)

- Line size
- Frequency
- Threshold
- Trend

Opening Saved Charts

To open saved charts:

1. Click the  button in the lower pane and select Open Charts from the drop-down menu.
2. Select a chart group and click **Open**. All chart observations and settings are restored, but with current data.

Custom period ranges are restored with the same period that was saved. Predefined periods (such as Last 24 hours) display data based on when the template is applied, not when it was saved.

If any charts have errors (such as a missing element), an error message is displayed. Click **OK** and the remaining charts in the saved chart group appear.

Deleting Chart Groups

To delete a chart group:

1. Click the  button in the lower pane, and select Delete Charts from the drop-down menu.
2. Select the chart group you want to delete, and click **Delete**. A confirmation dialog box appears.
3. Click **Yes**. The chart group is deleted.

Toolbars in Performance Manager

Performance Manager provides two toolbars, one in the upper pane and another in the lower pane. The toolbar in the upper pane is the same as the one in System Manager. See Toolbars in Performance Manager above for information about the toolbar in the upper pane.

Toolbar in Lower Pane of Performance Manager

Icon	Description
	Saves the graph in HTML, XLS, or CSV.
	Opens saved charts, saves chart settings, and deletes chart groups. For additional details, see Saving Chart Settings on previous page.
	Lets you print a graph.

Icon	Description
	Clears the graph of the elements you selected.
	Enables you to determine the type of graph displayed. Select one of the following options, and then click the  button: • **Bar** – Displays each data point as a bar. The data for the different elements is displayed side by side. • **Line** – Displays each data point as a dot with a line connected to the previous data points. The data for the different elements for a specific point in time is displayed in the same column. • **Summary** – Displays a single line that summarizes the values for a single statistic. Multiple statistics can be shown with multiple lines. See Viewing Summary Charts on page 618.
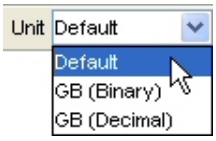	Enables you to determine the unit of measurement in the graph. Select one of the following options, and then click the  button: • **Default** – Displays the data in its default unit, which is usually megabytes. • **GB (Binary)** – Displays the data in binary gigabytes. (The computer handles the data in binary format. A gigabyte is equal to 1,073,741,824 bytes.) • **GB (Decimal)** – Displays the data in gigabytes. (A gigabyte is equal to 1,000,000,000 bytes.)

Icon	Description
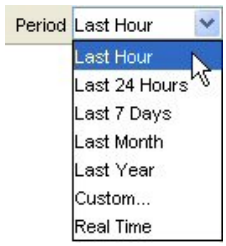	Enables you to format the graph to provide data collected during the time period specified. Select the option to the left of the **Period** combo box. Select one of the following options and click the  button: • **Last Hour** – Information collected during the last hour is reported. If you select Last Hour, the only frequency available is All. • **Last 24 Hours** – Information collected during the last 24 hours is reported. • **Last 7 Days**– Information collected during the last seven days is reported. • **Last Month**Information collected during the last 30 days is reported. “Last Month” does not mean the previous month on the calendar. • **Last Year**– Information collected during the last 365 days is reported. “Last Year” does not mean the previous calendar year. • **Custom...** – Lets you created custom settings for the graph. • **Real Time** – Displays the data as it is currently being gathered by the management server. If you select Real Time, the only option for frequency is All and unit is Default. To set a custom period, select **Custom**. See Setting a Custom Period on page 620 for more information.
	Zoom-out button enables you to return to the original data set.
	Applies the current period, frequency and trending information to the chart.
	This icon appears in the tool bar of the chart. It enables you to filter out the additional data series if the chart contains more than one series of data.

Icon	Description
	Enables you to modify the performance data displayed in the graph and change graph settings. When you select it, the following options are displayed: Performance data in the graph Data and Delete note pad – Displays elements on the chart (including statistics for the element) and allows you to delete some or all of them. Specify graph settings Chart: • **Type** – Allows you to select from the line, bar, or summary chart • **Unit** – Scales the y-axis by Gigabyte. You can select either decimal or binary value. • **Line size** – Affects the size of the line for line charts. • **Frequency** – Allows you to chart information at specified intervals. • **Threshold Visible** – Mandatory field to set a threshold or graph an existing threshold. • **Threshold Value** – Actual value of the threshold identified by a red line displayed on line and bar charts. • **Show Grid** – Check this field to display the Chart Grid. Period: • **Period bullet** – Lets you format the graph to provide data within the time period specified • **Custom** – Allows you to customize a time period value by selecting start and end times from a calendar pop-up. • **Trend**– Specifies the number of trending data points you want to see. Trending information only applies to line and bar charts with a single statistic. The time period of the trend is determined by the frequency setting.
	Creates a chart window.

Icon	Description
	Enables you to arrange the chart windows as follows: • **Tiled** – Displays the windows within their own tile. • **Cascade** – Displays the windows on top of each other, with the active window on top. • **Minimize All** – Minimizes all windows. • **Restore All** – Opens all windows that have been minimized.

Comparing the Performance of Different Elements

Use Performance Manager to compare the performance of different elements.

The following example describes how to compare the physical memory used on different hosts.

1. Access Performance Manager as described in Accessing Performance Manager on page 612.
2. Click the element in the topology whose performance you want to see. For example, click Cortez.
3. Select a component of Cortez.
4. Select the graph monitoring option you want to see on the Monitoring tab, located in the lower pane. For example, select Physical Memory Used (%).
5. Scroll to the bottom of the storage tree in the lower-left pane.
6. Expand the **All Elements** node by clicking the (+) symbol next to the node name.
7. Expand the node for the host you want to compare. For example, to compare a host name HostA against Cortez, expand the following nodes in order: **All Elements** > **Hosts** > **HostA** > **Element Paths** > **Application Path**.
8. Select the same type of component you selected previously for Cortez.
9. Select the same type of graph monitoring option in the left pane. For example, if you selected “Physical Memory Used (%)” for the first element, you must select the same option for the other elements. You cannot select “Physical Memory Used (%)” for the first element and then “Processor Utilization (%)” for the second element. This is because these two options measure different types of data. Performance Manager displays information for the different elements in the same graph.
10. Repeat steps 7 through 9 for each host you want to compare against Cortez.

Viewing Summary Charts

Performance Manager provides summary charts which display a single line that summarizes the values for a single statistic. Multiple statistics can be shown with multiple lines.

You cannot see more than one element at a time on a chart that has multiple elements.

The line has the following attributes:

- The vertical length of the line indicates the minimum and maximum value of the statistic within the selected data time frame.
- The green marker marks the median.
- The blue marker marks the average value.

Summary charts are currently not supported for real time performance display.

To view a summary chart:

1. Access Performance Manager as described in Accessing Performance Manager on page 612.
2. Create a chart as described in Creating Performance Charts on page 613.
3. Select **Summary** from the Chart combo box in the lower pane.
4. Click the  button in the lower pane for the change to take effect.

Viewing Trending Information for Performance

The management server can display trending information in its charts. For example, you can configure Performance Manager to display trending information for the next week. This information can give you an indication of an element's future performance based on past performance.

An element's performance can drastically change in the future. Keep in mind that the data trends are just assumptions and should not be treated as fact.

Trending requires at least two sets of data gathered within the frequency specified.

To view trending information:

1. Access Performance Manager as described in Accessing Performance Manager on page 612.
2. Click the element whose performance you want to see.
3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you might need to select an element's port, such as a switch.
4. In the lower-left pane on the Monitoring tab, click a performance-monitoring option.

   The performance monitoring options listed in this figure vary according to the type of element. The monitoring buttons in the figure are for a switch.
5. Click the **Edit Chart Property** () button in the lower pane.
6. In the Performance Graph Editing Dialog window, enter a number in the Trend box.

The number corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you enter 5 in the Trend box, the chart provides trending information for five frequency intervals, such as five weeks if weeks was selected from the Frequency menu. If the frequency is set to **All**, the trend interval is an hour.

7. Click **OK**. The trending information is displayed in the chart.

If there is not enough data to display, Performance Manager does not display the chart. For example, if you selected the weekly option from Frequency menu and you only have two days of data, a chart is not displayed, regardless of the value in the Trend box. Performance Manager does not display a chart if there is not enough data, and the trending number is ignored.

Removing Performance Data from a Graph

To remove multiple data from a performance graph:

1. Access Performance Manager as described in Accessing Performance Manager on page 612.
2. In Performance Manager, click the graph you want to modify in the bottom-right pane.
3. Click the **Edit Chart Property** () button in the lower pane.
4. In the Performance Graph Editing Dialog window, click the **Delete** () button corresponding with the data you want to remove from the graph.
5. When you are asked if you want to remove the data, click **Yes**.
6. Click **OK**.
7. Click the  button in the lower pane.

Setting a Custom Period

You can format the graph to provide data within a custom time period. This feature can be extremely useful for pinpointing performance changes. For example, assume you changed the firmware of a switch two weeks ago, and you want to compare the performance of that switch before and after you changed its firmware. You could create a graph that provides performance data two weeks before you changed the firmware. You could then create another graph that provides performance data two weeks after you changed the firmware.

To set a custom period:

1. Access Performance Manager as described in Accessing the Management Server on page 60.
2. Click the element whose performance you want to see.
3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you might need to select an element's component, such as a port on a switch.

4. In the lower-left pane on the Monitoring tab, click a performance monitoring option. The performance monitoring options vary according to the type of element.
5. Select **Custom** from the Period combo box in the lower pane.

   The Performance Graph Editing Dialog window appears when you select the Custom option and there are no previous custom settings.
6. In the Performance Graphic Editing Dialog window, select the Custom option near the lower-left corner.
7. Click the calendar button to the right of the Start box.
8. Enter the time in the time box. Make sure the time resembles a 24-hour clock; for example, 22:00 for 10 p.m.
9. Click the date. The date is highlighted in pink.
10. You can navigate the calendar as follows:

    - Displays the same month in the previous year

    - Displays the previous month

    - Displays the next month

    - Displays the same month in the following year
11. When you are done, click **Set**. The start time and date are displayed in the Start box.
12. To set the end date, click the calendar button to the right of the End box. Repeat steps 8 through 10. The ending time and date are displayed in the End box.
13. Click **OK**.
14. Click the button.

Monitoring with Direct Attached Storage

A port on a storage system that is directly attached to a host does not appear in the left panel for monitoring. To monitor the port, connect the port to a switch.

Supported Host Configurations for Monitoring

The table in this section describes which host configurations the management server can monitor. Supported configurations are indicated with **Y**, and unsupported configurations are indicated with **N**.

Unsupported configurations can still obtain processor and memory statistics from the host, with the exception of Windows NT 4, which does not provide any monitoring information to the management server.

If the host has several configurations listed, and one of them cannot be monitored, then monitoring is not supported for any configuration on the host. For example, assume you have a host with Solaris 9 Sun Foundation Suite Leadville with MPXIO and Solstice Disksuite/Volume Manager. Even though the management server supports monitoring for Solstice Disksuite/Volume Manager, neither of those devices can be monitored because Solaris 9 Sun Foundation Suite Leadville with MPXIO is not supported, as the following formula shows:

```
A monitorable configuration (Y) + an unmonitorable configuration (N)=
unmonitorable configuration (N)
```

In all configurations, you cannot monitor a VCM database device.

The product does not differentiate slides that are allocated by Solstice DiskSuite,and it cannot detect the size of the disk partitions on a host using Solstice DiskSuite.

The following table is not a complete list of all products in the marketplace.

Host Monitoring Support

Host Configuration	Monitoring Supported?
AIX 5.3, 6.1	Y
AIX 5.3, 6.1 SDD	N
AIX 5.3, 6.1 SDDPCM	N
AIX 5.3. 6.1 with AMS	N
HP-UX 11.x Itanium	Y
HP-UX 11i	Y
HP-UX 11.0	Y
HP-UX 11i LVM	Y
HP-UX 11.0 LVM	Y
HP-UX 11i PV Link	Y
HP-UX 11.0 PV Link	Y
HP-UX 11i PV Link Volumes	Y
HP-UX 11.0 PV Link Volumes	Y
HP-UX 11.i with PowerPath	Y
HP-UX 11.0 with PowerPath	Y
HP-UX 11.0 with HDLM	Y

Host Configuration	Monitoring Supported?
HP-UX 11.i with HDLM	Y
HP-UX iv1 (11.11) PA-Risc	Y
HP-UX iv2 (11.23) PA-Risc, Itanium	Y
HP-UX iv3 (11.31) PA-Risc, Itanium	Y
Irix 6.5.x	Y
Irix 6.5.x XVM	Y
Irix 6.5.x CXFS	Y (only on node sending I/O)
Redhat 4.0	Y
Redhat 4.0 Itanium	Y
Redhat 5.0	Y
Redhat 5.0 Itanium	Y
Solaris 9,10	Y
Solaris 10 (x86_64)	Y
Solaris 9,10 PowerPath	Y
Solaris 9,10 HDLM	N
Solaris 9,10 RDAC	N
Solaris 9,10 Sun Foundation Suite Leadville	Y
Solaris 9,10 DAS	N
Solaris 9,10 Sun Foundation Suite Leadville + MPXIO	N
Solaris 9,10 Sun Foundation Suite Leadville + MPXIO + Solstice Disksuite/Volume Manager	N
SuSe LINUX Enterprise Server Release 9 SP4	Y
SuSe LINUX Enterprise Server Release 10, SP1, SP2	Y
VMware ESX 3.5	Y

Host Configuration	Monitoring Supported?
VMware ESX 3.5 Update 2	Y
VMware ESX 3.5 Update 3	Y
VMware ESX 3.5 Update 4	Y
VMware ESX 3.5 Update 5	Y
VMware ESXi 3.5	Y
VMware ESXi 3.5 Update 2	Y
VMware ESXi 3.5 Update 3	Y
VMware ESXi 3.5 Update 4	Y
VMware ESXi 3.5 Update 5	Y
VMware ESX 4.0	Y
VMware ESX 4.0 Update 1	Y
VMware ESX 4.0 Update 2	Y
VMware ESX 4.1	Y
VMware ESXi 4.0	Y
VMware ESXi 4.0 Update 1	Y
VMware ESXi 4.0 Update 2	Y
VMware Virtual Machines	Y (only if the virtual machine is discovered via CIM extension. Refer to the table entry for the guest operating system.)
Windows 2008	Y
Windows 2003 HDLM	N
Windows 2003 Volume Manager	N
Windows 2003	Y
Windows NT 4	N
Windows 2000	N

Sudden Dips in Charts in Performance Manager

In Performance Manager on the Monitoring tab, charts that display data gathered by certain volume and drive counters display their charts with the results increasing to the maximum value and then decreasing rapidly to a very low number that starts rising again.

The sudden dip in the charts occurs because the counters holding the values displayed in the charts revert back to zero after reaching the upper limit, and continue to go up again. The counters do not usually display zero in a chart because they are quickly gathering data again.

The following charts for individual drives are affected: ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, Unrecovered Errors, Recovered Errors, Timeouts, Retried Requests.

The following charts for individual volumes are affected: Bytes Read, Bytes Read Large, Bytes Written, Bytes Written Large, ReadIOs Large, WriteIOs Large, ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, ReadHitIOs.

Values Increase in Charts for Aggregated Drives and Aggregate Volumes

Values in charts for aggregate drives and aggregate volumes continue to rise smoothly in Performance Manager and on the Monitoring tab. The only time the values in the charts come back down to the values returned by the array is when the service for the management server restarts, and then the charts display the values returned by the array. The values in the charts continue to rise until the service for the management server is stopped.

The values continue to rise because they do not revert back to zero when an individual drive or volume counter that makes up the aggregation contains the maximum number of integers. Instead, they continue to rise smoothly.

The following charts for aggregate drives are affected: ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, Unrecovered Errors, Recovered Errors, Timeouts, Retried Requests. Charts for aggregated drives are also provided for arrays, controllers, and volume groups.

The following charts for aggregate volumes are affected: Bytes Read, Bytes Read Large, Bytes Written, Bytes Written Large, ReadIOs Large, WriteIOs Large, ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, ReadHitIOs. Charts for aggregated drives are also provided for arrays, controllers, and volume groups.

Using Performance Metrics

HP Storage Essentials provides performance management metrics for the following:

- Host performance metrics — See HP Storage Essentials Performance Management – Host-Related Metrics on next page.

- HP StorageWorks Enterprise Virtual Array (EVA) metrics — See HP Storage Essentials Performance Management – EVA Metrics on page 639.
- HP XP array metrics — See HP Storage Essentials Performance Management – XP Metrics on page 649

HP Storage Essentials Performance Management – Host-Related Metrics

HP Storage Essentials provides the following host-related performance metrics:

- Host Performance Metrics
- Host Disk Performance Metrics
- Host Bus Adapter (HBA) Performance Metrics
- ESX Server Performance Metrics
- Microsoft Exchange Server Performance Metrics
- Oracle Server Performance Metrics
- Microsoft SQL Server Performance Metrics
- Sybase Performance Metrics
- IBM DB2 Performance Metrics
- Informix Performance Metrics
- InterSystems Caché Database Performance Metrics
- Switch Performance Metrics

For more detailed information, see related topics for each metric.

Host Performance Metrics

The following metrics can be used to measure host performance.

Metric	Description	Units	Use To...
Free Physical Memory	Amount of physical memory available	KB	Measure available main memory for additional processes and threads.
Free Virtual	Amount of virtual (paged) memory available	KB	Measure physical memory optimization and availability over time.
Physical Memory Used	Percentage (%) of main memory used by processes running on the host	%	Track memory utilization trends.

Metric	Description	Units	Use To...
Processor Utilization	Shows total CPU utilization percentage (%) for all processes running on the host	%	Identify CPU bottlenecks.
Virtual Memory Used	Percentage (%) of paging file that is in use	%	Determine if a system has insufficient physical memory.

Host Disk Performance Metrics

Metric	Description	Units	Use To...
Disk Read	Average time in seconds to read data from disk	KB/s	Compare read times for a given application (for example, read compared to writes).
Disk Total	Total read and write requests in seconds	KB/s	Test maximum throughput.
Disk Utilization	Based on the IRP (I/O request packets) round trip times the Average Disk Sec/Transfer. Indicates how busy a physical disk is over time	%	Determine the average disk utilization for a given application or known number of processes. Utilization indicates how busy a disk is.
Disk Write	Average time in seconds to write data to disk	KB/s	Compare write times for a given application (for example, writes compared to reads).

Host Bus Adapter (HBA) Performance Metrics

Metric	Description	Units	Use To...
Bytes Received	Inbound traffic in megabytes per second at the HBA port	MB/s	Measure network traffic for load balancing, multi-pathing optimization and network performance.
Bytes Transmitted	Outbound traffic in megabytes per second at the HBA port	MB/s	Measure network traffic for load balancing, multi-pathing optimization and network performance.
CRC Errors	Cyclic Redundancy Errors over a period	errors	Isolate CRC errors on a specific initiator or between devices.
Link Failures	Link Failures over a period	failures	Isolate connection failures and the effect on performance.

ESX Server Performance Metrics

The following performance metrics are tracked for the ESX Server.

Metric	Description	Units	Use To...
Free Physical Memory	Amount of physical memory available	KB	Measure available main memory for additional processes and threads.
Physical Memory Used	Physical memory being consumed by all processes	%	Indicate physical memory optimization and availability over a period of time.
Processor Utilization	Total CPU utilization (%) for all processes running on the host	%	Identify of CPU bottlenecks.
Disk Read	Average time in seconds to read data from disk	KB/s	Compare read times for Virtual Machines and their applications.
Disk Total	Total read/writes in seconds	KB/s	Test maximum throughput.
Disk Write	Average time in seconds to write data to the disk	KB/s	Compare write times for Virtual Machines and their applications.
Bytes Transmitted***	Inbound traffic in megabytes per second through the HBA	MB/s	Measure network traffic for load balancing, multi-path optimization, and network performance.
Bytes Received***	Outbound traffic in megabytes per second through the HBA	MB/s	Measure network traffic for load balancing, multi-path optimization, and network performance.
CRC Errors***	Cyclic Redundancy Errors (CRC) over a period of time	errors	Isolate CRC errors on a specific initiator or between devices.
Link Failures***	Link Failures over a period of time	failures	Isolate connection failures and effect on performance.

Microsoft Exchange Server Performance Metrics

This following metrics measure performance for the Microsoft Exchange Server. An asterisk (*) marks a metric that does not collect historical data. All other metrics use collected historical data to provide statistics.

Metric	Description	Units	Use To...
Application Server			
Exchange Services	Services Running/Not Running (Site Replication, Routing Engine, POP3, Exchange Management, MTA Stacks, Info Store, IMAP4, Events)	count	Indicate whether or not a critical Microsoft Exchange service is running.
Final Destination currently unreachable Queue Size*	Final destination currently unreachable queue size. Contains the number of messages that cannot reach their intended destination	count	Set alerts. Possible causes: No route, Microsoft Exchange connector unreachable, no remote delivery queue available.
Messages awaiting directory lookup Queue Size*	Messages awaiting directory lookup queue size. These are messages sent to a Microsoft Exchange server.	count	Set alerts or troubleshoot Event ID 9035, 6004, 9003, 9004. Possible causes: Message Archiving, Insufficient Permissions, or global catalog server issues.
Messages to be routed Queue Size*	Number of messages to be routed	count	Set alerts or troubleshoot message routing issues in Microsoft Exchange.
Pre-submission Queue Size*	Pre-submission queue size	count	Measure the number of message waiting in the pre-submission queue. Possible causes are 3rd party applications, SMTP events, DNS.
SMTP Local Delivery Queue Size*	Number of messages in the local queue	count	Set alerts. Large numbers of messages indicate possible disk I/O issues or dismounted stores.
SMTP Server Queue Summary*	Summary of local and remote queues	count	Set alerts and notifications. Helps pinpoint disk, DNS, or network I/O issues.
Storage Groups			

Metric	Description	Units	Use To...
Log Size (MBytes)	Storage group log file size	MB	Help manage log file sizes in Microsoft Exchange.
Storage Group Size	Size of the Storage Group in MB. Contains one or more mailbox or public folders, transaction logs, and system files depending on the version	MB	Provide a facility to track the Microsoft Exchange group utilization from a higher level than only mailbox or folders sizes.
Message Stores and Public Folders			
Active Client Logons*	Number of logons that have been active (issued by any MAPI requests) within a specific time interval	count	Track the number of sessions/users accessing the Microsoft Exchange server at a given time.
Average Delivery Time*	Rate at which messages are delivered to all recipients	Msg/s	Measure long delivery times that could be caused by deferred delivery queues or stalled messages.
Mail Related Objects Count	Public folder mail related objects count	count	Count the number of mail messages, appointments, meeting requests, tasks, task requests, contacts, remote mail items, and notes.
Receive Queue Size*	Number of messages in the mailbox store's receive queue	count	Track MTA queues, IS queues, and SMTP queues. Possible issues are a related service, looping or corrupt message, or an issue with the source server.
Send Queue Size*	Number of Deferred Delivery or Submission messages	count	Track MTA queues, IS queues, and SMTP queues. Possible issues are a related service, looping or corrupt message, or an issue with the destination server.
Store Size	Mailbox store size	MB	Set alerts. Helps alert administrators when a store size is reaching its limits.

*Not supported with Microsoft 2007

Oracle Server Performance Metrics

The following performance metrics are available for Oracle Server.

Metric	Description	Units	Use To...
Buffer Hit Ratio	Percentage (%) of requests against data block buffer	%	Measure the effectiveness of the Oracle data block buffer. Used with a database that has an undersized DB Cache size where the “working set” of frequently-referenced data has not been cached
Dictionary Hit Ratio	Ratio of logical reads to physical disk reads	%	Monitor ratio balance between logical and physical disk reads. As the hit ratio approaches 100 percent, more data blocks are found in memory, resulting in fewer disk I/Os and faster overall database performance.
File Read Percent	Percentage (%) of sequential reads	%	Measure the performance and efficiency of indexes.
File Total I/O Percent	Percentage (%) of sequential reads and writes	%	Capture overall sequential database I/Os.
File Write Percent	Percentage (%) of sequential writes	%	Identify other bottlenecks like redo waits, bind variables, bulk operations, and index contention.
In Memory Sort Ratio	Percentage (%) of sorts (from ORDER by clauses or index building) that are done to disk compared to in-memory	%	Differentiate and identify the number of disk sorts performed in TEMP tablespace versus those performed in-memory (RAM sorts).
Library Cache Hit Ratio (%)	Monitors the percentage (%) of entries in the library cache that were parsed more than once (reloads) over the lifetime of the instance	%	Set the shared_pool_size large enough to prevent excessive re-parsing of SQL
Parse CPU to Total CPU Ratio	Percentage (%) of reentrant SQL statements compared to re-parsed SQL statements; also referred to as the Parse to Execute ratio	%	Show the number of unique incoming SQL statements or that SQL statements are NOT reentrant.

Metric	Description	Units	Use To...
Redo Buffer Allocation Retries	Total number of retries needed to allocate space in the redo buffer	%	Monitor the buffer allocation retries over a period of time while the application is running. If the Redo Buffer Allocation Retries value is continuously increasing, then increase the LOG_ BUFFER value.
System Event Time Waited	Provides wait state details for ongoing Oracle transactions	ms	Analyze time-based events, and system-wide and session wait events.
Tablespace Read Percent	Percentage (%) of tablespace reads over an interval	%	Monitor tablespace usage and growth.
Tablespace Total I/O Percent	Percentage (%) of total tablespace I/O over an interval	%	Monitor tablespace usage and growth.
Tablespace Write Percent	Percentage (%) of tablespace writes over an interval	%	Monitor tablespace usage and growth.

Microsoft SQL Server Performance Metrics

This following performance metrics are available for Microsoft SQL Server 2005 and 2008.

For a list of supported software versions for Microsoft SQL Server, see the HP Storage Essentials Support Matrix located in any of the top-level directories on the StorageEssentialsDVD.

Metric	Description	Units	Use To...
Blocked Processes	Number of blocked processes in your Microsoft SQL Server database	count	Flag the database administrator (DBA) that a number of blocked processes are present. The DBA can then use a SELECT statement to identify the offending processes.
Buffer Cache Hit Rate	Percentage (%) of pages found in the buffer cache. The % is calculated as the total number of successful cache hits divided by the total number of requested cache lookups.	%	Monitor the buffer cache hit ratio. After a system maintains a steady state of operation, this metric should achieve rates of 90 percent or higher. The buffer cache hit ratio can be increased by increasing the amount of memory available to SQL Server.

Metric	Description	Units	Use To...
CPU Usage Percentage	Percentage (%) of the total available CPU time that Microsoft SQL Server uses during the current interval	%	Identify potential performance issues. A high CPU usage rate can indicate performance bottlenecks. Options for resolving the bottleneck are to add multiple CPUs, allocate resources more efficiently, identify resource intensive applications, and reduce workloads, or upgrade CPUs.
Cache Memory	Total amount of dynamic memory the server is using for the dynamic SQL cache	KB	Determine if the current cache memory is meeting the needs of the SQL Server application.
Dead Locks	Number of lock requests per second that resulted in a deadlock. A deadlock is a situation where two or more competing actions are waiting for the other to finish.	Req/s	Generate an alert for a deadlock condition. You can then identify the SPIDs and the resources that are involved in a deadlock using the database logs.
Lock Requests	Number of new locks and lock conversions per second requested from the lock manager.	Req/s	Identify issues related to data retrieval. A high number of lock requests with low request rate is an indicator that SQL Server must do table scans when retrieving data.
Lock Time Outs	Number of lock requests per second that timed out, including internal requests for NOWAIT locks.	Req/s	Identify areas of lock contention and congestion in the database.
Lock Waits	Lock wait time is the time that a process spends waiting for another process to release a lock	Wait/s	Generate an alert to the DBA when the number exceeds preset thresholds;also provides a database congestion indicator.
Memory Usage Percentage	Number of pages in the procedure cache that are currently allocated to a process.	%	Monitor allocation of memory to processes. A negative number indicates that a process is freeing memory allocated by another process.

Metric	Description	Units	Use To...
Physical I/O Percentage	Percentage of elapsed time that the disk drive was busy servicing read or write requests.	%	Identify disk I/O problems in the SQL server instance. A value greater than 50 percent may indicate an I/O bottleneck.
Plan Cache Hits Ratio	Ratio between plan cache hits and lookups. The plan contains stored procedures, ad hoc and prepared Transact-SQL statements, and triggers.	%	Monitor the database plan cache performance.
Target Server Memory	Amount of dynamic memory the server can consume	KB	Manage database server memory.
Total Server Memory	Committed memory from the buffer pool. This is NOT the total memory used by the SQL Server.	KB	Manage database server memory.
Transactions	Total number of database transactions	count	Monitor and manage load balancing.
User Connections	Total number of database user connections	count	Identify the number of users who are accessing the database.

Sybase Performance Metrics

The following performance metrics are available for Sybase systems.

Metric	Description	Units	Use To...
CPU Usage Percentage	Percentage of the total available CPU time that Sybase instance uses during the current interval	%	Monitor the CPU usage of the Sybase instance. A high CPU usage rate can indicate performance bottlenecks. If you identify an issue, some options are to add multiple CPUs, allocate resources more efficiently, identify resource intensive applications, reduce workload, or upgrade CPUs.

Metric	Description	Units	Use To...
Memory Usage Percentage	Number of pages in the procedure cache that are currently allocated to a process	%	Monitor memory allocation to active processes. A negative number indicates that the process is freeing memory allocated by another process.
Physical I/O Percentage	Percentage of elapsed time that the disk drive was busy servicing read or write requests	%	Identify disk I/O problems in the Sybase instance. A value greater than 50 percent may indicate an I/O bottleneck.

IBM DB2 Performance Metrics

The following performance metrics are tracked for IBM DB2 databases.

Metric	Description	Units	Use To...
Instance Private Sort Memory	Private sort memory is calculated using the following formula: `(db2.sort_heap_ allocated / sheapthres)x 100` In this instance, sort_heap_ allocated is the system monitor element, and sheapthres is a DBM configuration parameter.	%	Track private sort memory utilization. Used to check that there is sufficient heap space to perform sorting and that sorts do not overflow.
Instance Monitor Heap Utilization	Utilization is calculated using the following formula for the Memory Pool Identifier SQLM_HEAP_MONITOR: `db2.pool_cur_size / db2.pool_max_size) x 100`	%	Track the consumption of the monitor heap memory, If this percentage reaches the maximum 100%, monitor operations might fail.

Metric	Description	Units	Use To...
Database Catalog Cache Hit Ratio (%)	Hit ratio indicates (as a percentage %) how well the catalog cache is working to avoid actual accesses to the catalog on disk	%	Monitor the ratio between catalog cache and physical disk reads. As the catalog cache hit ratio approaches 100 percent, more catalog blocks are found in memory, resulting in fewer disk I/Os and faster overall database performance.
Database Lock List Utilization	There is one lock list per database, and it contains the locks held by all applications concurrently connected to the database. The indicator is calculated using the following formula: `(db.lock_list_in_use / (locklist x 4096)) x 100`	%	Track the amount of lock list memory that is being used.
Database Deadlocks	Number of lock requests that resulted in a deadlock. A deadlock is a situation where two or more competing actions are waiting for the other to finish.	lock requests	Track deadlock conditions. The DBA will then identify the SPIDs and the resources that are involved in a deadlock using the database logs.
Database Shared Sort Memory Utilization	Shared sort memory is calculated using the following formula: `(db.sort_shrheap_ allocated / sheapthres_shr)x 100` In this instance, sheapthres_ shr is a atabase configuration parameter.	%	Track shared sort memory utilization for the database . Can be used to determine an appropriate value for the shared sort memory threshold.

Informix Performance Metrics

The following performance metrics are available for Informix.

Metric	Description	Units
Buffer Read	Number of reads from the buffer cache by the database server	count
Buffer Write	Number of writes to the buffer cache by the database server	count
Chunk Read	Total number of reads from the chunk. Use in combination with Chunk Write metric to monitor chunk usage.	count
Chunk Write	Total number of writes to the chunk. Use in combination with Chunk Read metric to monitor chunk usage.	count
DBSpace Reads	Total number of read calls that involve this dbspace. Use in combination with DBSpace Writes metric to monitor dbspace usage.	count
DBSpace Writes	Total number of write calls that involve this dbspace. Use in combination with DBSpace Reads metric to monitor dbspace usage.	count
Disk Reads	Total number of read operations from disk by the database server	count
Disk Writes	Total number of write operations to disk by the database server	count
Page Reads	Number of pages read from disk by the database server	count
Page Writes	Number of pages transferred to disk by the database server	count
Read Cached	Percentage of all read operations that are read from the buffer cache without requiring a disk read by the database server, calculated as follows: `100 x ((buffer_reads – disk_reads) / (buffer_ reads))` As the hit ratio approaches 100%, more data blocks are found in memory. This results in fewer disk I/Os and faster overall database performance.	%
Write Cached	Percentage of all write operations that are buffer writes by the database server, calculated as follows: `100 x ((buffer_writes – disk_writes) / (buffer_ writes))`	%

InterSystems Caché Database Performance Metrics

The following performance metrics are available for InterSystems Caché databases.

Metric	Description	Units
Application Errors	Number of application errors logged Not supported for Caché 5.0.	error count
Global Kills	Number of global kills since startup Only supported for Caché 5.0.	count
Cache Efficiency	Most recently measured cache efficiency. This is measured as Global references / (physical reads + writes).	GlobalRef/IO
Global Refs	Number of Global references since system startup	count
Global Sets	Number of Global Sets and Kills since system startup	count
Globals Per Second	Most recently measured number of Global references per second Not supported for Caché 5.0.	GlobalRef/s
Disk Reads	Number of physical block read operations since system startup	count
Disk Writes	Number of physical block write operations since system startup	count
Journal Entries	Number of entries written to the system journal	count
Logical Requests	Number of logical block requests since system startup	count
Routine Loads	Number of routine loads since system startup Only supported for Caché 5.0.	count
Routine Refs	Number of routine loads and saves since system startup Not supported for Caché 5.0.	count
Routine Saves	Number of routine saves since system startup Only supported for Caché 5.0.	count

Switch Performance Metrics

The following metrics are used by HP Storage Essentials to monitor switch performance.

Metric	Description	Units	Common Use
Aggregated Port Bytes Received	Sum of bytes received for all ports in a switch over an interval	MB/s	Measure inbound traffic for all ports on the switch.
Aggregated Port Bytes Transmitted	Sum of bytes transmitted for all ports in a switch over an interval	MB/s	Measure outbound traffic for all ports on the switch.
Bytes Received	Number of bytes received over a given interval	MB/s	Measure inbound traffic for specific ports on the switch.
Bytes Transmitted	Number of bytes transmitted over a given interval	MB/s	Measure outbound traffic for specific ports on the switch.
CRC Errors	Number of Cyclic Redundancy Check errors over a period of time	errors	Isolate CRC errors on a specific initiator or between devices
Link Failures	Number of link Failures over a period of time	failures	Isolate connection failures and the effect on performance

HP Storage Essentials Performance Management – EVA Metrics

HP Storage Essentials provides these performance metrics for HP StorageWorks Enterprise Virtual Arrays (EVA):

- EVA Storage System Metrics below
- EVA Controller Metrics on page 642
- EVA Disk Group Metrics
- EVA Host FC Port Metrics on page 645
- EVA Volume Metrics on page 647

For more detailed information, see related topics for each metric.

EVA Storage System Metrics

The following metrics are used to track performance for the HP StorageWorks Enterprise Virtual Array (EVA). Storage system metrics are stored in the HPEVASTORAGESYSTEMSTATS table.

Metric	Description	Units	Formula
Total Data Rate*	Rate data can be transmitted between devices for the entire storage system	Bytes/s	(Δ KBytesTransferred x 1024) / Δ Time
Total I/O Rate*	Average number of I/O operations in requests per second for both sequential and non-sequential reads and writes for the entire storage system	Req /s	Δ TotalIOs / Δ Time
Total Volume Avg Read Hit Latency	Average time to complete a read request (from initiation to information receipt) from the array cache memory for all volumes in the array	sec	(Δ ReadHitLatency / 1000) / Δ ReadHitIOs
Total Volume Avg Read Miss Latency	Average time to complete a read request (from initiation to information receipt) from the physical disks for all volumes	sec	(Δ ReadMissLatency / 1000) / Δ ReadMissIOs
Total Volume Avg Read Size	Average data read size for all volumes	Bytes	(Δ KBytesRead x 1024) / Δ ReadIOs
Total Volume Avg Write Latency	Average time to complete a write request (from initiation to receipt of write completion) for all volumes	sec	(Δ KBytesTransferred x 1024) / Δ Time
Total Volume Avg Write Size	Average write size for all volumes	Bytes	(Δ KBytesWritten x 1024) / Δ WriteIOs
Total Volume Data Rate	Rate data can be transmitted between devices for all volumes	Bytes/s	(Δ KBytesTransferred x 1024) / Δ Time
Total Volume Flush Data Rate	Rate at which data is written to physical disks in array	Bytes/s	(Δ FlushKBytes x 1024) / Δ Time)

Metric	Description	Units	Formula
Total Volume Flush Rate	Aggregate of all flush counters: mirror flush, cache flush, host writes to snapshots and snapclones	Bytes/s	Δ FlushRequests / Δ Time
Total Volume I/O Rate	Average number of I/O operations per second for both sequential and non-sequential read and write operations for all volumes	Req/s	Δ TotalIOs / Δ Time
Total Volume Mirror Data Rate	Rate at which data travels across the mirror port to complete read and write requests to all virtual disks	Bytes/s	(Δ MirrorKBytes x 1024) / Δ Time
Total Volume Pct Read IOs	Percentage (%) of read I/O operations per second for both sequential and non-sequential reads for all volumes	%	100 x (Δ ReadIOs / Δ TotalIOs)
Total Volume Pct Write IOs	Percentage (%) of write I/O operations per second for both sequential and non-sequential writes for all volumes	%	100 x (Δ WriteIOs / Δ TotalIOs)
Total Volume Prefetch Data Rate	Rate data is read from the physical disk to cache in anticipation of subsequent reads when a sequential stream is detected	Bytes/s	(Δ PrefetchKBytes x 1024) / Δ Time)
Total Volume Read Data Rate	Rate data is read from the virtual disk by all hosts and includes transfers from the source array to the destination array	Bytes/s	(Δ KBytesRead x 1024) / Δ Time
Total Volume Read Hit Data Rate	Rate at which data is read from the array cache memory because of read hit requests	Bytes/s	(Δ ReadHitKBytes x 1024) / Δ Time)
Total Volume Read Hit Rate	Number of read requests per second completed from the array cache memory	Req/s	Δ ReadHitIOs / Δ Time
Total Volume Read Miss Data Rate	Rate at which data is read from physical disks because the data was not present in the array cache memory	Bytes/s	(Δ ReadMissKBytes x 1024) / Δ Time

Metric	Description	Units	Formula
Total Volume Read Miss Rate	Number of read requests (per second) that were not available from cache memory and therefore were completed from the physical disks instead	Req/s	Δ ReadMissRequests / Δ Time
Total Volume Read Rate	Number of read requests per second completed from a virtual disk that were sent to all hosts	Req/s	Δ ReadIOs / Δ StatisticTime
Total Volume Write Data Rate	Rate data is written to the virtual disk by all hosts; includes transfers from the source array to the destination array	Bytes/s	Δ KBytesWritten x 1024) / Δ Time
Total Volume Write	Number of write requests per second completed to a virtual disk that were received from all hosts	Req/s	Δ WriteIOs / Δ Time

*This metric is a duplicate of an SMI-S statistic.

EVA Controller Metrics

The following metrics describe the performance of the EVA Controller.

Metric	Description	Units	Formula
Average Read Size*	Amount of data read (per second) from physical disk	Bytes	(Δ KBytesRead x 1024) / Δ ReadIOs
Average Write Size*	Amount of data written (per second) to physical disks	Bytes	(Δ KBytesWritten x 1024) / Δ WriteIOs
CPU Percent**	Percentage (%) of time that the central processing unit on the controller is active. A completely idle controller shows 0%. A controller saturated with activity shows 100%.	%	100 x (Δ CpuBusyCounter / Δ StatisticsTime)
Data Transfer Percent**	Similar to % Processor Time except that it does not include time for internal processes not related to host-initiated data transfers	%	100 x (Δ DataTxCounter / Δ StatisticsTime)
Percent Reads*	Percentage (%) of CPU time dedicated to reads.	%	100 x (Δ ReadIOs / Δ TotalIOs)

Metric	Description	Units	Formula
Percent Writes*	Percentage (%) of CPU time dedicated to writes	%	100 x (Δ WriteIOs / Δ TotalIOs)
Read Data Rate*	Rate at which data is read from the controller by all disks	Bytes/s	(Δ KBytesRead x 1024) / Δ Time
Read Latency	Average time it takes to complete a read request (from initiation to receipt of write completion) through the controller	sec	(Δ ReadLatency / 1000) / Δ ReadIOs
Reads*	Rate at which data is read from each host port	Req /s	Δ ReadIOs / Δ Time
Total Data Rate*	Rate at which data can be transmitted between devices for the controller	Bytes /s	(Δ KBytesTransferred x 1024) / Δ Time
Total I/O Rate*	Average number of I/O operations as requests per second for both sequential and non-sequential reads and writes for the controller	Req/s	Δ TotalIOs / Δ Time
Write Data Rate*	Rate at which data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array	Bytes/s	(Δ KBytesWritten x 1024) / Δ Time
Write Latency	Average time it takes to complete a write request (from initiation to receipt of write completion).	sec	(Δ WriteLatency / 1000) / Δ WriteIOs
Writes*	Number of write requests per second completed to a virtual disk that were received from all hosts	Req /s	Δ WriteIOs / Δ Time

*This metric is a duplicate of an SMI-S statistic.

**This metric is collected as both an HP Storage Essentials and an EVAPerf performance statistic.

EVA Physical Disk Metrics

This following metrics track performance statistics for EVA physical disks.

Metric	Description	Units	Formula
Average Read Size*	Amount of data read from physical disk	Bytes	(Δ KBytesRead x1024) / Δ ReadIOs

Metric	Description	Units	Formula
Average Write Size*	Amount of data written to physical disk	Bytes	(Δ KBytesWritten x 1024) / Δ WriteIOs
Drive Latency metric**	Average time to complete read/write requests from the physical disk drive	Sec	(Δ DriveLatency / 1000) / Δ TotalIOs
Percent Reads*	Percentage (%) of CPU time dedicated to reads	%	100 x (Δ ReadIOs / Δ TotalIOs)
Percent Writes*	Percentage (%) of CPU time dedicated to writes	%	100 x (Δ WriteIOs / Δ TotalIOs)
Queue**	Average number of outstanding requests against the physical disk	Req	Δ DriveQueueDepth / Δ Statistic Time
Read Data Rate***	Rate at which data is read from the virtual disk by all hosts, including transfers from the source array to the destination array	Bytes/s	(Δ KBytesRead x 1024) / Δ Time
Read Latency**	Average time to complete a read request (from initiation to information receipt) from the array volume	Sec	(Δ ReadLatency / 1000) / Δ ReadIOs
Reads***	Rate at which data is read from each host port	Req/s	Δ ReadIOs / Δ Time
Total Data Rate*	Rate at which data can be transmitted between devices for the host port	Bytes/s	(Δ KBytesTransferred x 1024) / Δ Time
Total I/O Rate*	Average number of I/O operations (requests per second) for both sequential and non-sequential reads and writes for the host port	Req/s	Δ TotalIOs / Δ Time
Write Data Rate***	Rate at which data is written to the virtual disk by all hosts, including transfers from the source array to the destination array	Bytes/s	(Δ KBytesWritten x 1024) / Δ Time
Write Latency**	Average time to complete a write request (from initiation to receipt of write completion)	Sec	(Δ WriteLatency / 1000) / Δ WriteIOs
Writes***	Number of write requests per second completed to a virtual disk that were received from all hosts	Req/s	(Δ KBytesRead x 1024) / Δ Time

*This metric is a duplicate of an SMI-S statistic.

**This metric is collected as an HP Storage Essentials and an EVAPerf performance statistic.

***This metric is collected by HP Storage Essentials and EVAPerf and is a duplicate of an SMI-S statistic.

EVA Host FC Port Metrics

For EVA storage systems in which HP Storage Essentials supports performance data collection, the following metrics are collected for EVA Host FC

Metric	Description	Units	Formula
Average Read Size	Amount of data read (per second) from physical disks	Bytes	(Δ KBytesRead x1024) / Δ ReadIOs
Average Write Size	Amount of data written (per second) to physical disks	Bytes	(Δ KBytesWritten x 1024) / Δ WriteIOs
Bad CRC	Number of bad CRC errors. Indicates that the Cyclic Redundancy Check (CRC), which compares a data stream against a stored checksum, has found the data stream changed and therefore no longer reliable. Use to help the transmitter detect errors in the frame that are caused by bad writes, bad media, damaged links/hardware, excessive link errors, and transfer rates.	count	–
Bad Receive Characters	Number of bad receive characters in the bit stream. Use to determine the number of bad frames associated with the Bad CRC metric above.	count	–
Discard Frames	Number of frames discarded due to Bad CRCs. Frames are the basic unit of communication between two N_ports, and are composed of a starting delimiter, header, payload, CRC, and end delimiter.	count	–
Link Fail	Number of link failures. Use to find issues with the fiber optic cable or transceiver or the SAN infrastructure.	count	–
Loss of Signal	Number of times the receiver reports loss of signal. Indicator that fiber optic signal no longer exists. Use to assist in troubleshooting signal loss.	count	–

Metric	Description	Units	Formula
Loss of Sync	Number of times the receiver logic reports loss of sync has timed-out. Use to determine the number of times an intermittent loss of synchronization in communication signals was received by an enclosure connected to a Fibre Channel (FC) loop.	count	–
Percent Reads	Percentage (%) of CPU time dedicated to reads.	%	100 x (Δ ReadIOs / Δ TotalIOs)
Percent Writes	Percentage (%) of CPU time dedicated to writes	%	100 x (Δ WriteIOs / Δ TotalIOs)
Protocol Error	Number of errors in the protocol between the channel and the control unit. Use to differentiate between protocol errors and link errors.	count	–
Queue Depth*	Average number of outstanding host requests against all virtual disks accessed through this host port	Req	Δ QDepth / Δ Time
Read Data Rate	Rate at which data is read from the controller by all disks	Bytes/s	(Δ KBytesRead x 1024) / Δ Time
Read Latency**	Average time to complete a read request (from initiation to receipt of write completion) through the controller	sec	(Δ ReadLatency / 1000) / Δ ReadIOs
Reads**	Rate at which data is read from each host port	Req /s	Δ ReadIOs / Δ Time
Receive Abnormal End of Frame	Number of times a bad frame was detected during data transmission	count	–
Total Data Rate*	Rate in which data can be transmitted between devices for the host port	Bytes/s	(Δ KBytesTransferred x 1024) / Δ Time
Total I/O Rate*	Average number of I/O operations as requests per second for both sequential and non-sequential reads and writes for the host port	Req /s	Δ TotalIOs / Δ Time
Write Data Rate**	Rate at which data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array	Bytes/s	(Δ KBytesWritten x 1024) / Δ Time

Metric	Description	Units	Formula
Write Latency**	Average time to complete a write request (from initiation to receipt of write completion)	sec	(Δ WriteLatency / 1000) / Δ WriteIOs
Writes**	Number of write requests per second completed to a virtual disk that were received from all hosts	Req /s	Δ WriteIOs / Δ Time

*This metric is a duplicate of an SMI-S statistic.

**This metric is collected as an HP Storage Essentials and EVAPerf performance statistic.

EVA Volume Metrics

The following metrics are used to track performance statistics for HP StorageWorks Enterprise Virtual Array (EVA) volumes.

Metric	Description	Units	Formula
Average Read Size*	Amount of data read (per second) from physical disks	Bytes	(Δ KBytesRead x 1024) / Δ ReadIOs
Average Write Size*	Amount of data written (per second) to physical disks	Bytes	(Δ KBytesWritten x 1024) / Δ WriteIOs
Flush Data Rate	Rate at which data is written to a physical disk for the associated virtual disk	Bytes/s	(Δ FlushKBytes x 1024) / Δ Time)
Flushes	Number of cache flush requests per second	Req /s	Δ FlushRequests / Δ Time
Mirror Data Rate	Rate at which data travels across the mirror port to complete read and write requests for the associated virtual disk	Bytes/s	(Δ MirrorKBytes x 1024) / Δ Time
Percent Reads*	Percentage (%) of CPU time dedicated to reads.	%	100 x (Δ ReadIOs / Δ TotalIOs)
Percent Writes*	Percentage (%) of CPU time dedicated to writes	%	100 x (Δ WriteIOs / Δ TotalIOs)
Prefetch Data Rate	Rate data is read from the physical disk to cache in anticipation of subsequent reads when a sequential stream is detected.	Bytes/s	(Δ PrefetchKBytes x 1024) / Δ Time)

Metric	Description	Units	Formula
Read Data Rate*	Rate data is read from the virtual disk by all hosts and includes transfers from the source array to the destination array.	Bytes/s	(Δ KBytesRead x 1024) / Δ Time
Read Hit Data Rate**	Rate at which data is read from the array cache memory because of read hit requests.	Bytes/s	(Δ ReadHitKBytes x 1024) / Δ Time)
Read Hit Latency**	Average time to complete a read request (from initiation to information receipt) from the array volume	sec	(Δ ReadHitLatency / 1000) / Δ ReadHitIOs
Read Hits***	Number of read requests (per second) completed from the array cache memory	Req /s	Δ ReadHitIOs / Δ Time
Read Miss Data Rate**	Rate data is read from physical disks because the data was not present in the array cache memory	Bytes/s	(Δ ReadMissKBytes x 1024) / Δ Time
Read Miss Latency**	Average time it takes to complete a read request (from initiation to information receipt) from the physical disks for all volumes	sec	(Δ ReadMissLatency / 1000) / Δ ReadMissIOs
Read Misses**	Number of read requests (per second) that failed to complete from the array cache memory and were completed from physical disks instead	Req /s	Δ ReadMissRequests / Δ Time
Reads*	Rate data is read from each host host port	Req /s	Δ ReadIOs / Δ Time
Total Data Rate*	Rate data can be transmitted between devices for the hostport	Bytes/s	(Δ KBytesTransferred x 1024) / Δ Time
Total I/O Rate***	Average number of I/O operations in requests per second for both sequential and non-sequential reads and writes for the hostport	Req /s	Δ TotalIOs / Δ Time
Write Data Rate***	Rate data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array	Bytes/s	(Δ KBytesWritten x 1024) / Δ Time
Write Latency**	Average time to complete a write request (from initiation to receipt of write completion)	sec	(Δ WriteLatency / 1000) / Δ WriteIOs
Writes***	Number of write requests per second completed to a virtual disk that were received from all hosts	Req /s	Δ WriteIOs / Δ Time

*This metric is a duplicate of an SMI-S statistic.

**This metric is collected as both an HP Storage Essentials and an EVAPerf performance statistic.

***This metric is a duplicate of an SMI-S statistic and is collected as both an HP Storage Essentials and an EVAPerf performance statistic.

HP Storage Essentials Performance Management – XP Metrics

HP Storage Essentials provides the following performance metrics for HP XP arrays:

- XP System Metrics below
- XP Array Group Metrics on page 653
- XP Array Volume Metrics on page 654
- XP Front-End Controller (CHA) Metrics on next page
- XP Front-end Controller CLPR Metrics on page 651
- XP Back-end Controller (DKA) Metric on page 651
- XP MPB Controller Metrics on page 652

For more detailed information, see related topics for each metric.

XP System Metrics

The following metrics[1] measure the performance of HP XP or HDS arrays, including the P9500/XP storage arrays. See footnotes for exceptions.

Metric	Description	Units
CHA Cache Memory Busy Rate*	Rate at which the CHA cache is busy de-staging data to the DKA	%
CHA Shared Memory Busy Rate*	Rate at which the CHA shared memory is busy	%
Cache Usage	Cache utilization in megabytes	MB
DKA Cache Memory Busy Rate*	Rate at which the DKA cache is busy de-staging data to the disk	%
DKA Shared Memory Busy Rate*	Rate at which the DKA shared memory is busy	%
Percent Cache Usage**	Cache utilization percent	%

Metric	Description	Units
Percent Sidefile Usage**	Percent utilization of the sidefile. A sidefile is an internal buffer that saves a copy of the data to be transmitted to a remote XP array. Use to track continuous access (CA) sidefile cache utilization and the potential impact of DR activities	%
Percent Write Pending Data**	Percentage (%) of pending writes based on the percentage of cache being used to buffer writes on the selected controller. Use to determine if a CLPR is needed or if attention needs to be directed towards journal parity groups	%
Read Hits**	Read I/O requests per second satisfied from cache	Req/s
Sidefile Usage**	Sidefile cache utilization in megabytes	MB
Write Pending Data**	Indicator of pending writes based on cache in megabytes used to buffer writes on the selected controller.	MB

[1]Performance metrics in this table are not available for XP P9500 storage arrays discovered using Command View Advanced Edition (CVAE). For more information, see the Licensing chapter in the *Installation Guide*.

*Metric not available for XP and HDS RAID700 arrays. XP and HDS RAID700 arrays use MPB controllers.

**The default Cache Logical Partition (CLPR) used for obtaining performance information is CLPR 0. The default MP for the XP controller is MP 0.

XP Front-End Controller (CHA) Metrics

Front-end controllers route I/O to/from hosts to their disk cache slots. The front-end view shows all channel adapters (CHA) and their constituent ports. Selecting the CHA allows you to view the processor (MP) utilization.

Note: The Micro Processor Blade (MPB) controllers for XP P9500 storage arrays have a separate set of metrics to measure performance. This front-end controller (CHA) metric is therefore not used to measure performance for XP P9500 arrays.

Best Practices

Graph CHA MP Utilization and Port data rate simultaneously to determine if the port is the bottleneck. If MP utilization is less than 90%, and the port data rate is near its maximum, then you know that data is being de-staged to disk normally assuming the host(s) is running at its peak. If not, check the DKA and RAID Group utilization.

Controller CHA Metrics

This table describes the available performance metric for the front-end controller.

Metric	Description	Units	Formula
Processor Utilization	Processor utilization rate on the selected CHA controller	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter

[1] Metric not available for XP and HDS RAID700 arrays. XP and HDS RAID700 arrays utilize MPB controllers.

XP Front-end Controller CLPR Metrics

The following Cache Logical Partition (CLPR) metrics are available for the HP XP front-end controller.

Note: The MPB (Micro Processor Blade) controllers for XP P9500 storage arrays have a separate set of metrics to measure performance. These front-end controller CLPR metrics are therefore not used to measure the performance for XP P9500 arrays.

Metric	Description	Units	Formulas
Total Data Rate	Rate at which data can be transmitted between devices for the selected CLPR	Bytes\s	ΔTotalIDataTransferred / Δ StatisticTime
Total I/O Rate	Total number of read or write operations taking place per second for the selected front-end controller CLPR	Req\s	Δ TotalIOs / Δ StatisticTime

XP Back-end Controller (DKA) Metric

Back-end controllers route I/O from cache slots to the disk. The back-end view for XP arrays shows all the disk controller adapters (DKA). Selecting the DKA allows you to view processor utilization (MP).

Note: The Micro Processor Blade (MPB) controllers for XP P9500 storage arrays have a separate set of metrics to measure performance. This back-end controller (DKA) metric is therefore not used to measure performance for XP P9500 arrays.

Best Practices

You can graph DKA MP Utilization and RAID Group(s) IOPS to determine if the controller is acting as a bottleneck. If MP utilization is around 90% and the RAID Group data rate is much lower than expected for the RAID level and disk type, there might be a background process such as replication that is over-utilizing the DKA. Typically, you want your back-up disks on a different RAID Group than your production disks, and the back-up disks should be connected to a different DKA.

Controller DKA Metrics

The following table describes the performance metric available for the XP back-end controller.

Table 1 XP Back-end Controller (DKA) Metric

Metric	Description	Units	Formula
Processor Utilization [1]	Processor utilization rate on the selected DKA controller	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter

[1] Metric not available for XP and HDS RAID700 arrays. XP and HDS RAID700 arrays utilize MPB controllers.

XP MPB Controller Metrics

The following table describes the performance metrics collected for Micro Processor Blade (MPB) controllers for P9500/XP storage arrays (see footnote for exception). MBP processor statistics are obtained by HP Storage Essentials CIM extensions.

Table 2 MPB Controller Metrics[1]

Metric	Description	Units	Formula
Back End Utilization	Processor utilization for the back-end processes (back-end activities for target I/O requests) as a percentage of total processor time.	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Mainframe External Initiator Utilization	Processor utilization for the mainframe external initiator processes as a percentage of total processor time.	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Mainframe Target Utilization	Processor utilization rate for the mainframe target processes (front-end activities for processing mainframe I/O requests) as a percentage of total processor time.	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Open External Initiator Utilization	Processor utilization rate for the open external initiator processes (external storage access activities) as a percentage of total processor time.	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Open Initiator Utilization	Processor utilization rate for the open initiator processes (continuous access replication activities) as a percentage of total processor time.	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter
Open Target Utilization	Processor utilization rate for open target processes (front-end activities) as a percentage of total processor time.	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter

Metric	Description	Units	Formula
Processor Utilization	Processor utilization rate on the selected MPB controller. This rate is the sum of the other seven metrics listed in this table.	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter
System Utilization	Processor utilization rate of the array system processes as a percentage of total processor time.	%	Δ BusyTimeCounter / Δ ElapsedTimeCounter

[1] The performance metrics in this table are not available for XP P9500 storage arrays discovered using Command View Advanced Edition (CVAE). For more information, see the Licensing chapter in the *Installation Guide*.

XP Array Group Metrics

The following performance metrics[1] are available for HP XP and HDS array groups and includes the P9500/XP storage arrays (see footnote for exceptions).

Metric	Description	Units
Average Read Size	Average number of reads in bytes to the array group.	Bytes
Average Write Size	Average number of writes in bytes to the array group	Bytes
Percent Read Hits	Percentage (%) of cache reads for the array group	%
Percent Read Hits Random	Percentage (%) of random read I/O's that were de-staged from cache	%
Percent Read Hits Seq	Percentage (%) of sequential read I/O's that were de-staged from cache	%
Percent Reads	Percentage (%) of reads from cache	%
Percent Writes	Percentage (%) of writes from cache	%
Read Data Rate	Rate data is read from the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s

Metric	Description	Units
Read Data Rate Random	Rate at which random data is read from the array group by all hosts	Bytes/s
Read Data Rate Seq	Rate at which sequential data is read from the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s
Read Hits	Cache read hits in requests per second	Req/s
Read Response Time	Time required to complete a read I/O in seconds	sec
Reads	Number of read I/O's in requests per second	Req/s
Total Data Rate	Rate at which data can be transmitted between devices for the selected array group	Req/s
Total I/O Rate	Total number of read or write operations taking place per second for the selected array group	Req/s
Utilization	Percentage of time disks in the array group are busy	%
Write Data Rate	Rate at which data is written to the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s
Write Data Rate Random	Rate at which random data is written to the array group by all hosts	Bytes/s
Write Data Rate Seq	Rate at which sequential data is written to the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s
Write Response Time	Time required to complete a write I/O in seconds	sec
Writes	Number of write I/O's in requests per second	Req/s

[1]The performance metrics in this table are not available for XP P9500 storage arrays discovered using Command View Advanced Edition (CVAE). For more information, see the Licensing chapter in the *Installation Guide*.

XP Array Volume Metrics

The following metrics[1] are available for HP XP and HDS array volumes, including the P9500/XP storage array volumes (see footnotes for exceptions).

Metric	Description	Units	Formulas
Average Read Size	Average number of reads in bytes to the array group	Bytes	(Δ RandomReadData + Δ SequentialReadData) / (Δ RandomReadIOs + Δ SequentialReadIOs)
Average Write Size	Average number of writes in bytes to the array group	Bytes	(Δ RandomWriteData + Δ SequentialWriteData) / (Δ RandomWriteIOs + Δ SequentialWriteIOs)
Percent Read Hits	Percentage (%) of cache reads for the array group	%	100 x (Δ RandomReadHitIOs + Δ SequentialReadHitIOs) / (Δ RandomReadIOs + Δ SequentialReadIOs)
Percent Read Hits Random	Percentage (%) of random read I/O's that were de-staged from cache	%	100 x Δ RandomReadHitIOs / Δ RandomReadIOs
Percent Read Hits Seq	Percentage (%) of sequential read I/O's that were de-staged from cache	%	100 x Δ SequentialReadHitIOs / Δ SequentialReadIOs
Percent Reads	Percentage (%) of reads from cache	%	100 x (Δ RandomReadIOs + Δ SequentialReadIOs)/ (Δ RandomReadIOs + Δ RandomWriteIOs + Δ SequentialReadIOs + Δ SequentialWriteIOs)
Percent Writes	Percentage (%) of writes from cache	%	100 x (Δ RandomWriteIOs + Δ SequentialWriteIOs)/ Δ RandomReadIOs + Δ RandomWriteIOs + Δ SequentialReadIOs + Δ SequentialWriteIOs)
Read Data Rate	Rate at which data is read from the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s	(Δ RandomReadData + Δ SequentialReadData) / Δ StatisticTime
Read Data Rate Random	Rate at which random data is read from the array group by all hosts	Bytes/s	Δ RandomReadData / Δ StatisticTime

Metric	Description	Units	Formulas
Read Data Rate Seq	Rate at which sequential data is read from the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s	Δ SequentialReadData/ Δ StatisticTime
Read Hits	Cache read hits in requests per second	Req/s	(ΔRandomReadHitIOs + Δ SequentialReadHitIOs) / Δ StatisticTime
Read Response	Time required to complete a read I/O in seconds	sec	Δ ReadResponseTimeCounter / (Δ RandomReadIOs + Δ SequentialReadIOs)
Reads	Number of read I/O's in requests per second	Req/s	(Δ RandomReadIOs + Δ SequentialReadIOs) / Δ StatisticTime
Total Data Rate	Rate in which data can be transmitted between devices for the selected array group	Req/s	(Δ RandomReadData + Δ RandomWriteData + Δ SequentialReadData + Δ SequentialWriteData) / Δ StatisticTime
Total I/O Rate	Total number of read or write operations taking place per second for the selected array group	Req/s	(Δ RandomReadIOs + Δ RandomWriteIOs + Δ SequentialReadIOs + Δ SequentialWriteIOs) / Δ StatisticTime
Utilization	Percentage (%) of time that disks in the array group are busy	%	100 x Δ ActiveTime fields / Δ StatisticTime
Write Data Rate	Rate at which data is written to the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s	(Δ RandomWriteData + Δ SequentialWriteData) / Δ StatisticTime
Write Data Rate Random	Rate at which random data is written to the array group by all hosts	Bytes/s	Δ RandomWriteData / Δ StatisticTime

Metric	Description	Units	Formulas
Write Data Rate Seq	Rate at which sequential data is written to the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s	Δ SequentialWriteData/ Δ StatisticTime
Write Response Time	Time required to complete a write I/O in seconds	sec	Δ WriteResponseTimeCounter / (Δ RandomWriteIOs + Δ SequentialWriteIOs)
Writes	Number of write I/O's in requests per second	Req/s	(Δ RandomWriteIOs + Δ SequentialWriteIOs) / Δ StatisticTime

[1]Performance metrics in this table are not available for XP P9500 storage arrays discovered using Command View Advanced Edition (CVAE). For more information, see the Licensing chapter in the *Installation Guide*.

HP Storage Essentials Performance Management – NetApp Metrics

HP Storage Essentials provides the following information about NetApp performance metrics:

- NetApp Performance Metrics below
- NetApp Raw Statistics on page 667

See the HP Storage Essentials *Storage Performance Management Guide* for information about monitoring NetApp device performance.

NetApp Performance Metrics

The following performance metrics are provided for NetApp storage systems.

NetApp System Performance Metrics

The following performance metrics are provided for NetApp storage systems and include iSCSI, CIFS, DAFS, FCP, HTTP, and NFSv3 operations.

Metric	Description	Unit	Formula
Cache			
Buffer Cache Hits Count	Buffer Cache or system memory read cache hits per second. Use to determine if access latency is contributing to performance issues.	Req/s	Δ Buf_ Load_Cnt / Δ Time

Metric	Description	Unit	Formula
Buffer Cache Misses	Buffer cache miss count per second (rate). A cache miss is simply data that is not in the cache. The result is that the system must fetch the data from the disk. Use to determine PAM deployment requirements and configuration.	Req/s	Δ Buff_ Miss_Cnt / Δ Time
Inode Cache Hits Count	Number of hits for inodes read from disk that are cached and subsequently accessed. Use to determine cause of increase in file system performance.	Req/s	Δ Inode_ Cache_ Hit / Δ Time
Inode Cache Misses Count	Number of inode cache misses per second (rate). A cache miss is simply data that is not in the cache. The result is that the system must fetch the data from the disk. The inode cache behaves in the same fashion. Use to determine if the inode cache needs to be increased.	Req/s	Δ Inode_ Cache_ Miss / Δ Time
Name Cache Hits	Number of name cache hits per second (rate). Use to determine frequency of name cache hits. The name cache improves file lookup in the file system.	Req/s	Δ Name_ Cache_ Hit / Δ Time
Name Cache Misses	Number of name cache misses per second (rate). A cache miss is simply data that is not in the cache. The result is that the system must fetch the data from the disk. The name cache behaves in the same fashion. Use to determine if the name cache needs to be increased.	Req/s	Δ Name_ Cache_ Miss / Δ Time
Common Internet File System (CIFS) I/O Operations			
CIFS Latency	Average latency for Common Internet File System (CIFS) operations in milliseconds	ms	Δ CIFS_ Latency / Δ CIFS_ Latency_ Base
CIFS Operations	Number of Common Internet File System (CIFS) operations per second	Req/s	Δ CIFS_ Ops / Δ Time
Direct Access File System (DAFS) I/O Operations			
DAFS Operations	Number of Direct Access File System (DAFS) operations per second	Req/s	Δ DAFS_ Ops / Δ Time

Metric	Description	Unit	Formula
Fibre Channel Protocol (FCP) I/O Operations			
FCP Operations	Number of Fibre Channel Protocol (FCP) operations per second	Req/s	Δ FCP_ Ops / Δ Time
FCP Read Data	FCP bytes read per second	Bytes/s	Δ FCP_ Read_ Data / Δ Time
FCP Read Latency	Average latency for read operations observed over all LUNs in the system accessed over FCP in milliseconds	ms	Δ FCP_ Read _ Latency / Δ FCP_ Read_ Ops
FCP Read Operations	Total number of read operations per second observed over all the LUNS in the system accessed over FCP	Req/s	Δ FCP_ Read_ Ops / Δ Time
FCP Write Data	FCP bytes written per second	Bytes/s	Δ FCP_ Write_ Data / Δ Time
FCP Write Latency	Average latency for write operations observed over all LUNs in the system accessed over FCP in milliseconds	ms	Δ FCP_ Write _ Latency / Δ FCP_ Write_ Ops
FCP Write Operations	Total number of write operations per second observed over all the LUNs in the system accessed over FCP	Req/s	Δ FCP_ Write_ Ops / Δ Time
Hypertext Transfer Protocol (HTTP) I/O Operations			
HTTP Operations	Number of Hypertext Transfer Protocol (HTTP) operations per second	Req/s	Δ HTTP_ Ops / Δ Time
Internet Small Computer System Interface (iSCSI) I/O Operations			

Metric	Description	Unit	Formula
ISCSI Operations	Number of Internet Small Computer System Interface (iSCSI) operations per second	Req/s	Δ ISCSI_ Ops / Δ Time
ISCSI Read Data	iSCSI bytes read per second	Bytes/s	Δ ISCSI_ Read_ Data / Δ Time
ISCSI Read Latency	Average latency of read operations observed over all LUNs in the system accessed over iSCSI in milliseconds	ms	Δ ISCSI_ Read _ Latency / Δ ISCI_ Read_ Ops
ISCSI Read Operations	Total number of read operations per second observed over all the LUNs in the system accessed over iSCSI	Req/s	Δ ISCSI_ Read_ Ops / Δ Time
ISCSI Write Data	iSCSI bytes written per second	Bytes/s	Δ ISCSI_ Write_ Data / Δ Time
ISCSI Write Latency	Average latency of write operations observed over all LUNs in the system accessed over iSCSI in milliseconds	ms	Δ ISCSI_ Write_ Latency / Δ ISCI_ Write_ Ops
ISCSI Write Operations	Total number of write operations per second observed overall the LUNs in the system accessed over iSCSI	Req/s	Δ ISCSI_ Write_ Ops / Δ Time
Network File System (NFS) Operations			
NFSV3 Average Ops Latency	Average latency of the NFS v3 operations	ms	Δ NFSv3_ Avg_Op _ Latency / Δ NFSv3_ Avg_Op_ Latency_ Base

Metric	Description	Unit	Formula
NFS Operations	Number of Network File System (NFS) operations per second	Req/s	Δ NFS_ Ops / Δ Time
NFSV3 Read Latency	Average latency for NFS v3 read operations in milliseconds	ms	Δ NFSv3_ Read _ Latency / Δ NFSv3_ Avg_ Read_ Latency_ Base
NFSV3 Write Latency	Average latency for NFS v3 write operations in milliseconds	ms	Δ NFSv3_ Write _ Latency / Δ NFSv3_ Avg_ Write_ Latency_ Base

NetApp Host Processor (CPU) Performance Metrics

The following performance metrics are available for NetApp host processors (CPU).

NetApp does not expose memory utilization statistics because all memory that is not used for initialization is allocated to the file system for caching. In other words, all memory is used all of the time. How it is allocated is what really matters. Processor utilization and cache hit ratio are better indicators of how utilized a filer is. If the cache hit ratio is low and CPU utilization is high, the filer might be reaching maximum utilization.

Metric	Description	Unit	Formula
Processor Utilization	Total CPU utilization (%) by all the processes running on the filer. Indicates the percentage (%) of time that the processor is active. A completely idle processor shows 0%. A processor saturated with activity shows 100%. Use to identify CPU bottlenecks.	%	100 x (Δ Processor_ Busy / Δ Processor_ Elapsed_ Time)

NetApp Aggregate Performance Metrics

The following metrics are available for NetApp aggregate transfer rates.

Metric	Description	Unit	Formula
Total Transfers	Total number of transfers per second serviced by the aggregate.	Req/s	Δ Total_Transfers / Δ Time
User Read Block	Number of blocks read per second on the aggregate.	Blocks/s	Δ User_Read_ Blocks / Δ Time
User Reads	Number of user reads per second to the aggregate.	Req/s	Δ User_Reads / Δ Time
User Write Blocks	Number of blocks written per second to the aggregate.	Blocks/s	Δ User_Write_ Blocks / Δ Time
User Writes	Number of user writes per second to the aggregate.	Req/s	Δ User_Writes / Δ Time

NetApp File System Performance Metrics

The following metrics are available for NetApp file system performance.

Metric	Description	Unit	Formula
Reserved Inodes	Provides the count of reserved inodes in a file system. Use to count the reserved inodes. The first ten (10) inodes on a file system are special inodes. Inodes 7-10 are reserved and usually not used.	counter	–
Total Inodes	Total number of inodes. Inodes are file system data structures or metadata used to store basic file data like ownership and file permissions. Use to view the inode limit and determine if more are needed.	counter	–
Used Inodes	Total number of inodes that are currently used. Use to show the number of inodes in use. Use to alert the admin when inode utilization is approaching total inodes available.	counter	–

NetApp Front-End Port Performance Metrics

The following front-end port performance metrics are available.

Metric	Description	Unit	Formula
Other Operations	Number of other operations per second	Req/s	Δ Other_Ops / Δ Time
Queue Full Responses	SCSI queue full responses per second	Req/s	Δ Queue_Full / Δ Time
Read Data	Number of bytes read from the filer per second	Bytes/s	Δ Read_Data / Δ Time
Read Operations	Number of read operations per second	Req/s	Δ Read_Ops / Δ Time
Write Data	Number of bytes written to the filer per second	Bytes/s	Δ Write_Data / Δ Time
Write Operations	Number of write operations per second	Req/s	Δ Write_Ops / Δ Time

NetApp IP Ports Performance Metrics

This following metrics are available for NetApp IP port performance.

Metric	Description	Unit	Formula
Bytes Received	Inbound traffic in megabytes per second through the filer network interface controller (NIC). Use to assess network traffic for load balancing, multi-path optimization, and network performance.	MB/s	(Δ Recv_ Data x 1024) / Δ Time
Bytes Transmitted	Outbound traffic in megabytes per second through the filer network interface controller (NIC). Use to assess network traffic for load balancing, multi-path optimization, and network performance.	MB/s	(Δ Send_ Data x 1024) / Δ Time

Metric	Description	Unit	Formula
Packets Received	Inbound traffic in packets per second through the filer network interface controller (NIC). Network packets contain data headers, address source and destination, payload and CRC fields. Use to measure network traffic for load balancing, multi-path optimization, and network performance. Contrary to bytes received testing, packet testing is actually a better test because inbound packets have either arrived or not. Byte testing does not indicate whether or not a packet transmission completed.	Pkt/s	Δ Recv_ Packets / Δ Time
Packets Transmitted	Outbound traffic in packets per second through the filer network interface controller (NIC). Network packets contain data headers, address source and destination, payload and CRC fields. Use to measure network traffic for load balancing, multi-pathing optimization and network performance. Contrary to bytes transmitted testing, packet testing is actually a better test because outbound packets have been either sent or not. Byte testing does not indicate whether or not a packet transmission completed.	Pkt/s	Δ Send_ Packets / Δ Time
Receive Errors	Errors per second while receiving packets	Errors/s	Δ Recv_ Errors / Δ Time
Send Errors	Errors per second while sending packets	Errors/s	Δ Send_ Errors / Δ Time

NetApp Volume Performance Metrics

The following table volume performance metrics are available.

Metric	Description	Unit	Formula
Average Latency	Average latency in milliseconds for all operations on the volume	ms	Δ Avg_Latency / Δ Total_Ops
Other Latency	Average latency time for other writes to the volume in milliseconds	ms	Δ Other_Latency / Δ Other_Ops
Other Operations	Number of other operations per second to the volume.	Req/s	Δ Other_Ops / Δ Time

Metric	Description	Unit	Formula
Read Data	Bytes read per second from the volume	Bytes/s	Δ Read_Data / Δ Time
Read Latency	Average latency time for reads to the volume in milliseconds	ms	Δ Read_Latency / Δ Read_Ops
Read Operations	Number of reads per second to the volume	Req/s	Δ Read_Ops / Δ Time
Total Operations	Total number of operations per second serviced by the volume	Req/s	Δ Total_Ops / Δ Time
Write Data	Bytes written per second to the volume	Bytes/s	Δ Write_Data / Δ Time
Write Latency	Average latency time for writes to the volume in milliseconds	ms	Δ Write_Latency / Δ Write_Ops
Write Operations	Number of writes per second to the volume	Req/s	Δ Write_Ops / Δ Time

NetApp LUN Performance Metrics

The following metrics are available for NetApp LUN performance.

Metric	Description	Unit	Formula
Average Latency	Average latency in milliseconds for all operations on the LUN	ms	Δ Avg_Latency / Δ Total_Ops
Other Operations	Number of other operations per second	Req/s	Δ Other_Ops / Δ Time
Queue Full Responses	Number of queue full responses per second	Req/s	Δ Queue_Full / Δ Time
Read Data	Number of bytes read per second	Bytes/s	Δ Read_Data / Δ Time
Read Operations	Number of read operations per second	Req/s	Δ Read_Ops / Δ Time
Total Operations	Total number of operations on the LUN per second	Req/s	Δ Total_Ops / Δ Time

Metric	Description	Unit	Formula
Write Data	Number of bytes written per second	Bytes/s	Δ Write_Data / Δ Time
Write Operations	Number of write operations per second	Req/s	Δ Write_Ops / Δ Time

NetApp Disk Performance Metrics

The following performance measurements are available for NetApp disk drives.

Metric	Description	Unit	Formula
Disk Busy	Percentage of time there was at least one outstanding request to the disk	%	100 x (Δ Disk_ Busy / Δ Disk_ Busy_Base)
Total Transfers	Total number of disk operations involving data transfer initiated per second	Req/s	Δ Total_ Transfers / Δ Time
User Read Blocks	Number of blocks transferred for user read operations per second	Blocks/s	Δ User_Read_ Blocks / Δ Time
User Read Latency	Average latency per block in milliseconds for user read operations	ms	Δ User_Read_ Latency / Δ User_Read_ Blocks
User Reads	Number of disk read operations initiated each second for retrieving data or metadata associated with user request	Req/s	Δ User_Reads / Δ Time
User Write Blocks	Number of blocks transferred for user write operations per second	Blocks/s	Δ User_Write_ Blocks / Δ Time
User Write Latency	Average latency per block in milliseconds for user write operations	ms	Δ User_Write_ Latency / Δ User_Write_ Blocks
User Writes	Number of disk write operations initiated each second for retrieving data or metadata associated with user requests	Req/s	Δ User_Writes / Δ Time

NetApp Qtree Performance Metrics

The following performance metrics are available for the NetApp QTree.

Metric	Description	Unit	Formula
CIFS Operations	Number of Common Internet File System (CIFS) operations per second to the qtree	Req/s	Δ CIFS_ Ops / Δ Time
Internal Operations	Number of internal operations generated by activities such as snapmirror and backup per second to the qtree	Req/s	Δ Internal_ Ops / Δ Time
NFS Operations	Number of NFS operations per second to the qtree	Req/s	Δ NFS_ Ops / Δ Time

NetApp Raw Statistics

The following raw values are used to calculate the performance metric data provided in the Performance Manager for NetApp system devices. All raw metrics are counters derived from the NetApp device.

For a list of the Performance Manager metrics that use these raw statistics, see the NetApp Performance Metrics on page 657.

The following values are derived from the NetApp device and stored in the NAS_SYSTEM_ STATS database table. This database contains the iSCSI, CIFS, FCP, NFSv3 and general system statistics.

Raw Statistic	Description	Unit
buf_load_cnt	Number of buffer cache or system memory read cache hits	Hit count
buff_miss_cnt	Number of buffer cache misses	Miss count
cifs_latency	Average latency for Common Internet File System (CIFS) operations	Time count in milliseconds
cifs_latency_ base	Total observed CIFS operations to be used as a base counter for CIFS average latency calculation	Time count in milliseconds
cifs_ops	Number of Common Internet File System (CIFS) operations	Request count

Raw Statistic	Description	Unit
dafs_ops	Number of Direct Access File System (DAFS) operations	Request count
fcp_ops	Number of Fibre Channel Protocol (FCP) operations	Request count
fcp_read_data	Number of FCP bytes read	Byte count
fcp_read_ latency	Average latency for read operations observed over all LUNs in the system accessed over FCP	Time count in milliseconds
fcp_read_ops	Total number of read operations observed over all the LUNS in the system accessed over FCP	Request count
fcp_write_ data	Number of FCP bytes written	Byte count
fcp_write_ latency	Average latency for write operations observed over all LUNs in the system accessed over FCP	Time count in milliseconds
fcp_write_ops	Total number of write operations observed over all the LUNs in the system accessed over FCP	Request count
http_ops	Number of Hypertext Transfer Protocol (HTTP) operations	Request count
inode_cache_ hit	Number of hits for inodes read from disk that are cached	Hit count
inode_cache_ miss	Number of inode cache misses	Miss count
iscsi_ops	Number of Internet Small Computer System Interface (iSCSI) operations	Request count
iscsi_read_ data	iSCSI bytes read	Byte count
iscsi_read_ latency	Average latency of read operations observed over all LUNs in the system accessed over iSCSI	Time count in milliseconds
iscsi_read_ ops	Total number of read operations per second observed over all the LUNs in the system accessed over iSCSI	Request count
iscsi_write_ data	Number of iSCSI bytes written	Byte count
iscsi_write_ latency	Average latency of write operations observed over all LUNs in the system accessed over iSCSI	Time count in milliseconds

Raw Statistic	Description	Unit
iscsi_write_ops	Total number of write operations observed on all LUNs in the system accessed over iSCSI	Request count
name_cache_hit	Number of name cache hits	Hit count
name_cache_miss	Number of name cache misses	Miss count
nfsv3_avg_op_latency	Average latency of the NFS v3 operations	Time count in milliseconds
nfsv3_avg_op_latency_base	Array of select NFS v3 operation counts for latency calculation	Time count in milliseconds
nfs_ops	Number of Network File System (NFS) operations	Request count
nfsv3_read_latency	Average latency for NFS v3 read operations	Time count in milliseconds
nfsv3_write_latency	Average latency for NFS v3 write operations	Time count in milliseconds

NetApp Aggregate Raw Statistics

The following values are derived from the NetApp device and stored in the NAS_AGGREGATE_STATS database table.

Raw Statistic	Description	Unit
total_transfers	Number of transfers	Request count
user_read_blocks	Number of blocks read	Data block count
user_reads	Number of user reads	Request count
user_write_blocks	Number of blocks written	Data block count
user_writes	Number of user writes	Request count

NetApp Volume Raw Statistics

The following values are derived from the NetApp device and stored in the NAS_FILESYSTEM_STATS database table. This database contains the volume and file system statistics.

Raw Statistic	Description	Unit
avg_latency	Average latency all operations on the volume	Time count in milliseconds
other_latency	Latency time for other writes to the volume	Time count in milliseconds
other_ops	Number of other operations to the volume	Request count
read_data	Number of bytes read from the volume	Byte count
read_latency	Average latency time for reads to the volume	Time count in milliseconds
read_ops	Number of read operations to the volume	Request count
total_ops	Total number of operationsserviced by the volume	Request count
write_data	Number of bytes written to the volume	Byte count
write_latency	Average latency time for writes to the volume	Time count in milliseconds
write_ops	Number of write operations to the volume	Request count
wv_fsinfo_inos_reserve	Reserved inodes in the volume	Inode count
wv_fsinfo_inos_total	Total number of inodes in the volume	Inode count
wv_fsinfo_inos_used	Used inodes in the volume	Inode count

NetApp Processor (CPU) Raw Statistics

The following values are derived from the NetApp device and stored in the HOSTCPUSTAT database table.

Metric	Description	Unit
processor_ busy	Time CPU is being utilized. Uses the wall-clock time since boot and is used for calculating processor utilization.	Time count in seconds
processor_ elapsed_ time	Time CPU is running, including busy and idle periods	Time count in seconds

NetApp Front-end Port Raw Statistics

The following values are derived from the NetApp device and stored in the NAS_FCPORT_ STATS database table.

Raw Statistic	Description	Unit
other_ops	Number of other operations to the volume	Request count
read_data	Number of bytes read from the filer	Byte count
read_ops	Number of read operations to the volume	Request count
write_data	Number of bytes written to the filer	Byte count
write_ops	Number of write operations to the volume	Request count
queue_full	Number of responses that SCSI queue is full	Response count

NetApp IP Port Raw Statistics

The following values are derived from the NetApp device and stored in the IP_PORTSTATS database table.

Raw Statistic	Description	Unit
recv_ data	Amount of inbound data in bytes received through the filer network interface controller (NIC)	Byte count
recv_ errors	Number of errors which occurred while receiving packets	Error count

Raw Statistic	Description	Unit
recv_ packets	Number of inbound packets received through the filer network interface controller (NIC).	Packet count
send_ data	Amount of outbound data in bytes sent through the filer network interface controller (NIC)	Byte count
send_ errors	Number of errors which occurred while sending packets	Error count
send_ packets	Number of outbound packets sent through the filer network interface controller (NIC).	Packet count

NetApp LUN Raw Statistics

The following values are derived from the NetApp device and stored in the NAS_LUN_STATS database table.

Raw Statistic	Description	Unit
avg_ latency	Average latency time count for all operations on the LUN	Time count in milliseconds
other_ops	Number of other operations	Request count
queue_full	Number of queue full responses	Response count
read_data	Number of bytes read	Byte count
read_ops	Number of read operations from the LUN	Request count
total_ops	Number of total operations on the LUN	Request count
write_data	Number of bytes written to the LUN	Byte count
write_ops	Number of write operations to the LUN	Request count

NetApp Disk Raw Statistics

The following values are derived from the NetApp device and stored in the NAS_DISKDRIVE_ STATS database table.

Raw Statistic	Description	Unit
disk_busy	Time count for when there was at least one outstanding request to the disk	Time count
disk_busy_ base		Time count
total_ transfers	Total number of disk operations involving data transfers	Request count
user_read_ blocks	Number of blocks transferred for user read operations	Block count
user_read_ latency	Latency for user read operations	Time count in microseconds
user_reads	Number of disk read operations associated with a user request	Request count
user_write_ blocks	Number of blocks transferred for user write operations	Block count
user_write_ latency	Latency for user write operations	Time count in microseconds
user_writes	Number of disk write operations associated with user requests	Request count

NetApp QTree Raw Statistics

The following values are derived from the NetApp device and stored in the NAS_QTREE_STATS database table.

Raw Statistic	Description	Unit
cifs_ops	Number of Common Internet File System (CIFS) operations to the qtree	Request count
internal_ ops	Number of internal operations generated by activities (such as snapmirror and backup) to the qtree	Request count
nfs_ops	Number of NFS operations to the qtree	Request count

16 Provisioning Manager

Provisioning Manager tools in your management software vary according to your software and licenses. See the List of Features to determine if you have access to Provisioning Manager wizards. The List of Features is accessible from the Documentation Center (**Help** > **Documentation Center**).

Your account must belong to a role that has "System Manager" selected for you to be able to perform SAN zoning operations, such as creating zone aliases, zones, and zone sets.

This section contains the following topics:

About the Provisioning Manager

HP Storage Essentials has built-in element provisioning capabilities (for SAN zoning and storage system provisioning), as well as passive management capabilities.

With HP Storage Essentials, you also can add the optional (value added) GUI-based Provisioning Manager for “path provisioning” capabilities, as well as for zoning and scheduling.

When licensed, you can use Provisioning Manager to perform path provisioning, zoning, and scheduling tasks.

Provisioning Manager does not provide steps for creating a new LUN (or pool of LUNs).

You access Provisioning Manager from the management software home page, as well as from other management server menus. When accessed, the tab choices enable you to select **Path Provisioning**, **SAN Zoning**, or **Storage System Provisioning**. Each choice opens a wizard-like display from which to perform an activity.

Provisioning tools are accessible by clicking **Provisioning Manager** ().

- **Path Provisioning tool –** When licensed, this tool is available to provide mapping, masking, and zoning capabilities, and for scheduling provisioning tasks so they take place when the network traffic is light.

  This tool does not provide a step for creating a new LUN (or a pool of LUNs). If needed, create them before opening the tool by using the native management software of the array or the external tools feature. These are described later in this section. You can also create new LUNs by using the element provisioning capabilities of the management server.

- **SAN Zoning tool –** Enables you to create and modify zones, zone aliases and zone sets. The SAN Zoning tool is accessible by clicking the **Provisioning** button next to the fabric on which you want to do provisioning. You can view the properties of a fabric or switch by clicking

its link in the table. (If Provisioning Manager is licensed for your system, you can also access the SAN Zoning tool by clicking the **SAN Zoning** tab in the Provisioning Manager window.). For more information, see SAN Zoning Overview on page 679.

- **Storage System Provisioning tool –** Enables you to manage your storage system. Depending on the particular array, you might be able to create and manage storage pools, volumes, and host security groups. For more information, see About Path Provisioning on page 757. Also, in this section, see the important note about creating subfolders for virtual disks and hosts for an EVA from the management server, and the considerations regarding XP arrays.

  You access the Storage System Provisioning tool by clicking the **Provisioning** button next to the storage system on which you want to do provisioning. You can also view the properties of a storage system by clicking its link in the table. (If Provisioning Manager is licensed for your system, you can also access the Storage System Provisioning tool by opening the Provisioning Manager window, then clicking the **Storage Systems** tab.) For more information, see About Setting Up Storage Partitioning on page 694.

You cannot create a subfolder for virtual disks and hosts for an EVA from the management server. However, you can create an EVA Array volume within a subfolder of virtual disks which was already created, if the folder name does not contain a space. If the folder name contains a space, then you cannot create a volume within that folder. Both EVA Array virtual disk subfolders containing volumes, and an EVA Array volume having snapshots, will look alike in the management server; however, you must make sure you are creating the EVA Array volume within a folder.

Ports designated as an Initiator on a storage system belonging to the HDS Freedom Storage Lightning 9900 Series or the Freedom Storage Lightning 9900V Series cannot be used for provisioning. If you select one of these ports, you receive a message saying that provisioning failed because the HiCommand Database was not refreshed.

Considerations for XP Arrays

For XP arrays, management server storage “pools” are defined as sets of LDEVs which are compatible with similar ones for “logical unit size expansion” (LUSE). For more information, see About Volumes on HDS Storage Systems on page 706. However, V-VOLs, POOL LDEVs, and LUNs are also included in pools, and when viewed in Performance Manager, they are separated into a separate column labeled “Reserved.” The management server does not perform replication provisioning in this management server version, but it does protect a user from attempting replication-related or other illegal operations on such LDEVs.

To create a LUSE for XP arrays, some volumes need to be deleted in order to free up their extents to the storage pool. Volumes that cannot be deleted (because they have LUNs, P-Vols, V-Vols, etc) can therefore never be LUSE candidates, and their respective **Delete** icon is not available in their Volume list.

You can use filtering to facilitate selection of LDEVs as LUSE candidates. An ArrayGroup (RaidGroup) attribute for each LDEV is provided, and the Show Volume Filter function is able to filter on ArrayGroup. The location of the LDEV is shown within ArrayGroup using the row number of the LDEV in the arrayGroup table, and you can sort the volume table on that column. Also, an Array Control Processor (ACP) attribute for each LDEV is provided, and the Show Volume Filter function filters on the ACP. These functions are also available in CLI commands.

The "Add LUNs to the Host Security Group" page contains an additional column showing LDEV characteristics such as whether the volume is a P-VOL, S-VOL, V-VOL, External Vol, Command Device, External Command Device, and so forth. This data is also available using CLI commands.

External Tools and Web-based Native Vendor Device Managers

The management server detects management URLs for most external tools automatically upon discovery. Examples include Brocade Web Tools for Brocade switches, and Remote Web Console for XP arrays.

You can launch web-based native vendor device managers from the management server topology screen. To do so, right-click the device in the topology to access its external tool links and links to web-based native device management. (Do not double-click the device because this drills down into the object property pages.)

You can also add custom "External Tools" links to any web-based tool. You must add the link in the form of a URL. An example of an External Tool is the Emulex HBAnywhere Utility, a centralized host bus adapter (HBA) management suite that provides control and management of all Emulex HBAs in a SAN from a single console. After adding this custom link, the utility would be available upon right-clicking any managed host.

External Tools Feature

The External Tools feature that ships with the management server enables you to:

- **Browse the element**—Access a host or a switch through its main Web page. The software assumes the host or switch has a Web page at http://<hsIPaddress>. In this instance, <hsIPaddress> is the IP address of the host or switch. To access the main Web page of the host or switch, right-click the element in **System Manager** and select **External Tools** > **Browse to 192.168.1.2**. In this instance, 192.168.1.2 is the IP address of the host or switch.
- **Telnet to the element**—Access a host or a switch through the telnet utility. Telnet must be already enabled on the element. The command uses telnet://<hsIPaddress>. In this instance, <hsIPaddress> is the IP address of the host or switch. To telnet to a host or switch, right-click the element in **System Manager** and select **External Tools** > **Telnet to 192.168.1.2**. In this instance, 192.168.1.2 is the IP address of the host or switch.
- **Set up external tools**—Enables you to add a URL for accessing management software, such as Emulex HBAnyware, Hitachi HiCommand Device Manager and EMC ControlCenter Navisphere. See Setting up External Tools on next page for more information.
- **Access the management tool for the storage system**—In some instances, the management tool for the storage system is accessible from this menu. For example,

HiCommand for HDS storage systems and Command View for HP XP storage systems are accessible from the External Tools menu.

Setting up External Tools

You can add URLs for accessing external tools used for managing an element, such as Emulex HBAnyware, Hitachi HiCommand Device Manager and EMC ControlCenter Navisphere for storage systems.

Note: When you add a URL, it applies only to the element you originally right-clicked.

To add a URL for accessing external tools:

1. Access **System Manager**.
2. Right-click the element, and select **External Tools** > **Set Up External Tools**.
3. Click **Add New Management URL**.
4. In the Description box, enter the name of the product you plan to access.
5. In the URL box, enter the URL that is used to access the product.
6. Click **OK**.

When you right-click the element and select **External Tools**, the external tool is listed.

To delete the URL for an external tool, click the corresponding **Delete** button in the External Tools window.

After you become adept at provisioning, you can access the provisioning screens using one of the following methods:

- Double-click a storage system or switch in System Manager and then click the **Provisioning Manager** tab.
- Right-click a storage system or switch in the Access tab in System Manager. For more information, see Access Tab on page 440.

About Provisioning Brocade Switches after Upgrading

After you upgrade the management server, perform Get Details for any subset of elements that includes the Brocade switch before performing any provisioning operations that involve that switch. For more information, see Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries on page 71.

Managing Zones

This section contains the following topics:

SAN Zoning Overview

Use SAN zoning to control what can be seen in the storage area network (SAN). SAN zoning enables you to group elements into zones, which can then be grouped into active and inactive zone sets. Only elements in an active zone set can be seen. A switch fabric can have multiple zone sets, but only one zone set can be active.

Uses of Zones

Zones are an excellent way to split hardware resources because they work by exclusion. For example, you can set up your switch ports so that elements connected to some of the ports appear in one zone and the rest appear in another zone. Members of a zone can only communicate with other members of the same zone. If two elements are not in the same zone, they cannot communicate.

Zones are usually created for a particular task, such as controlling access between devices or groups. You might create zones based on an application or an operating system. For example, some network administrators prefer to put all of the Microsoft Windows computers in one zone and all of the Sun Solaris computers in another; or you might create zones according to an application. For example, you might want to create a zone for Production and another zone for Finance. This way the users in the Finance department are not even aware of the disks and ports available for Production, and vice versa.

Only elements in an active zone set can communicate with each other. If you do not want users in the Production and Finance zones to have access to the same storage, the two zones must be in two different zone sets, both of which must be active. Since you can only have one active zone set to a fabric, the Production zone belongs to a zone set in one fabric and the Finance zone belongs to another zone set in another fabric.

A zone can be in more than one zone set, which allows for flexibility. For instance, in our example, the Finance zone could be in both an active zone set and an inactive zone set. Assume that the Finance zone is a member of an active zone set named Zone Set One and also a member of an inactive zone set named Zone Set Two, and that Zone Set Two contains additional zones. If you activate Zone Set Two, users can be aware of those additional elements and still have access to the Finance zone (because it is a member of Zone Set Two). Only one zone set in a fabric can be active at a time.

You can create zone aliases to keep track of your zones easily. Instead of having to remember a port's name, you can assign a name that is easy to remember. As a best practice, a zone should contain either zone aliases or ports, but not both.

Types of Zoning

The SAN Zoning tool is able to manage the two types of zoning:

- **Switch Port Zoning (also known as hard zoning)** – A hard zone is created by assigning a domain/port to a zone. Any device attached to the port is automatically in the zone.
- **WWN Zoning (also known as soft zoning)** – A soft zone is created by assigning a world wide name (WWN) of a device port to a zone.

The following shows an example of hard zoning. Ports 1 through 5 on the switch are assigned to a zone for Production and ports 4 through 8 are assigned to the zone for Finance. Users in Finance can access storage systems B and C but not storage system A. Likewise, users in Production can access storage systems A and C, but not storage system B.

Resources in Two Zones

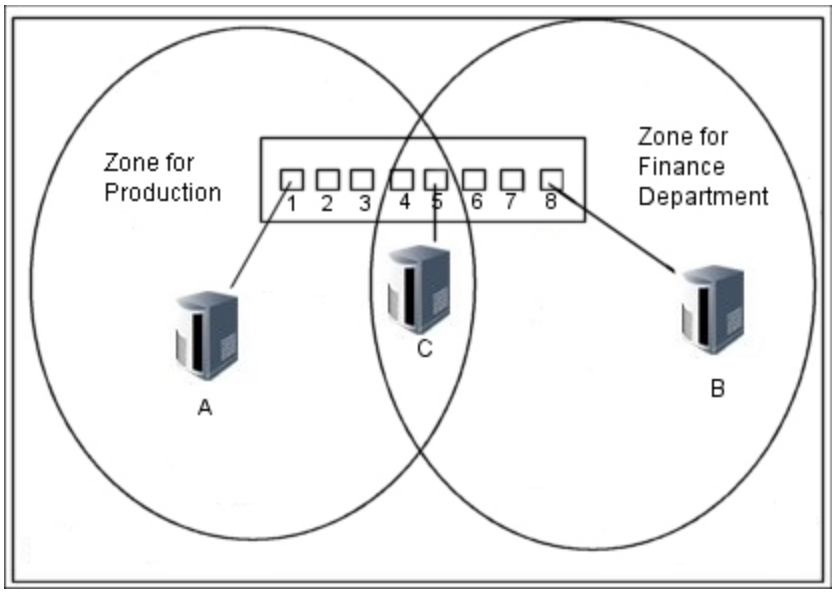

The following shows an overview of zoning structure.

Overview of Zoning Structure

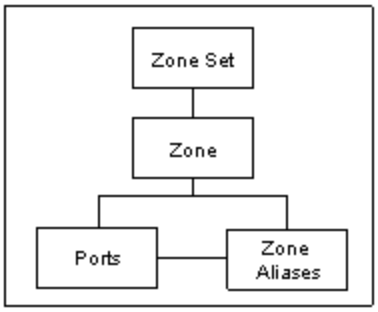

Zoning Structure

- **Zone Sets** – A zone set is a collection of zones. You can have only one zone set active at a time in a fabric; however, you can have a zone in more than one zone set. Zones sets are usually created for a particular task.
- **Zone** – A collection of zone aliases and ports.
- **Ports** – The WWN of the port to which an element is connected. The WWN of the port can be either the WWN of a switch port or the WWN of the connected element.

Setting Up Zoning

Use the following information as a guideline for setting up zoning.

Step	Description
1	Create a zone alias — *Optional*: Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports. See Creating a Zone Alias on page 683.
2	Create zones — Zoning is the primary tool to constrain groups of SAN members. A zone defines a logical SAN that contains limited element membership. The only elements visible to members of a zone are other members of that zone. See Creating a Zone in a Fabric on page 686.
3	Create zone sets — A zone set contains multiple zones. See Creating a Zone Set on page 688.
4	Activate a zone set — A switch fabric can have multiple zone sets defined, but only one zone set can be active. See Activating a Zone Set on page 691

Keep in mind the following:

- If you use another product to make zoning changes (such as adding a zone) you must perform Get Details in order to make the management server aware of these changes.
- The management server creates a zone by finding the port or WWN of discovered elements. The management server cannot create a zone using Fibre Channel addresses. If you use a

third-party software to create a zone using Fibre Channel addresses, the active zone will appear empty in the user interface of the management server.

Activities Supported by Zoning

Zoning Support by Switch Type below discusses the options supported for each type of switch.

Zoning Support by Switch Type

Options	Brocade	McDATA (BNA Discovery)	McDATA (Host Agent Discovery)	Cisco SNMP	Cisco SMI-S *
View Active and Interactive Zones/Zone Sets	Y	Y	Y	Y	Y
View Aliases	Y	Y	N	Y	Y
Zone Alias Provisioning	Y	Y	N	Y	Y
Zone/Zone Set Provisioning	Y	Y	Y	Y	Y
Zone Set Copying	N	Y	Y	N	N
Port Statistics	Y	Y	Y	Y	Y

* Soft zoning is only supported for Cisco SMI-S switches. For more information, see Glossary on page 913 in the Glossary. Hard zoning is not supported on Cisco SMI-S switches.

Issues for McDATA switches

Only one client at a time can provision on a McDATA fabric. However, since each fabric has a separate lock, you can perform simultaneous provisioning on two different fabrics. For example, you could perform provisioning by using the user interface and the CLI at the same time on two different fabrics. Simultaneous provisioning on the same fabric is not supported.

The management server does not support enabled default zones on McDATA switches. When a default zone is enabled on a McDATA switch, it is not listed as part of the active zone set.

Accessing Information about Zone Aliases

The software provides a listing of zone aliases in a fabric. You can view the properties of the zone alias and its port from this page.

For more information about which zoning features are supported for your switches, see Zoning Support by Switch Type above.

To access information about zone aliases in a fabric:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric on which you want to do provisioning.
4. Click **Step 1 Zone Alias**.

   This page lists the zone aliases and their ports under the following columns:

   - **Name** – Click the name of the zone alias to view its properties.
   - **Ports** – In some instance, you might be able to click the link of a port to view its properties.
   - **Active** – A check mark appears in the Active column if the zone alias is included in an active zone set.
   - **Edit** – Click the **Edit** ( ) button to edit an alias. For more information, see Modifying a Zone Alias on next page.
   - **Delete** – Click the **Delete** ( ) button for the zone alias you want to delete. For more information, see Deleting a Zone Alias on page 685.

   To create a zone alias, click **New Zone Alias**. For more information, see Creating a Zone Alias below.

Creating a Zone Alias

Zone aliases are used to give meaningful names to switch ports, HBA ports, or storage system ports.

To create a zone alias:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to create a zone alias.
4. Click **Step 1 Zone Alias**.
5. Click **New Zone Alias**.
6. From the Fabric drop-down menu, select the fabric in which you want to create the zone alias.

   Only the ports in the fabric you selected are displayed in the Potential Ports pane.
7. From the Switch drop-down menu, select the switch on which you want to create the zone alias.

8. In the Zone Alias Name box, enter a name for the zone alias. For more information, see Naming Conventions for Brocade and McDATA Switches on page 795.
9. Add ports to the zone alias by selecting a port in the **Potential Ports** pane.

   A port is not in the fabric if the  icon is next to it.
10. Remove ports from the zone by selecting them in the **Ports in the Zone Alias** pane and clicking **Remove From Zone**.
11. Click **OK**. The zone alias is created.

Zone Naming Conventions

The following naming conventions apply to zones, zone sets, and zone aliases:

- **Naming Conventions for Brocade Switches**:
  - Names can have a maximum of 64 characters.
  - Names must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
  - Names are case-sensitive. For example, Zone1 and zone1 are considered to be different zones.
  - You cannot create a zone with the same name as an existing zone, zone alias or zone set. For example, if you create a zone alias named “new”, you cannot give a zone, zone alias, or zone set the same name.
  - The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign ($).
- **Naming Conventions for McDATA Switches**:
  - Names can have a maximum of 64 characters.
  - Names must begin with a letter.
  - Names cannot contain spaces.
  - Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign ($).
  - All names must be unique and must not differ by case. For example, myzone and MyZone are considered to be the same zone.

Modifying a Zone Alias

To modify a zone alias:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to modify the ports.

4. Click **Step 1 Zone Alias**.
5. Click the **Edit** ( ) button.
6. Take one of the following actions:
   - To add ports, select a port in the **Potential Ports** pane.
   - To remove ports, select the ports in the **Ports in the Zone Alias** pane you want to remove, and then click **Remove From Zone**.
   - To select all of the ports, select the check box next to the Port heading.
7. Click **OK**.

Deleting a Zone Alias

To delete a zone alias:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to delete a zone alias.
4. Click **Step 1 Zone Alias**.
5. Click the **Delete** ( ) button for the zone alias you want to delete.
6. When you are asked if you want to delete the zone alias, click **OK**.

   You cannot delete a zone alias if it is the only member in a zone.

Accessing Information About Zoning

Keep in mind the following:

- For more information about which zoning features are supported for your switches, see Zoning Support by Switch Type on page 682.

To access information about zones and manage them:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to access information about zoning.
4. Click **Step 2 Zone**.

   This page lists the zones, their aliases and ports under the following columns:
   - **Name** – Click the name of the zone to view its properties.
   - **Zone Aliases** – Click the name of the zone alias to view its properties.

- **Ports** – In some instance, you might be able to click the link of a port to view its properties.
- **Active** – A check mark appears in the Active column if the zone is included in an active zone set.
- **Edit** – Click the **Edit** ( ) button to edit a zone. For more information, see Adding and Removing Zone Members on the facing page.
- **Delete** – Click the **Delete** ( ) button for the zone you want to delete. For more information, see Deleting a Zone on the facing page.

5. To create a zone, click **New Zone**. For more information, see Creating a Zone in a Fabric below.

Creating a Zone in a Fabric

To learn why zones are so important, see SAN Zoning Overview on page 679. A zone must have at least one member.

Keep in mind the following:

- Soft zoning is only supported for Cisco SMI-S switches. For more information, see the definition of soft zoneGlossary on page 913 in the Glossary. Hard zoning is not supported on Cisco SMI-S switches.
- For more information on zoning features supported for your switches, see Zoning Support by Switch Type on page 682.

To create a zone:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to create a zone.
4. Click **Step 2 Zone**.
5. Click **New Zone**.
6. From the Fabric drop-down menu, select the fabric in which you want to create the zone. Only the ports in the fabric you selected are displayed in the Potential Ports pane.
7. From the Switch drop-down menu, select the switch on which you want to create the zone.
8. In the Zone Name box, enter a name for the zone. For more information, see Zone Naming Conventions on page 684.
9. Add members to the zone by selecting a member in the **Potential Members** pane.

   Keep in mind the following:

- A zone member can be a port attached to a switch, WWN, or zone alias.
- As a best practice, a zone should contain zone aliases only, and there should be a zone alias for each port/WWN.
- You cannot create a zone with an existing name.
- A port is not in the fabric if the  icon is next to it.

10. Click **OK**.

Adding and Removing Zone Members

A zone must have at least one member.

To add and remove zone members:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to modify a zone.
4. Click **Step 2 Zone**.
5. Click the **Edit** () button.
6. Take one of the following actions:
   - To add a member to the zone, select a member in the Potential Members pane. A zone member can be a port attached to a switch, a WWN or a zone alias.

     Or

   - To remove from members from the zone, select the members in the **Zone Members** pane and click **Remove From Zone**.
7. Click **OK**.

Deleting a Zone

You cannot delete a zone if it is the only member in one of the zone sets or if it is a member of an active zone set. To delete a zone in an active zone set, first move the zone to an inactive zone set, and then delete it.

If you are using EFC Manager to delete zones, see Changes in EFC Manager Requiring Get Details on page 898.

To delete a zone:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to delete a zone.

4. Click **Step 2 Zone**.
5. Click the **Delete** ( ) button for the zone you want to delete.
6. When you are asked if you want to delete the zone, click **OK**.

Accessing Information about Zone Sets

For more information about which zoning features are supported for your switches, see Zoning Support by Switch Type on page 682.

To access information about zone sets and manage them:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to access information about a zone set.
4. Click **Step 3 Zone**.

   This page lists information about zone sets under the following columns:

   - **Name** – Click the name of the zone set to view its properties.
   - **Zones** – Click the name of the zone to view its properties.
   - **Active** – To make a zone set active, select its corresponding **Active** option. When you select a zone set, you make elements outside of the zone set inaccessible. For more information, see Activating a Zone Set on page 691.
   - **Edit** – Click the **Edit** ( ) button to edit a zone set. For more information, see Modifying a Zone Set on the facing page.
   - **Delete** – Click the **Delete** ( ) button for the zone set you want to delete. For more information, see Deleting a Zone Set on page 690.
5. To create a zone set, click **New Zone Set**. For more information, see Creating a Zone Set below.

Creating a Zone Set

To learn why zone sets are so important, see SAN Zoning Overview on page 679.

For more information about which zoning features are supported for your switches, see Zoning Support by Switch Type on page 682.

To create a zone set in a fabric:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to create a zone set.

4. Click **Step 3 Zone Set**.
5. Click **New Zone Set**.
6. From the Fabric drop-down menu, select the fabric in which you want to create the zone set.

   Only the ports in the fabric you selected are displayed in the Potential Members pane
7. From the Switch drop-down menu, select the switch on which you want to create the zone set.
8. In the Zone Set Name box, enter a unique name for the new zone set. For more information, see Zone Naming Conventions on page 684.
9. Take the desired action:
   - To make a zone set active, select the **Activate this Zone** option. Only one zone set can be active at a time. When you make a zone set active, the previous zone set becomes inactive.
   - To add zones to the zone set, select a zone in the **Zones Not in this Zone Set** list and click the greater than sign (**>**). A zone can be in multiple zone sets.
   - To remove zones from the zone set, select a zone in the **Zones in this Zone Set** list and click the less than sign (**<**).
10. Click **OK**.

Modifying a Zone Set

To modify a zone set:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to modify a zone set.
4. Do one of the following:
   - Click **Step 3 Zone Set**.

     Or
   - Click **Step 4 Activate this Zone Set**.
5. Click the **Edit** () button.
6. Take the desired action:
   - To add zones to the zone set, select a zone in the **Zones Not in this Zone Set** list and clicking the greater than sign (**>**). A zone can be in multiple zone sets.
   - To remove zones from the zone set, select a zone in the **Zones in this Zone Set** list, and click the less than sign (**<**).
7. Click **OK**.

Deleting a Zone Set

The software does not display all elements in a zone set, such as quick loop and fabric assist elements. When you delete a zone set, all elements, including quick loop and fabric assist, which are not viewable in the software, are deleted.

Only the zone set is deleted, not the zones contained in the zone set. For example, assume Zone A is contained in two zone sets, one named Zone_Set_One and another named Zone_Set_Two. If you delete Zone_Set_One, the zone has not been deleted so it is still in Zone_Set_Two.

If you are using EFC Manager to delete zone sets, see Changes in EFC Manager Requiring Get Details on page 898.

To delete a zone set:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric which contains the zone set you want to delete.
4. Do one of the following:
   - Click **Step 3 Zone Set**.

     Or
   - Click **Step 4 Activate Zone Set**.
5. Click the **Delete** () button for the zone set you want to delete.
6. When you are asked if you want to delete the zone set, click **OK**.

Copying a Zone Set

This feature copies a zone set and all of its members, such as zones and zone aliases. You can use this feature to copy inactive and active zone sets. The newly created zone set is inactive.

The management server stops the copying process of an active zone set if it finds one of the following:

- An inactive zone set that has the same name as the name entered for the copy.
- An inactive zone with the same name as an active zone, but they do not have the same content.

Active zones in a zone set have corresponding inactive zones for redundancy. If you attempt to copy an active zone set containing a zone that does not have a corresponding inactive zone, the management server creates an inactive zone with the same name as the active zone. The inactive zone is used as a backup for the active zone.

This feature is supported only for switches that support zone set copying. For information on which switches support zone set copying, see Zoning Support by Switch Type on page 682. Aliases in the zone set are not copied over for McDATA switches.

To copy a zone set:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to copy a zone set.
4. Click **Step 3 Zone Set**.
5. Click the corresponding button for the zone set you want to copy.
6. Enter the name of the new zone set. If you are copying an active zone set, do not enter the name of a pre-existing inactive zone set. For more information, see Zone Naming Conventions on page 684.
7. Click **OK**. The zone is copied.

Activating a Zone Set

You can only have one zone set in a fabric active at a time. When you make a zone set active, the previous active zone set becomes inactive. However, you could have a zone in more than one zone set.

For more information about which zoning features are supported for your switches, see Zoning Support by Switch Type on page 682.

To make a zone set active:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **SAN Zoning** tab.
3. Click the **Provision** button for the fabric in which you want to activate a zone set.
4. Click **Step 4 Activate Zone Set**.
5. Select the corresponding **Active** option.
6. *McDATA switches*: The management server enables you to create a backup copy of the zone set you want to activate. To create a backup of the zone set that will become active:
   a. Select the option **Make a backup copy of the active zone set after activation**.
   b. (*Optional*) In the **Name** box, modify the name that has been assigned to the backup zone set. The management server assigns the name by appending the date and time of the zone set you have selected to become active, as shown in the following example:

      ```
      zone_name_2005-05-17_13-41-05
      ```

      In this instance:

- zone_name is the name of the zone you are making active.
- 2005-05-17 is the date you made the zone active in the format yyyy-mm-dd. In the example, the date is May 17, 2005.
- 13-41-05 is the time the copy was made in the format hh-mm-ss, using 24 hour notation. The time is formatted as hour-minute-second, and it uses the 24-hour notation. In the example, the time is 1:41:05 p.m.

The name of the backup zone set must follow the naming conventions described in Zone Naming Conventions on page 684.

The management server truncates the name of the backup zone set if it is more than 44 characters long so that it can fit the date and time into the zone set name, which cannot be more than 64 characters. For example, if you have a zone set named McDATA_Switches_Burlington_Massachusetts_United_States, The management server truncates the name to McDATA_Switches_Burlington_Massachusetts_Un_2005-05-17_13-41-05 when it creates the backup zone set.

7. To activate the zone, click **OK**.

Zones and Zone Sets Listed Twice

Sometimes the Navigation tab and the Provisioning Manager pages list the zones and zone sets twice for McDATA switches. This is because EFC Manager contains an offline Zoning Library. This Zoning Library holds all zone sets, zones, and zone members. When you activate a zone set, the zone set, along with its zones and zone members, is copied to the McDATA switches and activated. This creates an active copy in addition to the saved copy that already exits in the Zoning Library. If you edit the saved copy of the zone or zone set in the Zoning Library, you must reactivate the saved copy of the active zone set.

For example, assume you have the following information in the EFC Zoning Library:

```
ZoneSet A

        Zone A1

                ZoneMember A1a

ZoneSet B

        Zone B1

                ZoneMember B1b
```

```
If a user activates ZoneSetB, the existing information in ZoneSetB is
copied to the switch and activated. The Zoning Library, however,
still contains the older information. The following is what you would
see in the Zoning Library:

ZoneSets:

    ZoneSet A

    ZoneSet B   (this is the inactive ZoneSet in the Zoning Library)
```

```
    ZoneSet B   (this is the active ZoneSet which is not the same as
the one above)

Zones:

    Zone A1

    Zone B1   (this is the inactive one)

    Zone B1   (this is the active one)

ZoneMembers:

    ZoneMember A1a

    ZoneMember B1b   (inactive)

    ZoneMember B1b   (active)
```

ZoneSetB and its members are listed twice in the previous example. They are also displayed twice on the Navigation tabs and Provisioning Manager pages. For example, If you click **Zone A1** in EFC Manager, or if you remove a zone from the Zoning library, the zone is shown only once in the user interface of the management server. The zone is not displayed at all in the management server if it is an inactive zone. Each time you make changes in the Zoning Library using EFC Manager, you must perform Get Details for the management server to obtain the latest information.

About the Messages Displayed in the Brocade Console

You can ignore the following messages in the Brocade Console. It is just the management server talking to the Brocade switch and subscribing to the event service.

```
wwn: 10:0:0:60:69:12:1a:5, cnt: 1, ipclnts: 0

total number of api event proxy servers: 0

total number of api event proxy servers: 0

total number of api event proxy servers: 1

proxy server 0 Did: 16776194

wwn: 10:0:0:60:69:10:56:fb, cnt: 1, ipclnts: 0

total number of api event proxy servers: 1

proxy server 0 Did: 16776194

wwn: 10:0:0:60:69:10:56:fb, cnt: 1, ipclnts: 0

wwn: 10:0:0:60:69:12:1a:5, cnt: 1, ipclnts: 0
```

To view the status of the switch, access the Brocade Console by typing the user name and password for the switch in a Telnet window, and then enter `switchshow` at the prompt.

Managing Storage

This section contains the following topics:

About Setting Up Storage Partitioning

Each storage vendor treats storage partitioning differently. For example, Hitachi and EMC ship their storage systems with the volumes already created. Other storage vendors, such as LSI, ship their storage systems as empty arrays.

Despite the differences among storage systems, you can still use this product to manage your provisioning. Some tasks, such as volume creation, might create different results depending on the type of storage system.

To learn how host security groups are created on your storage systems, see Rules for Creating Host Security Groups on page 712.

Provisioning and Pool Support

Storage System	Storage Provisioning	Create/ Delete Pool	Create Pool Using Settings*	Additional Information
3PAR	N	N	N	

Storage System	Storage Provisioning	Create/ Delete Pool	Create Pool Using Settings*	Additional Information
CLARiiON	Y	Y	Y	The EMC Navisphere CLI is required to communicate with a CLARiiON storage system. And the CLARiiON storage system must be configured to recognize the management server as a privileged user. For more information, see Discovering EMC CLARiiON Storage Systems on page 113.
EMC Symmetrix	Y	N	N	
HDS	Y	N	N	

Storage System	Storage Provisioning	Create/ Delete Pool	Create Pool Using Settings*	Additional Information
HP EVA	Y	Y	Y	To determine provisioning support for HP StorageWorks Arrays, see "Provisioning and Pool Support" and "Volume and Host Security Group Support" in the *User Guide*. EVA arrays can only be provisioned if they are actively managed by the Command View server that they are discovered through. When an EVA is discovered by the legacy built-in EVA provider, a cache is created and populated with the current array configuration. Each subsequent cache refresh will start 30 minutes after completion of the previous cache refresh. The time the cache refresh takes depends on factors such as the EVA configuration, model, and SAN traffic. When you perform a provisioning operation—for example, create, delete, or modify a pool or volume—the cache information about provisioning is immediately updated. If you provision an EVA using Command View EVA or a different management station, the cached information about the EVA will not be accurate until the cache is refreshed.
HP MSA 1000/1500	Y	N	N	Selective Storage Presentation (SSP) must be enabled for provisioning to work.
HP MSA P2000 G2 (2312fc/2324fc)	N	N	N	

Storage System	Storage Provisioning	Create/ Delete Pool	Create Pool Using Settings*	Additional Information
HP MSA P2000 G3 FC	N	N	N	
HP XP	Y	N	N	
IBM DS	Y	N	N	IBM and other storage systems that use external providers show raw capacity for a storage pool, instead of formatted capacity.
IBM XIV	N	N	N	
LSI	Y	Y	Y	
X9000	N	N	N	
Xiotech	Y	N	N	

*Refers to the functionality that enables you to choose the type of pool, usually RAID level.

Volume and Host Security Group Support

Storage System	Create/ Delete Volume	Create Volume Using Settings*	Create/ Delete Meta Volume	Host Security Group Provisioning Supported	Additional Information
3PAR	N	N	N	N	
CLARiiON	Y	Y	Y	Y	RAID level can be specified for first volume in a pool, subsequent volumes inherit this setting.
EMC Symmetrix	Y	N	Y	Y	Creating a volume means marking an existing device as accessible to host security group management. Deleting a volume means returning the device to the free device pool.

Storage System	Create/ Delete Volume	Create Volume Using Settings*	Create/ Delete Meta Volume	Host Security Group Provisioning Supported	Additional Information
HDS	Y	N	Y	Y	Creating a volume means marking an existing device as accessible to host security group management. Deleting a volume means returning the device to the free device pool.
HP EVA	Y	Y	N	Y	
HP MSA 1000/1500	Y	N	N	Y	
HP MSA P2000 G2 (2312fc/2324fc)	N	N	N	N	
HP MSA P2000 G3 FC	N	N	N	N	
HP XP	Y	N	Y	Y	Creating a volume means marking an existing device as accessible to host security group management. Deleting a volume means returning the device to the free device pool.

Storage System	Create/ Delete Volume	Create Volume Using Settings*	Create/ Delete Meta Volume	Host Security Group Provisioning Supported	Additional Information
IBM DS	Y	N	N	Y	Volume groups are not represented in the provisioning or properties of IBM storage systems. The host security groups are based on a host connection, which does not indicate which volume groups each of them belongs to. A volume group can have many host connections, but only one host connection can be part of a volume group.
IBM XIV	N	N	N	N	
LSI	Y	Y	N	Y	
X9000	N	N	N	N	
Xiotech	Y	Y	N	Y	During a provisioning operating on Xiotech, the Provisioning Manager interface remains busy and eventually times out. You can close the provisioning page after a few minutes and check on the array, which should be updated properly.

*Refers to the provisioning capability on the management server that enables you to create different kinds of volumes from a pool depending on whether you want it optimized, such as for streaming, random access,or high availability.

For more information about how storage provisioning works on your storage system, see:

- Issues Specific to EMC Symmetrix Storage Systems on page 724
- Issues Specific to HDS Storage Systems on page 725
- Issues Specific to LSI Storage Systems on page 728

Setting Up Storage Partitioning

To set up storage partitioning:

1. *LSI and CLARiiON only*: Create a storage pool (sometimes referred to as a volume group or RAID group). For more information about storage pools, see About Storage Pools below
2. Create a volume. For more information about volumes, especially in CLARiiON storage systems, see About Volumes below.
3. Create a host security group. For more information about host security groups, see About Host Security Groups below.

About Storage Pools

A storage pool is a group of disks associated together through a RAID configuration. The pool's capabilities define the level of protection for the associated volumes and LUNs. EMC Symmetrix and HDS storage systems have storage pools that are predefined. LSI and CLARiiON storage systems require that a storage pool be created and volumes be allocated from the storage pool. For more information, see Creating a Storage Pool (HP EVA, LSI, and CLARiiON) on page 702.

About Volumes

A volume is a virtual disk. Volumes are created in sizes that are desirable for being shown as a LUN. A volume can be associated with more than one fibre channel port, creating multiple LUNs corresponding to the same volume. (The defining characteristics of a LUN are the volume, port, and LUN number.)

On CLARiiON storage systems a volume is owned by one of the storage processors. Creating a volume also creates a LUN for this volume and for each port of the storage processor that owns the volume. A LUN mapped to a port is visible to all the ports on that controller. Mapping a volume to a port on a CLARiiON storage system also maps that volume to all ports that reside on the same storage processor as the selected port. It also causes the volume to be unmapped from all the ports of the other storage processor. Some storage systems have their volumes fully configured during install. For these storage systems, Users can concatenate multiple volumes together to create a new volume.

For more information, see Creating a Storage Volume on page 707.

About Host Security Groups

Host security groups define which initiators (HBA ports) have access to specified storage volumes. Host security groups are associated with a fibre-channel port and contain a list of HBA port initiators and the volumes they can see.

For more information, see Creating Host Security Groups on page 718 and Rules for Creating Host Security Groups on page 712.

Modifying the Cache Settings (LSI)

To modify the cache settings:

1. Click the **Edit** () button for the volume you want to modify.
2. Enter the cache read ahead multiplier (0 to 65535 bytes) in the Cache read-ahead multiplier box.

   A cache read ahead multiplier copies additional data blocks into the cache while it is reading and copying host-requested data blocks from disk to cache. To disable this option, enter 0.
3. (*Optional*): Select **Read Caching**.

   Use this option to store the host's operations in controller cache memory.
4. *(Optional)*: Select **Write Caching**.

   Use this option to write data to the cache memory of a controller.
5. (*Optional*): Select **Write Caching with Mirroring**.

   Use this option to preserve data if a controller or the cache fails. When this option is enabled, the data is written to two redundant controllers of the same cache size. This configuration provides redundancy in case a controller fails. One controller performs uncompleted write operations when the other controller fails.

For information about changing the owner assigned to the volume, see Changing the Owner of a Volume (LSI and CLARiiON) below.

Changing the Owner of a Volume (LSI and CLARiiON)

When a volume is created, the management server automatically assigns a controller to be the owner of the volume. You can change the owning controller if you want to use a different controller for LUN masking.

If the owner becomes unreachable as a result of a network failure or the owner itself fails, the other controller in the pair automatically becomes the owner of the volume.

To change the owner:

1. Click the **Edit** () button for the volume you want to modify.
2. Select the owner for the volume from the **Current Owner** menu.
3. Click **OK**.

Managing Storage Pools

This section contains the following topics:

Creating a Storage Pool (HP EVA, LSI, and CLARiiON)

A storage pool is a group of disks associated together through a RAID configuration. The pool's capabilities define the level of protection for the associated volumes and LUNs. You should create at least one storage pool before provisioning a volume.

Note:When provisioning a storage pool from HP Storage Essentials for HP EVA, the size of the pool created will be always more than the size specified during provisioning. This is because Command View EVA does not support creation of disk groups by size but by the number of disks. The SMI-S standard APIs only allow disk group creation by size. Hence, due to the size conversion HP Storage Essentials always uses more space than the lowest possible number to create a storage pool.

To create a storage pool:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system in which you want to create a storage pool.
4. Click **Step 1 Storage Pool**.
5. Click **New Storage Pool**.
6. Select a setting and size for the storage pool. The sizes displayed depends on the RAID level you want. For example, RAID 0 does not require additional drives, so you can assign more space to the pool.

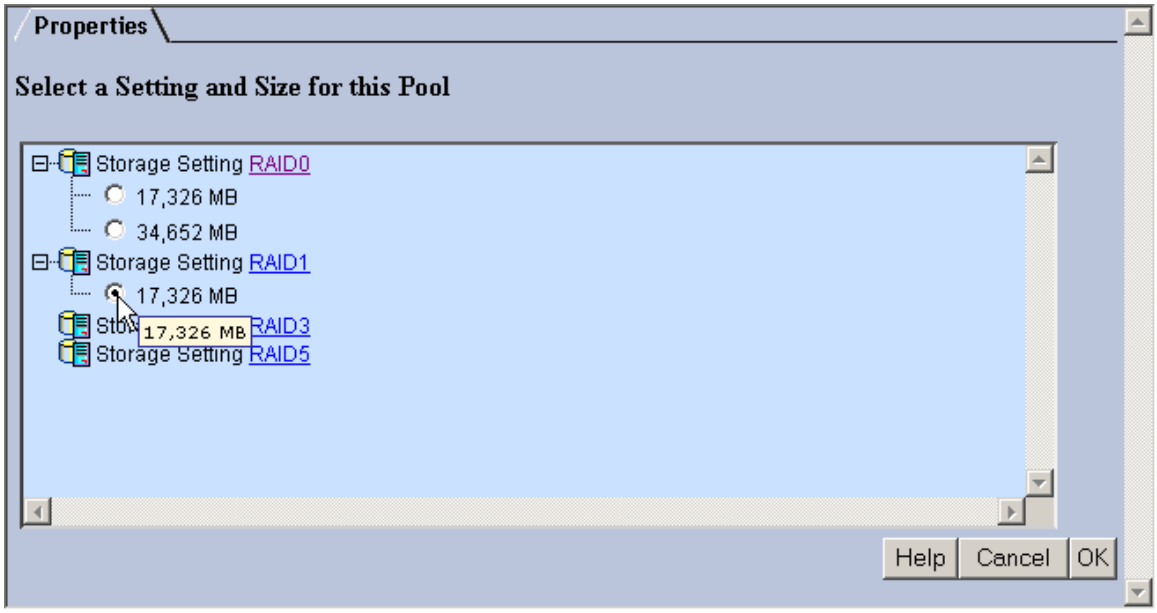

7. Click **OK**. The storage pool is created.

When you create a pool on an LSI storage system, a placeholder volume is created inside the new volume group. The name of the placeholder volume starts with “Required - do not delete.” The placeholder volume is required because the storage pool cannot not exist without it. The management server does not display the placeholder volume, but other monitoring products could display this volume.

Accessing Information about Storage Pools

If you used another product to make provisioning changes, perform the Get Details operations prior to viewing information about storage pools. The Get Details operation enables the management server to be aware of your changes.

To access information about storage pools:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system in which you want to access information about storage pools.
4. Click **Step 1 Storage Pool**.

   The following information about the storage pools is displayed:

   - **Pool Name** – Click the name of the storage pool to view its properties.
   - **Size** – Displays the amount of space assigned to the storage pool.
   - **Available** – Displays the amount of space available in the storage pool.
   - **Used** – Displays the amount of space used in the storage pool.
   - **Volumes** – Click the name of the volume to view its properties. If the storage system has a large number of volumes, not all the volumes are displayed. To display all the volumes, select the **Show All** option.
   - **Capabilities** – Click the RAID level or name to view its properties.
   - **Mainframe** – Displays whether the storage pools have volumes that are on a mainframe. (Available to only HDS storage systems).
   - **Delete** – Click the **Delete** () button for the storage pool you want to delete. For more information, see Deleting a Storage Pool (HP EVA, LSI and CLARiiON Only) on next page. (Available only to LSI and CLARiiON storage systems.)
5. To create a storage pool, click the **New Storage Pool** button in the upper-right corner of the page. For more information, see Creating a Storage Pool (HP EVA, LSI, and CLARiiON) on previous page.
6. If the table contains more than ten entries, the following buttons for navigating through the table are enabled:

- – Move to the first page.
- – Move back one page.
- – Move forward one page.
- – Move to the last page.

Deleting a Storage Pool (HP EVA, LSI and CLARiiON Only)

When you delete a storage pool on an LSI storage system, all the volumes in the volume group are deleted, including the placeholder volume.

To delete a storage pool:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system in which you want to delete a storage pool.
4. Click **Step 1 Storage Pool**.
5. Click the **Delete** ( ) button for the storage pool you want to delete. A message warns you about the other volumes that will be deleted.
6. Click **OK**. The storage pool and its volumes are deleted.

Managing Volumes

This section contains the following topics:

Accessing Information about Volumes

Some storage vendors require a password to access the storage system. If the correct password is not entered, an authentication error message is displayed. Refer to Get Details to determine which user account was used to access the storage system during discovery.

If you use another product to make provisioning changes, you must perform Get Details for the management server to be made aware of these changes.

To access information about volumes:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system you want to provision.
4. Click **Step 2 Volume**. All volumes are displayed.

   To filter the volumes displayed, see Filtering Volumes on next page.

   The following columns list information about the volumes:

   - Volume – Click the name of the volume to view its properties.
   - Size – Displays the size of the volume in megabytes (MB).
   - Ports – Click a ports link to view its properties.
   - Pool – Click the name of the pool to view its properties.
   - Mainframe – Displays whether the storage pools have volumes that are on a mainframe. (This option is available to only HDS storage systems).
   - *Default Owner – The controller which owns the storage system when it is rebooted.
   - *Current Owner – The controller which currently owns the volume.
   - *Segment Size – Displays the amount of space assigned to a volume in megabytes (MB).
   - *Read ahead – Displays the cache read ahead multiplier.
   - *Edit – Click the **Edit** () button for the volume you want to edit. For more information, see Modifying the Cache Settings (LSI) on page 700.
   - Delete – Click the **Delete** () button for the volume you want to delete. For more information, see Deleting a Storage Volume on page 710.

     *Not accessible to all storage systems.
5. To create a volume, click the **New Volume** button in the upper-right corner of the page. To delete several volumes at once, select the volumes you want to delete and then click **Delete Selected Volumes**.

If you have an HDS storage system, see About Volumes on HDS Storage Systems on next page.

If the table contains more than 10 entries, the following buttons for navigating through the table are enabled:

-  – Move to the first page.
-  – Move back one page.
-  – Move forward one page.
-  – Move to the last page.

Filtering Volumes

To filter the list of volumes displayed:

1. Click the **Show Volume Filter** link to display the filtering options.

   If the volume filter is turned on, the link appears as **Hide Volume Filter**.
2. Enter all or part of a volume name in the Name Contains box.
3. Enter a size in megabytes in the Size (>=) MB box.

   Volumes that are greater than or equal to the size specified appear on the page.
4. Select one of the following:
   - **All Volumes** to display all existing volumes.
   - **Unmapped Volumes** to display only unmapped volumes are displayed.
   - **Mapped Volumes** to display only mapped volumes are displayed.
5. Select the port you want to display from the **Ports** menu.

   If you want to display all ports, select **All** from the Ports menu.
6. Select the storage pools you want displayed from the **Storage Pools** menu.
7. Click **Filter**. The table is updated to display only the elements that meet the filter criteria.

To reset the filter criteria, click **Reset**.

About Volumes on HDS Storage Systems

Volumes from single LDEVs are shown as LDEV:0 on HDS storage systems. When volumes made up of multiple LDEVs are first created, they are not mapped to a target port on the storage system. The software remembers that these LDEVs constitute a single volume, but it does not make changes to the storage system until the volume is mapped to a port. As a result, they are referred to as Groups; for example:

```
Group:0(LDEV:0, LDEV:1)
```

In this instance, 0 in Group:0 is the volume identifier and LDEV:0 and LDEV:1 are the LDEVs that make up this volume.

After you create a storage volume on an HDS storage system, map the volume to a target port on the storage system, using the storage system Provisioning tool. In the tool, click **Step 3, LUNs**.

Once the volume is mapped, it is displayed as a logical unit size expansion (LUSE); for example:

```
LUSE:0(LDEV:0, LDEV:1)
```

In this instance, 0 in Group:0 is the volume identifier and LDEV:0 and LDEV:1 are the LDEVs that make up this volume.

Creating a Storage Volume

Some storage vendors require a password to access the storage system. If the correct password is not entered, an authentication error message is displayed. Refer to **Discovery** > **Details** to determine which user account was used to access the storage system during discovery.

When you create a storage volume, you can set its size, volume capabilities, and storage pool.

On HDS and Symmetrix storage systems, volumes are shipped already created. When you create a volume in the management server on these storage systems, you are defining the volume as being allocated.

For XP arrays, you do not have option to select the SLPR while creating a volume. HP Storage Essentials automatically assigns the SLPR for the volume creation, even though SLPR might not be required.

Rounding Volume Size

Some vendor's tools for HDS might round off the volume size, so that a 6.87-GB volume appears as 7 GB (7168 MB) in the tool. The management server displays the size of the volume without rounding. For example, assume you want to create a 14-GB (14336 MB) LUSE volume, and according to the storage tool, you have two 7-GB LDEVs, which are really 6.87 GB (7034.88 MB). If you look at the native tool, it would be logical to assume only two LDEVs would be required to create the 14-GB LUSE volume. The management server would use three LDEVs because each LDEV is only 6.87 GB.

If you are creating a volume on an HP EVA storage system, its external SMI-S provider might round the specified number of megabytes to the nearest whole gigabyte.

Support for PvLinks

PvLinks based on HP disk partitions are not supported. Any volumes created on such PvLink meta devices are shown as local. If you partition a regular (non-PvLink) external disk and create volumes based on it, then volumes are recognized as external volumes.

For HDS Storage Systems

A LUSE volume on an HDS storage system is not created until you map that volume to a target array port. In the management server, the create volume and LUN creation tasks are two different operations. So if you want to create a LUSE volume and then perform LUN creation on the HDS box, it is a two-step process. First, use the management server to create LUSE volumes. Then create a LUN and map a volume to a target port, by creating a host security group. For more information, see Creating Host Security Groups on page 718 and Rules for Creating Host Security Groups on page 712.

HDS ships some of its storage systems with volumes already created. When the software first discovers an HDS storage system, it detects the volumes created by HDS. When you use the software's “create a volume” feature, you are assigning the already-created volume. For more information, see About Provisioning on HDS Storage Systems on page 725.

LUSE made up of volumes from different RAID levels are not supported. You cannot use this product to provision this type of LUSE. Existing LUSEs of this type could be incorrectly reported.

It is not possible to extend LUSEs using the management server software. You must perform this operation using the native tools.

Note: In a Hitachi Freedom Storage Thunder 9500 V Series storage system, it is not possible to create a LUSE from a group of RAID-0 LDEVs.

For LSI Storage Systems

Make sure you select a volume group that can accommodate the requested size for the new volume.

You can create volumes from existing free extent areas within a volume group.

No volume-to-LUN masking is done by default, except for CLARiiON storage systems. For more information, see Creating Host Security Groups on page 718 and Rules for Creating Host Security Groups on page 712.

Creating a Storage Volume

To create a storage volume:

1. Access the Create Storage Volume wizard:
   a. Click **Provisioning Manager** ().
   b. In the right pane, click the **Storage Systems** tab.
   c. Click the **Provision** button for the storage system in which you want to access information about volumes.
   d. Click **Step 2 Volume**.
   e. Select the desired number of LDEVs for the LUSE volume then click **Delete Selected Volumes**. Note the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.
2. Click **New Volume**.

   You can also access the Create Storage Volume wizard from the Navigation tab in System Manager. To access the wizard from the Navigation tab, click the Volumes link for a storage system, then click the **New Volume** button at the bottom of the screen.
3. *LSI and CLARiiON*: In the Volume Name box, enter a name for the volume. If you do not provide a name, the software assigns one.

   If you enter a volume name, observe the following conventions:

- It cannot be more than 30 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or one of the following symbols: dollar sign ($), caret (^), or an underscore (_).
- The name is case-sensitive; for example, “StorageVolume1” and “storagevolume1” are different storage volumes.
- The name must be different from any other volume name on an LSI storage system.

4. In the Size box, enter the size of the volume in megabytes (MB) or gigabytes (GB). Select the appropriate unit of measurement from the menu to the right of the Size box.

   The management server creates a volume of at least the size specified. For example, assume you asked the management server to create a 15-MB volume, and you have only three free extents: 10 MB, 17 MB, and 100 MB. A 17-MB volume will be created instead of a 15-MB volume because that is the closest size available. Keep in mind that although the management server tries to find free extents that make the volume size as close to the requested size, there is no guarantee it will pick the optimal combination of free extents.

5. Select a storage pool for the volume.

   Keep in mind the following:

   - If you do not see space available in the storage pool, you must delete volumes. For example, assume you want to create an 8-MB volume, but you do not have space available. Each volume is made up of 4 MB, therefore you must delete two volumes from that storage pool. Make sure those volumes you delete are not being used. For more information, see Deleting a Storage Volume on next page.
   - If you do not see any storage pools, verify that you have obtained all element details from the storage system. For more information, see Discover Switches on page 84.

6. Click **Next**.

7. Select a volume capability. The volume capabilities listed depend on the type of storage system. For example, if an EMC Symmetrix storage system is selected, Pool default settings are displayed. If an LSI storage system is selected, the following information is displayed:

   - <Default> – Provides the default cache read ahead multiplier and the default segment size for the storage system.
   - File System (Typical) – Provides a cache read ahead multiplier of 1 with a segment size of 64 KB.
   - Database – Provides a cache read ahead multiplier of 0 with a segment size of 64 KB.
   - Multimedia – Provides a cache read ahead multiplier of 8 with a segment size of 128 KB.
   - Custom – Enables you to customize the cache read ahead multiplier and the segment size.

HDS only: Under the Volume Capabilities tab, keep the default selection and click **Finish**. Once the settings have been made, you will see a new volume called a Group Volume in the list of unmapped volumes. Technically this is not a LUSE yet, as it has not been assigned to a port; it is a logical grouping within the management server. Think of it as a place holder. From this point you can select the new group volume and assign it to a port. Once the volume has been assigned to a port, the management server makes the changes to the array and creates the LUSE. For more information about how to create a LUN, see Creating Host Security Groups on page 718 and Rules for Creating Host Security Groups on page 712.

8. *LSI only*: If you selected the **Custom** option, follow these steps:
   a. Enter the cache read ahead multiplier (0 to 65535 bytes) in the Cache read ahead multiplier box.

      A cache read ahead multiplier copies additional data blocks into the cache while it is reading and copying host-requested data blocks from disk to cache. Select the multiplier that maximizes performance for the way the volume will be utilized.

   b. Select a segment size from the menu.
9. Click **Finish**.

Deleting a Storage Volume

When you delete a storage volume on an HDS or Symmetrix storage system, the software marks the deleted volume as hidden in the CIM repository, making it unassigned, instead of being deleted. The software keeps track of the "deleted volumes."

Keep in mind the following:

- Some storage vendors require a password to access the storage system. If the correct password is not entered, an authentication error message is displayed. Refer to **Discovery** > **Details** to determine which user account was used to access te storage system during discovery.
- If you remove volumes from host storage groups that are command devices or are pair volumes on HDS storage systems, later modification of the pair volumes could be disabled.

To delete a storage volume:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system in which you want to access information about volumes.
4. Click **Step 2 Volume**.
5. To display only unmapped volumes, select the unmapped volumes node in the left pane.
6. Click the **Delete** () button for the volume you want to delete.
7. When you are asked if you want to delete the volume, click **OK**.

8. To delete several storage volumes at once, select the storage volumes you want to delete and then click **Delete Selected Volumes**.

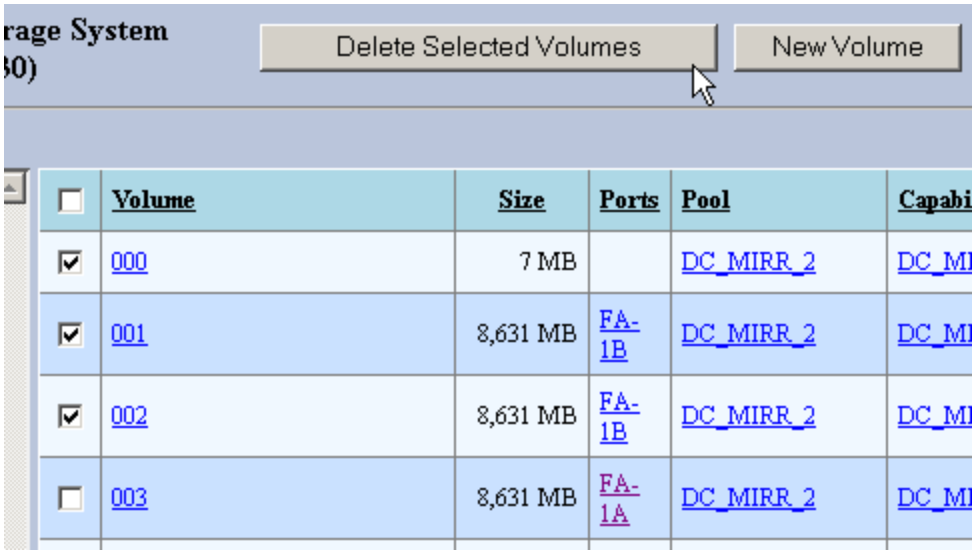

To select all volumes, select the check box next to the Volume heading, as shown in the following figure:

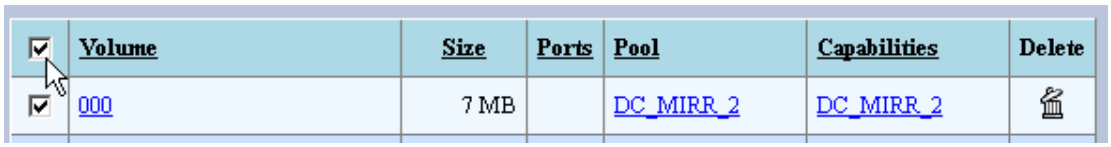

Changing the Cache Block Size for a Storage System (LSI)

To change the cache block size:

1. Do one of the following:
   - Click **Provisioning Manager** ( ), and then click the link for the storage system.

     Or

   - Click **System Manager** ( ) in the left pane.
2. Scroll to the bottom of the Navigation page.
3. Click **Change**.
4. Select the cache block size from the menu. Set a higher cache size for applications that require a lot of input and output, such as multimedia.
5. Click **OK**.

Modifying the Cache Settings (LSI)

To modify the cache settings:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system in which you want to access information about volumes.
4. Click **Step 2 Volume**.
5. Click the **Edit** () button for the volume you want to modify.
6. Enter the cache read ahead multiplier (0 to 65535 bytes) in the Cache read-ahead multiplier box. To disable this option, enter 0.
7. (*Optional*) Select **Read Caching**.

   When this option is enabled, the host's operations are stored in controller cache memory.
8. (*Optional*) Select **Write Caching**.

   When this option is selected, data is written to the cache memory of a controller
9. (*Optional*) Select **Write Caching with Mirroring**.

   Use this option to preserve data if a controller or the cache fails. When this option is enabled, the data is written to two controllers of the same cache size, providing redundancy, so that if one controller fails, the other controller performs the uncompleted write operations.
10. To change the current owner, select a new owner from the **Current Owner** menu.
11. Click **OK**.

Rules for Creating Host Security Groups

The management server now uses host security groups instead of LUN masking and LUN mapping. With the introduction of host security groups, the management server has a new definition of *mapped* for this release. *Mapped* refers to capacity that is accessible by one or more hosts external to the array (aggregated capacity of volumes that are accessible from hosts external to the subsystem).

The number of initiators allowed in the host security group depends on the type of storage system. For more information about how each storage system treats host security groups, see the topics referenced in the following table.

Allowed Initiators in Host Security Groups

Storage System	Allowed Initiators in Host Security Groups	For more information, see...
EMC CLARiiON		Host Security Groups on EMC CLARiiON Storage Systems on the facing page
EMC Symmetrix	1 initiator for host security masking	Host Security Groups on EMC Symmetrix Storage Systems on page 714

Storage System	Allowed Initiators in Host Security Groups	For more information, see...
LSI		Host Security Groups on LSI Storage Systems on next page
HDS		Host Security Groups on HDS Storage Systems on page 715
HP EVA		Host Security Groups on HP EVA Storage Systems on page 715
HP MSA	1 initiator	Host Security Groups on HP MSA 1000/1500 Storage Systems on page 715
IBM	Only 1 initiator per host security group	Host Security Groups on IBM Storage Systems on page 715
Xiotech	1 initiator	Host Security Groups on Xiotech Storage Systems on page 716

For the Volume Creation and LUN Security option in Path Provisioning, the All Ports node is not shown because volumes cannot be placed inside host security groups for All Ports.

Host Security Groups on EMC CLARiiON Storage Systems

Keep in mind the following rules for host security groups on EMC CLARiiONstorage systems

- When a volume is created, by default, it is assigned to one of the two controllers. Even though this volume is mapped to a controller, it is not visible from the outside world by a host. The management server reports this volume as unmapped since it is not visible by a host initiator.
- Volumes can be only on SPA or SPB because CLARiiON is active/passive storage, which means it can have only one active path to a volume. Addition of initiators to any of the ports on a storage processor is listed for all ports of that storage processor.
- The host security group is created on all ports of the processor you select unless you select an initiator that uses a different processor and does not belong to a host security group. For example, assume you select processor SPA, and then you select an initiator that belongs to SPB buts does not belong to a host security group. The host security group is created for all ports on SPB.
- Host security groups can consist of initiators (WWNs) only. You do not need to specify volumes. The initiator is shown in both host security groups SPA and SPB.
- Host security groups can consist of volumes (LUNs) only. You do not need to specify initiators.
- When you select an initiator for the host security group, the initiator has to be registered with the CLARiiON storage system.

- You can have more than one initiator in a security group if you have the proper multipathing software installed on the particular host where the initiator is located.

Host Security Groups on LSI Storage Systems

Keep in mind the following rules for host security groups on LSI storage systems.

- When you create a host security group using the management server, the host security group appears as a host, with its volumes and ports displayed underneath the tree in SANtricity.
- When you create a host security group using the management server, you can specify the controller, but not the port, for the host security group.
- In SANtricity, an initiator is equivalent to a host port.
- You can have multiple volumes and initiators in the host security group.
- If you create a volume on a host that does not have multipathing, make sure the volume is on the preferred path. Use SANtricity to make sure the volume is on the preferred path.
- The ID for a host security group changes when you rename the host security group.
- When creating a host security group, if you provide a volume, but not an initiator, the host security group is created, but the volume is part of the Default Group, not the group you created. This occurs because until a volume is assigned to a host security group with an initiator, it is visible to all initiators and is assigned to the Default Group. To assign a volume to a specific host security group, you must add an initiator to that host security group.
- If you create an host security group for a Unix host on an LSI array, you must use SANtricity to change the host mode of the host security group to Linux in order to be compatible with the Unix host. If you do not, the Unix host will show a large number of LUN numbers available. You must use SANtricity to change the host mode because it is not possible to select the host mode when creating an host security group on an LSI array. The host security group will be created with the Windows NT host mode by default if it is not manually changed.

Host Security Groups on EMC Symmetrix Storage Systems

Keep in mind the following rules for host security groups on EMC Symmetrix Storage Systems.

- If LUN security is not turned on for an FA port, all volumes assigned to the FC port are visible to hosts that are on the SAN and have been zoned by the SAN. All volumes assigned to the FC port appear in the mapped category.
- When you create a host security group on a Symmetrix storage system, you are creating LUN mapping and masking in one step. In the native tools for Symmetrix storage systems, you will not see the host security group you created by using the management server. Instead you will see a volume bound to a port and a masked LUN bound to a host in the native tools.
- Host security group is associated with individual ports.
- Host security groups only allow one initiator for host security masking.
- To create a host security group, you must specify a port, initiator, and a volume.

- Every port has a LUN host security group, even if no LUNs are defined for that port. To bind a LUN to a port, edit the host security group and add the desired LUN to a port.
- You can also add LUNs to a Mask host security group. To add initiators, you must create the host security group.

Host Security Groups on HDS Storage Systems

Keep in mind the following rules for host security groups on HDS storage systems.

- FC port contains only volumes but no initiators (HBA WWN) assignment, the management server displays these volumes as unmapped since no external host can see these volumes yet.
- You can have zero to multiple initiators in a host security group.
- A host security group can be on only one port on the array. You can have host security groups with the same name, as long as they are on different ports.
- Host security groups appear in the native tool for HDS storage systems. In the logical view, the host security groups are listed by LDEV; in the physical view, they are listed by port.
- In the native tool for HDS storage systems, host security groups are referred to as a host security domain.

Host Security Groups on HP MSA 1000/1500 Storage Systems

Keep in mind the following rules for host security groups on HP MSA 1000/1500 storage systems.

Provisioning is not supported for MSA P2000 G2 or G3 storage systems.

- When you create a host security group on the MSA 1000/1500, the host mode is set to the value defined by the smi.ProvisioningHpMsa.hostConnectionProfile property (the default is Windows). To change the value, follow the instructions in Customizing Properties on page 344, and then restart the service.
- You can have one initiator per host security group.
- You can have one or more volumes in a host security group.
- A host security group spans all ports on the array.

Host Security Groups on HP EVA Storage Systems

Keep in mind the following rules for host security groups on HP EVA storage systems.

- You can have multiple initiators per host security group.
- You can have zero to multiple volumes in a host security group.
- A host security group spans all ports on the array.

Host Security Groups on IBM Storage Systems

Keep in mind the following rules for host security groups on IBM storage systems.

- You can name a host security group on IBM storage systems. The host security group will be given the name of the initiator you select for the host security group.
- To assign a host mode to a host security group, you must modify a property, as described in Setting the Host Mode for IBM Storage Systems on page 723.
- The management server can read the names of host security groups created by the native tool.
- There can be only one initiator per host security group.
- A volume can be assigned to more than one initiator.
- You can select any number of ports from *one* to *all* when creating the host security group.
- You can create the host security group with or without LUNs.
- You can add mapped and unmapped volumes to a host security group, but they should have the same host mode.
- The default is that all ports are in the host security group. If no ports are selected, the default is used.

Host Security Groups on Xiotech Storage Systems

Keep in mind the following rules for host security groups on Xiotech storage systems.

- The initiators used in a host security group must be registered with the array.
- Any given host security group is assigned to only one storage port.
- Host security groups must have one initiator and at least one volume.

Managing Host Security Groups

This section contains the following topics:

Accessing Information about Host Security Groups

Host Security Groups define which initiators (HBA ports) have access to specified storage volumes. They are associated with a Fiber Channel port and contain a list of HBA port initiators and the volumes they can detect.

Keep in mind the following:

- Each type of storage system treats host security groups differently. For more information, see General Provisioning Issues on page 723.

- Not all HDS storage systems support host security groups. Refer to the documentation accompanying the HDS storage system.

If you use another product to make provisioning changes, you must perform Get Details (**Discovery** > **Details**) for the management server to be made aware of these changes.

You can access information about host security groups from the Provisioning Manager wizard or from System Manager.

To access host security groups from the Provisioning Manager wizard:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system for which you want to access host security group information.
4. Click **Step 3 Host Security Group**.

To access host security groups from System Manager:

1. Double-click a storage system in System Manager.
2. Click **Host Security Groups** in the Navigation tab.

This page lists the following information about host security groups:

The following two features are only available when accessed from the Provisioning Manager wizard.

- View all host security groups – Click the **All** category in the tree. All the host security groups appear in the right pane.
- View only host security groups assigned to a certain port – Click a port in the tree. The host security group assigned to the port appears in the right pane.

The following information is available on this page:

- Port (If applicable) – Displays the port associated with the host security group.

  You can filter the list of ports so that the information for only one port is displayed. To filter the list of ports: Click **+Filter**, select the port that you would like to view, and click **Apply**. Only the information for the selected port is displayed in the table below.
- Name – Displays the name of the host security group.
- Initiators – Displays one of the following:
  - The caption for the discovered port if the port has been discovered by the management server (for example Columbia:Adapter0 Port 0). A tool tip over the caption gives the full WWN.
  - The WWN if the port has not been discovered.
- Volumes – Displays the volumes in the host security group.

- Host Mode – Displays the port settings for your operational environment. The settings for the host mode vary by the model of the storage system. With some hardware, you must select a special host mode on the port for the storage system to enable certain servers and HBAs to "see" the LUNs on the port. Refer to your documentation for the storage system.
- Host Mode 2 (If applicable) – Displays optional settings on the port that describe how the host accesses the port. Multiple options exist. For more information, see you documentation for storage system.

If the table contains more than 10 entries, the following buttons for navigating through the table are enabled:

-  – Move to the first page.
-  – Move back one page.
-  – Move forward one page.
-  – Move to the last page.

You can also create, edit and delete host security groups from this page. The following topics provide more information:

- Creating Host Security Groups below
- Editing a Host Security Group on page 720
- Deleting a Host Security Group on page 722
- Setting the Host Mode for IBM Storage Systems on page 723

Creating Host Security Groups

Host Security Groups define which initiators (HBA ports) have access to specified storage volumes. They are associated with a Fiber Channel port and contain a list of HBA port initiators and the volumes they can see.

Keep in mind the following:

- You cannot use the management server to add a host to a host group. For example, you cannot have nested host groups.
- Each type of storage system treats host security groups differently. For more information, see General Provisioning Issues on page 723.

To create a host security group, you must assign a name to the host security group, assign a host mode (if applicable to your system), and add initiators and volumes to the group. Follow these steps:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **Storage Systems** tab.

3. Click the **Provision** button for the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.
5. Click the **New Host Security Group** button in the upper-right corner of the screen.

Task A – Add Details for the Host

1. Enter a unique name for the host security group in the Name box.

   You cannot name a host security group on IBM storage systems. The host security group will be given the name of the initiator you select for the host security group.

   Keep in mind the following:

   - The name must contain the following number of characters. If you enter no characters, you are given the option of using a default name.
     - CLARiiON storage systems – 1 to 64 characters
     - LSI storage systems – 1 to 30 characters
     - All other storage systems – 1 to 50 characters
   - The first and last letter cannot be spaces.
   - You cannot have the following characters in the name:

     `<, >, ., :, ,, |, *, ?, \, \\, \t, \n, \b`

2. (If applicable) Select the port you want associated with the host security group. This port should contain your LUNs.

   Each type of storage system handles ports for host security groups differently. For more information, see Rules for Creating Host Security Groups on page 712.

3. *HDS storage system*: Click the **Options** button to the right of the Host Mode box, select a host mode resembling the port settings for your environment, and then click **OK**. If your host mode is not listed, enter it in the Host Mode box.
4. *HDS storage systems*: If your storage system supports a second host mode, enter the second host mode in the Second Host Mode box.

   A second host mode is an optional setting on the port that describes how the host accesses the port (not applicable to all storage systems).

5. *IBM storage systems*: You cannot assign the host mode for an IBM storage system in the user interface. You must modify an internal property to set the host mode. For more information, see Setting the Host Mode for IBM Storage Systems on page 723.
6. Click **Next**.

Task B – Add Initiators to the Host Security Group

1. Add an initiator to the host security group by clicking the **Add** button in the upper-right corner.
2. Do one of the following:

- Enter the WWN of the port you want to add to the host security group.

  Or

- Select the initiator you want to add to the host security group.

  Notice that when the mouse hovers over the port, you are shown additional information, such as the name and WWN of the port on the switch that the host uses.

3. Click the **Add** button at the bottom of the window.
4. When you are finished with adding initiators, click **Close**.
5. To remove an initiator from the host security group, click the **Delete** ( ) button. To remove multiple HBA initiators from the list, select the HBA ports you want to remove and then click **Remove Selected**.
6. Click **Next**.

Task C – Add Volumes to the Host Security Group

1. Add a volume to the host security group by clicking the **Add** button in the upper-right corner of the window.
2. Select a volume. Then do one of the following:

   - If you want the unit number to be selected automatically by the server, leave the **Auto-Select** option selected.

     Or

   - If you want to choose a unit number, deselect the **Auto-Select** option and enter the unit number in the Unit Number box at the top of the window.

     For LSI storage systems LUN numbers cannot be duplicated, and that the management server can use an existing LUN number if the access mode for the created LUN is “No Access.”

3. Click **Add** at the bottom of the window.

   The volume is added to host security group.

4. When you are done adding volumes, click **Close**.
5. To remove a volume from the host security group, click the **Delete** ( ) button. To remove multiple volumes from the list, select the volumes you want to remove and then click **Remove Selected**.
6. Click **Finish**.

Editing a Host Security Group

To edit a host security group:

1. Click **Provisioning Manager** ().
2. In the right pane, click the **Storage Systems** tab.

3. Click the **Provision** button for the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.
5. Click the **Edit** () button for the host security group you want to edit.

Task A – Edit Details for the Host

1. Modify the name of the host security group.

   The ID for a host security group changes when you rename the host security group.
2. Change the port you want associated with the host security group. This port should contain your LUNs.

   Each type of storage system handles ports for host security groups differently. For more information, see Rules for Creating Host Security Groups on page 712.
3. *HDS storage system*: Click the **Options** button to the right of the Host Mode box. Select a host mode resembling the port settings for your environment and then click **OK**. If your host mode is not listed, enter it in the Host Mode box.
4. *HDS storage systems*: If your storage system supports a second host mode, enter the second host mode in the Second Host Mode box.

   A second host mode is an optional settings on the port that describes how the host accesses the port (not applicable to all storage systems).
5. *IBM storage systems*: You cannot assign the host mode for an IBM storage system in the user interface. You must modify an internal property to set the host mode. For more information, see Setting the Host Mode for IBM Storage Systems on page 723.
6. Click **Next**.

Task B – Change the Initiators Assigned to the Host Security Group

1. Change the initiators assigned to the host security group, by doing one or more of the following:

   Add an initiator as follows:

   a. Click **Add**, and then do one of the following:
      - Enter the WWN of the port you want added to the host security group.

        Or
      - Select the initiator you want added to the host security group.

   b. Click **Add** and then click **Close** to exit the window.

   c. Delete an initiator by clicking the **Delete** () button for the initiator you want removed.

   d. Delete multiple initiators. Select the initiators you want removed and click **Remove**

Selected.

Caution: Removing an HBA can cause hosts that are using it to lose access to their storage. This can result in the loss of data.

2. Click **Next**.

Task C – Change the Volumes Assigned to the Host Security Group

You cannot delete the default host security group.

1. Change the volumes assigned to the host security group:

   Add a volume as follows:

   a. Click **Add**, select a volume, and then do one of the following:

      - To have the unit number be selected automatically by the server, leave the **Auto-Select** option selected.

        Or

      - To choose a unit number, deselect the **Auto-Select** option and enter the unit number in the Unit Number box at the top of the window.

   b. Click **Add**.

   c. Click **Close** to exit the window.

      For LSI storage systems, LUN numbers cannot be duplicated. The management server can use an existing LUN number if the access mode for the created LUN is "No Access."

      Caution: Removing an HBA can cause hosts that are using it to lose access to their storage. This could result in the loss of data.

      Remove a volume: Click the **Delete** ( ) button for the volume you want removed from the host security group.

      Remove multiple volumes: Select the volumes and then click **Remove Selected**.

2. Click **Finish**.

Deleting a Host Security Group

You cannot delete the default host security group.

To delete a host security group:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click the **Storage Systems** tab.
3. Click the **Provision** button for the storage system in which you want to access information about host security groups.
4. Click **Step 3 Host Security Group**.

5. Click the **Delete** ( ) button for the host security group you want to delete.

   The host security group is removed.

6. To remove multiple host security groups, select the host security groups you want to remove. Then click **Delete Selected**.

Setting the Host Mode for IBM Storage Systems

The host mode for an IBM storage system cannot be set in the user interface, but it can be set by assigning the host mode in a property. The property you modify depends on the type of IBM storage system you are using and the version of the IBM CIM Agent.

To set the host mode:

1. Click **Show Default Properties** at the bottom of the page.
2. Do one of the following based on your configuration:
   - **DS Storage Systems using version 5.2.0 of the CIM Agent (or earlier)** – Uncomment smi.ProvisioningIbmEss.hostConnectionProfile=AIX by removing the hash symbol (#) in front of it.
   - **DS Storage Systems using version 5.2.1 or later of the CIM Agent** – Uncomment smi.ProvisioningIbmDs.clientType=Microsoft Windows by removing the hash symbol (#) in front of it.
3. Assign a host mode to the property you uncommented in the previous step; for example:

   ```
   smi.ProvisioningIbmEss.hostConnectionProfile=AIX
   ```

   In this instance, AIX is the host mode.

   A list of supported host modes is provided above the property.

4. Click **Save**.

General Provisioning Issues

This section contains the following topics:

Provisioning Can Make a Device Inaccessible

Provisioning can break a connection between an array and a host. When you rezone a device, make sure no users or applications are using the device. For example, assume a port for a disk drive is a member of zone set A, which is active. If you make zone set A inactive and this port is not a member of the new active zone set, the disk drive will become unavailable.

Provisioning Does Not Make an Operating System Aware of a Device

When a port in a zone set becomes active, you must take steps to make it available to the operating system. For example, assume a port for a disk drive is a member of zone set A. If you make the zone set active, the host will not automatically recognize the disk drive. You will need to configure the operating system so that it becomes aware of the device. Refer to the documentation that accompanies the operating system for the host.

Provisioning Issues by Vendor

This section contains the following topics:

Issues Specific to EMC Symmetrix Storage Systems

About Provisioning on EMC Symmetrix Storage Systems

EMC ships its Symmetrix storage system with volumes already created. When the software first discovers an EMC Symmetrix storage system, it assumes the devices on the Symmetrix storage system are volumes.

This software refers to the term “device” to define a piece of hardware in the storage network. EMC uses the term “device” to refer to a volume on one of its storage systems. In this section, the term “device” is used in the context of EMC storage systems.

When you use the software's “create a volume” feature, you are assigning the already created volume. If necessary the software will create a meta device, which is a device that is a concatenation of several devices.

The software does not delete the volumes created by EMC. When you use the software's “delete a volume” feature, the software marks the volume as hidden in its repository. These “hidden volumes” are stored in the “Free Device” list. If you use a device in the “Free Device” list when you create a volume, that device is removed from the “Free Device” list.

Process Has an Exclusive Lock Message

The Symmetrix storage system can become locked for many reasons. For example, the storage system becomes locked when it performs LUN mapping, LUN masking or Get Details. The Symmetrix storage system might also remain locked after a provisioning operation has failed. You will receive a message resembling the one shown below if a process has already locked the EMC Symmetrix storage system and you attempt a process that requires a lock on the storage system.

```
SYMAPI routine SymDevMaskSessionStart failed with error code 188: The
operation failed because another process has an exclusive lock on the
local Symmetrix
```

After the management server has detected the lock on the Symmetrix storage system, it tries to access the storage system for 15 minutes and logs the errors.

If you receive the error message, determine whether someone is performing an operation that requires a lock, such as LUN mapping, LUN masking, or Get Details. This applies even if one of the processes is being used by a third-party product, such as for LUN masking. If this is the case, wait until the process is complete. Remove the lock manually only if you are certain that no other processes are happening on the storage system. To learn how to remove the lock, refer to the documentation for the Symmetrix storage system.

If a provisioning failure causes the Symmetrix storage system to remain locked, you are alerted to this situation in Event Manager and on the Properties tab. You might receive a message similar to the following:

```
Unable to end device masking session. Symmetrix '000001835005700' may
be locked.
```

Some EMC Volumes, Their LUNs and LUN Maskings Are Hidden

EMC volumes, their LUNs, and LUN maskings for the volumes that play a special role (such as holding device masking information, attached BCVs, and gate-keeper devices) are hidden. On the host side, the software shows LUN maskings for these volumes. So you might see a LUN masking, but not its volume, residing on the EMC storage system.

Issues Specific to HDS Storage Systems

About Provisioning on HDS Storage Systems

The management server does not allow LUSE of different RAID levels.

HDS ships some of its storage systems with volumes already created. When the software first discovers an HDS storage system, it detects the volumes created by HDS. When you use the software's create a volume feature, you are assigning already-created volumes.

The software does not delete volumes created by HDS. When you use the software's delete a volume feature, the software marks the volume as hidden in its repository. These hidden volumes are stored in the Free LDEVs list.

If you use an LDEV from the Free LDEVs list when you create a volume, the LDEV is removed from the list because it is now assigned.

HDS cannot create a LUSE volume (made up of multiple LDEVs) without mapping it to a target port (that is, without creating a LUN). In the software, creating a volume and creating a LUN are two different operations. Therefore, the software keeps the volumes, made from multiple LDEVs in the Grouped LDEVs list in the repository. Once these volumes are mapped to the target port and a LUN is created for them, they are removed from the repository and a real LUSE volume is created on the HDS box.

For example, assume you have several 2-GB Free LDEVs and you want to create a 4-GB volume. Since the requested 4-GB volume is larger than one LDEV, two of the Free LDEVs will be used for the 4-GB volume.

Unable to Provision When HDS CruiseControl Is Enabled

When HDS CruiseControl is enabled on an HDS array, such as an HDS Lightning 9980V, you are unable to do provisioning. You might also receive the following error message:

```
An error was encountered during this operation. Some of the operation
may have been applied to the storage subsystem. A refresh of the
storage subsystem is recommended. "The LDEV is HIHSM reserved; cannot
be used in a LUSE".
```

To use the provisioning tool, disable HDS CruiseControl. Refer to the HDS CruiseControl product documentation for more information.

Increasing the Wait Time for the Management Server

By default, the management server waits 20 minutes for a response from HiCommand Device Manager after sending a provisioning command. If the management server does not receive a response from HiCommand Device Manager after 20 minutes, the management server assumes the provisioning command did not go through. It then tries to contact HiCommand Device Manager again while the previous command is still active. After two retries, the management server stops attempting to contact HiCommand Device Manager.

For example, assume you initiated a provisioning command to add the host security group. The management server waits 20 minutes for a response. If it receives no response during that time, it sends another “Add” command and waits 20 minutes for a response. If no response is received, it sends another command. After the second retry, the management server stops attempting to contact HiCommand Device Manager. Multiple host security groups with the same configuration of LUNs and WWNs were created when the management server attempted to contact HiCommand Device Manager.

The management server has a similar behavior when a delete command is initiated. It sends the delete command to HiCommand Device Manager, waits 20 minutes, and sends another delete command if it receives no response. The second command tries to delete the same host security group, but the target host security group is deleted when the first command is completed, and the second command returns an error.

If you need more time for HiCommand Device Manager to respond, you can increase the amount of time the management server waits, by modifying the cimom.provider.hds.ProvisioningTimeout property.

To change the provisioning timeout property:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the **Disk Space** tree on the management server.
2. Click **Show Default Properties** at the bottom of the page.

3. Copy the following text:

   ```
   cimom.provider.hds.ProvisioningTimeout
   ```

4. Repeat step 1 to return to the Advanced page.
5. Paste the copied text into the Custom Properties box. How you paste the text depends on your Web browser. If you are using Microsoft Explorer or Netscape Navigator, right-click the box and select **Paste**.
6. Make any necessary changes in the Custom Properties box. Change the value of the cimom.provider.hds.ProvisioningTimeout property. (The value is in milliseconds.) For example, assume you want the management server to wait an hour. You would assign 3600000 to cimom.provider.hds.ProvisioningTimeout, since 3600000 milliseconds is one hour; for example:

   ```
   cimom.provider.hds.ProvisioningTimeout=3600000
   ```

7. Click **Save**.
8. The product notifies you if a restart of the AppStorManager service is required.

Initiator Ports Cannot Be Used for Provisioning

Ports designated as an initiator on a storage system belonging to the HDS Freedom Storage Lightning 9900 Series or Freedom Storage Lightning 9900V Series cannot be used for provisioning. If you select one of these ports, you receive a message saying that provisioning failed because the HiCommand Database was not refreshed. The management server does not support provisioning for ports designated as initiators on these storage systems.

Mapping Issue on HDS 9900V Storage Systems

On HDS 9900V storage systems, if a host is already mapped to a volume and you try to map the same host to a volume in another host storage domain, corresponding to the same port, it will fail. However for HDS 9900, the host can be mapped to a volume in another host storage domain corresponding to the same port.

A Default LUN Number Is Used Instead of a User-Specified One

When you create a LUN, the user-specified LUN is ignored. The LUN is created with the next available default number.

Issues Specific to HP Storage Systems

Cannot Always Delete Selected Volume on MSA 1000/1500 Arrays

Volumes on MSA 1000/1500 Arrays must be deleted in the reverse order of their creation. For example, if you have six volumes, and you want to delete the second one you created, you must delete the volumes one at a time, starting with the volume created sixth and continuing with the fifth, fourth, third, and then the second. Attempting to delete other volumes will return a generic error code 4.

If the volume deleted is the last volume of the storage array (also known as the storage pool), the storage array is also deleted.

Selective Storage Presentation Must Be Enabled on MSA 1000/1500 Arrays

Selective Storage Presentation (SSP) for the MSA 1000/1500 arrays must be enabled for provisioning to work.

Issues Specific to LSI Storage Systems

Creating and Deleting Storage Pools

For LSI, a storage pool is the same as a volume group. Create at least one storage pool before provisioning a volume. When you delete a storage pool on an LSI storage system, all the volumes for the volume group are deleted, including the placeholder volume.

Creating and Deleting Storage Volumes

Keep in mind the following when you create a volume on an LSI storage system:

- Make sure you select a volume group that can accommodate the requested size for the new volume.
- You can create volumes from existing free extent areas within a volume group.
- The volume capabilities, their cache read ahead multiplier, and segment size are shown in the following table.

No volume-to-LUN masking is done by default.

Volume Usage

Volume Capability	Cache Read Ahead Multiplier	Segment Size
File and Default	1	64 KB
Database	0	64 KB
Multimedia	8	128 KB
Custom	0 to 65,535	8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB

The management server creates a placeholder volume when a storage pool is created. This placeholder volume is not viewable in the management server, but it might be viewable in other storage tools; do not delete it.

17 Using Backup Manager

This section contains the following topics:

About Backup Manager

The Backup Manager feature enables you to:

- Monitor the overall status of the backup process
- Visualize the backup configuration and recoverability of a file, directory, volume, or server
- View the status of the physical infrastructure supporting the backup process, backup application, backup server, network, tape library, and media
- Obtain information on reasons for backup failures and advisory information for configuring new backup schedules

Only a single click on an element in the topology is required to obtain more information about the element.

Backup Manager monitors the backup applications running on discovered hosts. To determine which backup applications are supported, see the support matrix, which is accessible from the Documentation Center.

Caution: The management server is able to detect the presence of the following after you obtain backup details:

- **Backup Application** – A backup application (such as NetBackup, HP Data Protector, EMC NetWorker, or IBM Tivoli Storage Manager) serving as the master in a backup hierarchy. A backup application is responsible for managing other media managers.

- **Backup Manager Host** – A managed host that is running the backup application. The IP address of the Backup Manager host must be specified in Step 1 of discovery before the backup application can be discovered.
- **Media Manager Application** – A backup application functioning as a server to control the media in a backup hierarchy. A media manager application can be responsible for managing different types of hardware, such as tapes and drives.
- **Media Manager Host** – A host that has the backup application running as the media manager application. A media manager application and its host can be discovered through the Backup Manager host (similar to the way that hosts are detected through a switch). If a media manager host is discovered through a backup manager host, the media manager host is considered to be "unmanaged," meaning that the management server has discovered it, but cannot obtain additional information about the element. If the IP address of the media manager host is specified in Step 1 of discovery, the media manager host will be considered "managed."
- **Media** – Any device that is used to store backup data, such as tape
- **Media Pool** – A logical grouping of the backup media
- **Sessions** – Scheduled and executed backup sessions
- **Tape Library** – A device hosting a collection of tapes
- **Robot** – An automated device inside the tape library; responsible for manipulating the tapes
- **Backup Client** – A host that is being backed up by a backup application. A backup client can be managed as a non-generic element if its IP address appears in the discovery list. Otherwise, backup clients that are identified through the backup application are considered generic elements.

Master servers without the tape library connectivity are not shown in the drive monitoring page.

When the tape library name provided by Data Protector and the name provided by the SMI-S agent of the tape library do not match, a tape library is displayed as two different tape libraries, once as a generic tape library and again as a discovered tape library.

Requirements for Backup Manager

The CIM extension supports only one backup solution on a host. For this reason, Backup Manager does not support both EMC NetWorker or NetBackup Master Server, HP Data Protector Cell Manager, and IBM Tivoli Storage Manager on the same host. If EMC NetWorker or NetBackup Master Server, Data Protector Cell Manager, and IBM Tivoli Storage Manager are installed on the same host, by default only Data Protector Cell Manager is discovered. EMC NetWorker or NetBackup Master Server and IBM Tivoli Storage Manager are ignored by the CIM extension.

Before you can use the Backup Manager feature, you must follow these steps:

1. Install a CIM extension on the host running the backup application. For information about installing CIM extensions, see the *Installation Guide*.

2. Discover the host running the backup application, selecting the **Include backup details** option (see Step 1 – Discovering Your Hosts and Backup Manager Hosts on page 179). HP recommends that you also select the **Include infrastructure detail** option, so you can also monitor and manage the host itself.

   Make sure you have at least 500 MB available if you are using the host as a Backup Manager host in a large environment; for example, 300 clients, 25,000 sessions, and 500,000 images.

3. Schedule backup collection for your Backup Manager hosts as described in Scheduling Backup Collection for Backup Managers on page 357.

Determining if You Have Enough Media to Run a Backup

If you are performing many, or large, backups, you should make sure you have enough media available for the backup. Backup Manager provides several methods for determining this:

- **Media tab** – Provides information about the discovered media, including its usage count.
- **Media Pool tab** – Provides information about media in the pools, such as whether it is Available, Allocated, Frozen, or Suspended.
- **Media Summary reports** – Provides information about all discovered media over a defined time period.

To use these methods:

1. Click **Backup Manager** ().
2. Click the **Topology** tab on the right side of the window.
3. Expand the Backup Applications node in the left pane, and then select a Backup Manager host.
4. If necessary, expand the lower pane so you can view the Media and Media Pool tabs.
5. Click the **Media** tab in the lower pane. The following information is displayed:
   - **Media ID** – The identification number of the media
   - **Media Pool** – The name of the media pool
   - **Usage Count** – How often the media has been used
   - **Barcode** – The barcode associated with the media
   - **Retention** – How long the media is retained
   - **State** – Whether the media is Full, Available, or Active

   To learn more about a specific medium, select its row. Additional information is displayed in the lower-right pane.
6. To learn more about the media pool that contains the media, click the **Media Pool** tab. The following information is displayed:

- **Media Pool** – The name of the media pool
- **Backup Manager** – The name of the Backup Manager host to which the media pool belongs
- **Media Manager** – The name of the media manager to which the media pool belongs
- **Library** – The name of the library to which the media pool belongs
- **Available** – The media is available for backup.
- **Allocated** – The media is currently either actively being used or has a valid backup on it.
- **Frozen** – The media will never become available for backup again, but it is still available for restores.
- **Suspended** – The media will not be used again until all backups written to it expire. It is still available for installations however.

You can also set up a policy that will notify you when the number of available media for a storage pool is running low. For example, you could create a policy that sends you an e-mail when the number of available media for a storage pool is less than two. See Creating a Performance, Backup or Utilization Policy on page 582.

Viewing Running Sessions

For information about configuring session monitoring, see Session Monitoring on page 358.

To view sessions that are running:

1. Click **Backup Manager** ().
2. Click the **Topology** tab on the right side of the window.
3. Right-click a Backup Manager application, and then select **Show Running Sessions**. The Running Sessions page is displayed.
4. If desired, select filter criteria, and then click **Apply Filter**. The table is updated to display only the sessions that meet the filter criteria you entered. The following information about each session is displayed:
   - **Job ID** – The identifier assigned to the session
   - **Backup Manager** – The name of the backup manager
   - **Media Manager** – The name of the media manager
   - **Clients** – The names of the clients
   - **Backup Policies** – The names of the backup policies
   - **Schedules** – The names of the schedules
   - **Session Status** – The session status: Success or Failure
   - **Session State** – The session state: Done, Queued, or Active

- **Media Used** – The type of media used for the backup

  Media Used information is not available for IBM Tivoli Storage Manager.
- **Start Time** – The starting time and date of the session
- **End Time** – The end time and date of the session
- **Size** – The size of the session in kilobytes (KB)
- **Files** – The number of files

After Caché is removed and the CIM extensions are installed, the first session monitoring collection collects all of the sessions for the past week. One hour after a fresh installation of Caché, these sessions are removed from the Running Sessions page.

The Running Sessions page is automatically refreshed every two minutes. To manually refresh the page, click **Getting Latest Sessions**.

For IBM Tivoli Storage Manager,the sessions information is collected using the activity log.

Determining Last Successful Backup

Backup Manager provides several tools to help you determine if the last scheduled backup was successful. The quickest way to do this is from the Summary tab in Backup Manager. The following topics provide more information:

- Viewing the Summary Backup Charts below
- Viewing Backup Results for a Backup Manager Host below
- Viewing Backup Results for a Client on next page
- Viewing Backup Information for a Client on page 735

Viewing the Summary Backup Charts

To access summary information about last night's backup from the backup charts:

1. Click **Backup Manager** ().
2. Click the **Summary** tab on the right side of the window.

By default, the Backup SLA Performance chart is displayed in the upper-left pane on the Summary page. This chart includes the overall results of the backups made in the past 14 days. It tells you if the overall backup was successful, partially successful, or failed.

For more information, see About the Summary Backup Charts on page 745 and Modifying Summary Backup Charts on page 751.

Viewing Backup Results for a Backup Manager Host

To quickly view the results of the backup sessions performed by a backup application:

1. Click **Backup Manager**().
2. Click the tab on the right side of the window.
3. In the left pane, expand the Backup Applications node, and then select the Backup Manager host.
4. If necessary, expand the lower pane so you can see the tabs.
5. Click the Sessions tab. Backup Manager displays the following information for each session the backup application performed:
   - **Session ID** – The identifier assigned to the session
   - **Client** – The name of the client
   - **Backup Policy** – The name of the backup policy
   - **Schedule** – The name of the schedule
   - **Status** – The status of the backup: Success, Partial, or Failure
   - **Start Time** – The starting time and date of the backup
   - **Duration** – The amount of time in seconds it took for the backup to be displayed
   - **Size** – The size of the backup in gigabytes (GB)
   - **Files** – The number of files backed up
6. To learn more about a session, select the session's row in the table and see the Session Detail, Policy Detail, and Schedule Detail tabs in the lower-right pane.

Viewing Backup Results for a Client

To view the results of the last backup for a client:

1. Click **Backup Manager** ().
2. Click the **Topology** tab on the right side of the window.
3. In the left pane, expand the Clients node, and then select the client.
4. In the topology pane, check the color of the check mark above the icon for the client. The icon colors have the following meaning:
   - **Green** – The last backup on the client was successful.
   - **Yellow** – The backup on the client was partially successful.
   - **Red** – The backup on the client failed.

To view detailed backup information for a client, see Viewing Backup Information for a Client on the facing page.

Viewing Backup Information for a Client

Backup Manager tracks backup information for a client for the past 30 days. The Backup tab gives an at-a-glance view of the backup coverage for a selected element. For example, you can select the period for the coverage and review the policy, schedule, and results for the backups executed for that period.

To obtain detailed information about the backup sessions for a client:

1. Click **Backup Manager** ().
2. Click the **Topology** tab on the right side of the window.
3. Select the client from the topology pane or from the tree in the left pane.
4. If necessary, expand the lower pane so you can view the Properties and Backup tabs.
5. Click the **Backup** tab to view the following information for each backup policy:
   - **Backup Policy** – The name of the backup policy
   - **Schedule** – The name of the schedule
   - **Date** – The date, end time, and status
6. To learn more about a policy or schedule, select a cell in the table. Additional information is provided on the Policy Detail and Schedule Detail tabs in the lower-right pane.

About the User Interface for Backup Manager

Backup Manager has an easy-to-use interface that provides the following options:

- **Toolbar** – Provides buttons and menus to help you modify the topology and charts in Backup Manager (see About the Toolbars in Backup Manager on page 737)
- **Summary and Topology tabs**:
  - **Summary** tab – Displays summary charts for backup elements (see About the Summary Backup Charts on page 745 and Modifying Summary Backup Charts on page 751).
  - **Topology** tab – Displays the topology of the backup elements

    The lower pane on the Topology tab becomes visible when you select a discovered element. The tabs are determined by the backup element selected (see About the Tabs in the Topology Lower Pane on page 747).

    Access to Navigation, Events, Collectors, Policies, and Chargeback – When you click an element on the Topology tab, the following links in the lower-right corner are enabled if that feature is supported for the selected element:

    - **Navigation** – Displays the navigation information for an element, such as which storage systems are connected to the element (see About the Navigation Tab on page 489).
    - **Events** – Displays the events for the element (see About the Events Tab on page 511).

- **Collectors** – Provides links to data collectors and reports about the element (see About the Collectors Tab on page 514).
- **Policies** – Enables you to set up policies for the element (see About the Policies Tab on page 515).
- **Chargeback** – Enables you to provide chargeback information about the element (see Asset Attributes of an Element on page 512).

You can also select an element and use the right-click menu options to obtain additional information (see Right-Click Menu Options on the Topology Tab on page 741).

About the Topology Icons in Backup Manager

The following table describes the icons that appear in the topology in Backup Manager.

Topology Icons in Backup Manager

Icon	Description
	When a client computers is shown with a green check mark, the backup on the computer was successful.
	When a client computers is shown with a yellow question mark, the backup on the computer was partial.
	When a client computers is shown with a red X, the backup on the computer failed.
	Host
	Master backup server (media). This image is an example of a master backup server for NetBackup.
	Backup server (media), such as NetBackup
	Tape Library

Icon	Description
	Tape Drive

About the Toolbars in Backup Manager

Backup Manager has two toolbars:

- The main toolbar that appears at the top of its screen. See Main Toolbar for Backup Manager below.
- The toolbar for charts that appears in the middle of its screen. See Toolbar for Charts on page 739.

Main Toolbar for Backup Manager

The following table provides a brief description of the buttons and menus on the main toolbar in Backup Manager. This toolbar is available at the top of the Backup Manager screen.

Toolbar in Backup Manager

Button	Description
	Saves the current topology or summary page, so that when you return to Backup Manager, the saved layout or summary is restored. This option can be especially useful if you want to keep the new location of elements you have moved. When you click the button, you are asked if you want the layout to apply to all users. • **Yes** – All users who log into the management server can view the topology or summary you created. Only users with system configuration capability can save their layout for all other users • **No** – No other users can view the topology or summary you saved.
	Prints the topology or the summary.
	Enables you to view information from past backups. For example, if you want to view a backup from last March, just click the Calendar icon, then select the date.
	Enables you to modify the summary charts. Enabled when the Summary tab is active.
	Magnifies the view. Enabled when the Topology tab is active.

Button	Description
	Decreases the magnification. Enabled when the Topology tab is active.
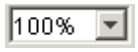	Enables you to set the magnification to a percentage of the default magnification. Enabled when the Topology tab is active.
	Opens a smaller pane, which provides a global view of the topology. This enables you to position the main view to a certain section of the topology (see Using the Global View on page 458). Enabled when the Topology tab is active.
	Fits the topology to the window, so you can view the entire topology. Enabled when the Topology tab is active.
	Enables you to move an element in the topology (see Arranging Elements in the Topology on page 456). Enabled when the Topology tab is active.
	Enables you to move the entire topology at once. Click the **Pan** () button, click any place in the topology, and then drag the mouse to a new location. Enabled when the Topology tab is active.
	Opens a new window, containing the topology, which enables you to view different domains of the topology at one time (see About the New Window Option on page 507). Enabled when the Topology tab is active.
	Enables you to change the topology layout. Enabled when the Topology tab is active.
	Restores the topology layout to the last saved version. Enabled when the Topology tab is active.

Button	Description
Find	Enables you to find an element by name or by Worldwide Name (WWN) in the topology. Enabled when the Topology tab is active. To find an element, enter the name or part of the name in the Find box, and then click the **Find Next** ( ) button. The management server highlights the elements that match in the topology and in the tree. If the management server has found multiple elements matching your search criteria, click the **Find Next** ( ) button to find the next element that matches your search criteria. To view the previous element that matches the search criteria, click the **Find Previous** ( ) button. The Find Previous ( ) button is disabled when only one element meets your search criteria.
	Exports the topology to an XML file that can be viewed in Microsoft Visio (see Exporting the Topology to Visio on page 741). Enabled when the Topology tab is active.
	Displays links between shared libraries. Additional connections between media servers and tape libraries, and media servers and disk drives are displayed. If the additional links between shared libraries are currently displayed, clicking the Show MultiPath button a second time will hide the links.
	Change Observer button – Monitors changes in the database status on the server. When changes are detected, the button turns amber. Click on the amber button and a pop-up window displays the elements that have changed on the server. When no changes are detected, the button is grayed out.
	Reloads the Change Observer button to display the latest changes to elements on the server.

Toolbar for Charts

To view a chart on the Topology tab, click an element in the topology. The following toolbar options are available.

Toolbar for Charts

Option	Description
	Converts the data in the chart to a list in a separate browser.

Option	Description
	Click to print a chart. See Printing Summary Charts on page 752 for more information.
	Click to switch the X and Y axes in a chart.
	To change the chart displayed, select another chart from the Title menu (see About the Summary Backup Charts on page 745).
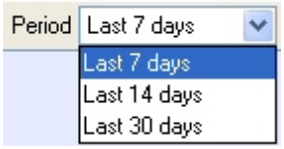	To change the period displayed in the chart, select a period from the Period menu.
	In Service Level Agreement charts, Backup Manager provides a green line that serves as a baseline. Use the Average Service Level menu to change the location of this baseline. The default baseline value is 95% of usage.

Changing the Topology Settings

The **Display Layout Settings Dialog** () button enables you to modify the following properties of the topology in Backup Manager:

- **Direction** – Horizontal or Vertical. The direction of the topology is Horizontal by default, with multiple elements of the same type displayed in a row. If you select Vertical, multiple elements of the same type are displayed in a column.
- **Alignment** – Left, Right, or Center. The default alignment of the topology is Center. You can change the alignment of the topology to be left- or right-justified. For example, if you select the Left, the backup clients are aligned along the left side of the topology window.
- **Horizontal Spacing** – The number of spaces in pixels between elements in a row
- **Vertical Spacing** – The number of spaces in pixels between elements in a column

To restore the layout to the default settings, click **Defaults**.

To change the layout settings:

1. Click **Backup Manager** ().
2. Click the **Display Layout Settings Dialog** button.

3. Select one of the following directions:
   - Horizontal
   - Vertical
4. Select one of the following alignments:
   - Left
   - Right
   - Center
5. To change the horizontal spacing, enter a new number in the Horizontal Spacing box.
6. To change the vertical spacing, enter a new number in the Vertical Spacing box.
7. Click **OK**.

You might need to use the scroll buttons to see the rearranged topology.

Exporting the Topology to Visio

To export the topology to an XML file that can be viewed in Microsoft Visio:

1. Click **Backup Manager** ().
2. Click **Export to Visio**.
3. Name the file, and then select the directory in which you want the file to be saved.
4. Click **Save**. The XML file is saved to the directory you selected.

For information about configuring Visio and viewing the exported file, see Viewing the Topology in Microsoft Visio on page 461.

Right-Click Menu Options on the Topology Tab

When you right-click an element on the Topology tab or in the Backup Applications tree, a list of options is displayed. The options displayed in the menu depend on the type of element you selected.

The following table describes the menu options displayed when an element is right-clicked on the Topology tab or in the Backup Applications tree.

Right-Click Menu Options on the Topology Tab

Right-Click Menu Option	Description
Go To Navigation Details	Directs you to the Navigation page. If the element is labeled Discovered, you are shown the Properties page. An element is labeled unmanaged when the management server has become aware of it, but cannot obtain additional information about it (see About the Navigation Tab on page 489).
Show Events	Displays events for the selected element (see About the Events Tab on page 511).
Show Policies	Displays the backup policies for the selected element (see About the Policies Tab on page 515).
Show Collectors	Displays the report collectors for the selected element (see About the Collectors Tab on page 514).
Show Chargeback	Displays the chargebacks for the selected element (see About the Monitoring Tab on page 515).
Update Element Data	The management server gathers new and changed details from the element and then redraws the topology with the updated information. **Important**: The Update Element Data functionality does not detect element components that have been removed, such as ports and LUNs. For example, assume you removed several LUNs from an array. If you right-click the storage system, and then select **Update Element Data**, the deleted LUNs still appear in the user interface. You must perform Get Details for the deleted LUNs to be removed from the user interface. For more information, see Right-Click Menu Options on the Topology Tab on previous page.

Right-Click Menu Option	Description
External Tools	Provides several ways to access an element: • **Telnet** – Enables you to access a host or a switch through the telnet utility. The Telnet feature is only accessible to Web browsers on Microsoft Windows operating systems. • **Browse** – Enables you to access the main Web page for a host or a switch. • **Set Up External Tools** – Enables you to add URLs for accessing the management tools for the storage system. In some instances, the management tool for the storage system is directly accessible from this menu (for example, HiCommand for HDS storage systems and Command View for HP XP storage systems). See Using External Tools on page 488
Add Virtual Application	Enables you to add an unsupported application so you can monitor it. For example, you might want to add a virtual application so you can monitor software that was created uniquely for your company. See Creating a Virtual Application on page 510.
Custom Commands	Enables you to run a custom command on an element; for example to start an executable or a script. See Setting Up Custom Commands on page 479.
Go to System Manager	Provides a topology that enables you to view how the devices in your network are connected. See About System Manager on page 413.
Show MultiPath/Remove MultiPath	Displays links between shared libraries. Additional links between media servers and tape libraries, and media servers and disk drives are displayed. If the additional links between shared libraries are currently displayed, the menu option becomes Remove MultiPath. Selecting Remove MultiPath hides the links between shared libraries.
Delete Element	This option enables you to delete a backup element in the topology.

The charts in Backup Manager provide a wealth of information about your backups. You can obtain detailed information about a data point displayed in a chart by right-clicking the data point. For example, assume you are looking at a Service Level Agreement (SLA) chart on the Summary tab and you want to obtain more information about a backup performed yesterday. You could right-click the bar **yesterday's date**, and then select **Show Details** to display the Sessions tab showing the additional details of that data point, such as the backup status of each client, in addition to the start and end time of the backup on that client.

Right-Click Menu Options on the Summary Tab

Right-Click Menu Option	Description
Summary Page Settings	The settings for the charts provided on the Summary tab are displayed. See About the Summary Backup Charts on the facing page and Modifying Summary Backup Charts on page 751.
Go To Topology	A graphical representation of the path of an element is displayed. This also includes multipathing. See Viewing Element Topology on page 501.
Show Details	The additional information about the data point you right-clicked is displayed. See the following table.

The following table explains what is displayed when you click **Show Details** on the Summary tab's right-click menu option.

When additional information is not available for a data point, Show Details is disabled.

Additional Information from Charts on the Summary Tab

When you right-click a bar and select Show Details in the following chart...	The Sessions tab displays the following information:
Service Level Agreement	Clients that were backed up, whether successfully or not. The failures are displayed first. To obtain details about a session, select the session in the Session tab and then expand the View the Details pane on the far right. The status you right-clicked is highlighted in the Sessions tab.
Backup Volume	Clients in the backup, sorted by size.
Windows Utilization Chart	The time span represented by the bar you clicked is highlighted in the Sessions tab. For example, assume a bar in the Windows utilization tab shows a duration of seven hours. To determine which sessions were running during that time, right-click the bar, and then select **Show Details**.
Largest Sessions	Clients in the backup, sorted by the size of the session.
Longest Sessions	Clients in the backup, sorted by duration of the session.

You can also obtain additional information from some of the charts that are displayed on the bottom pane of the Topology tab.

Show Details for Tabs on the Lower Pane of the Topology Tab

Right-click...	Select Show Details to view...
Any data point on the Charts tab	The Sessions tab for the data point you right-clicked.
Any element on the Servers tab	The Sessions tab showing the sessions for the element you right-clicked.
Any element on the Resources tab	The Media tab for the element you right-clicked.
Any element on the Media Pools tab	The Media tab for the elements contained in the media pool you right-clicked.

About the Summary Backup Charts

Backup Manager displays six summary backup charts on the Summary tab by default and offers many other charts as well. To learn how to display the various charts and/or modify which charts display by default, see Modifying Summary Backup Charts on page 751.

Backup Manager Summary Charts

View	Description
Servers	Displays the servers Backup Manager monitors with the following information for each server: • Volume • Sessions • Failed • Partial • Successful

View	Description
Resources	Displays the resources Backup Manager monitors, showing the following information for each server: • Media Pools • Available media • Allocated • Frozen • Suspended In the Available Media, Allocated, Frozen, and Suspended columns, the first number shows the number of available online media, and the second number shows the number of available offline media. This information is available only to the backup manager host.
Service Level Agreements (SLAs)	Displays the performance of backup SLAs, showing the percentage of the following types of sessions for each SLA: • Successful sessions • Partial sessions • Failed sessions
Backup Volume	Displays the backup volume of all backup applications in gigabytes (GB). This chart can also display the backup volume of a backup manager host.
Window Utilization	Displays the number of hours it takes for all backup sessions on a server to run. Keep in mind this time might seem extended if you have overlapping sessions. For example, bsession1 starts at 11 p.m. on Monday. While it is running, bsession2 starts. At 2 a.m. on Tuesday, bsession1 stops, but bsession2 continues to run until 9 a.m. on Tuesday. The Windows Utilization report shows the backup sessions running for 10 hours – from the beginning of the bsession1 to the end of bsession2.
Backup Manager Hosts with Most Executed Sessions	Displays the five backup manager hosts with the most executed sessions. Only successful sessions are counted.

View	Description
Most Unsuccessful Backup Manager Hosts	Displays the five backup manager hosts with the most unsuccessful sessions. Unsuccessful sessions include failed and partially completed sessions.
Servers with Most Available Media	Displays the five backup manager hosts with the largest number of media in the Available state and for each of the displayed hosts, the chart shows media that is Allocated, Frozen, or Suspended.
Servers with Fewest Available Media	Displays the five Backup Manager hosts with the lowest number of media in the Available state. The chart displays the following states: • Allocated • Frozen • Suspended • Available
Five Largest Sessions	Displays the five largest sessions in gigabytes (GB).
Five Longest Sessions	Displays the five longest sessions in seconds.

About the Tabs in the Topology Lower Pane

The lower pane on the Topology tab is displayed when you select a discovered backup element. The following tabs are displayed according to the element type you selected.

Tabs in the Lower Pane of Backup Manager Topology

Tab	Element Type	Description
Properties	All elements	Provides property information for an element, including information about whether the element supports backup.
Backup	Clients	Provides information about the last time the client was backed up. See Viewing Charts for a Backup Manager Host on page 752 for more information.

Tab	Element Type	Description
Charts	• Backup Manager Hosts • Media Managers • Tape Libraries	Shows a chart for the selected element.
Servers	• Backup Manager Hosts • Media Managers	Displays the servers Backup Manager monitors with the following information for each server: • **Volume** – The size of the volume backed up, in kilobytes • **Sessions** – The number of backup sessions that have run in the specified time • **Failed** – The number of failed sessions during the specified time • **Partial** – The number of partial sessions during the specified time • **Successful** – The number of successful sessions within the specified time
Resources	• Backup Manager Hosts • Media Managers • Tape Libraries	Displays the resources Backup Manager monitors with the following for each server: • **Media Pools** – The number of media pools that the backup manager host can access • **Available Media** • **Allocated** • **Frozen** • **Suspended** In the Available Media, Allocated, Frozen, and Suspended columns, the first number shows the number of available online media; the second number shows the number of available offline media. Note that this information is available only to the backup manager host.

Tab	Element Type	Description
Sessions	• Backup Manager Hosts • Media Managers	Displays the following information for the sessions assigned to a backup server: • **Session ID** – The identifier for the session • **Client** – The DNS name of the computer on which the session is taking place • **Backup Policy** – The name of the backup policy • **Schedule** – The name of the schedule for the session • **Status** – The status of the session • **Start Time** – The time the session started • **End Time** – The time when the session ended • **Duration** – The amount of time in seconds the session ran • **Size** – The size of the session • **Files** – The number of files that were backed up
Media	• Backup Manager Hosts • Media Managers • Tape Libraries	Displays the following information for the media attached to a backup server or tape library: • **Media ID** – The identifier for the media • **Media Pool** – The media pool to which the media belongs • **Usage Count** – How often the media is used • **Retention** – How long the media is retained • **State** – Whether the media is Full, Available, or Active

Tab	Element Type	Description
Media Pool	• Backup Manager Hosts • Media Managers • Tape Libraries	Displays the following information for the media pools containing the selected element: • **Media Pool** – The media pool to which the media belongs • **Backup Manager –** The name of the backup manager host in the media pool • **Library** – The name of the library in the media pool • **Available Media** – The number of available media • **Allocated –** The number of allocated media • **Frozen** – The number of frozen media. • **Suspended** – The number of suspended media
Drive Utilization	• Tape Libraries • Drives	Displays the following information for the drives in a tape library: • **Library** – The name of the tape library that contains the drive • **Drive** – The name of the drive • **Media ID** – The media identifier • **Status** – The running status of the drive For information about configuring drive monitoring, see Drive Monitoring on page 358.

Sorting Information in the Lower Pane

You can sort the information displayed on the tabs in the lower pane by clicking the heading of a column. You can also sort more than one column at a time. The sorting feature for multiple columns can be extremely useful. For example, if you have several clients with failed backups, you would take the following steps to sort the table to show the clients according to their status:

1. Click the **Status** heading in the session column to sort the sessions according to status.
2. Press the **CTRL** key, and then click the **Client** heading.

The clients are sorted first according to their status and second according to their client name. You can now easily view all clients with failed sessions in alphabetical order.

You can sort as many columns as you want on a tab. The arrow indicates ascending or descending sort order. The arrow decreases in size for each additional column that is sorted. The largest arrow corresponds to the column that is sorted first, the second largest arrow corresponds to the second sort, and so on.

Modifying Summary Backup Charts

You can modify Backup Manager to display charts other than the default. To learn more about the data and options available in the Summary Backup charts, see About the Summary Backup Charts on page 745.

The Summary Settings page shown below displays a grid that lists the charts available on the Summary tab in Backup Manager. To change a chart, select a grid on the this page and change the settings as described in the steps in this section.

The top four grids on the Summary Settings page correspond to the top four charts on the Summary tab, and the lower four grids on the Summary Settings page correspond to the lower Summary tab. There are only two tables on the bottom half of the Summary tab because each of the tables spans two grids.

Summary Settings Page for Backup Manager Charts

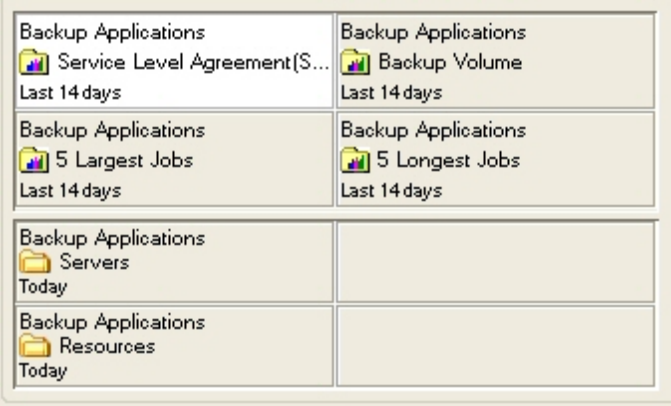

You can also modify the Summary tab instantly by clicking one of the buttons displayed in the Summary Settings page. The following table describes these buttons.

Buttons on the Summary Settings Page

Button	Description
Clear All	Clears the settings for all the charts.
Clear	Clears the settings for the selected chart.
Revert	Returns the Summary Settings page to the previous setting.
Defaults	Returns the Summary Settings page to the default setting.

To modify a chart displayed on the Summary tab in Backup Manager:

1. Click **Backup Manager** ().
2. Click the  icon.
3. To change the title for the summary page, enter a new title in the Title box.
4. Select the grid in which you want the chart to appear on the screen.
5. Select one of the following options from the Backup Element menu:
   - **Backup Applications** – The chart includes the results from all backup applications.

     Or
   - **A specific backup element** – The chart includes the results from only the backup application you selected.
6. Select the type of chart you want from the View menu. To learn more about the available charts, see About the Summary Backup Charts on page 745.
7. Select a period for coverage from the Period menu.
8. Click **OK**.

The Summary page is updated with your changes and the chart accessed from the selected grid will reflect your changes.

Viewing Charts for a Backup Manager Host

To quickly view charts for a backup manager host:

1. Click **Backup Manager**().
2. Select the backup manager host on the Topology tab.
3. Click the **Charts** tab in the lower pane.
4. Select a chart from the Title menu.
5. Select a period of coverage for the chart.

To learn more about the charts in Backup Manager, seeAbout the Summary Backup Charts on page 745.

See Toolbar for Charts on page 739 for information about the toolbar for charts.

Printing Summary Charts

To print a summary chart:

1. Access a backup summary chart by clicking an element on the Topology tab.
2. Scroll to the bottom of the screen.

3. Click the **Print** ( ) button.
4. Click **Landscape** at the top of the new window if you want the picture to be printed in landscape format. To revert to portrait format, click **Portrait**.
5. To change the magnification of the image on the printed page, select the desired percentage.
6. Click **Print** when you are ready to print the chart.

Changing Collection Times for Media and Session Collectors

For CIM extension hosts with Data Protector installed, the CIM extension queries the data from DataProtector for session and media details and puts the data into the cache database as soon as the CIM extension is started.

The background collection occurs every 15 minutes for sessions and every 24 hours for media.

Changing the Frequency of Collection Times

To change the collection time for the session and media collectors:

1. Check if the `dp-connector.properties` file exists in the following directory on the host with the CIM extension:
   - **Windows hosts**: `C:\Program Files\APPQcime\CimExtensions\tools`
   - **Linux, Solaris, and HPUX hosts**: `/opt/APPQcime/tools`
2. If the `dp-connector.properties` file does not exist, create it in the directory specified in the previous step.
3. Add the following properties to the `dp-connector.properties` file if they do not exist in the file:

   ```
   cachedb.sessions.waitBgCollection=300000

   cachedb.media.waitBgCollection=3600000
   ```

   The time shown in the example is in milliseconds.
4. Set the values.
   - `cachedb.sessions.waitBgCollection` Corresponds to the amount of time before the next collection for a session.
   - `cachedb.media.waitBgCollection` Corresponds to the amount of time before the next collection for a media.
5. Save the file.

Stopping Background Collections when a CIM Extension Starts

You can stop the background collection whenever a CIM Extension starts. In this case, the session and media collection will be started only for each get details that includes backup details.

To stop background collections when a CIM extension starts:

1. Check if the `dp-connector.properties` file exists in the following directory on the host with the CIM extension:
   - **Windows hosts**: `C:\Program Files\APPQcime\CimExtensions\tools`
   - **Linux, Solaris, and HPUX hosts**: `/opt/APPQcime/tools`
2. If the `dp-connector.properties` file does not exist, create it in the directory specified in the previous step.
3. Add the following properties to the `dp-connector.properties` file if they do not exist in the file:

   ```
   BUMODEL.startScheduler=no

   BUMODEL.forceStart=yes
   ```

4. Save the file.

Changing the Number of Days Session is Stored in the Cache Database

Seven days of session data is stored in the cached database by default.

To change the number of days session data is stored in the cached database:

1. Check if the `dp-connector.properties` file exists in the following directory on the host with the CIM extension:
   - **Windows hosts**: `C:\Program Files\APPQcime\CimExtensions\tools`
   - **Linux, Solaris, and HPUX hosts**: `/opt/APPQcime/tools`
2. If the `dp-connector.properties` file does not exist, create it in the directory specified in the previous step.
3. Add the following property to the `dp-connector.properties` file if it did not previously exist in the file:

   ```
   sessions.timeframe=2592000000
   ```

   The time in the example is in milliseconds. The value assigned to the sessions.timeframe property is in milliseconds (30 days).
4. Save the file.

Known Issues

- Display Limitation with a large number of Backup Sessions/Savesets. The management server user interface is limited by the number of savesets and/or sessions it may display. In testing when there is in excess of 100,000 items, the user interface will not display the information.
- Restore Sessions not Reported for EMC Networker. Backup Manager does not currently support the display of restore sessions for EMC Networker.
- Backup Details are not Gathered from Multi-homed Hosts if CIM Extension IP is Restricted. The CIM extensions have the capability of being restricted to listening on and using a particular IP address on multi-homed systems (hosts with multiple IP addresses). Although the host information will be returned, backup information will not be sent back to the management server while the CIM extension is in this type of configuration.
- Status not Displayed in Topology for Backup Master Server . The Session Status is not displayed in the Backup Manager Topology for the Backup Master Server Hosts.
- NetBackup Does Not Report Media Information for Disk-based Backups. NetBackup does not report media information for disk-based backups. Consequently, the Media Used field is blank in Backup Manager.
- Disk-based Backup Devices are Shown as Physical Tape Libraries in HP Data Protector. Disk-based backup devices are shown as physical tape libraries in HP Data Protector. In Veritas NetBackup, only robotic tape library devices are shown; other backup devices (simple tape drives or disks) are not shown.

18 Path Provisioning

Your license determines whether Path Provisioning is available. See the List of Features to determine if you have access to Path Provisioning. The List of Features is accessible from the Documentation Center (**Help** > **Documentation Center**).

This section contains the following topics:

About Path Provisioning

Path Provisioning enables you to schedule provisioning tasks to take place when the network traffic is light. For example, you could use Path Provisioning to schedule multiple provisioning tasks to take place at 1 a.m., and when you come in later that morning, review the status of the provisioning tasks.

You can also use Path Provisioning to identify host/storage dependencies so you can make informed decisions when deciding where new volumes, zones, or LUN security is needed.

To view the latest provisioning information in Path Provisioning, click **Refresh**. The Refresh button updates the following:

- The Path Provisioning screen with the following changes:
  - Changes made in Provisioning Manager. For example, if you use the wizards in Provisioning Manager to create a host security group, when you access Path Provisioning, your changes are not shown until you refresh.
  - Changes from executed jobs. After a job is executed in Path Provisioning, the Path Provisioning screen is not updated until you click Refresh or exit and re-enter Path Provisioning.

  Adding a volume to an existing host security group on an EMC Symmetrix or DMX array only performs the masking operation, and does not map the port on the array. The native tools for the array show that the masking to the host initiator took place, but the volume is still not mapped to a storage port.
- Other parts of the product:
  - Application Viewer
  - Capacity Manager
  - Performance Manager
  - Backup Manager
  - Provisioning Manager
  - System Manager

Keep in mind the following:

- Path Provisioning runs within a Java applet. If you receive “out of memory” messages when you view Path Provisioning, you might need to increase the amount of memory assigned to the Java plug-in on the client computer.
- (Windows only clients) If you do not have the Java plug-in already installed and you are running Firefox, you must use Microsoft Internet Explorer to install the plug-in. After you install the plug-in, you can use Firefox to run the plug-in.
- If you select a direct-attached host, the storage ports appear in the LUN pane with the  icon next to them, indicating they are unreachable. You can still select these storage ports and schedule the job. These storage ports are shown as unreachable in the user interface because the user interface uses switches to display the association between a host and a storage system. Therefore, if the management server cannot detect a switch, as with a DAS connection, the user interface assumes the storage ports are unreachable.
- A port designated as an Initiator on a storage system belonging to the HDS Freedom Storage Lightning 9900 Series or Freedom Storage Lightning 9900V Series cannot be used for provisioning. If you select one of these ports, you receive a message saying that provisioning failed because the HiCommand Database was not refreshed.
- The HBAs displayed might not have a connection to the selected storage system. This is done to provide flexibility. For example, you can select a disconnected HBA for a job you want to take place when the HBA is connected to the storage system.

- You can view zones, zone sets and zone aliases on a Cisco switch; however, you cannot use the management server to create, modify, or remove them from a Cisco switch.

McDATA Switches

- Zone aliases are not supported for McDATA switches.
- Only manageable fabrics will be displayed in the Path Provisioning. If no provisioning can be done on the fabric (any vendor), it is not displayed in Path Provisioning. It will be displayed in Provisioning Manager.
- Path Provisioning looks for the names of the active zone set and of the active zones and all of their saved counterparts in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

Brocade Switches

While configuring Path Provisioning for a Brocade switch, you might occasionally see the following message in the Message Console tab:

```
Failed Path Provisioning Job, <job information>, cause: CIM_ERR_
FAILED

Fabric Session is Locked: SessionState = 2

Try again later
```

Click **Refresh** to see the final status. This message generally means that the Brocade fabric is busy and the management server was unable to complete the operation. Try the operation again at a later time.

How Path Provisioning Works

When you select a storage system in Path Provisioning, the management server displays all the information relevant for provisioning with respect to the storage system selected. This information includes:

- Mapped, unmapped, and unmasked (mapped to one or more storage system ports but not associated to any host initiator port) volumes of the storage system.
- Already masked LUNs of the storage system and the front end ports of the storage system (possible candidates for LUN mapping).
- Hosts that are reachable (hosts belonging to the same fabric as the storage system) along with their HBA and host initiator ports. Single multipathing functionality is supported.
- Existing zones to which the ports of the storage system belong are displayed. If the Host and the Storage System belong to multiple fabrics, zones of all those fabrics are displayed.

The information is displayed in different panes: Storage System, Host, Volume, LUN and Zone. You can select relevant information from each of pane, such as which volumes are to be used for the LUN mapping task. Each selection can create a provisioning task to be performed. You can create a job that contains one of the following:

- A single task of one type, such as, Map LUN).
- A set of tasks of the same type, such as mapping multiple LUNs.
- Multiple LUN masking tasks or a job consisting of several tasks of different types (such as, Map LUN, zone required ports, mask LUN) that go together as a combination).
- Multiple jobs each consisting of multiple tasks of different types.

The jobs can be executed immediately or scheduled to start at a later time. A job with all its required details (such as, parameters to invoke Provider Service methods) is stored in a job queue. At the scheduled time, the scheduler retrieves the job from the job queue and performs the tasks of the job, using the details stored in the job.

The status of each job is displayed in the State column of the Provision Job section located in the lower pane of the screen. A job can have following status:

- **Created** – The job was created, but will not be executed. The job cannot be viewed by others and is deleted when the Web browser is closed. See Scheduling Provisioning Jobs on page 793 for information about changing the state of the job from “created” to “scheduled.”
- **Scheduled** – The job was tasked to execute at a specified time and date. Jobs are assigned a scheduled state after you select the job and click the **Execute Job** () button.
- **Started** – The job started. You cannot delete a job once it starts.
- **Failed** – The job failed.
- **Ended** – The job finished.

How to Use Path Provisioning

All the provisioning tasks are centralized on one screen. Provisioning enables you to either select a default template from the System Action combo-box or create a system action template consisting of the provisioning tasks best suited to your system, as described in Creating a System Action Template on page 789.

The default provisioning templates are the following:

- **Volume Creation + LUN Security + Zone Operation –** Lets you create a meta volume, map a volume to a Fibre Channel port and host HBAs (HSG), and then create a zone (see LUN Security and Zone Operation on page 781).
- **Meta Volume Creation** – Lets you create a meta volume (see Creating a Meta Volume on page 768).
- **LUN Security** – Lets you map a meta volume to Fibre Channel port and host HBAs (HSG) (see LUN Security on page 769).
- **Zone Operation** – Lets you perform a zone operation (see Zone Operation on page 774).
- **Volume Creation + LUN Security** – Lets you create a meta volume and then map the meta volume to a Fibre Channel port (see Volume Creation and LUN Security on page 778).

- **LUN Security and Zone Operation** – Lets you create a host security group with the host HBA WWN along with zoning operations (see LUN Security and Zone Operation on page 781).
- **Volume Assignment** – Lets you assign a volume to existing host security groups (see Volume Assignment on page 785).

Complete the steps in the various panes. A step that is not required for the action is disabled after all data is loaded.

Schedule the task as described in Scheduling Provisioning Jobs on page 793.

You can control which templates users can access (see Assigning a Template to a Role on page 801).

Overview for Path Provisioning

Step	Description	Where to Find Additional Information
1	Select a system action.	• Volume Creation, LUN Security, and Zone Operation on next page • Creating a Meta Volume on page 768 • LUN Security on page 769 • Zone Operation on page 774 • Volume Creation and LUN Security on page 778 • LUN Security and Zone Operation on page 781 • Volume Assignment on page 785
2	Complete the steps.	
3	*(Optional)* Schedule a provisioning job that you want to take place at a later time.	Scheduling Provisioning Jobs on page 793
4	Execute the job. The job does not run until you click the **Execute Job** ( ) button is clicked. Once this button is clicked, the job is saved in the management server database.	Executing Provisioning Jobs on page 794

About the User Interface

The Path Provisioning feature provides the following on its toolbar:

Feature Toolbar

Button	Description
	Change Observer button – Monitors changes in the database status on the server. When changes are detected, the button turns gold. Click the gold button and a pop-up window displays the elements that have changed on the server. When no changes are detected, the button is grayed out.
	Reloads the Change Observer button to display the latest changes to elements on the server.
Configure Templates	Lets you create and configure a System Action Template (see Creating a System Action Template on page 789 and Modifying a System Action Template on page 790).
Assign Templates	Enables you to assign a template to a role (see Assigning a Template to a Role on page 801).

Default System Action Templates

This section describes the contents of the default system action templates and the options included in each. It also provides instructions for using each of the options.

This section contains the following topics:

Volume Creation, LUN Security, and Zone Operation

You can use Path Provisioning to create a meta volume, map it to a Fibre Channel port and host HBA, and designate it to appear in a pre-existing zone or create your own.

If options are still selected from a previous job, clear the options you do not want in your next job. For example, assume you created a volume and now want to create a new volume on the same host. Clear the Volume pane by clicking the button.

To clear all the steps, except for the Step 1 (storage systems) action, select another option from the System Action combo-box.

(HDS storage systems only) Before you can create a volume, you must delete some unmapped LDEVs using the standard provisioning tool.

To delete LDEVs:

1. Click **Provisioning** ( ).
2. Click the Storage Systems tab and the **Provision** button for the storage system.
3. Click **Step 2 Volume**.
4. Select the desired number of LDEVs for the LUSE volume and click **Delete Selected Volumes**.
5. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

To access Path Provisioning:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select **Volume Creation + LUN Security + Zone Operation** from the System Action combo-box.

Step 1 – Select Storage System

To select a storage system:

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select the storage system on which you want to create the metavolume.

   The S column heading in the Storage Systems pane means that only a single selection is allowed.

   (*HDS only*) Select the storage system from which you want to create the LUSE volume.
3. Click the **Step 1** button below the pane.

The selected storage system's name is displayed below the Storage System pane. The Host pane is populated. Notice in the figure below that some hosts have a red X over their icon. This means the host is not accessible.

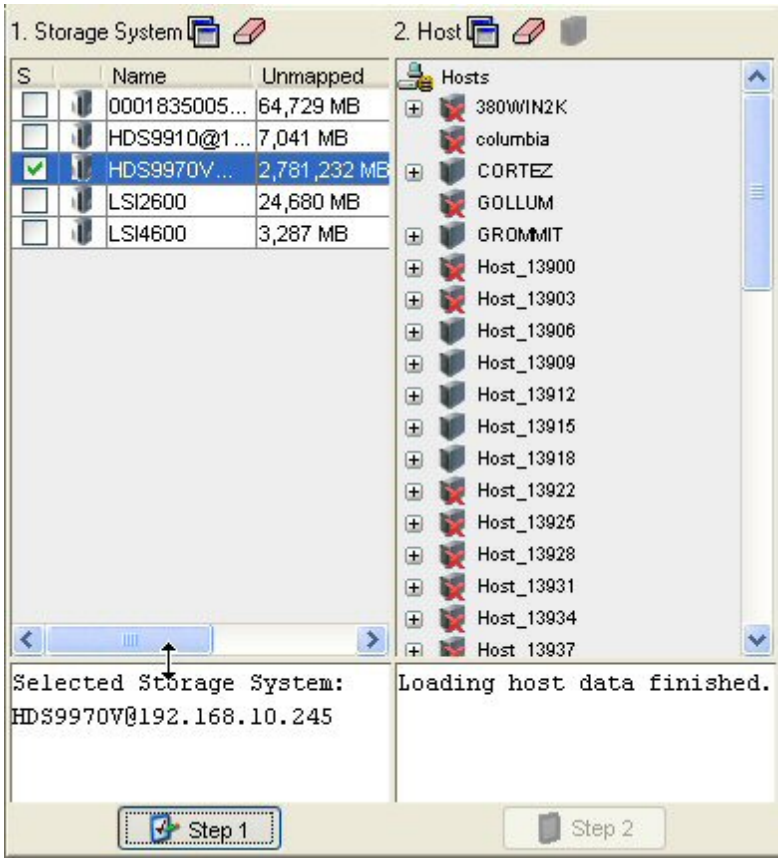

Step 2 – Select a Host

If you select host and storage ports that belong to an existing zone alias and you have the **Display Zone Aliases** option selected in Customize Zone Options dialog box, the existing zone alias is automatically selected and highlighted.

If you select hosts and storage ports that are not contained in an existing zone alias, the new hosts and storage ports are added into the existing zone alias after the provisioning job finishes successfully.

To select a host:

1. Wait for all data to be loaded. All data is loaded when you see the following messages:

   ```
   Host data loaded.

   Loading volume data finished.

   Loading HSG data finished.

   Loading zone data finished.
   ```

The Step 2 button is disabled until data is loaded.

2. Select a host that is accessible.

   To create a provisioning job for multipathing, click the button above the Host pane, and then select the option for multipathing. When you select this option, you must select the same host for both paths (see Host Customize Dialog Box on page 797).

   To automatically create a zone if a zone does not meet a preset criteria:

   a. Click the button above the Zone pane.

   b. Select the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.

   c. Set the criteria. See Customize Zone Options Dialog Box on page 799 for more information about setting the criteria.

   d. Select a host and click **Step 2**.

   If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane.

   If the management server does not find a zone that meets the criteria, it asks for a zone name. See Naming Conventions for Brocade and McDATA Switches on page 795 for more information about the naming requirements for a zone. After you enter a zone name, the new zone is displayed in the Zone pane, but it will not actually be created until the job runs.

   To configure zoning manually, click the button above the Zone pane and deselect the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, select a zone or create one manually.

3. Click **Step 2**.

   Information about the selected port, such as its WWN, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.

Step 3 – Select a Volume

To select a volume:

1. In the Volume pane select mapped and unmapped volumes. To select multiple volumes in Windows, hold down the **CTRL** key as you select the volumes.

   - **Mapped** – There are two types of mapped volumes:
     - **Masked** – The volume is exposed to the storage port and to the host.
     - **Unmasked** – The volume is exposed to the storage port, but not to the host.
   - **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
   - **Free Extents** – Lists available free extents that can be used to create a meta volume or

LUSE. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVS from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVS. When you select free extents, they must be of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the "unmapped" category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server (see Deleting a Storage Volume on page 710).

Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.

To narrow the type of volumes displayed in the Volumes pane by using the Customize Volume Options dialog box (), located above the Volumes pane (see Host Customize Dialog Box on page 797).

The Customize Volume Options dialog box is not available for the HP EVA.

If the LUN was already selected and Step 4 is clicked, skip this step or click the  button.

2. Click **Step 3**.
3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See Providing a LUN Number on page 789 for information about numbering LUNs.

Step 4 – Select a Host Security Group

1. Select a host security group. See Creating a Host Security Group on page 791 for information on how to create a host security group. See General Provisioning Issues on page 723 for information on how your storage system handles host security groups.
2. Click **Step 4**.

Step 5 – Select a Zone

If the zone was already selected and Step 5 is clicked, skip this step or click the  button to clear the selection.

If a zone is not selected or created yet, most likely the option **Automatically Configure Zoning** is not selected in the Customize Zone Options dialog box (). The management server assumes you want to select a pre-existing zone or create one manually.

- **To reuse a zone** – Select a zone in the Zone pane and then click **Step 5**, and expand a fabric node to view its zones.

- **To create a zone** – Select a fabric in the zone pane, click the button, and then enter a name for the zone (see Creating a Zone on page 777).

Zone Icon	Description
	Above Zone pane – Used to create zones. In the Zone pane – Represents a zone.
	Zone Alias
	Port
	The fabric cannot be reached.

The zoning filters and operations work only on zones that are part of the active zone set. *McDATA and Brocade switches only*: Path Provisioning looks for the names of the active zone set and of the active zones and verifies that all of their saved counterparts are matched in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

Creating a Zone

To create a zone:

1. Select a fabric in the Zone pane.

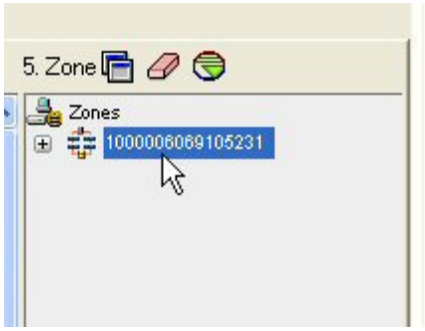

2. Click the button located above the Zone pane.
3. Enter a zone name in the dialog box. For naming conventions, see Naming Conventions for Brocade and McDATA Switches on page 795.
4. Click **OK**. The new zone is added to the Zone pane.
5. Click **Create Job**. The job is added to the bottom pane.
6. Take one of the following actions:
   - To execute the job now, click the **Execute Job** ( ) button.

     Or

- To execute the job at a later time, schedule the job as described in Scheduling Provisioning Jobs on page 793.

Creating a Meta Volume

If options are still selected from a previous job, clear the options you do not want in your next job. For example, assume you previously created a volume. To create a new volume on the same host, clear the Volume pane by clicking the button.

To clear the action taken in all Steps except Step 1, select another option from the System Action combo-box.

HDS only: Before you can create a volume, you must delete some unmapped LDEVs using the standard provisioning tool.

To delete LDEVs:

1. Click **Provisioning** ( ).
2. Click the **Storage Systems** tab and the **Provision** button for the storage system.
3. Click **Step 2 Volume**.
4. Select the desired number of LDEVs for the LUSE volume and click **Delete Selected Volumes**.
5. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

To use Path Provisioning to create meta volumes:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select **Meta Volume Creation** from the System Action combo-box.
4. Wait for the management server to load the storage systems into the Storage System panel.
5. Select the storage system on which you want to create the metavolume.

   HDS only: Select the storage system from which you want to create the LUSE volume.

   The S column heading in the Storage Systems pane means that only a single selection is allowed.
6. Click the **Step 1** button below the pane. The selected storage system's name is displayed below the Storage System pane.
7. In the Volume pane, select multiple free extents under the Free Extents node in the Volume pane.

   The meta volume containing the selected free extents is created when the job runs.

- **Mapped** – There are two types of mapped volumes:
  - **Masked** – The volume is exposed to the storage port and to the host.
  - **Unmasked** – The volume is exposed to the storage port, but not to the host.
- **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
- **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVS from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVS. When you select free extents, they must of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

  When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the "unmapped" category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server (see Deleting a Storage Volume on page 710).

  Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.

You can narrow the type of volumes displayed in the Volumes pane by using the Customize Volume Options dialog box (), located above the Volumes pane (see Host Customize Dialog Box on page 797).

The Customize Volume Options dialog box is not available for the HP EVA.

If the LUN was already selected and Step 4 is clicked, skip this step or click the  button.

8. Click **Step 3**.
9. Click **Create Job**. The job is added to the bottom pane.
10. Take one of the following actions:
    - To execute the job now, select the job and click the **Execute Job** () button.

      Or

    - To execute the job at a later time, schedule the job as described in Scheduling Provisioning Jobs on page 793.

LUN Security

If options are still selected from a previous job, clear the options you do not want in your next job. For example, assume you just mapped a volume to a port. Now you want to map the a different volume on the same host to another port. Clear the Volume and LUN panes. To clear a pane, click the  button.

To clear all the steps, except for the Step 1 (Storage Systems), select another option from the System Action combo-box.

To use Path Provisioning to designate subsystem LUN security:

1. Click **Provisioning Manager** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select **LUN Security** from the System Action combo-box.

Step 1 – Select Storage System

To select a storage system:

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select the storage system on which you want to create the metavolume.

   The S column heading in the Storage Systems pane means that only a single selection is allowed.
3. Click the **Step 1** button below the pane.

The selected storage system's name is displayed below the Storage System pane. The Host pane is populated. In the following example, some hosts have a red X over their icon. This means the host is not accessible.

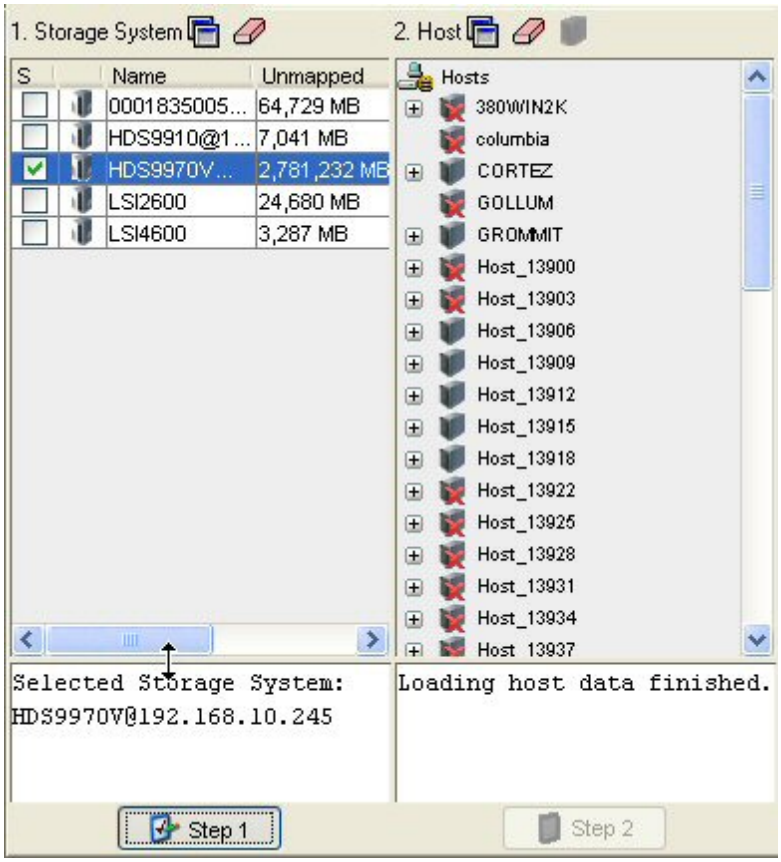

Step 2 – Select a Host

If you select host and storage ports that belong to an existing zone alias and you have the **Display Zone Aliases** option selected in Customize Zone Options dialog box, the existing zone alias is automatically selected and highlighted.

If you select hosts and storage ports that are not contained in an existing zone alias, the new hosts and storage ports are added into the existing zone alias after the provisioning job finishes successfully.

To select a host:

1. Wait for all data to be loaded. All data is loaded when you see the following messages:

   ```
   Host data Loaded.

   Loading volume data finished.

   Loading HSG data finished.
   ```

   The Step 2 button is disabled until data is loaded

2. Take one of the following actions:

- Select a host that is accessible.

  Or

- Add a host that is not currently connected to the network by clicking the button (see Adding a Host below).

To create a provisioning job for multipathing, click the button above the Host pane, and then select the option for multipathing. When you select this option, you must select the same host for both paths (see Host Customize Dialog Box on page 797).

To automatically create a zone if a zone does not meet a preset criteria:

a. Click the button above the Zone pane.

b. Select the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.

c. Set the criteria. See Customize Zone Options Dialog Box on page 799 for more information about setting the criteria.

d. Select a host and click **Step 2**.

If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane.

If the management server does not find a zone that meets the criteria, it asks for a zone name. See Naming Conventions for Brocade and McDATA Switches on page 795 for more information about the naming requirements for a zone. After you enter a zone name, the new zone is displayed in the Zone pane, but it will not actually be created until the job runs.

To configure zoning manually, click the button above the Zone pane, and then deselect the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, select a zone or create one manually.

3. Click the **Step 2** button. Information about the selected port, such as its WWN, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.

Adding a Host

The management server lets you add hosts that are not currently connected to the network. While you are creating a job, add the host.

Make sure the added host is physically connected to the network before the scheduled job runs.

1. Click the button.
2. Enter a name for the host in the Host Name box.
3. Enter a port name of the host in the Port WWN box.
4. Click the **Add** button.
5. Repeat Steps 2 and 3 for multiple ports.

6. When you are done with your changes, click **OK**. The host is added to the list of hosts.
7. Physically connect the host to the network before the job runs.

Step 3 – Select a Volume

To select a volume:

1. In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the **Ctrl** key as you select the volumes.
   - **Mapped** – There are two types of mapped volumes:
     - **Masked** – The volume is exposed to the storage port and to the host.
     - **Unmasked** – The volume is exposed to the storage port, but not to the host.
   - **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
   - **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. When you select free extents, they must of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

     When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the "unmapped" category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server (see Deleting a Storage Volume on page 710).

     Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.
2. Click **Step 3**.
3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See Providing a LUN Number on page 789 for information about numbering LUNs.

   You can narrow the type of volumes displayed in the Volumes pane by using the Customize Volume Options dialog box (), located above the Volumes pane (see Host Customize Dialog Box on page 797).

   The Customize Volume Options dialog box is not available for the HP EVA.

   If the LUN was already selected and Step 4 is clicked, skip this step or click the  button.

Step 4 – Select a Host Security Group

To select a host security group:

1. Select a host security group. See Creating a Host Security Group on page 791 for information on how to create a host security group. See General Provisioning Issues on page 723 for information on how your storage system handles host security groups.
2. Click **Step 4**.
3. Click **Create Job**. The job is added to the bottom pane.
4. Take one of the following actions:
   - To execute the job now, click the **Execute Job** ( ) button.

     Or
   - To execute the job at a later time, schedule the job as described in Scheduling Provisioning Jobs on page 793.

Zone Operation

Zoning filters and operations work only on zones that are part of the active zone set.

If options from a previous job are still selected, clear the ones you do not want in your next job. For example, assume you created a zone and you now want to create a new zone that includes the same host used previously. Clear the Zone pane. To clear a pane, click the button.

To clear all the steps, except for Step 1 (storage systems), select another option from the System Action combo-box.

To use Path Provisioning to perform zoning operations:

1. Click **Provisioning** ( ).
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. From the System Action combo-box, select **Zone Operation**.

Step 1 – Select Storage System

To select a storage system:

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select a storage system.

   The S column heading in the Storage Systems pane means that only a single selection is allowed.
3. Click the **Step 1** button below the pane.

The selected storage system's name is displayed below the Storage System pane. The Host pane is populated. In following example, some hosts have a red X over their icon. This means the host is not accessible.

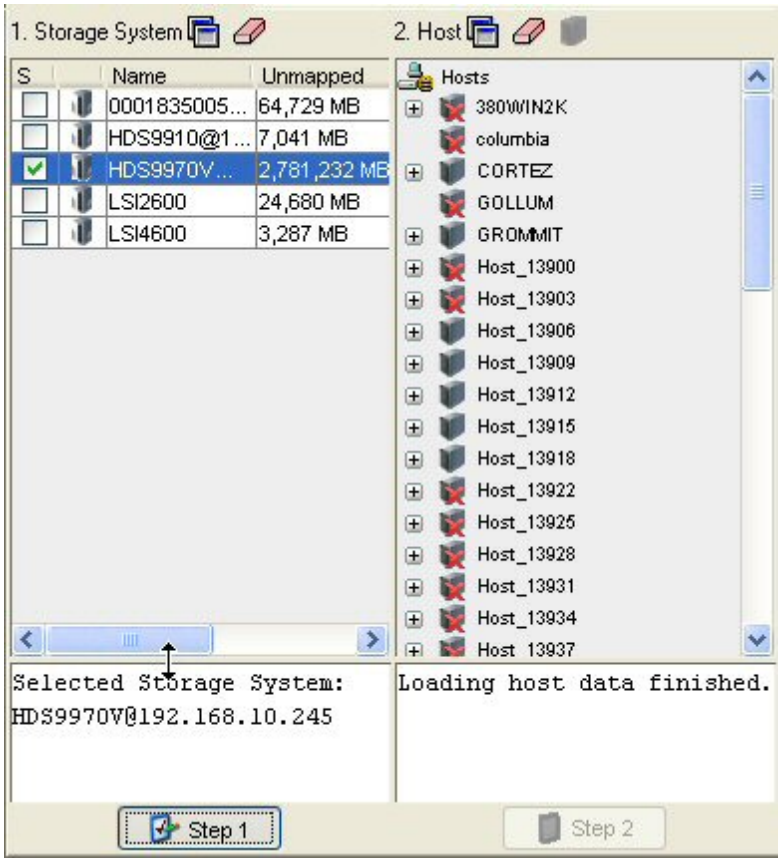

Step 2 – Select a Host

If you select host and storage ports that belong to an existing zone alias and you have the **Display Zone Aliases** option selected in Customize Zone Options dialog box, the existing zone alias is automatically selected and highlighted.

If you select hosts and storage ports that are not contained in an existing zone alias, the new hosts and storage ports are added into the existing zone alias after the provisioning job finishes successfully.

To select a host:

1. Wait for all data to be loaded. When all data is loaded, the following messages are displayed:

   ```
   Host data Loaded.

   Loading HSG data finished.

   Loading zone data finished.
   ```

   The Step 2 button is disabled until data is loaded.

6. Select a host that is accessible.

   Keep in mind the following:

- To create a provisioning job for multipathing, click the button above the Host pane, and then select the option for multipathing. When you select this option, you must select the same host for both paths (see Host Customize Dialog Box on page 797).
- To automatically create a zone if a zone does not meet a preset criteria:
  i. Click the button above the Zone pane.
  ii. Select the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.
  iii. Set the criteria. See Customize Zone Options Dialog Box on page 799 for more information about setting the criteria.
  iv. Select a host and click **Step 2**.
- If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane.
- If the management server does not find a zone that meets the criteria, it asks for a zone name. See Naming Conventions for Brocade and McDATA Switches on page 795 for more information about the naming requirements for a zone. After you enter a zone name, the new zone is displayed in the Zone pane, but it will not actually be created until the job runs.
- To configure zoning manually, click the button above the Zone pane, and then deselect the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, select a zone or create one manually.

2. Click **Step 2**. Information about the selected port, such as its Worldwide Name, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.
3. Select a host security group in the LUN pane. See Creating a Host Security Group on page 791 for information on how to create a host security group. See General Provisioning Issues on page 723 for information on how your storage system handles host security groups.
4. Click **Step 4**.

Step 3 – Select a Zone

If the zone is already selected and Step 5 is clicked, skip this step or click the button to clear the selection.

If a zone is not selected or created, most likely the option **Automatically Configure Zoning** is not selected in the Customize Zone Options dialog box ( ). The management server assumes you want to select a pre-existing zone or create one manually.

- **To reuse a zone** – Select a zone in the Zone pane, click **Step 5**, and expand a fabric node to view its zones.
- **To create a zone** – Select a fabric in the zone pane, click the button, and then enter a name for the zone (see Creating a Zone on the facing page).

Zone Icon	Description
	Above Zone pane – Used to create zones. In the Zone pane – Represents a zone.
	Zone Alias
	Port
	The fabric cannot be reached.

The zoning filters and operations work only on zones that are part of the active zone set.

McDATA and Brocade switches only: Path Provisioning looks for the names of the active zone set and of the active zones and verifies that all of their saved counterparts are matched in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

Creating a Zone

To create a zone:

1. Select a fabric in the zone pane.

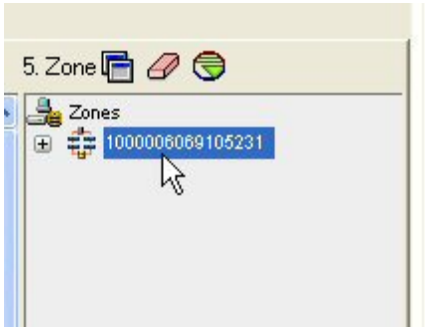

2. Click the button located above the Zone pane.
3. Enter a zone name in the dialog box. For naming conventions, see Naming Conventions for Brocade and McDATA Switches on page 795.
4. Click **OK**. The new zone is added to the Zone pane.
5. Click **Create Job**. The job is added to the bottom pane.
6. Take one of the following actions:
   - To execute the job now, click the **Execute Job** () button.

     Or
   - To execute the job at a later time, schedule the job as described in Scheduling Provisioning Jobs on page 793.

Volume Creation and LUN Security

If options are still selected from a previous job, clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Clear the Volume pane. To clear a pane, click the button.

To clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.

HDS only: Before you can create a volume, you must delete some unmapped LDEVs using the standard provisioning tool.

To delete LDEVs:

1. Click **Provisioning** ( ).
2. Click the **Storage Systems** tab and the **Provision** button for the storage system.
3. Click **Step 2 Volume**.
4. Select the desired number of LDEVs for the LUSE volume and click **Delete Selected Volumes**.
5. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

To create a meta volume and designate LUN security:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select **Volume Creation + LUN Security** from the System Action combo-box.

Step 1 – Select Storage System

1. Wait for the management server to load the storage system into the Storage System panel.
2. Select the storage system on which you want to create the metavolume.

   The S column heading in the Storage Systems pane means that only a single selection is allowed.

   (*HDS only*) Select the storage system from which you want to create the LUSE volume.
3. Click the **Step 1** button.

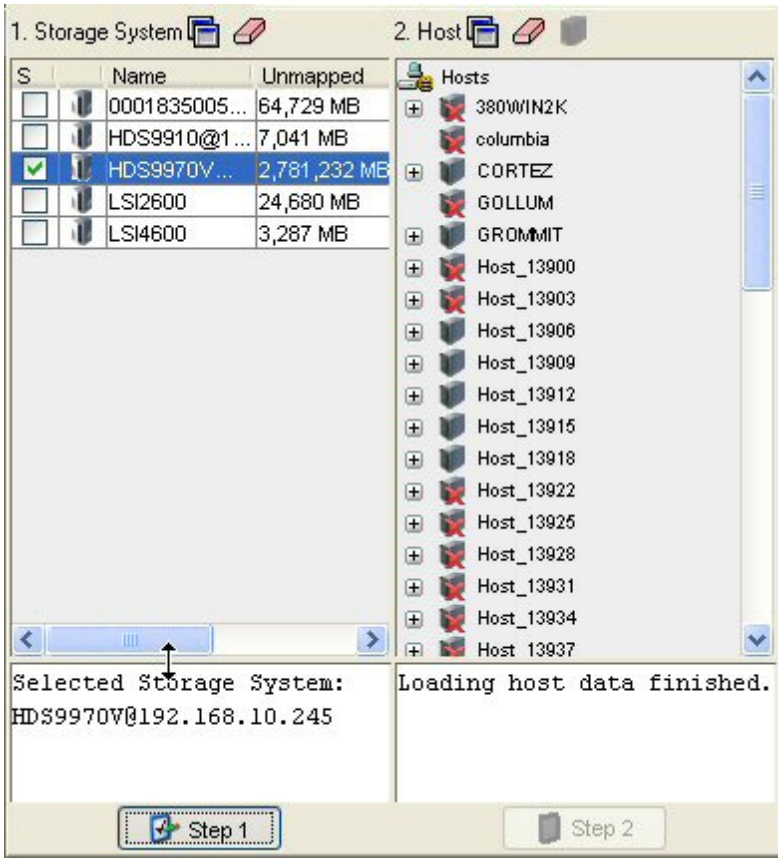

The name of the selected storage system is displayed.

Step 2 – Select a Volume

To select a volume:

1. In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the **Ctrl** key as you select the volumes.
   - **Mapped** – There are two types of mapped volumes:
     - **Masked** – The volume is exposed to the storage port and to the host.
     - **Unmasked** – The volume is exposed to the storage port, but not to the host.
   - **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
   - **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVS from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVS. When you select free extents, they

must of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server (see Deleting a Storage Volume on page 710).

Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.

Keep in mind the following:

- You can narrow the type of volumes displayed in the Volumes pane by using the Customize Volume Options dialog box ( ), located above the Volumes pane (see Host Customize Dialog Box on page 797).

  The Customize Volume Options dialog box is not available for the HP EVA.

- If the LUN was already selected and Step 4 is clicked, skip this step or click the button.

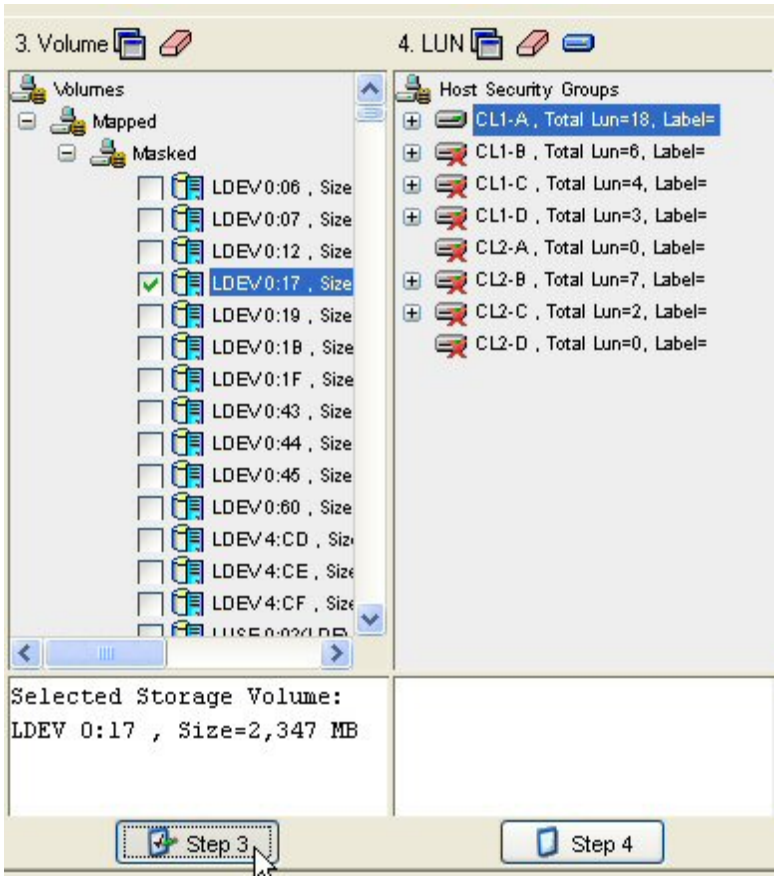

2. Click **Step 3**.
3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See Providing a LUN Number on page 789 for information about numbering LUNs.

Step 3 – Select a Host Security Group

1. Select a host security group in the LUN pane. See Creating a Host Security Group on page 791 for information on how to create a host security group. See General Provisioning Issues on page 723 for information on how your storage system handles host security groups.
2. Click **Step 4**.
3. Click **Create Job**. The job is added to the bottom pane.
4. Take one of the following actions:
   - To execute the job now, click the **Execute Job** ( ) button.

     Or
   - To execute the job at a later time, schedule the job as described in Scheduling Provisioning Jobs on page 793.

LUN Security and Zone Operation

If options are selected from a previous job, clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Clear the Volume pane. To clear a pane, click the button.

To clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.

This template does not create volumes or associate the volumes to HSG.

To use Path Provisioning to create a host security group (HSG) with the host HBA WWN along with zoning operations:

1. Click **Provisioning Manager** ( ).
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select the following from the System Action combo-box: LUN Security and Zone Operation.

Step 1 – Select Storage System

To select a storage system:

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select a storage system on which you want to create the metavolume.

The S column heading in the Storage Systems pane means that only a single selection is allowed.

3. Click the **Step 1** button below the pane.

   The selected storage system's name is displayed below the Storage System pane. The Host pane is populated. Notice in the following pane that some hosts have a red X over their icon. This means the host is not accessible.

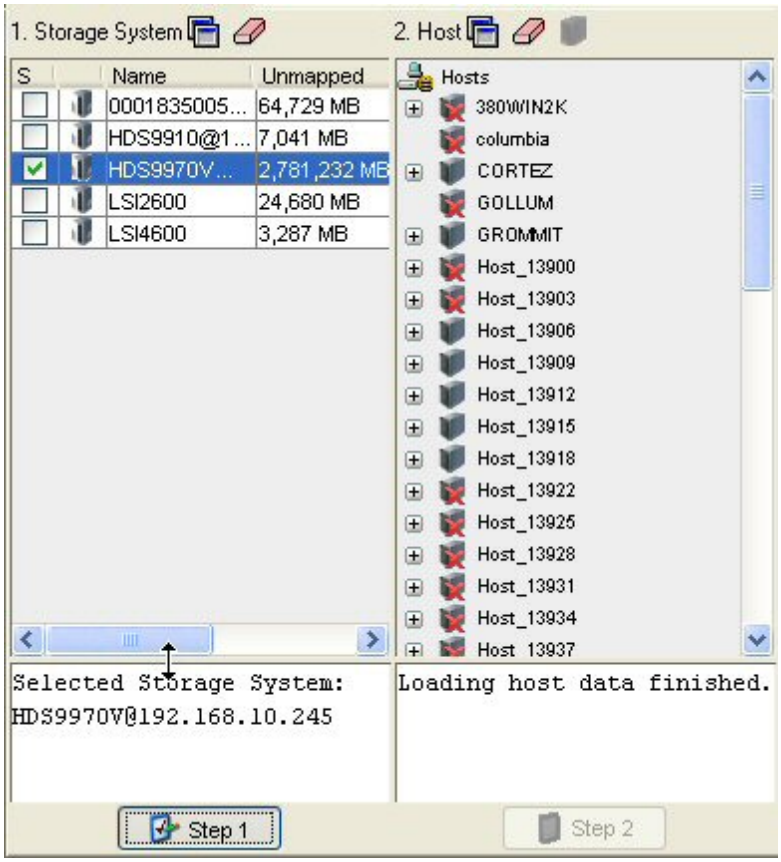

Step 2 – Select a Host

If you select host and storage ports that belong to an existing zone alias and you have the **Display Zone Aliases** option selected in Customize Zone Options dialog box, the existing zone alias is automatically selected and highlighted.

If you select hosts and storage ports that are not contained in an existing zone alias, the new hosts and storage ports are added into the existing zone alias after the provisioning job finishes successfully.

To select a host:

1. Wait for all data to be loaded. All data is loaded when you see the following messages:

   ```
   Host data Loaded.
   ```

```
Loading volume data finished.

Loading HSG data finished.
```

The Step 2 button is disabled until data is loaded.

2. Take one of the following actions:
   - Select a host that is accessible.

     Or

   - Add a host that is not currently connected to the network by clicking the button (see Adding a Host on page 772).

   To create a provisioning job for multipathing, click the button above the Host pane, and select the option for multipathing. You must select the same host for both paths (see Host Customize Dialog Box on page 797).

   To automatically create a zone if a zone does not meet a preset criteria:

   a. Click the button above the Zone pane.
   b. Select the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.
   c. Set the criteria. See Customize Zone Options Dialog Box on page 799 for more information about setting the criteria.
   d. Select a host and click **Step 2**.

   If the management server finds a zone that meets the criteria, it selects the zone in the Zone pane.

   If the management server does not find a zone that meets the criteria, it asks for a zone name. See Naming Conventions for Brocade and McDATA Switches on page 795 for more information about the naming requirements for a zone. After you enter a zone name, the new zone is displayed in the Zone pane, but it will not actually be created until the job runs.

   To configure zoning manually, click the button above the Zone pane, and then deselect the option **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**. In the zone pane, select a zone or create one manually.

3. Click **Step 2**. Information about the selected port, such as its Worldwide Name, is displayed below the Host pane. The volumes for that host are displayed in the Volume pane.

Adding a Host

The management server lets you add hosts that are not currently connected to the network. While you are creating a job, add the host.

Make sure the added host is physically connected to the network before the scheduled job runs.

1. Click the button.
2. Enter the name for the host in the Host Name box.

3. Enter the port name of the host in the Port WWN box.
4. Click **Add**.
5. Repeat Steps 2 and 3 for multiple ports.
6. To remove the host, click the button.
7. When you are done with your changes, click **OK**. The host is added to the list of hosts.
8. Physically connect the host to the network before the job runs.

Step 3 – Select a Host Security Group

To select a host security group:

1. Select a host security group. See Creating a Host Security Group on page 791 for information on how to create a host security group. See General Provisioning Issues on page 723 for information on how your storage system handles host security groups.
2. Click **Step 4**.
3. Click **Create Job**. The job is added to the bottom pane.
4. Take one of the following actions:
   - To execute the job now, click the **Execute Job** ( ) button.

     Or
   - To execute the job at a later time, schedule the job as described in Scheduling Provisioning Jobs on page 793.

Step 4 – Select a Zone

If the zone was selected and Step 5 is clicked, skip this step or click the button to clear the selection.

If a zone was not selected or created yet, most likely the option **Automatically Configure Zoning** is not selected in the Customize Zone Options dialog box ( ). The management server assumes you want to select a pre-existing zone or create one manually.

To reuse a zone – Select a zone in the Zone pane, click **Step 5**, and expand a fabric node to view its zones.

To create a zone – Select a fabric in the zone pane, click the button, and then enter a name for the zone. For more information, see Creating a Zone on the facing page.

Zone Icon	Description
	Above Zone pane – Used to create zones. In the Zone pane – Represents a zone.

Zone Icon	Description
	Zone Alias
	Port
	The fabric cannot be reached.

The zoning filters and operations work only on zones which are part of the active zone set. *McDATA and Brocade switches only*: Path Provisioning looks for the names of the active zone set and of the active zones and verifies that all of their saved counterparts are matched in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

Creating a Zone

To create a zone:

1. Select a fabric in the zone pane.

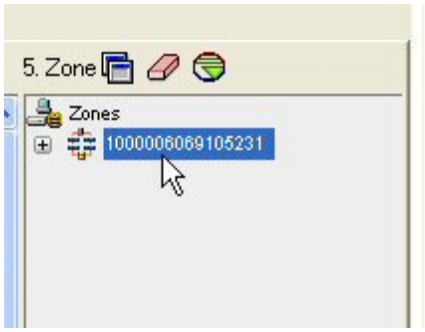

2. Click the button located above the Zone pane.
3. Enter a zone name in the dialog box. For naming conventions, see Naming Conventions for Brocade and McDATA Switches on page 795.
4. Click **OK**. The new zone is added to the Zone pane.
5. Click **Create Job**. The job is added to the bottom pane.
6. Take one of the following actions:
   - To execute the job now, click the **Execute Job** ( ) button.

     Or
   - To execute the job at a later time, schedule the job as described in Scheduling Provisioning Jobs on page 793.

Volume Assignment

If options are still selected from a previous job, clear the options you do not want in your next job. For example, assume you created a volume. Now you want to create a new volume on the same host used previously. Clear the Volume pane. To clear a pane, click the button.

To clear all the steps, except for the Step 1 (storage systems), select another option from the System Action combo-box.

HDS only: Before you can create a volume, you must delete some unmapped LDEVs using the standard provisioning tool.

To delete LDEVs:

1. Click **Provisioning** ().
2. Click the **Storage Systems** tab and the **Provision** button for the storage system.
3. Click **Step 2 Volume**.
4. Select the desired number of LDEVs for the LUSE volume and click **Delete Selected Volumes**.
5. Take note of the array group from which you deleted the LDEVs. You need this information to create the LUSE volume.

To assign a volume to existing host security groups:

1. Click **Provisioning Manager** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Select **Volume Assignment** from the System Action combo-box.

Step 1 – Select Storage System

To select a storage system:

1. Wait for the management server to load the storage systems into the Storage System panel.
2. Select a storage system on which you want to create the metavolume.

   The S column heading in the Storage Systems pane means that only a single selection is allowed.

   (HDS only) Select the storage system from which you want to create the LUSE volume.
3. Click the **Step 1** button below the pane. The selected storage system's name is displayed

below the Storage System pane.

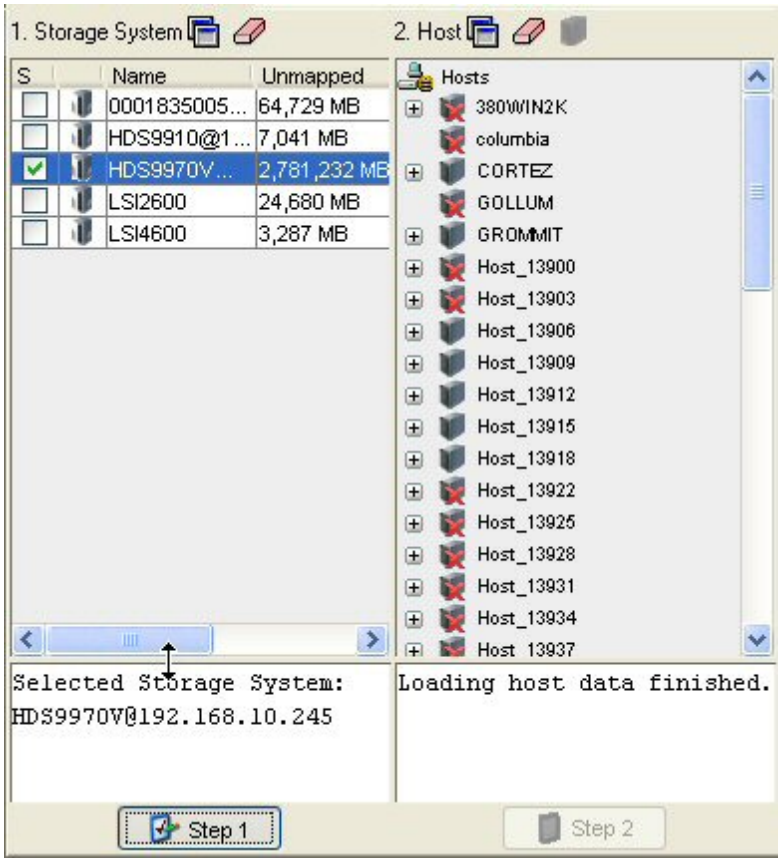

Step 2 – Select a Volume

To select a volume:

1. In the Volume pane select mapped and unmapped volumes. You can select multiple volumes on Windows computers by pressing the Ctrl key as you select the volumes.
   - **Mapped** – There are two types of mapped volumes:
     - **Masked** – The volume is exposed to the storage port and to the host.
     - **Unmasked** – The volume is exposed to the storage port, but not to the host.
   - **Unmapped** – The volume is not exposed to the storage port. The management server puts all unmapped volumes in this category when it first discovers an array.
   - **Free Extents** – Available free extents that can be used to create a meta volume. You can create meta volumes on EMC Symmetrix and LUSE on HDS storage systems. To create a meta volume or LUSE, select multiple free extents under the Free Extents node in the Volume pane. Select multiple LDEVS from the Free Extents menu by holding down the shift key on your keyboard and selecting free LDEVS. When you select free extents, they

must of the same type. For example, on Symmetrix, you cannot select a mirrored volume and a BCV (business continuous volume) to create a meta volume.

When you first discover a storage system, no free extents are displayed. This is because the management server puts all unmapped volumes into the “unmapped” category for the list of volumes by default. To move a volume to the free extent node, delete the unmapped volume. When the volume is deleted, it is moved to the free extents node. The free extents category is used internally by the management server (see Deleting a Storage Volume on page 710).

Make sure the free extents you select are not being used. Data on the free extents becomes unusable when a meta volume is created.

You can narrow the type of volumes displayed in the Volumes pane by using the Customize Volume Options dialog box ( ), located above the Volumes pane (see Host Customize Dialog Box on page 797).

The Customize Volume Options dialog box is not available for the HP EVA.

If the LUN was already selected and Step 4 is clicked, skip this step or click the button.

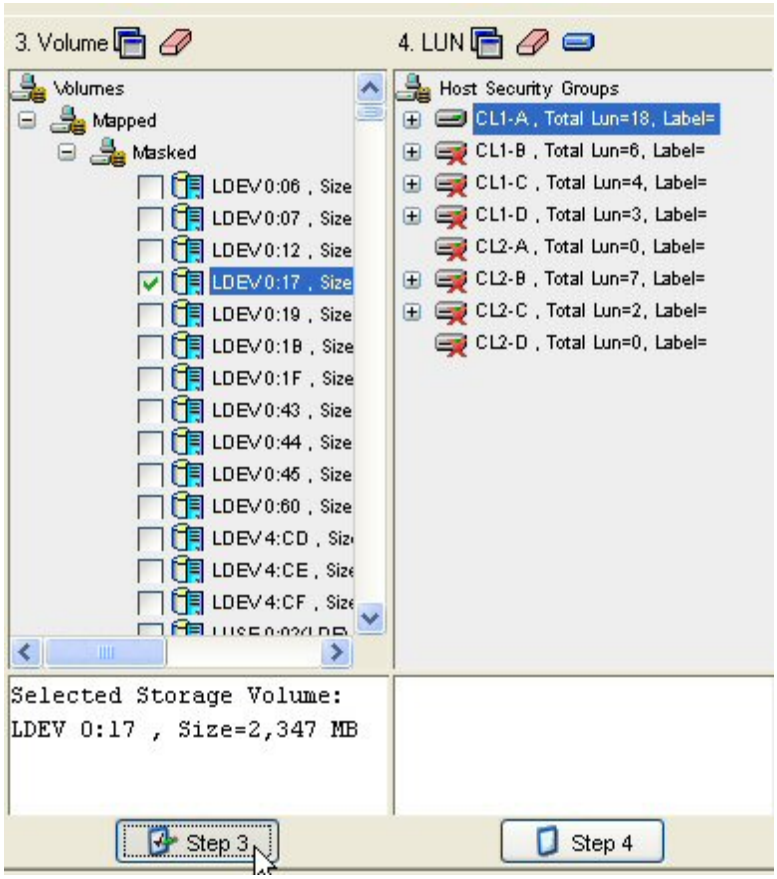

2. Click **Step 3**.

3. If you are asked to specify a LUN number, provide a LUN for each volume displayed. See Providing a LUN Number below for information about numbering LUNs.

Step 3 – Select a Host Security Group

To select a host security group:

1. Select a host security group in the LUN pane. See Creating a Host Security Group on page 791 for information on how to create a host security group. See General Provisioning Issues on page 723 for information on how your storage system handles host security groups.
2. Click **Step 4**.
3. Click **Create Job**. The job is added to the bottom pane.
4. Take one of the following actions:
   - To execute the job now, click the **Execute Job** ( ) button.

     Or
   - To execute the job at a later time, schedule the job as described in Scheduling Provisioning Jobs on page 793.

Providing a LUN Number

LUN numbers must be unique.

LSI storage systems: LUN numbers must be between 0 and 31.

Symmetrix storage systems: LUN numbers must be between 1 and 8190.

You can enter a LUN number for a volume by placing the cursor under the LUN Number column, as shown in the figure below. Click **OK** when you are done. This dialog box is displayed if your storage system requires you to provide a LUN number.

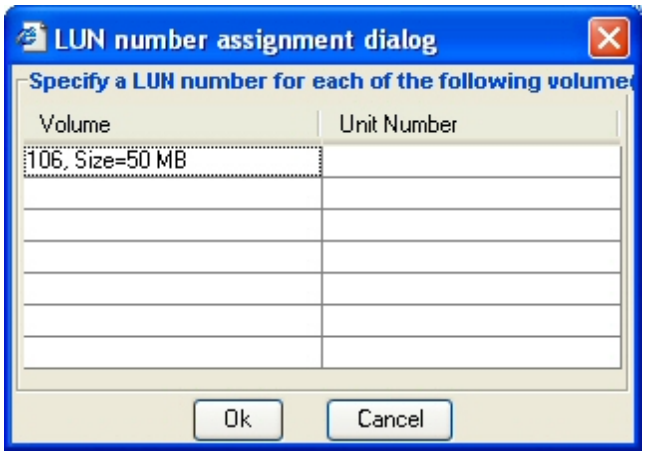

Creating a System Action Template

To create your own system action template:

1. Click **Provisioning Manager** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Click the **Configure Templates** button at the top of the screen.
4. Click the **New Template** button in the Provisioning Template Configuration dialog box.
5. To change the name that was assigned to the new template, enter the name you want for the template in the Template Name box and click **Apply**.
6. Select a master template on which you want to base your new template. For more information about the default master templates, see Default System Action Templates on page 762.
7. Click the following tabs and select the options you want for your template. Not all tabs are available for all default templates.
   - **Storage Options** – See Storage System Customize Dialog Box on page 797.
   - **Host Options** – See Host Customize Dialog Box on page 797.
   - **Volume Options** – See Customize Volume Options Dialog Box on page 798.
   - **Host Security Options** – See Customize HSG Options on page 798.
   - **Zone Options** – See Customize Zone Options Dialog Box on page 799.
8. When finished selecting all the options you want, take one of the following actions:
   - Click **Apply** to apply your changes and keep the Provisioning Template Configuration dialog box open.

     Or
   - Click **OK** to apply your changes and leave the Provisioning Template Configuration dialog box.

     Or
   - Click **Cancel** to leave the Provisioning Template dialog box without saving your changes. You will lose your changes unless you previously clicked **Apply**.

To use your new template, select it from the **System Actions** menu.

Modifying a System Action Template

You can only modify system action templates you created.

To modify a system action template:

1. Click **Provisioning Manager** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Click the **Configure Templates** button at the top of the screen.

4. Select the template you want to modify in the **Provisioning Templates** panel.
5. Make the necessary changes.
6. Click **Apply**.
7. When you are done, take one of the following actions:
   - Click **Apply** to apply your changes and keep the Provisioning Template Configuration dialog box open.

     Or

   - Click **OK** to apply your changes and leave the Provisioning Template Configuration dialog box.

     Or

   - Click **Cancel** to leave the Provisioning Template dialog box without saving your changes. You will lose your changes unless you previously clicked **Apply**.

Adding a Host

Make sure that the added host is physically connected to the network before the scheduled job runs.

The management server enables you add hosts that are not currently connected to the network while you are creating a job. This feature is only available when you select **LUN Security** from the System Action menu.

1. Click the button.
2. Enter a name for the host in the Host Name box.
3. Enter the port WWN of the host in the Port WWN box.
4. Click **Add**.
5. Repeat Steps 3 and 4 for multiple ports.
6. When done, click **OK**. The host is added to the list of hosts.
7. Physically connect the host to the network before the job runs.

Creating a Host Security Group

Keep in mind the following:

- Each storage system handles host security groups differently (see General Provisioning Issues on page 723).
- *For LSI storage systems*: When the **Volume Creation and LUN Security** option is selected on the System Action menu, you cannot add a host security group to an LSI storage system or to any storage system that LSI resells under a different brand.

- *For IBM storage systems:* You cannot assign the host mode in the user interface. You must modify an internal property (see Setting the Host Mode for IBM Storage Systems on page 723).

To create a host security group:

1. Select a storage system in the Storage System pane.
2. Click **Step 1**.
3. Select a host in the Host pane.
4. Click **Step 2**.
5. Select a port in the LUN pane.

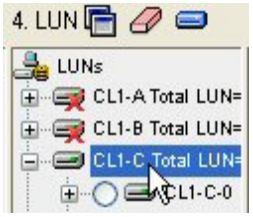

6. Click  at the top of the LUN pane.
7. When asked to provide a name for the new host security group, enter a unique name.

   For Symmetrix storage systems, you are not asked for the name of the host security group.

   If you have supported MSA P2000 G2 or G3 arrays, you might notice entries in the Initiators column with the value "FF:FF:FF:FF:FF:FF:FF:FF". Volumes shown there are LUNs on the MSA P2000 G2 or G3 arrays that were configured with Default Mapping.

   Keep in mind the following:

   - The name must contain 1 to 50 characters. If you enter no characters, you are given the option of using a default name.
   - The first and last letter cannot be spaces.
   - You cannot have the following characters in the name:

     <' > ; : , | / * ? \ \\ \t \n \b
8. *HDS only*:
   a. Select the host mode for the host security group.
   b. Provide a second host mode if applicable.
   c. Click **Create Host Security Group**.
9. For non-HDS storage systems, click **OK**. The host security group is created in the LUN pane.

Scheduling Provisioning Jobs

Keep in mind the following:

- You must create a provisioning job before you can schedule it (see About Path Provisioning on page 757).
- You cannot delete a job once it starts.
- You can deselect a job by clicking its check mark or by clicking the Clear Selection () button in the Provision Jobs pane.
- The User column in the Provision Jobs pane provides the name of the user who created the job. The admin user can see all of the created jobs. Other users can see only the jobs they created.

To determine the status of a job, look in the State column:

- **Created** – The job was created, but will not be executed. The job cannot be viewed by others, and it is deleted when the Web browser is closed. See Scheduling Provisioning Jobs above for information about changing the state of the job from “created” to “scheduled.”
- **Scheduled** – The job was tasked to execute at a specified time and date. Jobs are assigned a scheduled state after you select the job and click the **Execute Job** button () button.
- **Started** – The job started. You cannot delete a job once it starts.
- **Failed** – The job failed.
- **Ended** – The job finished.

To schedule a provisioning job:

1. Click the **Create Job** button in the lower pane. The Job Notes window opens. Enter a brief description of the job.

   To update the description, double-click it in the Description column of the Provision Jobs pane.

   The job is assigned a status of “created” but does not execute.

   Keep in mind the following:

   - When you close the Web browser window, all jobs with a status of “created” are erased.
   - Other users cannot see a job with a status of “created”.
   - Change the schedule of job only if its status is “created”. Once you click the Execute Job () icon, the job is saved in the database for the management server and its status changes from “created” to “scheduled”. Thus, it should not be modified.

If you are unable to click the **Create Job** button, verify that you clicked all the required Step buttons. If you are still unable to click the button, verify that the Host Customize dialog box ( ) is selected for multipathing. If it is, select a second path from the Path combo box and repeat the provisioning steps, except the steps for selecting a system action, storage system and host (see Host Customize Dialog Box on page 797).

2. Schedule the job by selecting it and clicking the button in the Provision Jobs pane.
   a. (*Optional*) In the Time box, change the time. The management server automatically displays a time 5 minutes from when you clicked the . Enter the time in 24-hour format, and separate hours and minutes with a colon. For example, enter 23:15 for 11:15 p.m.
   b. Select a date when you want the job to start.
   c. Click **Set**. If you click Set after the time has passed, you must reset the time.
3. To execute the job, select it and click the **Execute Job** ( ) button. The job will execute at the scheduled time.

   To execute a scheduled job, select it in the Provision Jobs pane. Detailed information appears in the Job Console pane, which is located in the lower-right corner of the window. The Message Console tab in the lower-right corner provides information such as whether the job ended or failed. If a job failed, the reason for the failure is provided on the Message Console tab.

   Keep in mind the following:

   - Jobs are executed according to the time set on the management server because they are centrally saved in the management server database.
   - When the management server is determining which job to perform first, it looks first for jobs requiring volume or zone creation. It does not perform the jobs in the order suggested by the Provision Job ID or according to the order of the jobs in the Provision Jobs table.
   - The management server can detect when the requested volume or zone was created. For example, assume you created a job to create a volume and, in the next few jobs, use this volume. Before creating the volume, these jobs first determine whether the volume was already created.

Executing Provisioning Jobs

To save and execute a job, click the **Execute Job** ( ) button. When you click that button, the job is saved on the management server. Other users can now see the job.

When the management server is determining which job to perform first, it looks first for jobs requiring volume or zone creation. It does not perform the jobs in the order suggested by the Provision Job ID, or the order of the jobs in the Provision Jobs table.

The management server can detect when the required volume or zone was created. For example, assume you created a job to create a volume and, in the next few jobs, use this volume. These jobs determine if the volume was created, before creating the volume.

You can schedule a job to start immediately or at a future time:

- To start a job immediately, click the **Execute Job** ( ) button.
- To start a job at a future time, first schedule the job and then click the **Execute Job** ( ) button. The job is executed according to the time set on the management server.

To view the latest information in Path Provisioning, click **Refresh**. For details about what the Refresh button updates, see About Path Provisioning on page 757.

Monitoring Provisioning Jobs

To view the latest status of the provisioning jobs listed, click the button in the Provision Jobs pane. The management server gathers information about the provisioning jobs listed to determine their latest status. Use this feature when you are not sure if a job ended.

Deleting Multiple Jobs

You cannot delete a job once it starts. A job starts when it has a state of “started.” To delete one or more jobs, select them in the M column and click the **Delete** ( ) button until they are all removed.

To deselect a job, click the check marks or click the **Clear Selection** ( ) button in the Provision Jobs pane.

Naming Conventions for Brocade and McDATA Switches

The following naming conventions apply to Brocade switches.

- The name must contain 1 to 64 characters.
- The name must begin with a letter. Any character other than the first character can be a letter, a number (0 to 9), or an underscore (_).
- The name is case sensitive. For example, “Zone1” and “zone1” are different zones.
- Two zones cannot have the same name as an existing zone, zone alias, or zone set. For example, if you create a zone named “new”, you cannot give that name to another zone, zone alias, or zone set.
- The following characters are invalid for Brocade switches: caret (^), dash (-), and dollar sign ($).

The following naming conventions apply to McDATA switches.

- The name can have a maximum of 64 characters.
- The first character of a zone name must be a letter (A-Z, AZ).
- A zone name cannot contain spaces.
- Valid characters are a-a, AA, 0-9, caret (^), dash (-), underscore (_), and dollar sign ($).

- All names must be unique and might not differ in case. For example, "myzone" and "MyZone" are considered the same name.

Using Multipathing with Path Provisioning

To use provisioning with multipathing, set the Host Customize dialog box to the multipathing option, and then repeat the provisioning steps for each path, as follows.

1. Select one of the following system actions:
   - **Volume Creation + LUN Security + Zoning** – Create a meta volume, map a volume to a Fibre Channel port and host HBAs (HSG), and then create a zone (see Volume Creation, LUN Security, and Zone Operation on page 762).
   - **LUN Security** – Map meta volume to Fibre Channel port and host HBAs (HSG) (see LUN Security on page 769).
   - **Zone Operation** – Create a new zone (see Zone Operation on page 774).
2. Select a storage system, and then click the **Step 1** button.
3. Click the  button above the Host pane.
4. Select the following option:

   ```
   Multipath: Select more than one port within a single server.
   ```

5. Select a port on a host, and then click **Step 2**.
6. Select a volume, host security group, or zone, as described in the following topics:
   - LUN Security on page 769
   - Zone Operation on page 774
7. Select the second path from the Path combo box.

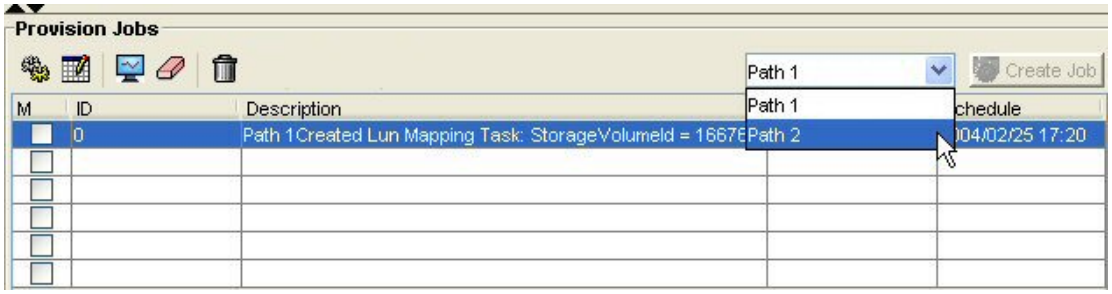

8. Repeat Step 6.

   You do not need to select a system action or storage system, and you must select the same host as the one used for the first path.
9. Set the schedule for the job as described in Scheduling Provisioning Jobs on page 793.

To narrow the types of volumes displayed in the Volume pan, set the Customize Volume Options dialog box. The Customize Volume Options dialog box is not available for the HP EVA.

Customizing Path Provisioning

This section contains the following topics:

Storage System Customize Dialog Box

The Storage System Customize dialog box lets you specify the type of storage systems displayed in the Storage System pane.

To specify the type of storage systems displayed in the Storage System pane:

1. Select one of the following:
   - **Show all available storage systems** – All storage systems are displayed.

     Or
   - **Show storage system with available raw storage** – Only storage systems with available raw storage are displayed.

     Or
   - **Show storage system with available unmapped storage** – Only storage systems with available unmapped storage are displayed.

     Or
   - **Show storage system with the following characteristics** – Select one or more of the storage tiers. Only storage systems belonging to the selected storage tiers are displayed. This menu is blank if no storage systems were assigned to storage tiers. Storage systems are assigned to storage tiers in Chargeback Manager. See Adding Asset Information on page 819 and Adding General Information on page 820 for information about assigning storage systems to a storage tier.
2. Click **OK**. The Storage System pane is updated.

Host Customize Dialog Box

The Host Customize dialog box enables you use multipathing with Path Provisioning.

To set multipathing:

1. Select one of the following options:
   - **Single path: Select one port of a host.** Select this option if you do not have multipathing or do not want to use multipathing with Path Provisioning.

     Or

   - **Multipath: Select more than one port within a single server.** Select this option if you want to use multipathing with Path Provisioning. For more information using multipathing with Path Provisioning, see Using Multipathing with Path Provisioning on page 796.
2. Click **OK** when finished making your selections.

Customize Volume Options Dialog Box

The Customize Volume Options dialog box is not available for the HP EVA.

1. Select one of the following options for metavolumes:
   - **Concatenating** – Only concatenating metavolumes are displayed.

     Or

   - **Striped** – Only striped metavolumes are displayed (*applies only to EMC storage systems*).
2. Click **OK**. The Volume pane is updated.

Customize HSG Options

1. To specify how LUN are mapped, select one or more of the following options in the Customize HSG Options dialog box:
   - **Perform number of path verification based on host selection** – The path verification is based on the host you selected.
   - **Automatically Assign Volume to Storage Port based on** – Select one of the following:
     - **Most unused ports (the number of LUNs assigned)** – Assigns a volume to a FA port based on how often a port is used. Unused ports have a better chance of having a volume assigned to them than frequently used ports. This option helps you spread out the traffic.
     - **Linked port if there is any** – A linked port is more likely to be assigned a volume than an unassigned port. The management server looks for any storage system port that is zoned to the selected HBA. If the management server cannot find a storage system port zoned to the selected HBA, it selects a port with the fewest LUNs.
   - **Verify that in a multipath configuration, storage system ports do not connect to the same switch** – This option makes sure the multipath configuration is preserved. Two storage system ports should not get connected to the same switch.

- **Assign a LUN number automatically** – Do not select this option if you want to assign a LUN number manually.

2. Click **OK** when done.

Customize Zone Options Dialog Box

When the option Automatically create new zone if no existing zone containing HBA and storage system ports is detected is selected in the Customize Zone Options dialog box (), the management server automatically selects a zone that meets its criteria. If the management server cannot find a zone that meets its criteria, it creates a zone on the fly. To set the criteria for automatically configuring zones, select one of the following options:

- **Create new zone if there is no zone containing ONLY the selected zone members (HBA and storage system ports)** – The management server checks to see if an identical zone exists. An identical zone contains only the same HBA and storage system ports you selected. If the zone contains additional members, it is not considered to be identical.
  - If an identical zone exists, it is selected in the Zone pane.
  - If an identical zone does not exist, the management server asks you to provide a name for the zone that will be created. The new zone appears in the Zone pane, even though it is not created until the job runs.
- **Create new zone if there is no zone containing AT LEAST the selected zone members –** The management server tries to find a zone that contains the HBA and storage system ports you selected.
  - If a zone contains additional members, the management server selects that zone in the Zone pane.
  - If the management cannot find a zone containing the HBA and storage system ports you selected, the management server asks you to name a zone that will be created. The new zone appears in the Zone pane, even though it is not created until the job runs.

The zoning filters and operations work only on zones which are part of the active zone set. *McDATA and Brocade switches only*: Path Provisioning looks for the names of the active zone set and of the active zones and verifies that all of their saved counterparts are matched in the zoning library in EFC Manager. The provisioning job only occurs if those names match.

About the Use Switch Port Zoning Mode Option

By default, Path Provisioning creates zones through WWNs, which means it looks for a port on a host and a port on a storage system to create the zone. Zoning through WWNs is not dependent upon the switch. This means you could change switches as long as the host and storage system are able to access each other through the network.

You can create zones through switch port zoning if you select the Use Switch Port Zoning Mode option. When you select the storage system and host, the corresponding switch ports are selected for the zone. For example, ports 1 and 2 on a switch are zoned when the host HBA port and storage system controller port are connected to either of the switch ports.

Display Zone Alias Option

Select the Display Zone Alias option to prevent you from creating duplicate zone aliases on the same port. If you have this option selected when path provisioning detects a Fabric support zone alias provisioning feature, all zone aliases in this Fabric will be also listed on the Path Provisioning Zone panel. If the selected ports are part of any existing zone alias provisioning feature, such zone aliases will be selected and highlighted to prevent you from creating duplicate zone aliases on the same port.

Automatically Configure Zoning

To automatically configure zoning:

1. Click the  button above the Zone pane before you select a host.
2. Select **Automatically create new zone if no existing zone containing HBA and storage system ports is detected**.
3. Select one of the options discussed in the previous bulleted list, and click **Apply**.
4. Select a storage system, and click **Step 1**.
5. Select a host and click **Step 2**. One of the following occurs:
   - If the management server cannot find a zone that meets the criteria set in Step 3, it asks for a zone name.
   - If the management server finds a zone that meets the criteria set in Step 3, it selects the zone in the Zone pane.
6. If you are asked for a zone name, enter the zone name and click **OK**. See Naming Conventions for Brocade and McDATA Switches on page 795 for restrictions on naming zones. The new zone is displayed in the Zone pane, but it will not be created until the job runs.

Manually Configure Zoning

The management server assumes you want to select a pre-existing zone or create one manually when you deselect the Automatically create new zone if no existing zone containing HBA and storage system ports is detected option.

- **To reuse a zone** – Select a zone in the Zone pane, and then click **Step 5**. Expand a fabric node to view its zones.
- **To create a zone** – Select a fabric in the zone pane, click the  button, and then enter a zone name.

Zone Icon	Description
	Location on the page determines its use: • **Displayed above Zone pane** – Button for creating zones. • **In Zone pane** – Icon for a zone.
	Zone Alias
	Port
	The fabric cannot be reached

To create a zone:

1. Select a fabric in the zone pane.

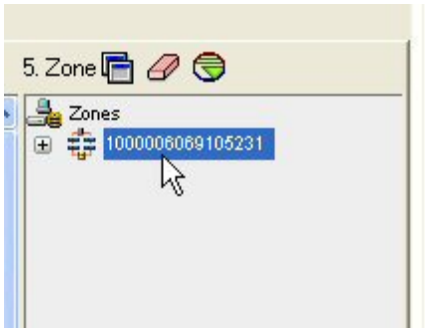

2. Click the button located above the Zone pane.
3. Enter a zone name in the dialog box. For naming conventions, see Naming Conventions for Brocade and McDATA Switches on page 795.
4. Click **OK**. The new zone is added to the Zone pane.

Assigning a Template to a Role

You can assign templates to a role to restrict a user's access to all templates. For example, you could specify that only users with administrator privileges can access the Volume Assignment template.

You must belong to a role that has "Provisioning Administration" privileges to be able to assign templates to a role. The domain administrator and storage administrators roles have these privileges by default. To determine if your role has "Provisioning Administration" privileges, your domain administrator will need to go to the **Security** > **Users** page, and click the name of the role to see which privileges are assigned to your role.

To assign templates to a role:

1. Click **Provisioning Manager** ().
2. In the right pane, click **Start Here** on the Path Provisioning tab.
3. Click **Assign Templates**.
4. Open the **Select Role** menu and select the role you want to have access to the Path Provisioning template.

   For a role to appear in the **Select Role** menu, it must have the following attributes:

   - Have provisioning privileges, meaning it has enough privileges to access Path Provisioning.
   - Does not have provisioning administration privileges. Roles that have Provisioning Administration privileges do not appear in the **Select Role** menu.
5. In the Provisioning Templates pane, select one of the following options for access to the templates for the users belonging to the role you selected in step 4:
   - Select the templates from the Assigned Provisioning Templates panel you want the users to be able to access and click **Add**.
   - Click **Add All >>** to make all the templates available to users assigned to the selected role.
   - Select the templates you do not want the users to be able to access and and click **<Remove**.
6. Click **<< Remove All** to make all the templates unavailable to users assigned to the selected role.

   The templates that the users are able to access appear in the Assigned Provisioning Templates panel.
7. When done, take one of the following actions:
   - Click **Apply** to apply your changes and keep the Provisioning Template Configuration dialog box open.

     Or
   - Click **OK** to apply your changes and leave the Provisioning Template Configuration dialog box.

     Or
   - Click **Cancel** to leave the Provisioning Template dialog box without saving your changes. You will lose your changes unless you had previously clicked **Apply**.

19 Chargeback Manager

Chargeback Manager helps you manage departmental ownership and track cost, which makes inquiries, such as audits and inventory reviews, easier.

This section contains the following topics:

About Chargeback Manager

The management server provides the following types of chargeback:

- **Asset-based –** Asset-based chargeback calculates chargeback based on the departmental ownership percentages and the depreciated value of the assets. Each piece of equipment is owned by a department or a set of departments. Each department has a percentage ownership of the equipment. Chargeback Manager provides asset information for virtual machines and virtual servers.
- **Storage-based** – Storage-based chargeback calculates charges based on the actual amount of storage used by an application, the type of storage it is using, and the ownership percentage assigned to each department. The chargeback number is further refined by an additional fixed infrastructure tax on a per-department basis.

After you add information about all of your assets, back up the database using the Database Admin Utility. Backing up saves your chargeback information. If the database fails, your asset information is restored when you restore the database (see Performing an RMAN Hot Backup on page 385).

First set up your chargeback as described in Setting Up Chargeback Manager below. When you are done, you can view chargeback as follows:

- **By element** – Displays chargeback for a single element (see Viewing Chargeback by Asset on page 840).
- **By department** – Displays chargeback for a department (see Viewing Chargeback by Department on page 840).
- **By owner** – Displays chargeback for an owner (see Viewing Chargeback by Owner on page 841).

Chargeback Manager helps you track the status of your elements. Elements that were recently discovered are automatically given the status of New. You can change the status of elements to In Use, Missing, or Repaired. Because the management server cannot determine what you plan to do with an element, you must change the status manually. You can easily change the status of a group of elements (see Setting the Status of an Asset on page 806).

Chargeback Manager also provides a brief listing of your assets by name, status, serial number, vendor/model, and description. You can filter elements by status and type for easy navigation. You can even create custom filters (see About Filtering Assets on page 844).

Setting Up Chargeback Manager

To view chargeback, you must first complete the following tasks.

1. Create an asset record (if it does not exist). You can use Chargeback Manager for applications and hardware that the management server does not detect. Create an asset record for each asset shown as an element in Chargeback Manager.

   Creating an Asset Record on the facing page

2. Set the status of the asset.

   Setting the Status of an Asset on page 806

3. *Optional*: Add asset information for asset management.

   Adding Asset Information on page 819

3. Add departments.

   Adding Departments on page 823

4. Set up the chargeback method —asset-based or storage-based.

   Setting Up Asset-Based Chargeback on page 824

   Setting Up Storage-Based Chargeback on page 829

Accessing Chargeback Manager

To access Chargeback Manager, click **Chargeback Manager** ().

Creating an Asset Record

You can use Chargeback Manager for any application or hardware, even those the management server does not detect. Create an asset record for the application or hardware the management server does not monitor, and then follow the steps for setting up Chargeback Manager, as described in Setting Up Chargeback Manager on previous page.

Only a user belonging to a role that has System Configuration selected on the Edit Role page (such as the Domain Administrator role) is allowed to create a record.

After you create a record, the element for which you created the record is treated as a discovered element. A discovered element is an element that is detected by the management server, but the management server cannot obtain detailed information about the element. If you create a record for an application, that application is treated as a virtual application.

You can easily remove an asset record by clicking the **Delete** () button. When you remove an asset record, the management server no longer monitors that asset element (see Deleting Elements from the Product on page 157).

To create a record:

1. Access Chargeback Manager, as described in Creating an Asset Record above.
2. Click **New**.
3. In the **Add Asset Record** window, enter the following:
   - **Name** – Name of the element.
   - **Vendor** – Vendor of the element.
   - **Model** – Model of the element.
4. Select the type of element:
   - **Application**
   - **Host**
   - **Storage System**
   - **Switch**
   - **Tape Library**
5. Click **OK**.

Setting the Status of an Asset

Chargeback Manager helps you track of the status of your assets. Asset elements that were recently discovered are automatically given the status of New. You can change the status of asset elements to In Use, Missing, or Repaired. Because the management server cannot determine what you plan to do with an asset, you must change the status manually. You can easily change the status of a group of assets instead of one at a time.

After you are done changing the status of your assets, save your settings (see Saving Chargeback Manager Information on the facing page).

To change status:

1. Access Chargeback Manager, as described in Accessing Chargeback Manager on previous page.
2. Under the Status column, select the status of the asset:
   - **New (Default)** – Status of the asset has not been set yet.
   - **Missing** – Asset cannot be found. It might have been taken off line.
   - **Repaired** – Asset is being repaired.
   - **In Use** – Asset is running.

   The status settings are set manually. For example, if the status of an asset element changes from In Use to Repaired, you must change this status manually. Refer to the Topology and Event Manager for the latest status of an element.
3. To change the status of multiple assets at one time:
   a. Select the asset elements you want to modify.
   b. Click **Set Status**.
   c. From the Asset Status menu, select the new status for the elements you selected.
   d. Click **OK**.

Saving Chargeback Manager Information

After you change the status of your elements, save your settings by clicking the **Save Listing to File** link. The following information is saved as comma-separated values, which can be viewed using a text editor, such as Notepad, or a spreadsheet program, such as Microsoft Excel.

- ID
- Name
- Status
- Category
- Serial Number
- Vendor
- Model

Viewing Assets

To obtain asset information about an element:

1. Access Chargeback Manager, as described in Accessing Chargeback Manager on page 805.

   The right pane displays the following:

   - **Name** – The name of the element.
   - **Status** – The status of the element. An element is automatically assigned the status of New when it is first discovered. You can change the status of an element to In Use, Repaired or Missing (see Setting the Status of an Asset on previous page).
   - **Chargeback Manager** – Click the  icon to view chargeback for an element. You must first set up Chargeback Manager before you can view its calculations (see Setting Up Chargeback Manager on page 804).
   - **Vendor/Model** – The vendor and/or model of an element.
   - **Type** – The type of element, such as an application. The following table shows the icons for the element types.
   - **Serial Number** – The serial number of the element.
   - **Description** – Description of the element.

Graphic	Element Type
	Application

Graphic	Element Type
	Host
	Switch
	Storage System
	Tape Library

2. To remove an asset record, click the **Delete** () button corresponding to the record you want to remove.

Managing Dynamic Storage Tiers

HP Storage Essentials provides flexible automated rules-based assignments for categorizing LUNs into tiers, so you can assign elements to tiers based on their tier definitions.

With HP Storage Essentials you can use a rule wizard to define a rule with a set of attributes for each tier. Based on the define rules for each tier, the wizard sets the priorities to each tier. You can also create a schedule to run defined tiers. Manual overrides are also supported.

Manual assignments can be overridden by the auto tier assignments. For example, assume you created a dynamic storage tier that requires its element to have a disk size of more than 900 GB. Then, you manually add an element that has a disk size of less than 900 GB to the storage tier. During the next refresh of rule-based membership, elements that do not fit the criteria for being a member of the storage tier are removed, except for the elements you manually added. The elements you manually added stay members of the storage tier even if they do not meet the criteria of the storage tier. To remove all elements you manually added quickly, use the clear all mappings option on the Dynamic Storage Tiers tab and then refresh the storage tiers.

Based on the tier definition and schedule, the elements are added automatically to the corresponding tier. A SAN administrator can easily identify which elements fall into which tier without manual intervention.

For example, assume you have the following environment:

- Tens of thousands of storage volumes, and each array was purchased with multiple tiers of storage.
- Your tier model is built on array type (enterprise, departmental), physical disk spindle size, and our offering (leveraged, dedicated).
- You have some volumes that fall out of your basic rules so you also need the ability to "override" a volume's rule generated tier to put our own on it.
- You must be able to differentiate virtual volumes from traditional volumes in our rules because they have a separate tier.

You could use dynamic storage tiers to have elements automatically added to the corresponding tier.

Overview of Setting Up Dynamic Storage Tiers

1. Create a dynamic storage tier, as described in Creating Dynamic Storage Tiers on next page. Create the storage tier to match the attributes of the elements you want to monitor. Elements that match the criteria will be automatically added to the storage tier.
2. Further refine the elements you want to add to the storage tier, such as volumes and LUNs. See Add or Remove Storage Elements from a Storage Tier on page 813 for more information.
3. Set up a refresh schedule for the storage tiers, as described in Setting a Schedule for Storage Tiers on page 815. To refresh the storage tiers now, click **Refresh Now**.

Dynamic Storage Tiers Tab

The table provides the following information to help you manage the tiers:

- **Priority**. Determines when the priority is kicked off. Storage tiers with a priority of 1 run first, and elements with a priority of 2 run second. Storage tiers with a priority of zero do not run. See Changing the Priority of a Storage Tier on page 817 for information about changing the priority for a storage tier.
- **Storage Tier Name**: The name of the storage tier
- **Monthly Cost per Gigabyte**: The cost of the storage tier
- **Total Elements in Tier**: The number of elements in the storage tier
- **Description**: A description of the storage tier
- **Has Attributes Defined?**: Shows a red checkmark if storage attributes are selected for the storage tier. Attributes must be defined for the storage tier to run. If you see no checkmark for a storage tier, the storage tier has no attributes defined and, therefore, that storage tier will not run.
- **Edit Elements**: Click to add or remove elements for a storage tier.
- **Edit Policy**: Click to modify the rules for a storage tier.
- **Delete**: Click for the storage tier you want to delete. To remove multiple storage tiers, select multiple storage tiers and click the Delete button above the table.

The tab also provides the following buttons:

- **Create New Storage Tier**: Brings up a wizard for creating a storage tier (see Creating Dynamic Storage Tiers on next page.)
- **Edit Priority**: Modifies the priority of a storage tier (see Changing the Priority of a Storage Tier on page 817).
- **Clear All Mappings**: Removes all elements from selected storage tiers (see Removing Mappings for a Storage Tier on page 817).
- **Delete**: Deletes selected storage tiers.

Setting Up Dynamic Storage Tiers for EMC Symmetrix

HP Storage Essentials obtains the architecture of array disk drives (Fibre, Flash, SATA, and so on) for the purpose of reporting and to create dynamic tier rules. However, not all EMC Symmetrix microcode and Solutions Enabler versions support the reporting of disk drive architecture (see your EMC documentation for more information). You can enable this reporting in your Symmetrix arrays by setting the cimom.emcDiskArchModelTypes property. This property serves as a supplemental lookup table that can be used by the internal EMC Symmetrix provider. The property must be formatted as a comma-separated list of DISK_MODEL:TYPE pairs. The following example shows two different disk drive models and their corresponding type, or architecture:

```
cimom.emcDiskArchModelTypes=SX373405FC:FC,SX3500071FC:SATA
```

The TYPE is displayed in the user interface. The DISK_MODEL is not case-sensitive.

Accessing the Dynamic Storage Tiers Tab

To access the Dynamic Storage Tiers tab:

1. Click **Chargeback Manager** ().
2. Click the **Dynamic Storage Tiers** tab.

Creating Dynamic Storage Tiers

Create storage tiers that match the attributes of the storage you want to monitor.

To create storage tiers:

1. Access the Dynamic Storage Tiers tab, as described in Accessing the Dynamic Storage Tiers Tab above.
2. Click the **Create New Storage Tier** button.
3. Change the storage tier name in the Storage Tier Name box.Provide the storage tier name in the Storage Tier Name box. You can provide up to 255 characters, but the name will be trimmed to 30 characters in the Dynamic Storage Tiers table. You will be able to see the full name of the storage tier when you place your mouse over the name of the storage tier, as shown in the following.

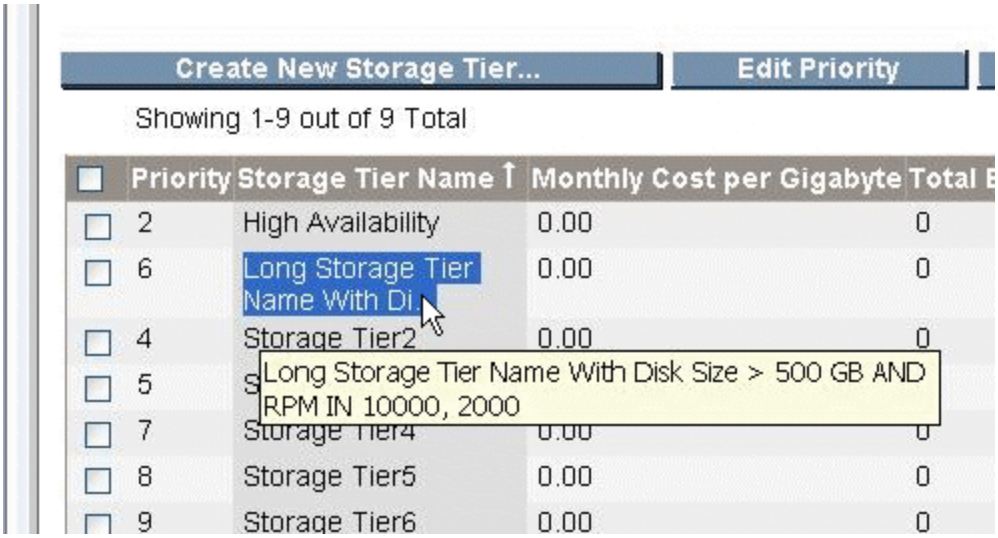

4. Provide the monthly cost per gigabyte.
5. Provide a description for the storage tier. You can provide up to 1023 characters, but the description will be trimmed to 30 characters in the Dynamic Storage Tiers table. You will be able to see the full description when you place your mouse over the description of the storage tier.
6. Select one of the following:
   - **All** if you want all storage systems to be eligible to belong to the storage tier.
   - **Selected** if you want only certain storage systems to be eligible to belong to the storage tier.

   Select **Vendors**, **Models**, or **Systems** from the Display box. If you select **Systems**, select the vendor and model of the storage system you want included in the tier.

   Select an item in the left pane and click the greater than sign (>) to move it to the right pane. To move all items in the left pane, click the double greater than sign (>>).

   The specified elements will be scanned during a refresh of the dynamic storage tiers. If they meet the rest of the criteria of the storage tier, they will be added as a member.
7. (*Optional*) Select a storage system type, offering, or both. Only storage systems that are of this type or offering are eligible to become a member of the storage tier. For example, if you select Enterprise and Leveraged, only storage systems that are set to Enterprise and Leverage are eligible to become members of the storage tier. They must also meet the criteria of the previous step in addition to the criteria for the storage attributes.
   You can change the options provided in the system type and offering menus. See Managing the Properties of Storage Systems on page 818 for more information. For information about editing the properties of a storage system, see Editing the Properties of Storage Systems on page 818.
8. You must select storage attributes for your storage tier to run. Select the disk storage attributes you want to be used for apply the rule. For example, if you want the storage tier to

apply to only disk sizes that are more than 500 GB, select Disk Size under Storage Disk Attributes and enter the following:

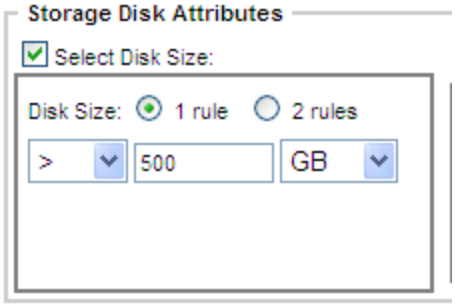

If you select multiple attributes, the element must meet all of the attributes to be included in the dynamic storage tier. For example, assume you set a storage tier to have the following attributes:

- Disk Size > 200 GB
- RAID5(5D+1P)

Only volumes that contain disk drives that are greater than 200 GB and are of RAID5(5D+1P) can become members of the storage tier, provided they also meet the storage system type and offering if those were set.

Make sure that the product can gather information about the storage attributes from your storage system. Otherwise, the storage system cannot be added to the storage tier, even though it meets the criteria.See Verify that HP Storage Essentials Can Obtain Storage Attributes from a Storage System below

The product is configured by default to not gather information from disk extents on EVA storage systems. You must modify a setting so the product can begin to gather information from disk extents on EVA storage systems.See Enabling the Product for EVA Disk Extents on the facing page.

Verify that HP Storage Essentials Can Obtain Storage Attributes from a Storage System

Make sure that the product can gather information about the storage attributes from your storage system. Otherwise, the storage system cannot be added to the storage tier, even though it meets the criteria.

The product is configured by default to not gather information from disk extents on EVA storage systems. You must modify a setting so the product can begin to gather information from disk extents on EVA storage systems. See Enabling the Product for EVA Disk Extents on the facing page.

To check if HP Storage Essentials can gather storage system attributes:

1. Click **Discovery** > **Details**.
2. In the elements column, click the storage system you want to verify. The Navigation tab appears.
3. Click the **Disk Drives** button. The disk drives on the storage system are displayed.
4. Click the name of a disk drive to view the properties and LDEVs that HP Storage Essentials detects that are on the drive.

   If you do not see the data for your storage systems, go to the Default Properties page (**Configuration** > **Product Health** > **Advanced** > **Custom Properties**) and verify that the following properties are not listed:

   ```
   synchronizerNoExtents=all
   ```

   ```
   synchronizerNoExtents=HDS,LSI,EMC,Clariion
   ```

   These properties are commented out by default. If someone enabled them, they would be listed in the Custom Properties box.
5. Make sure HP Storage Essentials can collect information from the disk extents from EVA storage system.

For more information about commenting out this property, see Enabling the Product for EVA Disk Extents below.

Enabling the Product for EVA Disk Extents

The product is configured by default to not gather information from disk extents on EVA storage systems. You must modify a setting so the product can begin to gather information from disk extents on EVA storage systems.
To enable the product to gather information from disks on EVA storage systems:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Add the following entry into the Custom properties box:

   ```
   synchronizerNoExtents=EVA
   ```
4. Save the settings and run Get Details (**Discovery** > **Details**) for the EVA array.

After the discovery, HP Storage Essentials displays the disk drives for the EVA array.

Add or Remove Storage Elements from a Storage Tier

You can further refine the elements that you want to be in the storage tier. For example, you can add specific storage volumes and LUNs to a storage tier.

To add storage elements to a storage tier:

1. Access the Dynamic Storage Tiers tab (see Accessing the Dynamic Storage Tiers Tab on page 810).
2. Click the **Edit** () button in the Edit Elements column.

3. Expand the Element Types heading:

4. Select an element type, such as Storage Systems or NAS Systems. The pane to the right populates when you select an element type.
5. Select the elements you want to add to the tier and click **Add Selected Elements to Storage Tier**. If you selected storage systems, click the Storage Volumes and Storage Pools tabs to add specific volumes and LUNs to the storage tier. The element is displayed in the Tier Members pane.

To remove elements from the storage tier:

1. Access the Dynamic Storage Tiers tab (see Accessing the Dynamic Storage Tiers Tab on page 810).
2. Click the **Edit** ( ) button in the Edit Elements column.
3. Select the elements you want to remove.
4. Click **Remove Selected Elements from Storage Tier**. The selected elements are removed from the storage tier.

Editing Dynamic Storage Tier Policies

You can always edit a dynamic storage tier policy so you can refine your information gathering. Manual assignments can be overridden by the auto tier assignments. For example, assume you created a dynamic storage tier that requires its element to have a disk size of more than 900 GB. Then, you manually add an element that has a disk size of less than 900 GB to the storage tier. During the next refresh of rule-based membership, elements that do not fit the criteria for being a member of the storage tier are removed, except for the elements you manually added. The elements you manually added stay members of the storage tier even if they do not meet the criteria of the storage tier. To remove all elements you manually added quickly, use the clear all mappings option on the Dynamic Storage Tiers tab and then refresh the storage tiers.

To add elements manually to a storage tier:

1. Access the Dynamic Storage Tiers tab, as described in Accessing the Dynamic Storage Tiers Tab on page 810.
2. Click the **Edit** ( ) button in the Edit Policy column.
3. Change the storage tier name in the Storage Tier Name box.Provide the storage tier name in the Storage Tier Name box. You can provide up to 255 characters, but the name will be trimmed to 30 characters in the Dynamic Storage Tiers table. You will be able to see the full

name of the storage tier when you place your mouse over the name of the storage tier, as shown in the following.

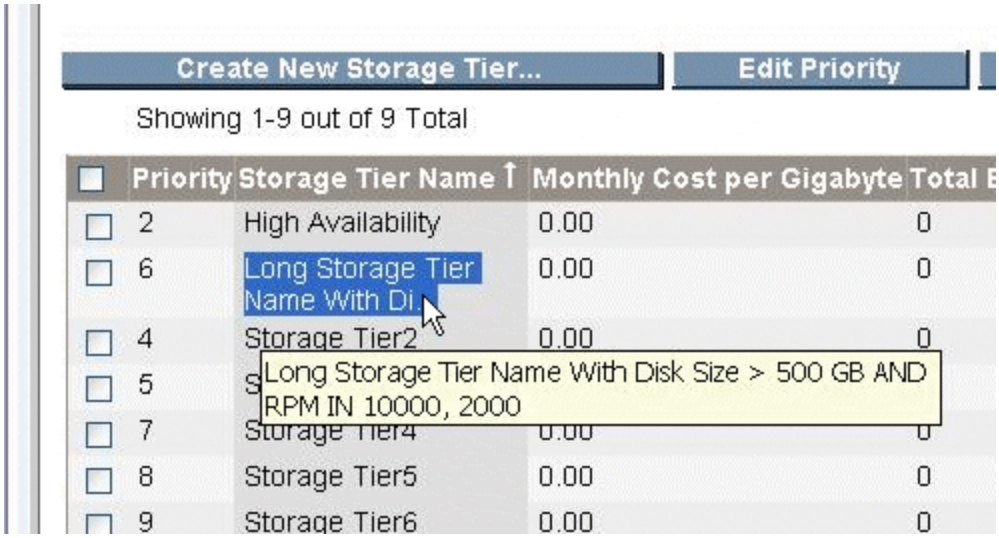

4. Change the monthly cost per gigabyte.
5. Change the description for the storage tier. You can provide up to 1023 characters, but the description will be trimmed to 30 characters in the Dynamic Storage Tiers table. You will be able to see the full description when you place your mouse over the description of the storage tier.
6. Select one of the following:
   - **All** if you want the storage tier to apply to all storage systems.
     - Select **Enterprise** or **Departmental** from the Storage System Type box.
   - **Selected** if you want the storage tier to apply to certain storage systems.
     - Select **Vendors**, **Models** or **Systems** from the Display box. If you select Systems from the Display box, select the vendor and model of the storage system you want included in the tier.
     - Select an item in the left pane and click the greater than sign (>) to move it to the right pane. To move all items in the left pane, click the double greater than sign (>>). To remove an item from the right pane, click the less than sign (<).
7. Select either **Leveraged** or **Dedicated** from the Offering box.
8. Select the disk storage attributes you want to be used for apply the rule.
9. Click **Finish**.

Setting a Schedule for Storage Tiers

During a refresh, a storage tier does the following:

- Verifies that its existing members still meet the criteria and removes any members that do not meet the criteria.
- Adds new elements, such as storage systems and volumes, that now meet the criteria of the storage tier.

Manual assignments can be overridden by the auto tier assignments. For example, assume you created a dynamic storage tier that requires its element to have a disk size of more than 900 GB. Then, you manually add an element that has a disk size of less than 900 GB to the storage tier. During the next refresh of rule-based membership, elements that do not fit the criteria for being a member of the storage tier are removed, except for the elements you manually added. The elements you manually added stay members of the storage tier even if they do not meet the criteria of the storage tier. To remove all elements you manually added quickly, use the clear all mappings option on the Dynamic Storage Tiers tab and then refresh the storage tiers.

For a storage tier to refresh its membership, the storage tier must have the following:

- A priority other than 0.
- Storage disk attributes selected. A red checkmark be showing in the Has Attributes Defined? column for the storage tier.

To set a schedule for refreshing the storage tiers:

1. Access the Dynamic Storage Tiers tab (see Accessing the Dynamic Storage Tiers Tab on page 810).
2. Select the **Enabled Schedule** option.
3. Click the calendar icon  .
4. In the Time box, enter the time, in 24-hour format, with the hour and minutes separated by a colon; for example, 22:15. Click the date on which you want the task to run. Today's date is highlighted in pink. Click **Set**.

   The date and time appear in the Next Scheduled Run box in yyyy-mm-dd hh min format.

   If you change the date in the box to a date that does not exist, the software automatically calculates the date to the next month. For example, if you enter 2010-11-31, the software calculates it as 2010-12-01.
5. In the Repeat Interval box, select one of the following units:
   - Second(s)
   - Minute(s)
   - Hour(s)
   - Day(s)
   - Week(s)
6. Click **Save Changes**.

Changing the Priority of a Storage Tier

Storage Tiers with a priority of 1 run first, and elements with a priority of 2 run second. Storage Tiers with a priority of zero do not run.

To change the priority of a storage tier:

1. Access the Dynamic Storage Tiers tab, as described in Accessing the Dynamic Storage Tiers Tab on page 810.
2. Click **Edit Priority**.
3. Modify the priority of a storage tier in the Priority column.
4. Click **Apply Priority**. The priority is set.

Removing Mappings for a Storage Tier

The Clear All Mappings button removes all elements from a storage tier. Use this option when you want to quickly remove all elements that were added manually to the storage tier. The next time you refresh the storage tiers, the elements that meet the criteria of a storage tier are automatically added.

To remove all elements from a storage tier:

1. Access the Dynamic Storage Tiers tab (see Accessing the Dynamic Storage Tiers Tab on page 810).
2. Select the storage tiers from which you want to remove the elements.
3. Click **Clear All Mappings**. All elements are removed from the storage tier.

Setting Custom Properties for Storage Systems

Use the Custom Properties tab to:

- Assign storage systems to a storage system type and/or storage system offering. See Managing the Properties of Storage Systems on next page
- Add or remove options from the Storage System Type and/or Storage System Offering menus. See Editing the Properties of Storage Systems on next page.

Accessing the Custom Properties Tab

To access the Custom Properties tab:

1. Click **Chargeback Manager** ().
2. Click the **Custom Properties** tab.

Managing the Properties of Storage Systems

You can change the options that are provided in the Storage System Type and Storage System Offerings menus.

To add an option to the Storage System Type or Storage System Offering menu:

1. Access the Custom Properties tab (see Accessing the Custom Properties Tab on previous page).
2. Click **Manage Properties**.
3. To add an item to one of the menus:
   - **Storage System Type**: Click **Add Storage System Type**.
   - **Storage Systems Offerings**: Click **Add Offerings**.
4. Enter a new menu option in the text box, and then click **Add**.
5. Click **OK**. The new menu option is added.

To remove an option from the Storage System Type or Storage System Offering menu:

1. Access the Custom Properties tab (see Accessing the Custom Properties Tab on previous page).
2. Click **Manage Properties**.
3. Click the Delete () button for the option you want to delete from the Storage System Type or Storage System Offering menu.

   If a storage system has the option you are deleting, the storage system will no longer be associated with the option. For example, assume you have a storage system that has Departmental for Storage System Type and you remove the Departmental option. The storage system will no longer have a Storage System Type associated with it.
4. Click **OK**. The menu option is removed from the product.

Editing the Properties of Storage Systems

The properties you select for a storage system determine its membership into a storage tier. For example, assume a storage system meets the storage requirements of a storage tier but the storage tier requires the storage system to be of storage system type Enterprise with an offering of leveraged. The storage system cannot be added as a member of the tier unless it also has a storage system type Enterprise with an offering of leveraged.

To edit the property of a storage system:

1. Access the Custom Properties tab, as described in Accessing the Custom Properties Tab on previous page.
2. Select the storage system you want to modify.
3. Click **Edit Storage System**.

4. (*Optional*) Select a storage system type from the Storage System Type menu.
5. (*Optional*) Select an offering from the Offering menu.
6. Click **OK**.

Adding Asset Information

This section contains the following topics:

Adding Asset Information

Chargeback Manager enables you to keep track of your asset information for an element. In addition to warranty and licensing information, you can also store contact information for the element. For example, assume a switch on the network is having some problems, and you want to contact the person in charge of that switch. You can use the element's asset record to find not only the contact information for that switch, but also the location of that switch.

After you add information about all of your assets, back up the database. Backing up the database saves your chargeback information. If the database fails, your asset information is restored when you restore the database (see Performing an RMAN Hot Backup on page 385).

To view and add asset information:

1. Access Chargeback Manager, as described in Accessing Chargeback Manager on page 805.
2. Do one of the following:
   a. Click the  icon corresponding to the element.

      Or

b. Click the link for the element in the right pane.

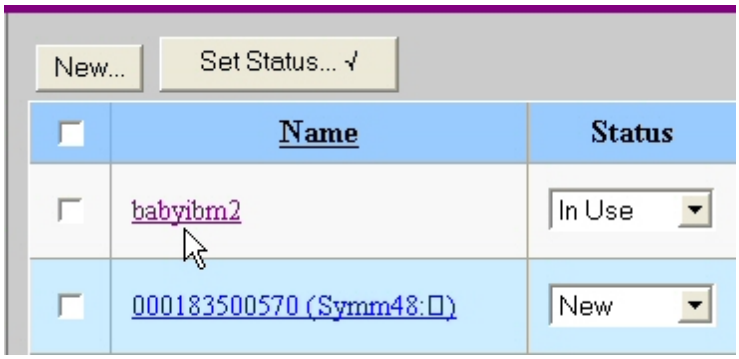

3. To access the different types of asset information, click the **Asset Record** node or one of the nodes under it.

   To view general information about an element, click the **Asset Record** node in the tree. To view specific asset information, such as ownership, click the **Staff** node.

   You can also access the tree for Chargeback Manager from Application Viewer and System Manager:

   a. To access the tree from Application Viewer, click the name of an application in the Application Viewer tree. In the right pane, click the **Asset Management** tab.

   b. To access the tree from System Manager, double-click an element in the topology. In the right pane, click the **Asset Management** tab.

Adding General Information

The management server provides a page for you to enter the following general information about an element. When you are done adding information, click the **Save Changes** button at the bottom of the page. To learn more about Chargeback Manager, see About Chargeback Manager on page 803 and Setting Up Chargeback Manager on page 804.

This page enforces the maximum number of characters you can enter in a box. When you can no longer add additional characters to a box, you have reached the maximum number of characters that can be entered for that box.

- **Custom Name –** A name you assign to the element (see Assigning a Custom Name on page 500).
- **Date Created** – The date the element was discovered.
- **Date Last Modified** – The date the record was last modified.
- **Description** – A description of the element.
- **Status** – The current status of the element. If the status of the element has changed, select the new status from the Status menu.

  - **New** – Default category for all detected elements.
  - **Missing** – No longer detectable through discovery.
  - **Repaired** – The element is being repaired. The software does not automatically select this status.
  - **In Use** – The element is in use.
- **Vendor** – The company that supplied the element.
- **Model** – The model of the element.
- **Serial Number** – The serial number of the element.
- **Barcode Number** – The barcode on the device.
- **Asset Code** – The asset code assigned to the element.
- **Asset Type** – The asset type assigned to the element.
- **Asset Tag** – The asset tag assigned to the element.
- **Asset Category** – The asset category assigned to the element.
- **Geographic Location** – The location of the element; for example, Boston, Massachusetts.
- **(Storage Systems Only) Storage Tier Classification** – Click the **Set Storage Tier Cost** link to set up storage tiers (see Managing Dynamic Storage Tiers on page 808).

Adding Staff Information

This page provides contact information about the element.

Keep this information up to date. Other users need this information to contact you about the element; for example, if it is having problems.

- **Administrator** – The person or department that maintains the element.
- **Staff Name** – The name of the person who maintains the element.
- **Staff Phone Number** – A phone number for the person who maintains the element.
- **Staff Department** – The department that maintains the element.
- **Staff E-Mail** – An e-mail address of the person who maintains the element.
- **Staff #2 Name** – The name of an additional person who maintains the element.
- **Staff #2 Phone Number** – A phone number for an additional person who maintains the element.
- **Staff #2 Department** – An additional department that maintains the element.
- **Staff #2 E-Mail** – An e-mail address of an additional person who maintains the element.

Adding Geographic Information

Use this page to add geographic information about the element. This page is helpful in keeping track of the locations of all your elements, especially if you have more than 100 elements. For example, assume you are told one of your servers is having problems and you need physical access to it. You can use this page to find where the server is located.

- **Rack Number** – The number of the rack that holds the element.
- **Floor** – The floor on which the element is located, for example third floor.
- **Data Center** – The name of the data center where the element is located.
- **Address** – The street address where the element is located.
- **City** – The city where the element is located, for example, Boston.
- **Region** – The region where the element is located, for example, New England.
- **Country** – The country where the element is located, for example, the United States.
- **Continent** – The continent where the element is located, for example, North America.
- **Zip Code** – The ZIP code for the town where the element is located. If your country does not use ZIP codes, leave this box blank.

Adding Licensing and Warranty Information

Use this page to provide licensing and warranty information:

- **License (maximum of 4000 characters)** – The license of the element.
- **Warranty Information (maximum of 4096 characters)** – Information about the warranty. In this box, you probably want to enter information such as how long the warranty lasts and what it covers.
- **Comments (maximum of 4000 characters)** – Any financial information you might want to add about the element

Adding Custom Information

To provide up to six custom properties:

1. In the **Name** field, assign a name for the box. Do not enter more than 50 characters; for example, Backup Contact.
2. In the **Value**field, provide the information for the box. Do not enter more than 255 characters; for example, Joe Smith.
3. Repeat steps 1 and 2 for each custom property you want to add.
4. When you are done, click **Save Changes**.

Managing Departments

This section contains the following topics:

Adding Departments

Before you can assign a department to an element, you must add it to the list of departments, as follows:

1. Access Chargeback Manager, as described in Accessing Chargeback Manager on page 805.
2. Click the **Departments** tab above the table.
3. Click **New**.
4. In the **Add Department** window, provide the following information:
   - Department Name (Required)
   - Department Number (Required)
   - E-mail
   - Phone
5. Click **OK**. The new department is added.

Editing a Department

To edit a department:

1. Access Chargeback Manager, as described in Accessing Chargeback Manager on page 805.
2. Click the **Departments** tab above the table.
3. Click the **Edit** () button corresponding to the department you want to edit.
4. In the Edit Department window, you can edit all boxes except the department number.
5. Click **OK**.

Removing a Department from Chargeback Manager

Over time, some departments in your company might merge, and others might be dissolved. To keep up with these changes, you might need to remove departments from your list. If an element is assigned only to the department that is removed, it no longer has an owner. However, if an element is assigned to this department and several others, it continues to be assigned to the other departments.

Say you want to delete a department called TooSmall. The TooSmall department owns 50 percent of a host, and the Server department owns 50 percent of the host. When you remove TooSmall, the host is owned by the Server department, but only by 50 percent.

To remove a department:

1. Access Chargeback Manager, as described in Accessing Chargeback Manager on page 805.
2. Click the **Departments** tab above the table.
3. Click the **Delete** ( ) button corresponding to the department you want to remove.

Setting Up Asset-Based and Storage-Based Chargeback

You can set up chargeback to be calculated based on the ownership of assets for each department (referred to as asset-based chargeback) or on the amount of storage used by a department or application (referred to as storage-based chargeback). The following topics explain how to set up chargeback.

About Asset-Based and Storage-Based Infrastructure Cost

Asset-based and storage-based infrastructure cost are optional chargeback calculations that you can specify on the Ownership page. Setting an infrastructure cost calculates a monthly infrastructure charge, which is identical for each department and is applied once each month on top of the department's total ownership cost. Modifying the infrastructure charge impacts the asset-based or storage-based chargeback result for all department owners. Asset-based infrastructure cost is added to the total asset-based chargeback calculation results. Storage-based infrastructure cost is added to the total storage-based chargeback calculation results.

Setting Up Asset-Based Chargeback

Asset-based chargeback calculates chargeback based on the departmental ownership percentages and the depreciated value of the assets. Each piece of equipment is owned by a department or a set of departments. Each department has a percentage ownership of the equipment.

The management server calculates monthly chargeback using the financial information you provide. You can then use these calculations to determine the cost impact on your enterprise on a monthly basis. For reporting purposes, you can also break the cost down by department. Also, if you have an infrastructure cost, you can add that into the calculations.

To set up asset-based chargeback:

1. Specify Financial information, as described in Step 1 – Specify Financial Information on next page.
2. Assign a Departmental Ownership Percentage, as described Step 2 – Assign Departmental Ownership Percentage on page 827.
3. Review the asset-based chargeback result, as described in Step 3 – Review Asset-Based Chargeback Result on page 828.

Note: You must have already added your departments, as described in Editing a Department on page 823.

To provide information for asset-based chargeback:

1. Access Chargeback Manager, as described in Accessing Chargeback Manager on page 805.
2. Do one of the following:
   - Click the  icon corresponding to the element.

     Or
   - Click the link for an element in the right pane, as shown in the following figure.

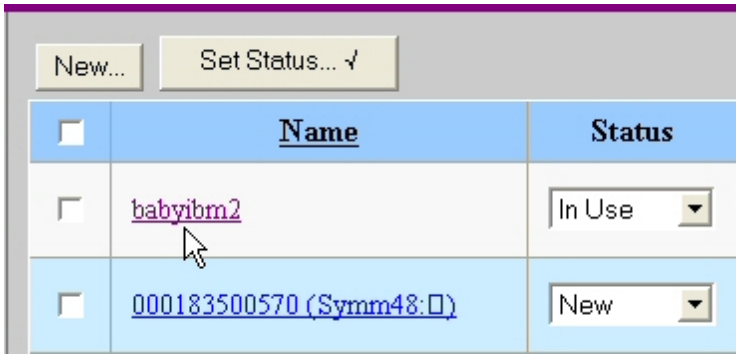

3. Click the **Asset-based** node under the Chargeback Manager node, as shown in the following figure.

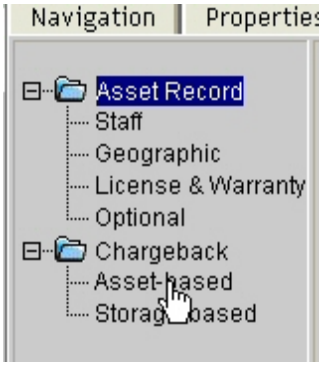

You can also access the tree from Application Viewer and System Manager.

- To access the tree from Application Viewer, click the name of an application in the Application Viewer tree. In the right pane, click the **Asset Management** tab.
- To access the tree from System Manager, double-click an element in the topology. In the right pane, click the **Asset Management** tab.

Step 1 – Specify Financial Information

1. Verify that **Step 1 – Specify Financial information** is selected at the top of the page.
2. Provide the following financial information:
   - **Purchase Order Number** – The purchase order of the element.
   - (*Required*) **Date Purchased** – The date the element was purchased. To select the date, click the calendar icon . If you select a future date, the purchase date is set to today when calculating depreciation. The management server only supports dates within the years 1900 through 3000.
   - **Reseller** – The company that directly sold you the element.
   - (*Required*) **Purchase Price** – The price of the element when it was bought.
   - **Salvage Value –** The amount of money an item is worth for salvage value. You cannot go below this number when depreciating an item.
   - (*Required*) **Depreciation Period (months)** – The time period in which you plan to keep the element.
   - (*Required*) **Depreciation Method** – How the depreciation is calculated. Select one of the following:
     - **Straight Line** – The device loses the same amount of value in each period. To learn more about how the management server calculates straight-line depreciation, see Calculating Straight Line Depreciation on page 834.
     - **Fixed Declining Balance** – Calculates depreciation based on the value of the asset each month instead of a fixed rate like straight line depreciation. To learn

more about how the management server calculates fixed declining balance, see Calculating Fixed Declining Balance on page 835.

- **Double Declining Balance** – Doubles the calculation of the Fixed Declining Balance method, thus doubling the speed at which a device depreciates. To learn more about this, see Calculating Double Declining Balance on page 837.

- **Value as of** – The value of the element as of the end of the previous month. For example, if you enter or view this chargeback information in the middle of March. the value is for the month of February, and does not include March. The value is calculated from the following boxes:
  - Date Purchased
  - Original Cost ($)
  - Depreciation Salvage Value ($)
  - Depreciation Period
  - Depreciation Method

3. Click **Save Changes**.

Step 2 – Assign Departmental Ownership Percentage

To assign the percentage of ownership to an element and the monthly infrastructure charge:

1. Select the option **Step 2 – Assign Departmental Ownership Percentage** at the top of the page.
2. Click **Add Ownership**.
3. Select a department from the **Department** menu.
4. Enter the percentage of ownership in the Ownership % box. If you do not see a department listed, add it to the list as described in Adding Departments on page 823. Click the **Manage Departments** link. After you have added the department, close the window you used to add the department and then refresh the page.
5. Click **OK**. The department with its percentage of ownership is added to the table.
6. If multiple departments own the element, repeat steps 1 through 5 for each department. You can have departments owning more than 100 percent of the element.
7. (*Optional*) Specify a monthly infrastructure charge for when asset-based calculation is being done. For details, see About Asset-Based and Storage-Based Infrastructure Cost on page 824.
   a. Click **Set Infrastructure Cost**.
   b. Enter the monthly infrastructure charge.
   c. Click **OK**.
8. When you are done with assigning the element to a department, click **Save Changes**.

Step 3 – Review Asset-Based Chargeback Result

The management server displays chargeback information up to the end of the previous month. For example, if you are viewing chargeback information in the middle of March, the calculations include the month of February, but not March.

To view the result of asset-based chargeback:

1. Select **Step 3 – Review Asset-based Chargeback Result** at the top of the page.
2. If you see empty values, verify that you provided the required values.

Ownership cost is determined by the following formula:

```
(Depreciation) x (Ownership %) = Ownership Cost
```

Ownership Cost reflects how much it costs a department to own the element. Depreciation is determined by the depreciation method you selected in Step 1 – Specify Financial Information on page 826.

Infrastructure cost is not included in ownership cost because the information is per asset. The asset-based infrastructure cost is a monthly charge that is applied to each departmental owner in addition to any ownership charges. The infrastructure cost is not included when you view the chargeback for individual elements, but is displayed when you view chargeback per department. The infrastructure cost is added to each department, regardless of whether the department owns one or 100 elements.

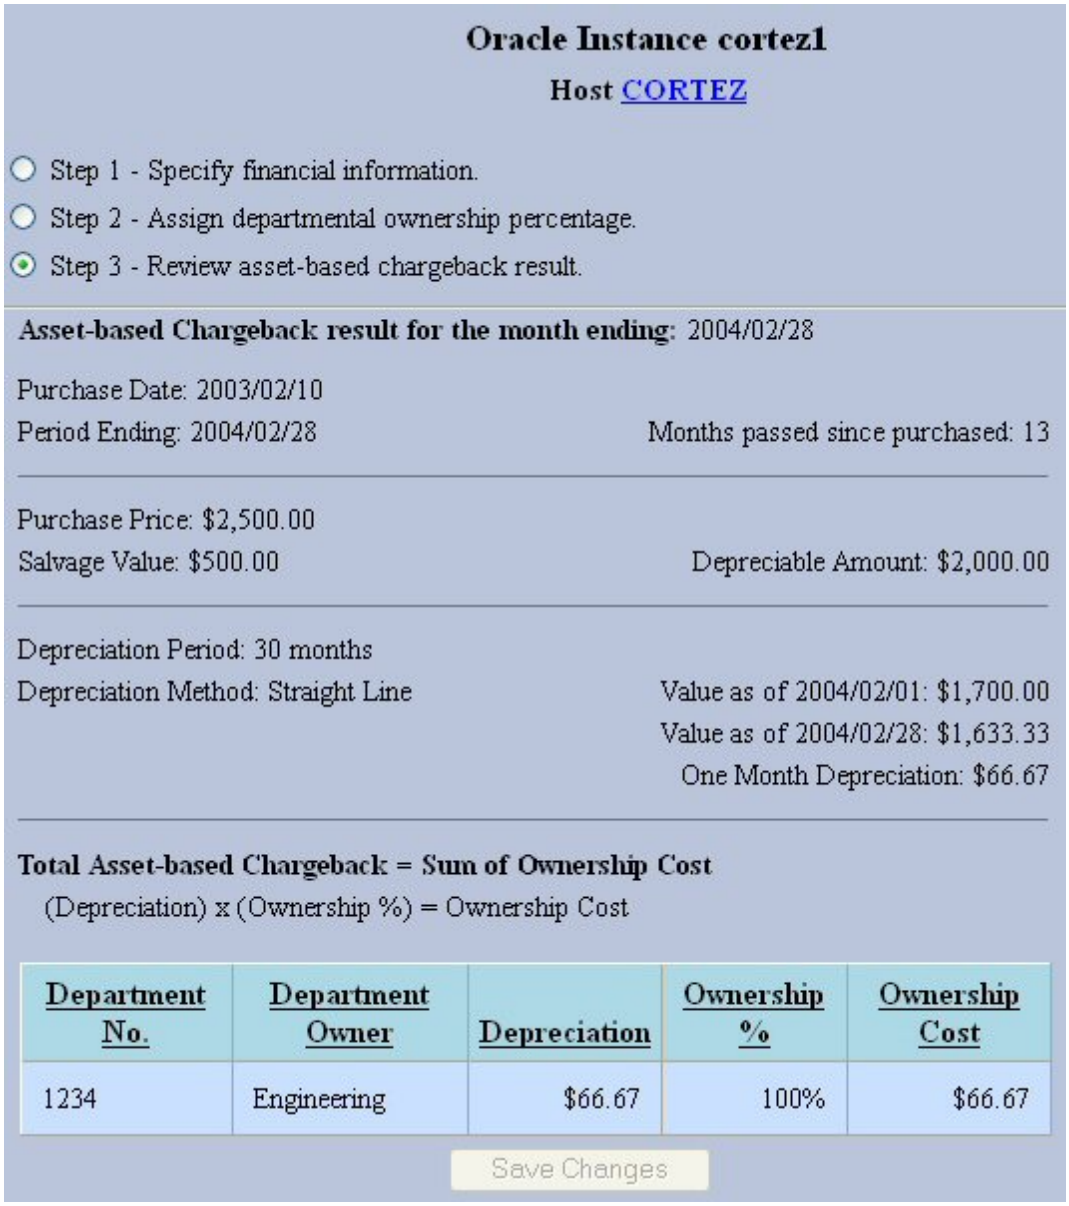

Setting Up Storage-Based Chargeback

Storage-based chargeback calculates charges according to the actual amount of storage used by an application on the storage system, the type of storage it is using, and the ownership percentage assigned to each department. The chargeback number is further refined by an additional fixed infrastructure tax on a per-department basis.

Obtaining storage-based chargeback requires completing the following tasks:

- Assign Departmental Ownership Percentage, as described in Step 1 – Assign Departmental Ownership Percentage on page 831.
- Review Storage Tier Cost, as described in Step 2 – Review Storage Tier Cost on page 832.

- Review Storage Dependency and Cost, as described in Step 3 – Review Storage Dependency and Cost on page 832.
- Review storage-based chargeback result, as described in Step 4 – Review Storage-Based Chargeback Result on page 832.

Keep in mind the following:

- You must have already added your departments, as described in Adding Departments on page 823.
- Not all applications use storage on storage systems. Storage-based chargeback is applicable only for those applications that use storage on storage systems.
- You must have access to the storage system the application uses. Verify that your organization and roles allow you access (see Security for the Management Server on page 287).
- Chargeback Manager provides only network capacity. If you look at the capacity of an application in Capacity Manager, the capacity differs. Capacity Manager displays the total capacity of an application, including the network drives (see How Capacity Differs in Chargeback Manager and Capacity Manager on page 833).

To provide information for storage-based chargeback:

1. Access Chargeback Manager.
2. Do one of the following:
   - Click the  icon corresponding to the element.

     Or
   - Click the link for a host running an application in the right pane, as shown in the following figure.

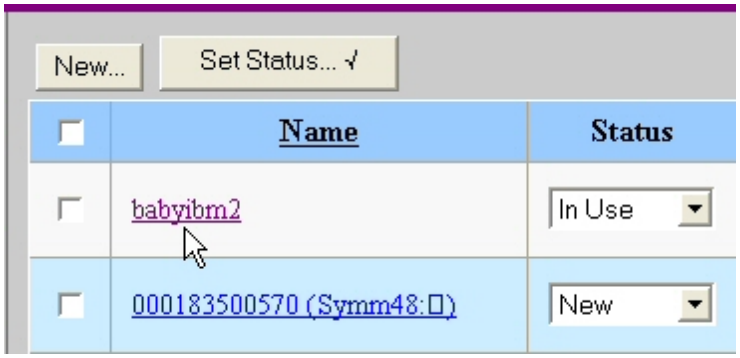

3. Click the **Storage-based** node under the Chargeback Manager node, as shown in the following figure.

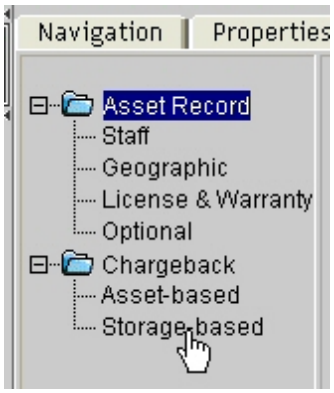

You can also access the tree from Application Viewer and System Manager.

- To access the tree from Application Viewer, click the name of an application in the Application Viewer tree. In the right pane, click the **Asset Management** tab.
- To access the tree from System Manager, double-click an element in the topology. In the right pane, click the **Asset Management** tab.

Step 1 – Assign Departmental Ownership Percentage

1. Select the option **Step 1 – Assign Departmental Ownership Percentage** at the top of the page.
2. Click **Add Ownership**.
3. Select department from the **Department** combo-box.

   If you do not see a department listed, add it to the list as described in Adding Departments on page 823. Click the **Manage Departments** link. After adding the department, close the window you used to add the department, and refresh the page.
4. Type the percentage of ownership in the Ownership % box.
5. Click **OK**. The department, with its percentage of ownership, is added to the table.
6. If multiple departments own the element, repeat the previous steps for each department. You can have departments owning more than 100 percent of the element.
7. *(Optional)* Specify a monthly infrastructure charge. For details, see About Asset-Based and Storage-Based Infrastructure Cost on page 824.
   a. Click **Set Infrastructure Cost**.
   b. Type the monthly infrastructure charge.
   c. Click **OK**.
8. When you are done with assigning the element to a department, click **Save Changes**.

Step 2 – Review Storage Tier Cost

Storage-based chargeback enables you to charge the application owners according to the amount of storage allocated to them. Each storage system is designated a storage tier classification. You can specify a charge for each storage tier. This charge is referred to as storage tier cost. The management server determines the storage cost of the application by multiplying the storage tier cost by the allotted storage. If the application uses more than one storage system, the storage cost from each storage system is added for the total storage cost.

Click the **Storage Tiers** link to review your storage tier costs.

Step 3 – Review Storage Dependency and Cost

The management server displays chargeback information up to the end of the previous month. For example, if you are viewing chargeback information in the middle of February, the calculations include the month of January, but not February.

Click the **Storage/volume dependency and cost details** link to view the details.

The Storage Dependency for Application table is displayed when both of the following conditions apply:

- The application depends on a storage system.
- The organizations to which you belong allow you to view the storage system.

If the table is empty and you know the application is dependent on a storage system, verify that you have access to the storage system; otherwise, data cannot be calculated.

The details are provided in a tree table. Expand the various nodes to drill down into the application cost and examine how the storage systems, storage volumes, storage pools, and assigned tiers are contributing to the total cost.

Step 4 – Review Storage-Based Chargeback Result

The management server displays chargeback information up to the end of the previous month. For example, if you are viewing chargeback information in the middle of February, the calculations include the month of January, but not February.

The total storage-based monthly chargeback is displayed. This number is calculated as follows:

```
(Total Storage Cost) x (Ownership %) = Ownership Cost
```

In this instance:

- Total Storage Cost is the total Storage Cost from (Step 3 – Review Storage Dependency and Cost).
- Ownership % is the percentage of ownership.

Editing Percentage of Ownership

To edit the department ownership of an asset:

1. Access Chargeback Manager.
2. Click the link for the element in the right pane.
3. Click **Asset-based** under the Chargeback node in the tree.
4. Verify that the option Step 2 – Assign Departmental Ownership Percentage is displayed in the right pane.
5. Click the **Edit** ( ) button corresponding to the percentage of ownership you want to modify.
6. In the Ownership % box, enter a new percentage of ownership.
7. Click **Save Changes**.

Removing Department Ownership of an Element

Sometimes you must remove ownership from an element; for example, when an element is moved from one department to another. When department ownership is removed from an element, the department is still accessible from the list of departments. To make the department inaccessible to all elements, remove it from the list of departments as described in Removing a Department from Chargeback Manager on page 823.

To remove ownership:

1. Access Chargeback Manager.
2. Click the link for the element in the right pane.
3. Click **Asset-based** under the Chargeback node in the tree.
4. Verify that the option Step 2 – Assign Departmental Ownership Percentage is displayed in the right pane.
5. Click the **Delete** ( ) button corresponding to the department you want to remove.
6. Click **Save Changes**. The department is removed.

How Capacity Differs in Chargeback Manager and Capacity Manager

The capacity displayed for an application in Chargeback Manager differs from the capacity displayed in Capacity Manager. The management server uses only network storage when calculating chargeback capacity. Local capacity is not included. The following figure shows the chargeback capacity for an Oracle instance named RETAIL. Notice that chargeback capacity is 0.89 GB.

Storage System	Mounted Storage	Unmounted Storage	Total Storage
000183500570 (Symm48:3830)	0.89 GB	0 GB	0.89 GB

If you view the Oracle instance RETAIL in Capacity Manager, you see local and network capacity, which is a total of 1,042 MB in the following example. Of the 1,042 MB, 133 MB is on a local drive. The rest (909 MB) is on a network drive. When you convert 909 MB to gigabytes (0.887 GB) and round the output (0.89 GB), the capacity in Capacity Manager matches the number in Chargeback Manager.

RETAIL

Database Instance Name	Total Capacity
RETAIL	1,042 MB

Database Files	Total Capacity
INDX	58 MB
RBS	520 MB
RETAILSPACE	5 MB
SYSTEM	264 MB
TEMP	72 MB
TOOLS	12 MB
USERS	108 MB
RedoGroup 1	1 MB
RedoGroup 2	1 MB
RedoGroup 3	1 MB

How a Depreciation Method Is Calculated

Depreciation is calculated in the following ways.

Calculating Straight Line Depreciation

When the management server calculates straight line depreciation, it calculates depreciation based on a fixed rate. These instructions describe how the management server performs the straight line depreciation calculation. An example is provided for each step, so that you can try the calculations for yourself.

The management server calculates straight line depreciation as follows:

1. The management server rolls back the purchase date to the beginning of the purchase month. If the purchase date is later than today (for example, a future purchase), then the purchase date is rolled back to today.

   Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchase date to January 1, 2003, when calculating months to depreciate.

2. It determines the period ending date. This is equivalent to the last day of the previous full month.

   Example: Assume today's date is January 9, 2004. The management server sets the period ending to December 31, 2003.

3. The management server calculates the delta between the purchase date and the period ending date. This determines how many month’s worth of depreciation amount the

management server need to take into account.

Example: Using the examples from the previous two steps, the delta is 12 months (January 1, 2003 through December 31, 2003).

4. It subtracts the salvage value from the purchase price. This is the depreciable amount.

   Example: Assume the purchase price for the element is $2500, and the Salvage Value is $100. The depreciable amount is $2400, which was calculated by subtracting the Salvage Value ($100) from the purchase price ($2500).

5. It takes the depreciable amount and divides it by the depreciation period (the number of months it takes for the asset to fully depreciate to either 0 or salvage value). This gives us the depreciation for a single month.

   Example: Let's use the depreciable amount ($2400) calculated in the previous step. Let's assume the depreciation period is 24 months. Divide $2400 by 24. The result is $100, which is the one month depreciation.

6. It multiplies the depreciation for a single month by the delta from step 3. This is the total depreciation.

   Example: To find the total depreciation, multiply the one-month depreciation from the previous step ($100) by the delta (12 months), which was calculated in Step 3. The result of $100 x 12 months is $1,200, which is the total depreciation.

7. To determine the value as of the end of last month, the management server simply subtracts the total depreciation from the purchase price.

   Example: Subtract the total depreciation ($1200), which was calculated in the last step, from the purchase price ($2500), which was provided in Step 4. The value as of the end of last month is $1300.

Calculating Fixed Declining Balance

The Fixed Declining Balance method calculates depreciation based on the value of the asset each month, instead of at a fixed rate like straight line depreciation.

These instructions describe how the management server performs the fixed declining balance calculation. An example is provided for each step, so that you can try the calculations for yourself.

The management server calculates fixed declining balances as follows:

Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchase date to January 1, 2003 when calculating months to depreciate.The management server rolls back the purchase date to the beginning of the purchase month. If the purchase date is later than today (for example, a future purchase), then the purchase date is rolled back to today.

1. It determines the period ending date. This is equivalent to the last day of the previous full month.

Example: Assume today's date is January 9, 2004. The management server sets the period ending to December 31, 2003.

2. The management server calculates the delta between purchase date and the period ending. This determines how many months worth of depreciation amount the management server needs to take into account.

   Example: Using the examples from the previous two steps, the delta is 12 months (January 1, 2003 – December 31, 2003).

3. The management server takes the user-specified depreciation period and use it as the life of the asset.

   Example: Let's assume the depreciation period is 24 months and that it is also the life of the asset.

4. The management server calculates the declining ratio using this formula: (1.0 / life). This determines the rate at which depreciation should occur each month.

   Example: Use the example from step 3 (24 months) in the following formula to find the rate of depreciate per month:

   ```
   1.0/24
   ```

   The depreciation ratio is 0.042.

5. For each month identified by delta from Step 2, the management server calculates the following:

   The example for the following steps can be found at the end of these instructions.

   a. Determine the “would-be” depreciation for the month. This means multiplying the asset value for the month by the declining ratio from step 4.

   b. Subtract the depreciation for the month from the asset value for the month. If the result is less than the salvage value, it means the asset value after depreciation would be less than the salvage. In this case, the management server simply depreciates the asset to the salvage value. Once the management server depreciates an asset down to its salvage value, the depreciation for that asset stops.

   c. If the management server subtracts the depreciation for the month from the asset value and the result is greater than the salvage value, then the management server knows it is safe to depreciate the asset by the depreciation amount calculated in step a. The depreciated asset value for the month would be asset value minus depreciation. The new asset value will be used to compute the depreciation for the next month. This process continues until one of the following occurs:

      - The management server has depreciated the asset value for the number of months equal to delta.
      - The asset value has depreciated down to the salvage cost. If no salvage cost is specified, then the asset value has depreciated down to 0.

Example: For Step 5, we complete Steps a through c for the first month and then repeat these steps for the second month.

Step 5a – We assume the asset value of the element is $2500. Calculate the “would-be” depreciation of the month by multiplying the asset value by the declining ratio from Step 4 (0.042):

```
$2500 x .042 = $105
```

Step 5b – Assume the salvage value is $100. Determine if the asset value after depreciation is less than the salvage value by using the following formula.

Asset value of the month ($2500) - Depreciation for the month ($105) = $2395

Since $2395 (the depreciated asset value) is greater than the salvage value ($100), the asset value after depreciation is $2395. Go to Step 6c.

Step 5c – The new asset value ($2395) is used to calculate the depreciation for the next month. Let's go through the calculations for the next month.

Step 5a – Assume the asset value of the element is $2395. Calculate the “would-be” depreciation of the month by multiplying the asset value by the declining ratio from Step 4 (0.042):

```
$2395 x .042 = $100.59
```

Step 5b – Assume the salvage value is $100. Determine if the asset value after depreciation is less than the salvage value by using the following formula:

Asset value of the month ($2395) - Depreciation for the month ($100.59) = $2294.41

Since the $2294.41 (the depreciated asset value) is greater than the salvage value ($100), the asset value for the month is $2294.41. Go to Step 5c. The management server repeats Steps 5a through 5c for 12 months (the delta from Step 2), unless the depreciated asset value reaches the salvage value, or 0 if the salvage value is not specified.

Calculating Double Declining Balance

The Double Declining Balance method and the Fixed Declining Balance are very similar. The difference is that instead of using the depreciation ratio determined by (1.0 / life), the management server doubles the ratio to increase the rate of depreciation. This provides for a more realistic depreciation when your asset tends to lose its value in the early part of its life. For instance, a new car’s blue book value decreases dramatically once it is sold and driven off the lot of the car dealership.

These instructions describe how the management server performs the double declining balance calculation. An example is provided for each step, so that you can try the calculations for yourself.

1. The management server rolls back the purchase date to the beginning of the purchase month. If the purchase date is later than today (for example, a future purchase), then the purchase date is rolled back to today.

Example: Assume the purchase date of an element is January 15, 2003. The management server adjusts the purchase date to January 1, 2003, when calculating months to depreciate.

2. It determines the period ending date. This is equivalent to the last day of the previous full month.

   Example: Assume today's date is January 9, 2004. The management server sets the period ending to December 31, 2003.

3. The management server calculates the delta between purchase date and the period ending. This determines how many months worth of depreciation amount the management server need to take into account.

   Example: Using the examples from the previous two steps, the delta is 12 months (January 1, 2003 – December 31, 2003).

4. The management server takes the user-specified depreciation period and uses it as the life of the asset.

   Example: We assume the depreciation period is 24 months and that it is also the life of the asset.

5. The management server calculates the declining ratio using this formula: (1.0 / life)*2. This determines the rate at which depreciation should occur each month.

   Example: Use the example from step 4 (24 months) in the following formula to find the rate of depreciation per month:

   ```
   (1.0/24)*2
   ```

   The depreciation ratio is 0.084.

6. For each month identified by delta from Step 3, the management server calculates the following:

   The example for the following steps can be found at the end of these instructions.

   a. Determine the “would-be” depreciation for the month. This means multiplying the asset value for the month by the declining ratio from step 5.

   b. Subtract the depreciation for the month from the asset value for the month. If the result is less than the salvage value, it means the asset value after depreciation would be less than the salvage. In this case, the management server simply depreciate the asset to the salvage value. Once the management server depreciates an asset down to its salvage value, the depreciation for that asset stops.

   c. If the management server subtracts the depreciation for the month from the asset value, and the result is greater than the salvage value, then the management server knows it is safe to depreciate the asset by the depreciation amount calculated in step a. The depreciated asset value for the month would be asset value minus depreciation. The new asset value will be used to compute the depreciation for next month. This process continues until one of the following occurs:

- The management server has depreciated the asset value for the number of months equal to delta.
- The asset value has depreciated down to the salvage cost. If no salvage cost is specified, then the asset value has depreciated down to 0.

Example: For Step 6, let's complete Steps a through c for the first month and then repeat these steps for the second month.

Step 6a – Let's assume the asset value of the element is $2500. Calculate the “would-be” depreciation of the month by multiplying the asset value by the declining ratio from Step 5 (0.084):

```
$2500 x 0.084 = $210
```

Step 6b – Assume the salvage value is $100. Determine if the asset value after depreciation is less than the salvage value by using the following formula.

Asset value of the month ($2500) - Depreciation for the month ($210) = $2290

Since $2290 (the depreciated asset value) is greater than the salvage value ($100), the asset value after depreciation is $2290. Go to Step 6c.

Step 6c – The new asset value ($2290) is used to calculate the depreciation for the next month. Let's go through the calculations for the next month.

Step 6a – Assume the asset value of the element is $2290. Calculate the “would-be” depreciation of the month by multiplying the asset value by the declining ratio from Step 5 (0.084):

```
$2290 x .084 = $192.36
```

Step 6b – Assume the salvage value is $100. Determine if the asset value after depreciation is less than the salvage value by using the following formula:

Asset value of the month ($2290) - Depreciation for the month ($192.36) = $2097.64

Since the $2097.64 (the depreciated asset value) is greater than the salvage value ($100), the asset value for the month is $2097.64. Go to Step 6c. The management server repeats Steps 6a through 6c for 12 months (the delta from Step 3), unless the depreciated asset value reaches the salvage value, or 0 if the salvage value is not specified.

Viewing Chargeback

If you see empty values, make sure chargeback was set up as described in Setting Up Asset-Based Chargeback on page 824 and Setting Up Storage-Based Chargeback on page 829.

This section contains the following topics:

Viewing Chargeback by Asset

You can view chargeback for an asset by clicking the icon next to an element listed on the Asset tab.

The management server displays chargeback information up to the end of the previous month. For example, if it is the second week of February and you want to view chargeback information, the calculations for chargeback would include the month of January, but not February.

To view chargeback by asset:

1. Access Chargeback Manager as described in Accessing Chargeback Manager on page 805.
2. Click the icon next corresponding to the element for which you want to view chargeback. Asset-based chargeback is displayed.
3. (Applications only) To view storage-based chargeback, click **Storage-based** under the Chargeback node to the left of the Asset-based chargeback information.

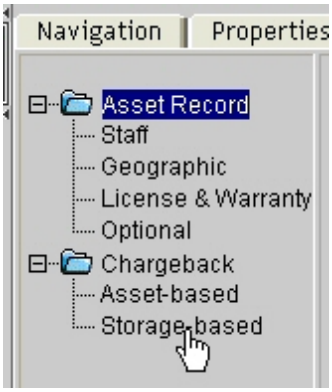

Chargeback information for the element is displayed.

Viewing Chargeback by Department

You can determine how much a department is being charged for equipment use by viewing chargeback by department. This feature lets you view the monthly costs associated with using hardware and applications.

The management server displays chargeback information up to the end of the previous month. Therefore, if you are viewing chargeback information in the middle of February, the calculations for chargeback include the month of January, but not February.

To view chargeback by department:

1. Access Chargeback Manager as described in Accessing Chargeback Manager on page 805.
2. Click the **Departments** tab in the right pane.
3. Click the icon corresponding to the element for which you want to view chargeback.

   You are shown the following information:

- **Department Name** – Provided when the department was added.
- **Department Number** – Provided when the department was added.
- **E-mail** – Can be blank if information was not provided.
- **Phone** – Can be blank if information was not provided.
- **Monthly Infrastructure Cost for Asset** – How much it costs to operate the element on a monthly basis.
- **Ownership Cost** – How much it costs the department in operating the element
- **Total Asset-based Chargeback** – How much it costs the department in operating the element. This number is based on the following formula:

  ```
  (Monthly Infrastructure cost) + ((Depreciation) x (Ownership %))
  ```

- **Monthly Infrastructure Cost for Storage-based Chargeback** – How much it costs for an application to use a specified amount of storage.
- **Total Storage-based Chargeback** – How much it costs the department for an application to use a specified amount of storage:

  ```
  Monthly Infrastructure Cost + Ownership Cost
  ```

  In this instance, Ownership Cost is (Ownership %) x (Storage Cost)

  This page also displays two tables:

  - Asset-Based Chargeback lists the asset, depreciation, ownership percentage, and ownership cost
  - Storage-Based Chargeback lists the application, storage allotted, storage used, storage cost, ownership %, and ownership cost. The storage allotted value includes mounted and unmounted storage. Any volumes the application can access are included in the storage calculations.

Viewing Chargeback by Owner

You can view chargeback for all elements by using the Ownership tab. The Ownership tab shows the ownership distribution across different departments and helps you to quickly identify the assets without a department owner.

To view chargeback by owner:

1. Access Chargeback Manager.
2. Click the **Ownership** tab in the right pane.
3. Select one of the following from the Chargeback Method menu:
   - **Asset-based** – Displays chargeback information for assets.

Or

- **Storage-based** – Displays chargeback information for storage (applications only).

The management server displays asset-based or storage-based chargeback information based on your selection. The management server displays chargeback information from the previous month. For example, assume you view chargeback information in the middle of February. The calculations for chargeback would include the month of January, but not February.

You can sort elements according to a column heading. Click a column heading in the table to sort the data. The arrow next to a column heading indicates whether the items are being sorted in ascending or descending order. If the arrow is pointing up, items are sorted in ascending and alphabetical order. If the arrow is pointing down, items are sorted in descending and reserved alphabetical order.

The following information is displayed:

- **Department Name** – The department that owns the element. This information was provided when the department was added.
- **Department Number** – The number of the department that owns the element. This information was provided when the department was added.
- **Application** or **Asset Name** – The name of the associated element or application.
- **Vendor** – The company that supplied the element.
- **Serial Number** – The serial number of the element.
- **Ownership Percentage** – The percentage of the element that the department owns.
- *Storage-Based Only:***Storage (GB)** – The amount in gigabytes that the application uses. This value includes mounted and unmounted storage. Any volumes the application can access are included in the storage calculations.
- *Storage-Based Only:***Storage Cost** – How much it costs to run the storage that the application uses.
- *Storage-Based Only:***Partition Size** – The partition size used.
- **Ownership Cost** – How much it costs the department to use the asset.

Chargeback Information for HP P4000 Devices

HP P4000 devices can be used with Chargeback Manager to track costs associated with applications and hosts using storage presented by a clustered device. To ease the association of a particular cost with a set of volumes, you can create dynamic storage tiers by RAID type. The following example shows a monthly cost of $200 GB being assigned to all volumes on devices with a model of “VSA” or “NSM2060G2”. This could be further refined by selecting a particular RAID level.

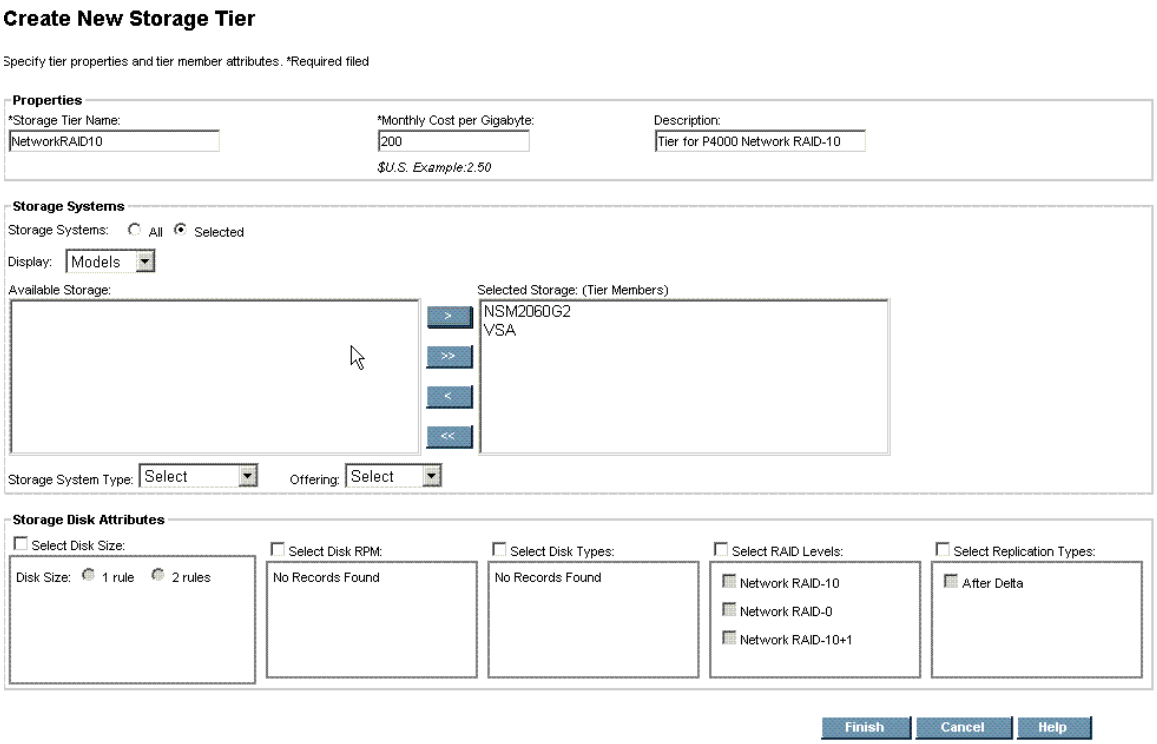

Once this is set, any new devices or volumes discovered through Get Details are added to the storage tier when the storage tiers are refreshed.

After creating storage tiers, the storage dependency of a particular application can be examined and billed on a monthly basis.

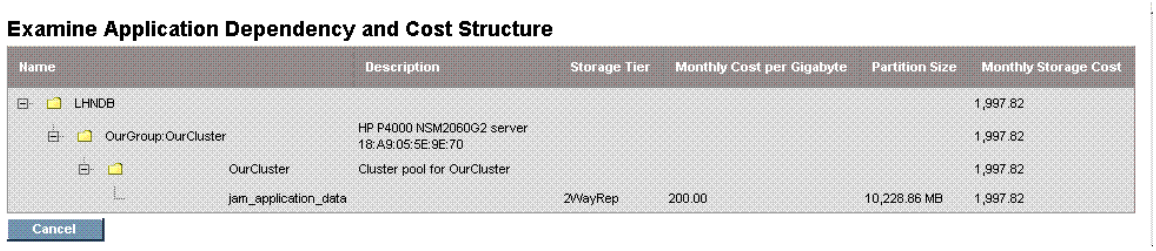

The example shows that Oracle database LHNDB consumes approximately 10 GB of two-way replicated storage (2WayRep) at a monthly cost of $1997.82.

Filtering Assets

This section contains the following topics:

- About Filtering Assets on next page
- Selecting an Element Type for Chargeback on next page
- Customizing the Element Type Filter on page 845

- Filtering Assets by Status on the facing page
- Customizing the Asset Status Filter on page 846
- Hiding Filters in Chargeback Manager on the facing page

About Filtering Assets

The management server provides several types of filters to specify which assets you want Chargeback Manager to display.

You can use all the filters at the same time, or only one of them. You can filter assets by:

- Status
- Element type

For example, if you need a host to install an application, but do not know which hosts are in use, you can set the filters so that only hosts with a status of In Use are displayed. You can then click the element to find the contact information for the owner.

You can also use the filters to find which elements are missing or repaired by doing the following:

1. Set the filter to display only hosts by selecting the **Host** option from the Show Element Type combo box.
2. Click **Custom**.
3. Verify that only the Missing and Repaired filters are selected.
4. Click **OK**.
5. After you set all your filters, click **Apply Filters**.

Selecting an Element Type for Chargeback

You can filter by element type, so that only certain types of elements are displayed. For example, you can specify that only hosts are displayed.

To filter by element type, select an option from the **Show Element Type** menu in Chargeback Manager. When asked if you want to apply your changes, if you want to apply them now, click **Yes**. To apply them at a later time, click **No** and **Apply Filters** when you are ready for your changes to take effect. Chargeback Manager displays only the elements you specified in your filter.

Element Type	Description
Application	Displays only applications, such as Microsoft Exchange and Oracle.
Host	Displays only hosts.
Switch	Displays only switches.
Storage System	Displays only storage systems.
All	Lists all elements.

Filtering Assets by Status

You can filter an asset by status, so that only certain assets of a specified status are displayed. For example, you can specify that only assets in use are displayed.

To filter by asset status:

1. Select an option from the **Show Element Type** menu in Chargeback Manager:
   - **All (Default)** – All assets are displayed.
   - **New –** Only assets with the status of New are displayed.
   - **Missing** – Only assets with the status of Missing are displayed.
   - **In Use** – Only assets with the status of In Use are displayed.
2. When asked if you want to apply your changes, and you want to apply them now, click **Yes**. To apply them at a later time, click **No** and then click **Apply Filters** for your changes to take effect. Chargeback Manager displays only the elements you specified in your filter.

Hiding Filters in Chargeback Manager

Hide the filters for additional screen space. When you hide the filters, the following features are hidden:

- Show Element Type
- Show Status

To hide the filters, click the **-Filters** link in the upper-left corner of Chargeback Manager.

To display the filters, click the **+Filters** link in the upper-left corner of Chargeback Manager.

Customizing Filters

You can customize your filters in Chargeback Managerso that you see only the elements and events you want.

For information about customizing filters, see:

- Customizing the Element Type Filter below
- Customizing the Asset Status Filter on next page

Customizing the Element Type Filter

You can customize the element type for your filter by clicking the **Customize** button next to the **Show Element Type** menu in Chargeback Manager.

For example, you can specify you want only hosts and switches displayed in Chargeback Manager.

To select more than one element for filtering:

1. Click the **Custom** button next to the **Show Element Type** menu in Chargeback Manager.
2. Select the element types you want displayed in Chargeback Manager.
3. Click **OK**.
4. When asked if you want to apply your changes, and you want to apply them now, click **Yes**. To apply them at a later time, click **No** and then click **Apply Filters** for your changes to take effect. Chargeback Manager displays only the elements you specified in your filter.

Customizing the Asset Status Filter

You can filter multiple assets by clicking the **Customize** button next to the **Show Status** menu in Chargeback Manager. For example, you can specify you want only assets that are missing displayed in Chargeback Manager.

To select more than one asset for filtering:

1. Click the **Custom** button next to the **Show Status** menu in Chargeback Manager.
2. Select the statuses you want displayed in Chargeback Manager.
3. Click **OK**.
4. When asked if you want to apply your changes, and you want to apply them now, click **Yes**. To apply them at a later time, click **No** and then click **Apply Filters** for your changes to take effect. Chargeback Manager displays only the elements you specified in your filter.

20 Troubleshooting

This section contains the following topics:

Troubleshooting the Web Browser

This section provides information about troubleshooting issues seen with the Web browser.

Receiving HTTP ERROR: 503 When Accessing the Management Server

If you receive a message resembling the following when you try to access the management server, make sure your database for the management server is running. If it is not, start the database.

```
Receiving HTTP ERROR: 503 javax.ejb.EJBException: null;
```

Windows

In the Services window, make sure the OracleOraHome11gR2TNSListener service has started and is set to automatic. For information on how to access the Services window, see the Windows documentation.

If the OracleOraHome11gR2TNSListener service has not started, but the AppStorManager service has started, start the OracleOraHome11gR2TNSListener service, and then restart AppStorManager.

UNIX

To verify that the Oracle service started, enter the following at the command prompt:

```
# ps -ef | grep ora
```

If the service started, output similar to the following is displayed:

```
/opt/oracle/product/9.2.0.1.0/bin/tnslsnr LISTENER -inherit

 ./appstormservice /opt/productname/ManagerData/conf/unix-wrapper.

  oracle   356     1  0   Jul 30 ?        0:01 ora_pmon_APPIQ

  oracle   358     1  0   Jul 30 ?        0:26 ora_dbw0_APPIQ

  oracle   360     1  0   Jul 30 ?        1:13 ora_lgwr_APPIQ

  oracle   362     1  0   Jul 30 ?        0:39 ora_ckpt_APPIQ

  oracle   364     1  0   Jul 30 ?        0:10 ora_smon_APPIQ

  oracle   366     1  0   Jul 30 ?        0:00 ora_reco_APPIQ

  oracle   368     1  0   Jul 30 ?
```

To start the service for Oracle, enter the following at the command prompt:

```
# /etc/rc3.d/S98dbora start
```

To stop the service for Oracle, enter the following at the command prompt:

```
# /etc/rc3.d/S98dbora stop
```

If you are starting the services manually, start the Oracle service before the service for the management server.

Security Alert Messages when Using HTTPS

To stop receiving a Security Alert message each time you use the HTTPS logon.

> **Note**: Enter the DNS name of the computer in the URL instead of localhost. If you use https://localhost to access the management server, you are shown a “Hostname Mismatch” error.

Installing the Certificate Using Microsoft Internet Explorer 6.0

To access the management server:

1. Type `https://machinename`

   In this instance, machinename is the name of the management server.
2. When the security alert message appears, click **OK**.
3. When you are told there is a problem with the site's security certificate, click **View**

Certificate.

4. When you are shown the certificate information, click the **Install Certificate** button at the bottom of the screen.
5. When you are shown the Certificate Import Wizard, click **Next** to continue the installation process.
6. Select one of the following:
   - **Automatically select the certificate store based on the type of certificate** – This option places the certificate automatically in the appropriate location.

     Or

   - **Place all certificates in the following store** – This option lets you pick the store where the certificate will be stored.
7. Click **Finish**.
8. When you are asked if you want to install the certificate, click **Yes**.

"Security certificate is invalid or does not match the name of the site," Message

If your users are shown a Security Alert window with the following message, you might want to modify the security certificate so users feel more comfortable with installing the certificate:

```
The name of the security certificate is invalid or does not match the
name of the site.
```

You can change the security certificate so that users receive the following message instead:

```
The security certificate has a valid name matching the name of the
page you are trying to view.
```

When you change the certificate, you must use the generateAppiqKeystore program to delete the original certificate, and then use the generateAppiqKeystore program to create a new certificate and to copy the new certificate to the management server.

Windows

To change the certificate on Windows:

1. Go to the %MGR_DIST%\Tools directory.
2. To delete the original certificate, enter the following at the command prompt:

   ```
   %MGR_DIST%\Tools> generateAppiqKeystore.bat del
   ```

   The original certificate is deleted.
3. To create a new certificate containing the DNS name of the management server, enter the following at the command prompt:

```
%MGR_DIST%\Tools> generateAppiqKeystore.bat
```

4. If the program is unable to detect a DNS name, enter the following at the command prompt:

   ```
   %MGR_DIST%\Tools> generateAppiqKeystore.bat mycomputername
   ```

   In this instance, mycomputername is the DNS name of the computer

5. To copy the new certificate to the management server, enter the following at the command prompt:

   ```
   %MGR_DIST%\Tools> generateAppiqKeystore.bat copy
   ```

   The new certificate is copied to the correct location.

Linux

To change the certificate on Sun Solaris and Linux:

1. Go to the [Install_Dir] directory and run the following command:

   ```
   eval `./usersvars.sh`
   ```

   The quotes must be entered as left single quotes as shown.

2. Go to the following directory:

   ```
   [Install_Dir]/Tools
   ```

   In this instance, [Install_Dir] is the directory into which you installed the management server.

3. To delete the original certificate, enter the following at the command prompt:

   ```
   perl generateAppIQKeyStore.pl del
   ```

   The original certificate is deleted.

   If you see an error message when you enter this command, a previous certificate might not have been created. You can ignore the error message.

4. To create a new certificate containing the DNS name of the management server, enter the following at the command prompt:

   ```
   perl generateAppIQKeyStore.pl
   ```

5. If the program is unable to detect a DNS name, enter the following at the command prompt:

   ```
   perl generateAppIQKeyStore.pl create mycomputername
   ```

   In this instance, mycomputername is the DNS name of the computer

6. To copy the new certificate to the management server, enter the following at the command prompt:

   ```
   perl generateAppIQKeyStore.pl copy
   ```

The new certificate is copied to the correct location.

"You Are About to Leave a Secure Connection" Message when Accessing Reporter

If you click the Reporter icon and you are running HP Storage Essentials from a secure website, you will be told you are leaving a secure Internet connection and asked if you want to continue.

If you do not want your users to see this message, follow these steps to change the SSLOnly property from false to true:

1. Log on to HP Storage Essentials.
2. Select **Configuration** > **Product Health**.
3. Click **Advanced** in the Disk Space tree.
4. Click **Show Default Properties** at the bottom of the page.
5. Copy the following line:

   ```
   #SSLonly=false
   ```

6. Return to the Advanced page.
7. Paste the copied text into the Custom Properties box. How you paste the text depends on your Web browser.
8. In the Custom Properties box, remove the hash (#) symbol in front of SSLonly property, and change false to true, so the line looks as follows:

   ```
   SSLonly=true
   ```

9. When you are done, click **Save**.

Client Unable to Access HP Storage Essentials

If the management server is behind a firewall, the firewall must be disabled if you want the client Web browser to be able to access HP Storage Essentials from outside of the firewall. Windows 2008 has a firewall enabled by default.

Configuring the Java Console

HP recommends that you configure your Java Console to the heap size to –Xmx320 for daily work. If it is absolutely necessary, you can increase the heap size to as high as –Xmx750m. Setting the heap size to –Xmx750m will, however, slow down the performance of the Web browser.

Refer to the documentation for your Java Console for more information on how to modify the Java heap size.

"Data is late or an error occurred" Message

If you see the message "Data is late or an error occurred" when you try to obtain information from a UNIX host, verify that you logged in as root when you started the CIM extension (./start). You must be logged in as root to use the ./start command, even if you are using the ./start -users username command, where username is a valid UNIX account.

The CIM extension only provides the information within the privileges of the user account that started the CIM extension. Therefore, you must use root to start the CIM extension. Only root has enough privileges to provide the information the management server needs.

If you continue to see the message, contact customer support.

appstorm.<timestamp>.log Filled with Connection Exceptions

When an Oracle redo log becomes corrupt, the management server is unable to connect to the database. Whenever this occurs, the management server writes to the appstorm.<timestamp>.log file. Many exceptions might cause the application log on Windows to become full.

To correct this problem, follow these steps to stop the management server and Oracle, and remove the corrupted redo log:

1. Stop the AppStorManager service, which is the service the management server uses.

   Note: While the service is stopped, the management server cannot monitor elements and users cannot access the management server.

2. To find the corrupt log file, look in the alert_appstorm.<timestamp>.log file, which can be found in one of the following locations:

   Windows: `\oracle\admin\APPIQ\bdump`

   UNIX: `$ORACLE_BASE/admin/APPIQ/bdump`

   You can verify if the redo log listed in the alert_appstorm.<timestamp>.log file is corrupt by looking for a "redo block corruption" error in the redo log.

3. On the management server, enter the following at the command prompt:

   ```
   Sqlplus /nolog
   ```

4. Enter the following:

   ```
   Sql> connect sys/change_on_install as sysdba
   ```

5. Enter the following:

   ```
   Sql> startup mount;
   ```

6. Enter the following:

```
Sql> ALTER DATABASE CLEAR UNARCHIVED LOGFILE
'C:\ORACLE\ORADATA\APPIQ\REDO02.LOG';
```

In this instance, C:\ORACLE\ORADATA\APPIQ\REDO02.LOG is the corrupted log file and its path.

7. Enter the following:

```
Sql> alter database open
```

8. Enter the following:

```
Sql> shutdown immediate;
```

9. Enter the following:

```
Sql> startup
```

Errors in the Logs

If you access the logs, you are shown messages resembling the following. To save space, the text has been shortened:

```
[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService]
Creating

[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService] Created

[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService]
Starting

[Aug 04 2004 11:59:07] INFO
[com.appiq.service.policyManager.policyService.PolicyService]
Starting Policy Factory

[Aug 04 2004 11:59:11] ERROR
[com.appiq.security.DatabaseSecurityManager] DatabaseSecurityManager
Error:

org.jboss.util.NestedSQLException: Could not create connection; -
nested throwable: (java.sql.SQLException: ORA-01033: ORACLE
initialization or shutdown in progress

); - nested throwable: (org.jboss.resource.ResourceException: Could
not create connection; - nested throwable: (java.sql.SQLException:
ORA-01033: ORACLE initialization or shutdown in progress

))
```

Volume Names from Ambiguous Automounts Are Not Displayed

Volume names from ambiguous automounts on Solaris hosts are not displayed on the Storage Volumes page or in Capacity Manager. Some Solaris hosts have autofs and NFS mounted through an automounter. The management server cannot display volume names from ambiguous automounts because it cannot determine if the comma-separated strings that are part of the mounted volume name are host names or part of the name of a remote volume.

The following example is a comma-separated string that is part of a mounted volume name. The management server cannot tell whether `test` and `three` are host names or part of the name of a remote volume. As a result, the management server does not display the volume name.

```
VolumeName = two:/ntlocal2,two:/comma,test,three,one:/ntlocal
```

Known Issues about Applications

This section provides information about known issues with applications.

- Oracle ACFS shown with Drive Type "Local" even if the file system is on an External Drive. The Drive Type on the Storage Volumes page is shown as "Local" for Oracle Automatic Cluster File System (ACFS) file systems even if the ACFS file system is on an external disk.
- Unmounted Databases not shown on Properties Page for InterSystem Cache Databases. On the Properties Page for InterSystem Cache Database instances, unmounted databases are not shown under Logical Elements.
- sblobspace Reported for an Informix Server even if the sblobspace is Removed. The sblobspace reported for an Informix installation continues to be reported by the management server even if the sblobspace is removed.
- Usernames to Discover Applications must be Unique. In the Setup->Applications tab, user names are unique. A single user name with different passwords cannot be used to discover databases on multiple hosts; the user interface will show only one entry for a particular user name.
- Redo Groups on Raw Devices shown only for one RAC Instance. Redo groups appear in the topology for only one RAC instance in an Oracle RAC configuration with raw devices.
- Capacity Charts for Informix Databases show dbspaces. Although databases are listed on the Capacity pages, the Capacity Manager Charts display data for dbspaces for Informix databases.
- Cannot Create a Virtual Application on an Oracle RAC Shared Volume on Solaris x86. At this time it is not possible to create a virtual application on a shared Oracle RAC volume on Solaris x86. You will see the following message: “java.lang.NullPointerException.”
- Update Element Data (Single Element Refresh) does not Update all Oracle Failover Information. Performing a single element refresh does not update the Oracle Failover

information about which node is active if there has been a failover. Get Details updates all the necessary information.

- Host Cluster Topology Does Not Show Oracle Database Instances as Shared. Oracle database instances on shared raw volumes in a cluster are not reported as shared on the Host Cluster Topology. The individual instances are shown as local to the host and not shared in the cluster. The Application Topology page shows the proper configuration.
- Status not Displayed for Oracle Database Instance Control Files. The status of the Oracle database instance's control files is not shown on the instance properties page.
- Exchange Services Statistics Chart Shows Raw Data. The Exchange Services Statistics Chart will report only the raw data available. It does not report on rolled-up data. This chart is being reconsidered, as a roll-up of a “service up” or “service down” value is not meaningful.

Troubleshooting CIM Extensions

This section describes how to troubleshoot issues with CIM extensions.

Configuring UNIX CIM Extensions to Run Behind Firewalls

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 IPs: 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 on how to start a CIM extension manually, see the *Installation Guide*
- The “If Mentioned in cim.extension.parameters” column provides information on modifying the cim.extension.parameters file (see Permanently Changing the Port a CIM Extension Uses (UNIX Only) on page 860).
- The “Step 1 Discovery (**Discovery** > **Setup**) and RMI Registry Port” column provides information about the IP addresses that are required for the discovery list. The CIM extention 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; for example, 192.168.1.1:1234. In this instance, 192.168.1.1 is the IP for 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
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
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

Configuration	Manual Start Parameters for CIM Extension	If mentioned in cim.extension.parameters	Step 1 Discovery and RMI Registry Port
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
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

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. 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
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. Dicover 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

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

AIX CIM Extension Does Not Start

In some cases, a CIM Extension installed on an AIX server does not start, and the `cxsw.out` file in `/opt/APPQcime/tools` shows an error message like the following:

```
[ Unable to mmap Java heap of requested size, perhaps the maxdata
value is too large - see Java README.HTML for more information. ]
```

To resolve this:

1. Open the `wrapper.conf` file in the `/opt/APPQcime/conf` directory in a text editor.
2. Set the wrapper.java.maxmemory property to 256, as follows:

   ```
   wrapper.java.maxmemory=256
   ```

3. Save the `wrapper.conf` file.

1. Locate and open the `wrapper.user-sample` file in the `conf` directory.
2. Copy your custom settings from the `wrapper.conf` file to the `wrapper.user-sample` file and save your changes.
3. Save or rename `wrapper.user-sample` as:

   ```
   wrapper.user
   ```

The CIM extension software retains and uses the `wrapper.user` file containing your custom settings after each future upgrade of the CIM extension.

Note: If further JVM custom settings are required, the changes should be added to and saved in `wrapper.user`.

Permanently Changing the Port a CIM Extension Uses (UNIX Only)

CIM extensions on UNIX use port 4673 by default. You can start a CIM extension on another port by entering `./start -port 1234`. In this instance, 1234 is the new port. With this method, you must always remember to provide the nondefault port when starting the CIM extension.

You can configure a CIM extension to remember the nondefault port, so you only need to enter `./start` to start the CIM extension:

1. Go to the /opt/APPQcime/conf directory.
2. Open the cim.extension.parameters file in a text editor, and provide the following:

   ```
   -credentials username:password

   -port 1234
   ```

   The values for -credentials and –port must be on separate lines, as shown in the example.

   In this instance:

   - username is the user that is used to discover the CIM extension. You will need to provide this user name and its password when you discover the host.
   - password is the password of username.
   - 1234 is the new port for the CIM extension.
3. Save the file.
4. Restart the CIM extension for your changes to take effect.

   The CIM extension looks for parameters in the cim.extension.parameters file whenever it starts, such as when it is started manually or when the host is rebooted.
5. The management server assumes the CIM extension is running on port 4673. If you change the port number, you must make the management server aware of the new port number.

   In the IP Address/DNS Name box in the Add Address for Discovery page (**Discovery** > **Setup** > **Add Address** on the HP SE Home page), enter a colon and then the port number after the IP address or DNS name, as shown in the following example:

   ```
   192.168.1.2:1234
   ```

   In this instance:

   - 192.168.1.2 is the IP address of the host
   - 1234 is the new port number

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

Troubleshooting Discovery and Get Details

This section contains the following topics:

Troubleshooting Mode

Troubleshooting Mode helps you identify and resolve host configuration issues during discovery. You can enable Troubleshooting Mode as follows:

- If errors occur during discovery, an error message appears at the top of the screen below the discovery step where the errors occurred. If you see an error message, enable Troubleshooting Mode by selecting the Enable Troubleshooting Mode check box located near the top of the page for each discovery step.

- A red icon appears in the Problems column for each host for which a problem was detected. When you click this icon for a particular host, a list of troubleshooting tips appears below the Enable Troubleshooting Mode check box. These tips enable you to resolve the configuration problems for that host.
- Click the link located in the error message for one of the discovery steps. For example, if you are on discovery step 3, click the “Discovery -> Setup in Troubleshooting mode” link located in the step 1 error message. Clicking this link brings you to the step 1 page with Troubleshooting Mode enabled.

  When Troubleshooting Mode is enabled during Get Details, the following additional information can help you identify configuration issues:

- Host Operating System
- CIM Extension Version
- HBA (Driver Version)
- Multipathing
- Volume Management

Unable to Discover Emulex Host Bus Adapters

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

CIMOM Service Not Starting After Trying to Discover Sybase or SQL Server Applications

If your management server is running on Linux, you cannot discover Sybase or SQL Server applications. If you already added a Sybase or SQL Server entry to be managed in the Discovery setup page and performed a Get All Element Details operation, entries for the Sybase or SQL server are added to the Oracle listener configuration file. On the next system reboot, or on the next restart of the Oracle service, the Oracle listener will error out, and the CIMOM service will not start.

To correct the issue:

1. Edit `ORA_HOME/network/admin/listener.ora` and remove the `SID_DESC` text blocks containing the `PROGRAM=hsodbc` string.

   In this instance, ORA_HOME is the Oracle home.

   If you have a SID_DESC block similar to the following text block, remove the entire block.

   ```
   SID_DESC =

   SID_NAME = SQLSERVERSID)

   ORACLE_HOME = /opt/oracle/product/9.2.0.4)
   ```

```
PROGRAM = hsodbc)
```

2. Restart Oracle with the following command:

   ```
   /etc/init.d/dbora restart
   ```

3. Restart the appstormanager service.
4. After the service starts, delete any Sybase or SQL entries from the Application tab in the discovery setup page. This is necessary to prevent them from being re-added to `listener.ora` on further discoveries.

NSK Host Managed by Multiple CMS Not Supported

A configuration of multiple CMS set up to manage the same NSK host is not supported. NSK does not support pre-emptive thread scheduling. Therefore, if the agent is running an `enumerateInstances` in response to a request from a CMS, it is not able to accept a connection request from a second CMS. When this happens, a `NO_CIMOM` exception is thrown in the CMS that initiated the connection request. The number of `syncronizerThreads` is limited to one for na NSK host; therefore, the same issue does not occur during GAED.

Super Group Users Discover NSK Hosts

Only users who are part of the super group should be configured (using the `-users` option) to discover the NSK host. A user who is *not* a member of the super group is not able to invoke HBA library calls; therefore, HBA details (adapter, port, and binding information) cannot be retrieved. This results in a failure to generate the NSK host topology.

Configuring E-mail Notification for Get Details

The management server enables you to send status reports about Get Details to users. These status reports can also be found in the `GAEDSummary.log` file in the `[Install_DIR]\logs` directory on the management server.

To configure the management server to send status reports on Get Details to an e-mail account:

1. Enable e-mail notification for the management server. See Enabling Email Notification on page 335.
2. Add or edit the e-mail address for the Admin account.

   The following status reports for Get Details are sent:

   - "gaedemail property is empty" – E-mail is sent to users whose roles have System Configuration selected.
   - "gaedemail property is populated" – E-mail is sent only to users whose e-mail is assigned to the gaedemail property.

To have additional users receive status reports for Get Details:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.

3. Copy the gaedemail property.
4. Return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Add the e-mail accounts that will receive the reports. For example, to enable user1@mycompany.com and user2@mycompany.com to receive reports, modify the gaedemail property in the Custom Properties box as follows:

   ```
   gaedemail=user1@mycompany.com;user2@mycompany.com
   ```

   Remove the hash (#) symbol from the gaedmail property.
7. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

"Connection to the Database Server Failed" Error

If you received an error message like the following after getting all element details, verify that the database instance is running:

```
The connection to the database server failed. Check that the Oracle
instance 'OIQ3 on host '192.168.1.162:1521 is running correctly and
has the management software for Oracle installed correctly.
```

If you receive such an error message, verify the following:

- Oracle instance OIQ3 on host 192.168.1.162 port 1521 is running.
- The management software for Oracle is installed on the server running the Oracle instance. One of the installation's tasks is to create an APPIQ_USER user account with enough privileges for the software to view statistics from the database.

After that, run Get Details again. If you continue to see the error message, contact customer support.

Using the Test Button to Troubleshoot Discovery

If you are having problems discovering an element, click the **Test** button on the Discovery setup page (**Discovery** > **Setup**). When you click the Test button, the management server attempts to ping the element, and then it runs a series of device-specific connectivity tests. The output of these tests can be viewed in the discovery log window.

The management server uses a provider to communicate with an element. A provider is software that communicates with the element and the management server. When you click the Test button, it checks every available provider against the element to see which one works. When this test is being performed, you might notice messages such as "Test provider not supported," "Connection Refused," or "Failed to Establish Connection." This means a provider was tested against the element and the provider was not the correct one.

When the correct provider is found, a message such as "ExampleComputer responds to a Wind32 system" or "Connection accepted" is displayed; for example:

```
Testing provider APPIQ_Win32Provider for: 192.168.1.2
```

```
ExampleComputer responds as a Win32 system with CIM Extensions
3.0.0.129
```

The success messages are intertwined with the other messages, so you need to scroll through the log messages. For example, the success message shown previously appeared in the middle of the log messages, as shown in the following example. The success message is underlined in the following example.

To make it easier to view the log messages, copy and paste the log messages from the log window to a text editor.

```
LOG MESSAGES


[2004/01/15 09:10]    Test Discovery Started

[2004/01/15 09:10]    Successfully pinged 192.168.1.2

[2004/01/15 09:10]

Testing provider APPIQ_SolarisProvider for: 192.168.1.2

Connection refused to host: 192.168.1.2; nested exception is:

java.net.ConnectException: Connection refused: connect

Testing provider APPIQ_CimProxyProvider for: 192.168.1.2

Test provider functionality not supported for APPIQ_CimProxyProvider

Testing provider APPIQ_McDataProvider for: 192.168.1.2

Can't connect.

No current SWAPI connection to host 192.168.1.2.  Cannot establish
connection

Testing provider APPIQ_AltixProvider for: 192.168.1.2

Connection refused to host: 192.168.1.2; nested exception is:

java.net.ConnectException: Connection refused: connect

Testing provider APPIQ_IrixProvider for: 192.168.1.2

Connection refused to host: 192.168.1.2; nested exception is:

java.net.ConnectException: Connection refused: connect

Testing provider APPIQ_Win32Provider for: 192.168.1.2

ExampleComputer responds as a Win32 system with CIM Extensions
3.0.0.129

Windows host does not support remote testing

VERITAS Volume Manager not available
```

```
HDLM Multipathing Software not available
Powerpath Multipathing Software not available
RDAC Multipathing Software not available
Testing provider APPIQ_EmcProvider for: 192.168.1
Can't connect
appiqSymInitialize() failed with error code 510
Testing provider APPIQ_AixProvider for: 192.168.1.2
Connection refused to host: 192.168.1.2; nested exception is:
java.net.ConnectException: Connection refused: connect
Testing provider APPIQ_HdsProvider for: 192.168.1.2
Cannot connect to Proxy
Cannot connect to Proxy
Testing provider APPIQ_BrocadeElementManager for: 192.168.1.2
Cannot connect
Cannot connect
Testing provider EngenioSSI_Provider for: 192.168.1.2
Failed to establish connection.
Testing provider APPIQ_ClariionProvider for: 192.168.1.2
NaviCLI not installed
No such file: C:\Program Files\EMC\Navisphere CLI\NaviCLI.exe
[2004/01/15 09:10]    Test Discovery Completed
TEST DISCOVERY COMPLETED in 5 seconds
```

By design, the Test button is not available when any of the discovery steps are occurring.

DCOM Unable to Communicate with Computer

Sometimes the following error message appears in the event log of the management server when the software is monitoring a Brocade switch:

```
DCOM was unable to communicate with the computer 192.168.10.21 using
any of the configured protocols
```

In this instance, 192.168.10.21 is the IP address of the Brocade switch.

Ignore this error message.

Duplicate Listings/Logs for Brocade Switches in Same Fabric

If you discover more than one Brocade switch in the same fabric, the Targets tab displays duplicate listings for the Brocade switches. Each Brocade switch is listed multiple times: with the IP address of the other switches and its own.

For example, if Brocade switches QBrocade2 and QBrocade5 are discovered in the same fabric, they are listed twice on the Targets tab. QBrocade2 appears once with its own IP address and then again with the IP address of QBrocade5, as follows:

```
192.168.10.22 Switch QBrocade2, QBrocade5 admin

192.168.10.25 Switch QBrocade2, QBrocade5 admin
```

Duplicate Entries for the Same Element on the Get Details Page

If an element is discovered through two different protocols, it might be listed twice on the Get Details page.

To change the protocol used to discover an element that has already been discovered, delete the element before attempting to perform Get Details again. See Deleting Elements from the Product on page 157.

For some elements, duplicate entries could result if a second protocol is available. For example, you could choose to discover an element through a supported API, but if the element supports SMI-S, and the SMI-S provider is also available, the element could be discovered again. In this example, you would then disable the SMI-S provider.

Element Logs Authentication Errors During Discovery

During discovery, you might see SNMP authentication errors on the element you are trying to discover. The management server is probing the element with an SNMP request. If the element does not know the management server, it logs authentication errors.

EMC Device Masking Database Does Not Appear in Topology (AIX Only)

An EMC device masking database attached to an AIX host does not appear in the Topology tree under the Application Path – Unmounted node on the Topology tab in System Manager.

If the EMC device masking database is attached to a host running Microsoft Windows or Sun Solaris, the masking database appears under the Application Path – Unmounted node.

Management Server Does Not Discover Another Management Server's Database

In some situations, the management server might not discover another management server's database. Make sure that the Oracle monitoring software (CreateOracleAct.bat for Microsoft Windows or CreateOracleAct.sh for UNIX) is installed on the management server to be discovered and that the Oracle instance is added to the discovery list.

Microsoft Exchange Drive Shown as a Local Drive

Microsoft Exchange Servers have a drive M. The software displays this drive as a local fixed disk, instead of a Microsoft Exchange Server special drive.

Unable to Discover Microsoft Exchange Servers

If DNS records for your Microsoft Exchange servers are outdated or missing, the discovery of Microsoft Exchange might fail because Microsoft Exchange is dependant on Active Directory, which is dependant on DNS. Since Active Directory is dependant on DNS, Active Directory replication and Active Directory lookups might fail or contain errors if DNS records are not accurate.

Nonexistent Oracle Instance Is Displayed

The software uses the Oracle Transparent Name Substrate (TNS) listener port to detect Oracle instances on a server. Sometimes an Oracle instance is removed from the server, but not from the TNS listener port. This results in the software detecting the nonexistent Oracle instance and displaying it in the topology. For information on how to remove the deleted Oracle instance from the TNS listener port, see the Oracle documentation.

Requirements for Discovering Oracle

To discover Oracle:

- The management software for Oracle must be installed. For information about installing the management software for Oracle, see the *Installation Guide*.
- By default, the software sets the TNS listener port to 1521. If you use another port, you can change the port number on the Discovery Targets tab.
- Oracle discovery relies on the TNS networking substrate on which Oracle is built (TNS is Oracle's proprietary protocol). The software does not use the TNS listener password. If you set a TNS listener password, the software is not able to discover the Oracle instances serviced by the listener.

Do Not Run Overlapping Discovery Schedules

If you are creating multiple discovery schedules, you must be careful to avoid scheduling conflicts; for example, concurrently scheduled Discovery tasks. Each scheduled task must have enough time to start and finish before the next Discovery task is scheduled to start. For example, if a scheduled Discovery is still in progress when another scheduled Discovery attempts to start, the Discovery task that attempts to start will not start, because the first discovery is still running. The discovery that is unable to start is rescheduled according to its recurring rule. If the discovery task is scheduled to run on a daily basis, the discovery would then start again on the next day. To check the status of scheduled discovery tasks, view the `appstorm.<timestamp>.log` file in the following directory:

```
[Install_Dir]\jbossandjetty\server\appiq\logs
```

Storage System Uses Unsupported Firmware

The following message is displayed when an LSI storage system is discovered, and is running unsupported firmware:

```
This storage system uses unsupported firmware. ManagementClassName:
class_name
```

In this instance, class_name is the management class name for the unsupported array.

The management class name for the unsupported array is displayed in the message.

New releases of storage system firmware are supported with each new release of this software. For the latest information on supported firmware, see the support matrix for your edition.

FC Port Total Request Rate and FC Port Total Throughput Reports Fail

The FC Port Total Request Rate and FC Port Total Throughput reports fail when attempting to retrieve data for RAID-450 class storage arrays (such as the HP XP128, HP XP512, and HP XP1024). To resolve this issue, run these reports on the attached switches by selecting the switch port that is connected to the array port you are interested in. Running reports on RAID-450 class storage array ports requires the discovery of the attached switches.

"CIM_ERR_FAILED: index out of bounds" During Step 1 Discovery

Step 1 Discovery produces the error "CIM_ERR_FAILED: index out of bounds" after discovering an ESX Server and attempting to probe the SMI-S provider on the subsequent IP address. This error is written to the management server logs and does not impact the Discovery operation.

An Event Might not Appear when a New Device is Discovered

When a new device is discovered, an event is generated in the management server. The event might not appear for all new devices that are discovered.

Discovery Logs Might Show ORA-01430 Error for the DATABASE_PORTS Table

The first Detail Discovery following an upgrade of the management server might show the following in the discovery logs:

```
Exception in alterTable batching for table: DATABASE_
PORTSerror occurred during batching: ORA-01430: column being
added already exists in table.
```

This error can be ignored.

Troubleshooting

This section contains the following topics:

Shown "Cannot initialize report engine" or "Invalid session WH 00013" Message

If you are shown one of the following messages and Report Optimizer is running on a 64-bit Linux system, the Oracle client might not have been installed correctly:

- Cannot initialize report engine
- Invalid session WH 00013

The workaround is to install the 11.1.0.6 Oracle client; however, before you install the Oracle client you must prepare the server for the installation, as described in the following steps.

To prepare the server for the installation of the Oracle database client:

1. Logon to the Linux server as root.
2. Make sure the X Window System can display.You can determine that the X Windows System is displaying properly by entering the xclock command. If the time is displayed, the X Windows System is working properly. You can press Ctrl+c to exit the clock. If you are running into issues with the X Windows System, refer to the documentation for X Window System for more information.

a. Logon as root.

b. Enter the following commands to enable the display for the Oracle client installer:

```
xhost +
export DISPLAY=:0.0
```

3. Create a 11.1.0.6 directory under the ora_11gR1_client directory by entering the following command:

```
mkdir -p /ora_11gR1_client/11.1.0.6
```

4. Change the owner of the new directory to oracle by entering the following command:

```
chown oracle:oinstall /ora_11gR1_client
```

5. Change the execution mode of the newly created directory to read, write, and execute for all by entering the following command:

```
chmod 777 /ora_11gR1_client
```

6. Download version 11.1.0.6 of the Oracle client from the following website:

   http://www.oracle.com/technetwork/database/enterprise/downloads/111060-linx8664soft-099033.html

   You must accept the license agreement on the website to download the software.

7. Save linux.x64_11gR1_client.zip to a directory where the user "oracle" has all privileges, for example `/tmp`.

8. Change to the directory where the zip file was downloaded, for example `/tmp`. Add execute permissions to the zip file by entering the following command:

```
chmod +x linux.x64_11gR1_client.zip
```

9. Logon as user oracle by entering the following command:

```
su oracle
```

10. Unzip linux.x64_11gR1_client.zip by entering the following command:

```
unzip linux.x64_11gR1_client.zip
```

To install the Oracle database client:

1. Change to the `<extracted file directory>/client` by entering the following command:

```
cd <extracted zip file directory>/client
```

   In this instance, `<extracted zip file directory>` is the directory containing the extracted files from linux.x64_11gR1_client.zip. For example, you would enter the following command if the `linux.x64_11gR1_client.zip` file was extracted to `/tmp`:

```
cd /tmp/client
```

2. Enter the following command to run the installation:

```
./runInstaller
```

3. On the Welcome page, click **Next**.
4. On the Select Installation Type page, click **Custom**, then **Next**.
5. On the Install Location page, enter the following in the Oracle Base field:

   ```
   /ora_11gR1_client
   ```

   The wizard finds the directory with the zip file, and it populates the Path field.
6. In the Name text box, change the value to the following:

   ```
   OraHome2
   ```
7. In the Path field, click the **Browse** button to set it to:

   ```
   /ora_11gR1_client/11.1.0.6
   ```

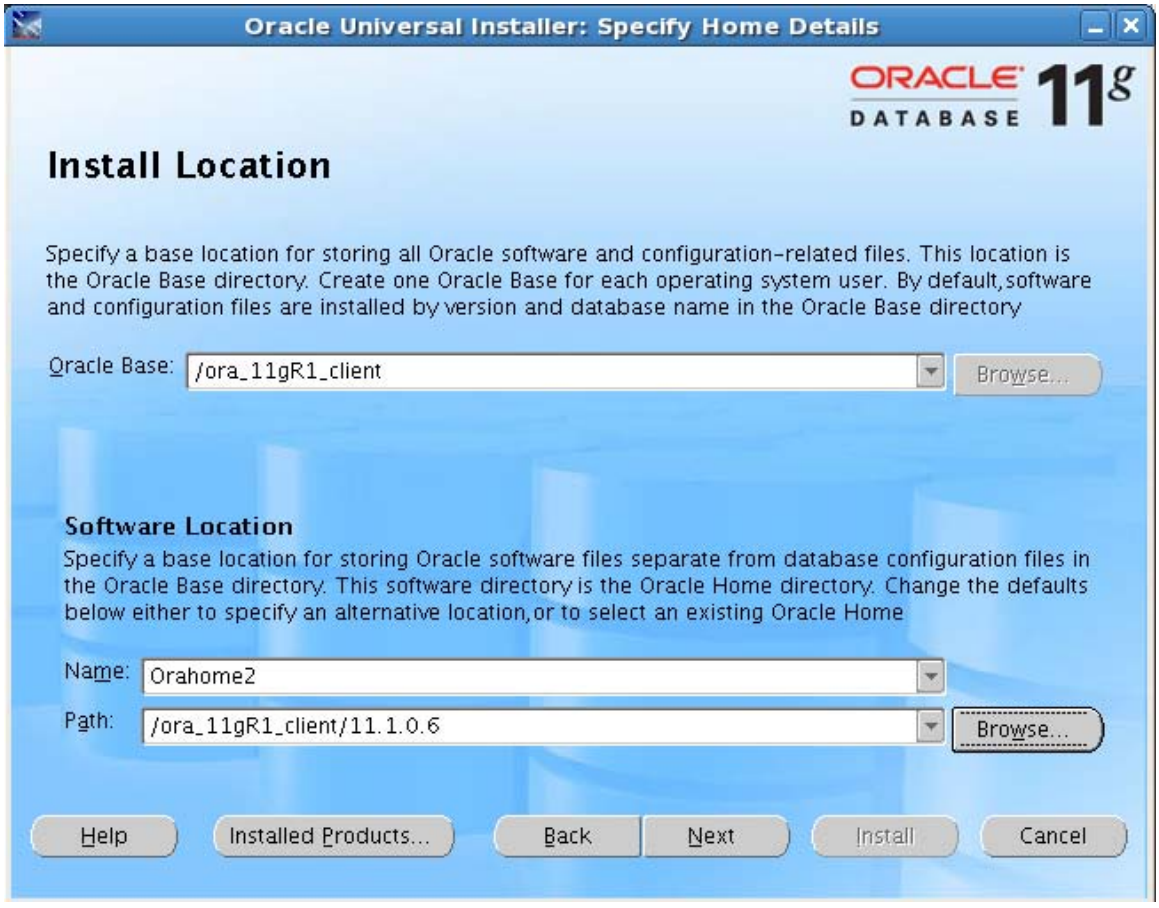

8. Click **Next** to submit your changes.
9. Wait for the installation to check for pre-requisites and then click **Next**.
10. Select the following and then click **Next**:

- SQL*Plus
- Oracle JDBC/THIN Interfaces
- Oracle Net
- Oracle ODBC Driver

11. On the Summary page, click **Install**.
12. On the Oracle Net Configuration Assistant Welcome page, select **Perform typical configuration**. Then, click **Next**.
13. Click **Next**.
14. Click **Finish**.
15. Refer to the configuration steps listed in the window. These configuration steps require a terminal window.
16. Open a terminal window to run the configuration steps.
17. Press Enter four times to accept the defaults for the configuration steps in the terminal window.
18. Type exit in the terminal window.
19. Click **OK** in the Execute Configuration Scripts window.
20. On the End of Installation page, click **Exit**.
21. Click **Yes** to exit.
22. To exit the installer background process press ctrl+c.
23. Return to root user by typing `exit` in the terminal window.
24. Logon to the Linux server as repadm:

    ```
    su - repadm
    ```

25. Edit the user profile (for the Bash UNIX shell it is `vi .bash_profile`) to ensure ORACLE_SID, ORACLE_HOME, LD_LIBRARY_PATH, and PATH environment variables are set correctly. Enter the following in the user profile or for the Bash UNIX shell in the .bash_profile:
    - `ORACLE_HOME=/ora_11gR1_client/11.1.0.6`
    - `export ORACLE_HOME`
26. Make sure the following environment variable is set in the .bash_profile:

    ```
    ORACLE_SID=REPORT
    ```

27. Prepend the path of the LD_LIBRARY_PATH variable with the following:

    ```
    /ora_11gR1_client/11.1.0.6/lib32:
    ```

28. Prepend the PATH variable so the following appears at the beginning:

```
/ora_11gR1_client/11.1.0.6/bin:
```

29. Make sure the environment variables are only listed once in PATH and LD_LIBRARY_PATH. If a variable is listed more than once, Linux will use the value that appears last.
30. Stop Report Optimizer by entering the following command:

    ```
    /etc/init.d/BobjEnterprise120 stop
    ```

31. Start Report Optimizer by entering the following:

    ```
    /etc/init.d/BobjEnterprise120 start
    ```

32. Run the Report Admin Utility to get the latest report data.
33. Run the reports.

"Connection failed." Message when Generating Reports

If you see the following message when you try to run reports in Report Optimizer, perform the steps in this section:

```
Connection failed. The server has reached the maximum number of
simultaneous connections. (Error: RWI 00239)
```

To resolve this:

1. Go to **CMC** > **Users** > **Administrator User** > **Properties** > **Change Connection**.
2. Select the **Named User** option.
3. Click **Save**.

Known Issues with Report Content

- Storage Details Report does not include Storage Pools that have no Volumes. The Storage Details Report omits Storage Pools that do not have any associated Storage Volumes. When a Storage Volume is discovered in the Storage Pool, the Storage Details Report shows the Storage Pool. If you would like to report on the details of the affected Storage Pools, do one of the following:
  - Use a different report, such as Storage Capacity Details
  - Provision a Storage Volume in the empty Storage Pool, then perform a Step 3 Get All Details and Report Cache Refresh. When the Reporter data is updated and the Storage Details Report data is refreshed, the Storage Pool will appear.
- Storage Pool Name not shown for LUSEs in Storage Details Report. The Storage Pool Names do not show in the Storage Details Report for LUSE storage extents on HDS devices.
- Storage Dependency Report does not show Virtual Storage Dependency if LUNs not Mapped to Hosts. The Storage Dependency Report for back end storage does not show Virtual Storage Dependency if the LUNs are not mapped to hosts.

- Stopped Oracle ASM Instances not Counted in Host Unused Capacity and Available White Space Reports. Disks that are part of an Oracle ASM disk group are removed from the Host Unused Capacity and Available White Space reports if the Oracle ASM instance is stopped and a Step 3 / Detailed Discovery is run. When ASM is active again, perform a Step 3 / Detailed Discovery operation to restore the expected information to the reports.
- Back-end Storage Dependency Report requires a LUN Mapped to a Host. The Storage Dependency Report for Back End Storage does not display external storage dependencies if there are no LUNs from the virtualizer mapped to a host. The report shows the dependencies from the host to the back end storage as long as a LUN is mapped from the storage virtualizer to the host. The management server user interface displays the external storage dependencies of a storage virtualizer even if no LUNs are mapped to hosts.
- Available White Space Report may show #MULTIVALUE for "White Space Size in GB". In multipath configurations where multipathed disks are not part of the same volume group, the "White Space Size in GB" will list "#MULTIVALUE" in the Available White Space Report.
- LUN Mount Report shows "Internal Volumes" for Storage Virtualizers. The LUN Mount Report shows storage for SVSP and IBM SAN Volume Controller as "internal volumes". The terminology used in the LUN Mount Report is being reviewed and may change in a future release.
- Host Unused Capacity Report does not show Source Array of EMC LUN masking disks. The Unused Capacity Report does not provide the source array of EMC LUN masking disks.
- Reports Concerning Storage in Oracle ASM Configurations. Oracle ASM configurations have not yet been fully modeled within the standard reports provided in the management server user interface. The standard reports do not report used capacity information in Oracle ASM configurations.
- Chargeback by Organization Report does not Contain Storage Tiers Configured on Storage Volumes. The Chargeback by Organization Report does not display Storage Tiers that are configured on Storage Volumes. Tiers created on Storage Systems and Storage Pools are reported correctly.
- Storage-Based Chargeback by Organization Report can Report Extra Storage. Creating an Organization that contains all storage volumes, another that contains all storage systems, and dividing the storage volumes and storage pools into separate tiers, can result in the Storage-Based Chargeback by Organization Report showing extra Total Capacity for the Organization that contains the storage volumes.
- Shared Raw Volumes, Shared ASM Disk Group Data Excluded from Total Capacity Chart for a Host Cluster. In the management server user interface, the Total Capacity Summary data reported in the Capacity Chart tab for the cluster excludes shared raw volume and shared ASM disk group information.
- Host Connectivity Report shows HSGs without Initiators. HSGs without initiators appear in the Host Connectivity Report even though the HSGs are not connected to the host.
- Capacities for Virtual Arrays Incorrect if Attached Storage is Discovered. The aggregated capacity reported for storage arrays is incorrect if virtual arrays, such as the IBM SAN Volume Controller and Hitachi Universal Storage Platform, are discovered by the management server

along with the storage arrays hosting the volumes served to the virtual arrays. The volumes are double-counted. This affects the following reports: Storage Array Capacity by Applications; Storage System Array Overhead Utilization; Storage System Array Utilization; Storage System Utilization.

- System Switch Reports Do Not Have Data if Only Switches Have Been Discovered. If you discover only switches, the System Switch Reports will not contain any data. When you discover a host or an array attached to those switches, the System Switch Reports will be populated properly.
- Oracle 10 RAC Shown Twice in OpenVMS Host Dependency Report. The Host Dependency Report lists Oracle 10 RAC dependencies twice for OpenVMS hosts that are part of a manually built cluster.
- Events from Tape Libraries are Not Shown in the Event Summary Report. Although events from tape libraries appear in Event Manager, such events are not displayed in the Event Summary Report.
- Information in some File System Viewer Reports does not include UNC Volumes. A number of File System Viewer Reports do not include information about UNC volumes: File Server Stale Files Summary; File Server Department; File Server Summary; File Server Summary by Operating System; TopN File Server Summary; TopN Volumes with Stale Files; TopN Volumes with Stale Files by File Server; Volume Details. UNC information is not shown in the Host Utilization Volume Details Report because mounted UNC shares are considered to have zero capacity.
- Application Viewer is Required to Generate Application Reports.Application Viewer is required to generate Application reports that include element and system-specific application data, even if a user has access to all elements in the organization.
- Report Data Might be Missing When Exported to Different Format. Report data might be missing when exported to different formats due to issues in the reporting engine used by the management server. For example on the Applications by Host report, the operating system is incorrectly in HTML format only and some report data is truncated. In the Dependency report for a host, the IP address might be truncated. In a Detail report for a host, the WWN and drive ID information might be missing the final character. These issues have been reported to the report engine development team.
- Task Dashboard and the Report Cache Refresh Time Stamp. On the Task Dashboard the time stamp for the last Report Cache Refresh is the last scheduled time for that operation. The manual Report Cache Refresh is not done with a task, so its results and time do not appear on the Task Dashboard.
- Capabilities Column in HP XP “Details” Report Displays a Text String. The Capabilities Column in the “Details” Report for HP XP arrays displays a placeholder text string because the details of storage pool capabilities are not reported by the Command View XP SMI software.
- Missing information in the Asset Details report. The **Asset Type** field is blank in the Asset Details report.

- Report Pack: HDS storage system pool details are missing in the Storage System Capacity report. HDS storage system pool details are not displayed in the Storage System Capacity report.
- Report Pack: "Last refresh date" is populated before the report initially runs. The **Last refresh date** field is populated before a report initially runs. You can ignore this value. The Last refresh date field should be blank until you click the **Refresh Data** button.
- Uninstalling Report Optimizer does not remove all folders. The uninstaller for Report Optimizer does not remove files and folders that were modified or created after the installation, such as the jre folder and the “Uninstall_HPSRMReportOptimizer" folder. You can safely leave the files and folders that were not removed by the uninstaller or you can manually remove them.
- Report Pack: The Prompt window has a number of usability issues . When some of the standard reports run, a Prompt window appears. This Prompt window is missing some field labels, and the Help button does not work correctly.
- Report Pack: An error message is not shown for the Library Utilization Report when the start date occurs after the end date. If you set the start date to occur after the end date for the Library Utilization report, you are not shown an error message and no data will be retrieved for this report.
- Report Pack: Run the Absolute Date Range filter for the Backup sessions report. The **Specification of relative date range** option does not work for the Backup sessions report. This report should always be run with the absolute date range filter. In the **Select Type** field type `IGNORE`. In the **Select Number** field, type `0`.
- Report Pack: Reports with many elements may not display properly. If you have many elements in a report, labels and legends in the graph of that report might not appear not properly. To workaround this problem, graphs can be enlarged in the edit mode of the report.
- Report Pack: In the Top N Aged Files report, text in a prompt window shows as "Top X File Name" instead of "Top N Aged Files". When you run the Top N Aged Files report from the Report Pack, a prompt window displays a field labeled **Top X File Name**. The label should read **Top N Aged Files**. The software will run a query for the Top N Aged Files based on the number entered in the Top X File Name field.
- Some reports do not let you navigate by year. You cannot navigate by year in the **Collection Time Range** filter in some reports. You are forced to navigate month by month.
- Start and end dates required for the Backup Sessions report when using the relative date range . Use n order to run the Backup Sessions report. By using the relative date range, you must provide dummy start and end dates; otherwise, the **Run Query** button is disabled.
- Top N Reports in Report Optimizer does not work the same way as in HP Storage Essentials. In Report Optimizer and in HP Storage Essentials, customers can use a filter called **Top N Reports**. However, this filter works differently in each product:
  - In HP Storage Essentials: The number of records displayed is based on the **N** value. For example, if you select N=10, the total number of records displayed is always less than or equal to 10 based on the number of files in that report criteria.

  - In Report Optimizer: The number of records displayed is based on rank and not the **N** value. If you select N=10, the total number of records displayed can vary from zero to many, based on the number of files present in a particular rank. For example: Assume you have four files of the following sizes: 5 GB, 2 GB, 2 GB, and 1 GB. The four files would be ranked as 1, 2, 2, and 4. The 5-GB file, which is the largest file in the group, is given the ranking of one; the two 2-GB files are given the ranking of two; and the 1-GB file is ranked last.

- Empty sections of reports overlap other data. Empty sections of reports sometimes overlap other data in the report. Save the report as an Excel or PDF file to view a properly formatted report.
- Host volume capacities are incorrect when filtered with the Select Statistics Type filter. Host volume capacities are incorrect when filtered with the Select Statistics Type filter. If you want to report on the last collection timestamp, none of the statistics type filters or objects need to be included in the query. Use the statistics type filters and objects only when reporting on historical data.
- The elements listed as other on the management server are missing in Report Optimizer. Events reported under the element type "OTHER" in HP Storage Essentials are not visible from the Universe. There are no reports based on events, hence the Report Pack is not affected. When generating event-based reports, HP Storage Essentials events reported under "ELEMENT TYPE = OTHER" are not visible through Report Optimizer.

Manually Importing the BIAR File

If the BIAR file import fails you must manually import the file.

To manually import the file:

1. Make sure that the Report Optimizer services are running:
    a. Open the Central Configuration Manager (**Start Menu** > **Programs** > **BusinessObjects XI Release 3.1** > **BusinessObjects Enterprise** > **Central Configuration Manager**).
    b. Make sure that the Apache Tomcat and Server Intelligence Agent services are running.
2. If you are upgrading from an expired evaluation license:
    a. Log on to the Central Management Console as described in Accessing the Central Management Console.
    b. In the Organize section, click **Servers**.
    c. Click **Servers List** in the left-hand pane, and then select all of the servers in the right-hand pane.
    d. Right-click the selected servers, and select **Enable Server** to turn on all of the servers in your system.
    e. Expand the **Service Categories** node in the left pane.
    f. Right-click the **Web Intelligence** node, and select **Enable Server**.

g. Click the **Core Services** node. Select **AdaptiveJobServer** and **AdaptiveProcessingServer**. Right-click your selection, and select **Enable Server**.

h. Open the Central Configuration Manager (**Start Menu** > **Programs** > **BusinessObjects XI Release 3.1** > **BusinessObjects Enterprise** > **Central Configuration Manager**).

i. Restart the Server Intelligence Agent service.

3. Change the password:

- Windows - Change the password in the ImportBiarFileWindows.properties file:

  i. Open the ImportBiarFile.properties file located in the installation directory:

  - For fresh installations, change
    `password=@password@`
    to
    `password=`
  - For upgrades, change
    `password=@password@`
    to
    `password=<youradministrator password>`

  ii. Save your changes.

- Linux - Change the password in the ImportBiarFileLinux.properties file:

  i. Open the ImportBiarFile.properties file located in the installation directory.

  - For fresh installations, change
    `password=@password@`
    to
    `password=`
  - For upgrades, change
    `password=@password@`
    to
    `password=<youradministrator password>`

  ii. Save your changes.

4. Enter the following command at the command line:

```
<Installation Directory>\ImportBiarFile.bat INSTALL <Installation Directory> >> <Name of log file>
```

5. After the BIAR file import is complete, change the password in the ImportBiarFile.properties file back to `password=@password@`.

Failed License Installation

If the license installation fails, you must manually install the license as follows:

1. Obtain the license key from the License.txt file on the installation DVD.
2. Launch the Central Management Console as described in Accessing the Central Management Console.
3. In the Manage section, click **License Keys**.
4. Remove the existing license keys by highlighting each key and clicking **Delete**.

   Remove all existing keycodes before adding new keycodes.
5. In the Add Key box, enter your new license key, and click **Add**.
6. Open the Central Configuration Manager (**Start Menu** > **Programs** > **BusinessObjects XI Release 3.1** > **BusinessObjects Enterprise** > **Central Configuration Manager**).
7. Make sure that the Apache Tomcat and Server Intelligence Agent services are running.

Error message: Account Information Not Recognized

If your license has expired, you will receive the following message on the Report Optimizer Log On page:

```
Account Information Not Recognized: Enterprise authentication could
not log you on. Please make sure your logon information is correct.
```

Contact your customer representative for an updated license.

Warning Message: The object named 'Root Folder' with id number '23' may never be modified or deleted

If this message appears in the installation log, you can ignore it.

Servers Disabled after License Expiration

If your license expires, the Report Optimizer servers are disabled even after you enter a valid key.

To enable the servers:

1. Verify that you created a server group as described in Creating a Server Group.
2. Log on to the Central Management Console as described in Accessing the Central Management Console.
3. In the Organizer section, click **Servers**.
4. Click **Server Groups List**.
5. Right-click the **Report Connector Services** group, and select **Enable Server**.

Resetting the Administrator Password

If you want to reset the Administrator password, you must know the password for “root” or “sa” user of MySQL.

To reset the Administrator password for Report Optimizer:

1. Go to the command prompt.
2. Browse to the install location of the MySQL bin folder. The default path is the following:
   - **Windows**: `<Report Optimizer install dir>\MySQL5\bin`
   - **Linux**: `<Report Optimizer install dir>/bobje/mysql/bin`

   In this instance `<Report Optimizer install dir>` is the installation directory for Report Optimizer.
3. Enter the following command at the command prompt:
   - **Windows**: `mysql -u sa -h your_ro_server_name -p boe120`
   - **Linux**: `./mysql -u sa -h your_ro_server_name -p BOE120`
4. Enter the MySQL password when prompted. The default password is the following: `Password123`
5. Enter the following command at the command prompt:
   `delete from CMS_InfoObjects6 where objectid=12;`
6. Enter the following command at the command prompt: `quit`
7. Restart Tomcat:
   - **Windows**: Right-click the **BOE120Tomcat** services in the Services Administration tool and select **Restart**.
   - **Linux**:
     i. Go to the following directory: `<Report Optimizer install dir>/bobje`
     ii. Verify that you are root user.
     iii. To stop Tomcat, enter the following command: `./tomcatshutdown.sh`
     iv. To start Tomcat, enter the following command: `./tomcatstartup.sh`
8. Restart Report Optimizer:
   - **Windows**. To restart Report Optimizer:
     i. Restart the MySQL service (BOE120MySQL) from Services, which is available from the Windows Control Panel. Refer to your Windows documentation for more information about restarting a service on Windows.
     ii. Click **Yes** when you are asked to restart the Server Intelligence Agent.
   - **Linux**. To restart Report Optimizer:
     i. To stop Report Optimizer enter the following command:

     ```
     /'etc/init.d/BobjEnterprise120 stop
     ```

     ii. To start Report Optimizer enter the following command:

     ```
     /etc/init.d/BobjEnterprise120 start
     ```

   The Administrator password is now empty.

Do Not Import a Windows BIAR File on Linux

Due to a limitation in the Business Objects software, it is not possible to import a Report Optimizer BIAR file created on Windows into Report Optimizer running on the Linux platform. You will see an error similar to the following: "The service container connected to the server with ID nnnn does not support the service with ID nnnn."

Uninstalling Reporter from Windows 64-bit Might be Slow

Due to an issue in a vendor-supplied utility, uninstalling Report Optimizerfrom a Windows 64-bit server may take two hours.

Cannot Launch Reporter with IE6 or IE7 if Larger or Largest Text Sizes are Specified

The reporting engine will not launch properly if the default text size set for the browser is “Larger” or “Largest”. Internet Explorer 6 and 7 exhibit this issue. As a workaround, set the default text size in the affected browser to be one of the other selections. Internet Explorer 8 does not exhibit this problem.

Installation Fails After Running the BusinessObjects Cleanup Scripts

You may be required to run the BusinessObjects cleanup scripts a second time to prepare the system for a reinstall of BusinessObjects. If the installation fails after you run the BusinessObjects cleanup scripts, run the cleanup scripts a second time.

Extra Directory is Added After a Failed Installation

After a failed installation, if you reinstall the product to a different directory, the original installation directory will still be added. It is safe to manually delete this directory.

“Windows DEP (Data Execution Prevention) can Occasionally Close WebIntelligence Report Server" Message

You can safely ignore the following message:

```
Windows DEP (Data Execution Prevention) can occasionally close
WebIntelligence Report Server.
```

The Email Address Object Provides Storage Group and User Information

The "email address" object located at **Application** > **exchange storage groups** > **exchange stores** > **exchange mail boxes** > **email address** returns user login information instead of an email address.

Troubleshooting Topology Issues

This section contains the following topics:

About the Topology

The software determines the topology by looking at the following:

- **Fibre Channel switch** – The Fibre Channel switch contains a list of all elements within the fabric. The software obtains a detailed listing of all elements connected to the switch fabric.
- **A host containing a Host Bus Adapter (HBA)** – All Fibre Channel host adapters look for available elements attached to the HBA. This information is gathered by CIM extensions and sent to the management server.
- **A proxy connected to the SAN** – Include a proxy that has a direct connection or a SAN connection to the management server. An example of a proxy is the EMC Solutions Enabler or Hitachi HiCommand Device Manager. LSI storage systems do not require a proxy, as they

can be accessed directly. Make sure the proxy service has started. On a computer running Windows, this can be determined by looking in the **Services** window.

About the Topology on previous page provides details about how to correct problems that might occur during discovery and data collection.

Troubleshooting Discovery and Get Details

Scenario	Description	What to Do
 The host appears discovered and it is connected to the switch.	The software is aware of the host, but it cannot obtain additional information about it.	Verify that a CIM extension is installed on the host. Try discovering the element again in HP SE, and then run Get Details.
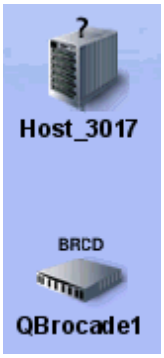 Host appears discovered and it is not connected to the switch.	The switch was previously made aware of the host, but it can no longer contact it. If the steps provided do not work, see Link Between a Brocade Switch and a Host Disappears from the Topology on page 887.	Verify that the host is on and the network cables are connected to it. Try discovering the element again in HP SE, and then run Get Details.
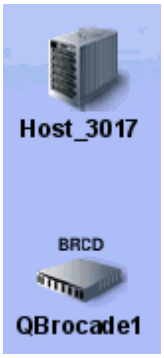 The host appears managed, but it is not connected to the switch.	There is a problem with Get Details from the host. If the steps provided do not work, see Link Between a Brocade Switch and a Host Disappears from the Topology on page 887.	Try getting the topology again: 1. Click the **Discovery** menu, and then click the **Topology** tab. 2. Verify the element is selected and click **Get Topology**.

Scenario	Description	What to Do
 The element appears discovered, but a connected switch does not appear.	The switch has not been discovered.	Try discovering the switch again. 1. Click the **Discovery** menu. 2. Click the **Setup** tab and the **Add Address** button on the IP Addresses tab. 3. Enter the IP address or DNS Name of the switch, and then enter its user name and password. Click **OK**. 4. Verify that the element is selected. 5. Click **Start Discovery**. 6. After discovery has completed, click the **Topology** tab. 7. Verify that the element is selected and click **Get Topology**.

Scenario	Description	What to Do
When discovering a Windows-based host, the correct IP address is entered, but the host does not appear in the topology. The following can be seen on the host: • In Windows Event Manager the WinMgmt.exe process is not running. This process starts WMI. The CIM extension for Microsoft Windows enhances Windows Management Instrumentation (WMI) so it can gather information from host bus adapters and make the information available to the management server. • In the Windows Event Log, DCOM error messages are shown.	An invalid user account was entered	Enter a valid user account that has administrative privileges so it can start WMI. *or* Enter credentials that were provided in the cxws.default.login file, as described in the section, "Creating Default Logins for Hosts," in the *Installation Guide*.

One way to determine what is happening is to look at the log messages during discovery and getting element details. For more information, see Viewing Discovery Logs on page 162.

Virtual Machine's Logical Disks Are Not Mapped to the Virtual Server

If a virtual machine is running Windows (and was discovered explicitly by using its IP address), and some of its disk drives do not have unique SCSI Target IDs, the disk drives will not be stitched to the virtual server. When this occurs, the topology is not able to map the logical disks to the virtual server. The path will stop at the level of the virtual machine.

Undiscovered Hosts Display as Storage Systems

On rare occasions, the management server displays undiscovered hosts as storage systems in System Manager.

To resolve this, follow these steps to provide the host's world wide name (WWN):

1. Determine the host's WWN. This information is available on the IEEE Standards Association web site at http://standards.ieee.org/regauth/oui/oui.txt.
2. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the `#hostPortWWNs=` property.
5. Return to the Advanced page.

6. Paste the copied text into the Custom Properties box.
7. Uncomment the `hostPortWWNs` property by removing the hash mark (#) in front of `hostPortWWNs`.
8. Enter the host's WWN in hexadecimal format. Multiple WWNs can be entered as a comma-separated list; for example:

   ```
   hostPortWWNs=00-01-C9,00-01-C8
   ```

9. Click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

No Stitching for Brocade Switches with Firmware 3.2.0

Stitching does not appear for hosts attached to Brocade switches running firmware 3.2.0. There is no stitching when the PID format is 0. The port setting must be the same for all Brocade switches in the fabric or the fabric will become segmented. The PID format should be set to 1 for all Brocade switches running firmware later than 2.6.0 and 3.0. The PID=0 setting is a legacy Port ID format that does not support the numbers of ports beyond 16.

Brocade SMI-A Switch Discovery

Brocade switches managed through SMI-A version 120.7.2 show only licensed ports when discovered through the management server. The embedded switch ports and ports without SFPs (Small Form-Factor Pluggable transceivers) are not shown. This is a permanent change in the behavior of the management server when discovering Brocade switches with SMI-A 120.7.2 software from Brocade.

Link Between a Brocade Switch and a Host Disappears from the Topology

If a link that used to work between a Brocade switch and a host disappears from the topology, you might need to run Get Details for the Brocade switch and the host. Also, confirm that both are online and there are no network connection issues. As a last resort, you might need to reboot the switch. In some instances, the API of the Brocade switch has been known to hang. Rebooting the switch clears the switch of the API hang.

Unable to Find Elements on the Network

The management server uses ping to find the devices on the network enabled for IP. Ping is a program that lets you verify that a particular IP address exists. Ping is not guaranteed to return a response from all devices. If discovery is not able to find a device automatically, enter the IP address for the device on the discovery Targets tab, which can be accessed by clicking the **Discovery** button at the top of the screen in the management server.

Sometimes ping cannot find the device if any of the following occurs:

- Network configuration does not support ping.
- Data center security (firewalls).

- Device has the ping responder turned off.
- Device does not support ping.

Unable to See Path Information

You will not be able to see path information if LUN masking information is missing. To view LUN masking information, follow the steps described in Accessing Information about Host Security Groups on page 716.

Device Locking Mechanism for Brocade Element Manager Query/Reconfiguration

The configuration for Brocade switches is locked while getting all details for elements in a zone. The software ensures that each CIM query locks out any reconfiguration. For example, if you are getting details for elements in all zones, you cannot add a new Brocade switch while you are doing it (the discovery or configuration process waits until the collection of details is finished before proceeding). However, simultaneous CIM queries do not lock each other out.

A Discovered Sun StorEdge A5000 JBOD Does Not Display Its WWN Properly

Although full monitoring and management support is available only to those devices for which there is a provider, the software's topology displays other devices found on your storage area network (SAN) to give you a more complete view. However, because these devices do not have a provider, only basic information is returned. In some cases, as with the Sun StorEdge A5000 JBOD (just a bunch of disks), the Worldwide Name (WWN) presented and reported to the management server might be different from the official WWN of the device, as the management server reports the WWN of the port connected to the fabric.

Unable to Monitor McDATA Switches

McDATA switches use the Fibre Channel Switch Application Programming Interface (SWAPI) to communicate with devices on the network. The McDATA switches allow only one SWAPI connection at a time. If the management server discovers the IP address of the McDATA switch, other management servers and third-party software are not able to communicate with the switch using SWAPI.

Use Enterprise Fabric Connectivity (EFC) Manager to communicate with the McDATA switch. EFC Manager versions 7.0 and later can communicate with the management server and the switch. This configuration enables multiple instances of the management server or other clients to contact EFC Manager, which in turn provides information about the switch. To communicate with the EFC Manager, discover the McDATA switches as described in Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries on page 71.

EFC Manager uses the SWAPI connection, preventing other third-party software from contacting the switch.

Unable to Detect a Host Bus Adapter

The software is unable to detect a host bus adapter if you install its driver before you completely install the Solaris operating system for the first time; for example, if you installed the HBA drives too early when you used JumpStart to install Solaris. The best way to install the HBA driver is to install it after Solaris is installed and running.

Navigation Tab Displays Removed Drives as Disk Drives

If you remove an internal disk from a Solaris host and do not enter the cfgadm command, the Navigation tab displays the empty slot as DiskDrives_XXXXX after getting element details. The cfgadmn command makes the software realize the drive has been removed. See the documentation that shipped with the Solaris operating system for more information about the cfgadm command.

Unable to Obtain Information from a CLARiiON Storage System

If you are having difficulty obtaining topology information or element details from a CLARiiON storage system, the NaviCLI might have timed out because the service processor is under a heavy load. The management server uses the NaviCLI to communicate with the CLARiiON storage system. This situation has been seen in the field when the service processor is running more than 35,000 IOs per second.

Try obtaining the topology and/or Get Details from a CLARiiON storage system when the service processor is not under such a heavy load.

Discovery Fails Too Slowly for a Nonexistent IP Address

If you enter a nonexistent IP address, the management server times out by default after 20 seconds on Windows or 3 minutes and 45 seconds on UNIX systems. To shorten the time-out period, modify the cimom.CimXmlClientHttpConnectTimeout property as described in this section.

The management server does not accept a period longer than its default setting. If you set the cimom.CimXmlClientHttpConnectTimeout property to more than 20 seconds on Windows or 3 minutes and 45 seconds on UNIX systems, the management server ignores the values of this property and reverts back to the default settings.

To modify the default time-out:

1. Access the management server.
2. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
3. Click **Show Default Properties** at the bottom of the page.
4. Copy the `cimom.CimXmlClientHttpConnectTimeout` property you want to modify.
5. Return to the Advanced page.
6. Paste the copied text into the Custom Properties box.

7. Make your changes in the Custom Properties box. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
8. To modify the time-out period, set the `cimom.CimXmlClientHttpConnectTimeout` property to the number of milliseconds you want. For example, to change the time-out period to 200 ms:

   ```
   cimom.CimXmlClientHttpConnectTimeout=200
   ```

9. When you are done, click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

SVSP Virtual Application Not Displayed in Topology

When discovering the HP StorageWorks SAN Virtualization Services Platform (SVSP), if the virtual application on a host does not show in the SVSP topology and is not listed as a dependency for SVSP, you might have an incorrectly configured system which requires the installation of MPIO and DSM software on the host. This additional software is a basic requirement for being able to mount the SVSP LUNs to an MS Windows server.

Switch Names Inconsistent

The naming convention for Cisco switches discovered for SVSP environments could be different in front-end and back-end topology diagrams. For example, the front-end Cisco switch name could be FCS104108, but the switch name could be 2001000DEC5F6941 in the back-end topology diagram.

"CIM_ERR_FAILED" Message

If you are in a McDATA environment where the EFC Manager Service Processor is managing multiple switches, it is possible that the management server will send SWAPI requests faster than the EFC Manager Service Processor can handle them. The management server might detect this as a failed connection and take corrective action. When this happens, you are shown a "CIM_ERR_FAILED" message whenever the management server tries to access the McDATA switches and directors.

The management server then attempts to reconnect to the EFCM by creating a new SWAPI connection. EFCM versions 8.x and later have five SWAPI connections available. EFCM versions 7.1.3 and later but before version 8.x have three SWAPI connections available. If the management server reconnects successfully, a reconnect event is generated, and no further action is necessary.

If the management server cannot reconnect to the EFCM, another event is generated with a severity of Major. If this happens, any Get Details operation the management server performs involving switches on that EFCM fails.

To prevent the "CIM_ERR_FAILED" messages, follow these steps to increase the delay between the management server's SWAPI calls to EFCM:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.

3. Copy `cimom.mcData.swapiThrottle=200`.
4. Return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Make your changes in the Custom Properties box by changing the value of `cimom.mcData.swapiThrottle`. Say the default is 200 ms and you want to change it to 800 ms. Enter the following:

   `cimom.mcData.swapiThrottle=800`

   If you want no delay, change the value to 0 for 0 milliseconds. The maximum delay you can have is 1,000 milliseconds (cimom.mcData.swapiThrottle=1000).
7. When you are done, click **Save**.
8. The product notifies you if a restart of the AppStorManager service is required.
9. Make sure that you can re-establish communication with EFCM by following the steps in Re-establishing Communication with EFCM below. You might have to change the value of the cimom.mcData.swapiThrottle property if you cannot re-establish communication with EFCM after following the steps in that section.

Re-establishing Communication with EFCM

To re-establish communication with EFCM:

1. To check the status of the connection, click the **Test** button on the Discovery Setup screen. If the McDATA provider reports that it can connect to EFCM, the connection has been restored. A provider is a component of the management server that is used to gather information about an element. In this case, the McDATA provider gathers information about McDATA switches for the management server. To ensure the management server does not have corrupt data as a result of the loss of communication, perform Get Details to obtain the latest information from the element.
2. If the ping to EFCM fails, a network problem exists and must be resolved. Once network connectivity is restored, click the **Test** button to verify that the McDATA provider can communicate with EFCM, and then do a Get Details.
3. If the Test button results from the management server indicate that it still cannot communicate with EFCM, wait approximately 3 minutes for the lost SWAPI connection to time out, and then click the **Test** button again. If this works, do a Get Details.
4. If, after 3 minutes, the Test button results continue to indicate a lost connection, perform the following steps to restore the connection. Note that these steps involve restarting services on the EFCM server. Any other applications using SWAPI to communicate with EFCM are affected by these actions.

   a. Open the EFCM client. Make sure that the EFCM is still actively managing at least one switch. If there are no switches under management, you will not be able to connect to this EFCM.

b. On the EFCM server, stop and restart the Bridge Agent service. Repeat Steps 1 through 3. If the connection is still down, proceed to step c.

c. On the EFCM server, stop and restart the EFCM services. On Windows, use the McDATA EFCM Manager options in the **Start** > **Programs** menu. Repeat Step 1 through 3. If the connection is still down, proceed to step d.

d. Reboot the EFCM server. Repeat Step 1 through 3. If the connection is still down, proceed to step e.

e. Stop and restart the service for the management server. Repeat Step 1 through 3. If the connection is still down, proceed to step f.

f. Reboot the management server. Repeat Step 1 through 3. If the connection is still down, proceed to step g.

g. If none of the previous steps restore the connection, see the support matrix for your edition to determine if the EFCM and switch versions are all supported. Contact technical support for further information.

CIM_ERR_FAILED When Trying to Activate a Zone Set Using McDATA SWAPI

When the user tries to activate a zone set using McDATA SWAPI, the operation might return CIM_ERR_FAILED with one of the following detailed messages:

```
Cannot activate zone set. SWAPI Handle is not valid for fabric
Cannot activate zone set. Active zone set information is out of
date for fabric
There is no active SWAPI connection for fabric
Fabric is not in the cache
```

These error messages indicate that the SWAPI connection to the EFCM managing the fabric is no longer valid, or the active zone information was changed on the fabric without using the management server. The management server does not activate a zone set under these conditions.

To fix this problem, click the **Test** button on the discovery screen to check the status of the SWAPI connection. If necessary, re-discover the EFCM to re-establish the SWAPI connection.

Once the connection is working, the provisioning operation should succeed. If it continues to fail because the active zone set information is out of date, run Get Details for this element to update the zoning information. See Get Details on page 151 for more information.

Communicating with HiCommand Device Manager over SSL

By default, the management server communicates with HiCommand Device Manager through a nonsecure connection. You can configure the management server so that it communicates with HiCommand Device Manager over a secure socket layer (SSL) connection by doing one of the following:

- **Use HTTPS in the discovery address**

  Prepend https:// to the discovery address to force the connection to HTTPS mode; for example, https://192.168.1.1. In this instance, 192.168.1.1 is the IP address of the host running HiCommand Device Manager. Use this option if you have one HiCommand Device Manager that you want to communicate through a secure connection (SSL) and another that you want to communicate through a nonsecure connection.

- **Modify an internal property**

  Change the value of the cimom.provider.hds.useSecureConnection to true, as described in the following steps. Use this option if you want all connections to HiCommand Device Manager to be SSL.

To set all connections with HiCommand Device Manager to SSL:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
2. Click **Show Default Properties** at the bottom of the page.
3. Copy the `cimom.provider.hds.useSecureConnection` property.
4. Return to the Advanced page.
5. Paste the copied text into the Custom Properties box.
6. Make your changes in the Custom Properties box. Make sure the property is not commented out by removing the hash (#) symbol in front of the property.
7. Change the value assigned to the cimom.provider.hds.useSecureConnection property to true, as shown in the following example:

   ```
   cimom.provider.hds.useSecureConnection=true
   ```

8. When you are done, click **Save**.

   To connect to another instance of HiCommand Device Manager using a nonsecure connection, prepend http:// to the discovery address to force the connection to nonsecure mode; for example, http://192.168.1.1. In this instance, 192.168.1.1 is the IP address of the host running HiCommand Device Manager.

9. The product notifies you if a restart of the AppStorManager service is required.

Unable to Discover a UNIX Host Because of DNS or Routing Issues

If the management server is unable to discover a UNIX host because of a DNS or routing issues, you must increase the amount of time that passes before the management server times out for that CIM extension. By default, the management server waits 1,000 ms before it times out. HP recommends increasing the time before the management server times out to 200000 ms (3.33 minutes), as described in the following steps. If you continue to see time-out issues, you can increase the time before the management server times out, but doing so will lengthen discovery.

To increase the time-out period:

1. Select **Configuration** > **Product Health**, and then click **Advanced** in the Disk Space tree.
2. Paste the following text into the Custom Properties box.

   ```
   cimom.cxws.agency.firstwait=200000
   cimom.cxws.agency.timeout=200000
   ```

   In this instance:

   `cimom.cxws.agency.firstwait` controls the amount of time required for the management server to wait after it first contacts the CIM extension on the host before the management server attempts to proceed with a username and password. The default value is 1,000 ms. You are modifying it to wait 200,000 ms or 3.33 minutes.

   `cimom.cxws.agency.timeout` controls the allowable interval of silence before either the CIM extension or the management server starts to question whether its partner is still alive. If one entity (management server or extension) does not receive a message from the other during the interval set by the timeout property, it sends an “are you there” message. If that message is not acknowledged during the interval set by the timeout property, the entity concludes that the connection is no longer functioning. The CIM extension stops attempting to make a connection. When this occurs on the side of the management server, the management server attempts to reconnect (and continues the attempt until the host becomes available). The default value is 1,000 ms. You are *modifying* it to wait 200,000 ms or 3.33 minutes.

3. Click **Save**.

The product notifies you if a restart of the AppStorManager service is required.

ERROR replicating APPIQ_EVAStorageVolume During Get Details for an EVA Array

Errors similar to ERROR replicating APPIQ_EVAStorageVolume might occur when an EVA-specific data cache is updated during a Get Details operation. For example, when Data Protector creates a snapshot, a new virtual disk is automatically created on the EVA array, and the EVA database used by the management server is updated to reflect this change.

If the EVA database is changed during a Get Details operation, small replication errors might be seen as a result. The array information will be updated with the correct information the next time Get Details runs.

Recalculating the Topology

When recalculating the topology or running Get Details, other tasks, using the management server can be delayed because recalculation is a resource-intensive operation. Recalculation occurs after a Get Details when provisioning is done and when you recalculate the topology manually.

During the recalculation period, you might not be able to log on to the application. If you are already logged into the application, navigation might not be possible until the topology recalculation is complete.

Display All Fabrics in Topology Cannot be Cleared

When you use the topology filter to “Display All Fabrics in Topology,” there is no option to clear the filter, and all fabrics continue to show. You must close down and restart your web browser to clear the filter.

Trunked ISL Label Appears Behind the Switch in Topology

The text label for an ISL trunk appears behind the switch and is partially hidden by the switch in the topology layout.

Brocade Fabrics Remain Connected in Topology even if the ISL Ports are Disabled

When you remove ISLs from a set of Brocade switches managed by SMI-S by disabling the ports, the management server topology will still show the ISLs as active.

Ensure that one switch per fabric is managed by the Brocade SMI-S proxy hosts, and perform a Step 3 Discovery on the management server to update the switch topology.

Troubleshooting the Java Plug-in

This section contains the following topics:

Incorrect Java Applets Cause Java Exceptions and User Interface Issues

In rare cases, the Java applets are not updated correctly. This can result in Java exceptions and user interface issues.

To resolve these issues:

1. Clear your web browser’s cache.
2. Restart the browser.
3. Clear the Java cache as follows:

a. Right-click the Java console, and select **Open Control Panel**.

b. On the General tab, click **Settings** in the Temporary Internet Files section.

c. Click **Delete Files**.

Unable to View Pages with the Java Plug-in on Linux and Solaris Clients

If your client is running Linux or Solaris, you will not be able to download the Java plug-in. You must manually install the Java plug-in as described in Installing the Java Plug-in on page 63.

Firefox on Windows is Unable to Download the Java Plug-in

(Windows only clients) If you do not have the Java plug-in already installed and you are running Firefox, you must use a web browser other than Firefox to install the plug-in. Firefox is unable to find the missing plug-in. After you install the plug-in, you can use Firefox to run the plug-in.

Java Applet Has Data from a Different Version of Management Server Software

If you attempt to monitor a host with old JAR (Java Archive) files, you might be unable to monitor the host, and you might see the following error message:

```
The Java applet has data from a different version of the management
server. Please close and re-start your browser.
```

The reason for this error message is that the client still has JARs from the previous version in its Java Plug-in cache. To remove the old JARs, clear the cache for the Java plug-in.

OutOfMemoryException Messages

In rare cases it might be necessary to increase the amount of memory for the Java plug-in on the client computer. This should only be done if you are seeing `OutOfMemoryException` messages in the Java console on the client side.

Unable to View System Manager after Upgrade

System Manager might not display if the Java applet plug-in for the Web browser is configured to use a proxy. This has been seen after the management server has been upgraded and the Web browser has cached Java class files. Clearing the cache does not correct this issue. The only known workaround is to disable the proxy.

Improving Reload Performance in System Manager

If your Java plug-in control panel cache is set at 50 MB, HP recommends increasing this setting to 150 MB or more. Increasing this setting improves the reloading performance of System Manager.

"The Java Runtime Environment cannot be loaded" Message

This error is caused when the Java Runtime Environment cannot allocate enough contiguous memory to start up with the requested settings. There are three workarounds for this problem. Attempt the workarounds in the order listed below. If the first workaround does not solve the problem, attempt the next listed workaround.

1. Access the product from a machine other than the one running the management server.
2. Use Firefox 2.0 or later with Java Runtime Environment 6 update 7:

   http://www.java.com/en/download/
3. Use Java Runtime Environment 6 update 10 beta:

   http://www.java.com/en/download/beta_6u10.jsp

Install the JRE Manually for 64-bit Clients

The product automatically downloads the correct JRE for clients on 32-bit operating system. If your client is on 64-bit operating system that is having difficulty rendering its applets to download, you should install the JRE manually to version 1.6.0_23 or later. Download the JRE from the following location:

http://www.oracle.com/technetwork/java/javase/downloads/index.html

HP Storage Essentials does not support the 64 bit web browser in this release. Use the 64-bit browser and JRE at your own risk.

Troubleshooting Provisioning

This section contains the following topics:

Cannot Access a Resource Owned by Another Controller

If you receive a message about not being able to access a resource owned by another controller, it is because you tried to access a controller that has not been discovered. You should discover all controllers on the LSI storage system.

For example, assume you discovered only one of the controllers on an LSI storage system with two controllers. To change a volume, such as add or delete a LUN, you will not be able to make the change to the volume associated with the controller that has not been discovered.

See Discovering Switches, Storage Systems, NAS Devices, and Tape Libraries on page 71 for more information on how to discover a controller.

Error -56

If you see error -56, the switch has network connection failures or problems. To solve the problem, make sure the switch is physically connected to the network, and then redo the task you were originally trying to complete.

If you now see -21(OBJECT_NOT_FOUND) errors, perform Get Details again for the switch.

"Can't delete this zone" Message

If you see the following message when you try to delete a zone, move the zone to an inactive zone set, and then delete the zone.

```
Can't delete this zone, it is member of an Active Zoneset
```

Changes in EFC Manager Requiring Get Details

If you use EFC Manager to delete zones or zone sets, perform Get Details on the management server afterwards. The changes are not reflected by the management server until Get Details is done.

Provisioning with Invalid HostMode2 Setting Partially Completes the Provisioning Operation

When provisioning a Host Storage Group (HSG) on an XP or HDS array, the provisioning operation will create the HSG if an invalid HostMode2 value is specified, but HostMode2 will not be set. You will see the following error: "java.lang.NumberFormatException: For input string: D." The HostMode2 input field does not properly check for the values entered. Edit the HSG to assign an appropriate HostMode2.

LUN Security Mode Sometimes Set to True Automatically

If a new host security group (HSG) is created on an HDS port that has LUN security disabled, the port LUN security mode is set to true automatically. The management server reflects this change when a full Get Details is performed on the device.

Custom Organizations

If you belong to a custom organization and you create a metavolume using volumes to which you have access, the resulting metavolume will not be available in your organization. The management server administrator or other suitably privileged account will need to move the metavolume into your organization.

Setting HSG Name

When you provision a host security group (HSG) on an array such as the Sun 6920 which is managed by an SMI-S provider, it is not possible to set the name of the HSG.

Incorrect Message About IBM ESS-800 Storage Systems Provisioning

When you create a volume on the IBM ESS-800, the management server returns right away, implying the volume creation is complete. However it might take a while for the ESS-800 to fulfill a request to create a volume. The ESS-800 does not indicate to the management server that creation of the volume has not finished. IBM is aware of this issue.

"You have selected a host that does not have a physical connection to the selected storage system!" Message

The message "You have selected a host that does not have a physical connection to the selected storage system!" is displayed if you attempt to create a zone on a switch that is not managed by the management server. This can happen if you select an initiator on a managed array that is physically connected to the unmanaged switch.

Troubleshooting Chargeback Manager

This section contains the following information:

"Name Contains" Filter in NAS Chargeback Returns Validation Error

Attempting to perform a filtering operation while editing NAS information in a storage tier will not work properly if Names Contained is specified. You will see a Validation Error on the page, and the items will not be filtered.

Creating Virtual Applications on the Host in Topology is the Preferred Method

Although it is possible to create an Application in the Chargeback Manager feature, these Applications might not be tied to a particular switch fabric and might not be shown or available in some parts of the management server user interface. Create Virtual Applications in the Topology by adding them directly to a host, and use the Virtual Applications in Chargeback Manager.

Business Cost Per Hour Field does Not Validate, Needs Refresh

In Chargeback Manager the Business Cost per Hour field does not validate the entry. No information is saved, and no error message is given if an inappropriate value is entered into that field. When the values are changed on that page, after you click to save the changes, you must do a page refresh (F5) in order to see the saved values.

Chargeback and Backup Applications

Disk-based backup media is not taken into account when calculating storage-based Chargeback Manager for a backup application.

Roles with Only Chargeback Manager Access

Roles with only Chargeback Manager access do not permit access to elements, even if the user has access to the Everything organization. Add the System Manager role to enable access to the various elements. See the user guide for details on creating and editing roles.

Incorrect Salvage Cost

Double Declining Balance and Fixed Declining Balance depreciation methods do not result in the correct Salvage Cost when the asset is fully depreciated.

Troubleshooting Host Virtualization

This section contains the following information:

Display of hdisks on IBM VIO Clients

When an IBM VIO client uses an hdisk directly, the management server displays the hdisk identifier with additional unnecessary characters. For example, in the Topology "hdisk0" will be shown as something similar to "hdisk0/Pseudo_lpar7".

ESX Servers with Non-Standard (All Zero) or Duplicate UUIDs

In some environments, ESX Servers have non-standard (all zero) or duplicate UUIDs. In these environments, the first ESX Server the management server discovers remains, but other ESX Servers with all zero or duplicate UUIDs are not shown by the management server. An error message is placed in the management server logs during Discovery when such an ESX Server is encountered. Ensure that all ESX Servers have an distinct non-zero UUID should there be difficulty discovering all the ESX Servers in the environment.

Copied VMware VMs Have the Same UUID Key

If you create a VMware VM and copy it to multiple machines, the virtual machines will report the same UUID key. In order to manage and report on these VMs properly, the management server software requires that the UUID be unique (as was intended by the VM software producer). When you deploy VMs, make sure they have distinct UUIDs.

VMware Size on Datastore is Inconsistent with Allocated Size

The Size on Datastore is reported by the VMware software as less than the Allocated Size and is displayed as such by the management server. This inconsistency has been reported to VMware.

Product Displays Unmanaged VMware Hosts

Although unmanaged VMware Hosts appear in Policy Manager, policies might not be created for any unmanaged host, including unmanaged VMware Hosts (VMware hosts without VMtools running or CIM extensions).

Backup Applications are not Supported on VMware Hosts

The management server does not support backup applications running on VMware Hosts. The Show Backup Topology button is disabled in the System Manager for VMware hosts.

Troubleshooting Hardware

This section contains the following topics:

About Swapping Host Bus Adapters

Swapping brands of host bus adapters (HBA) on a Microsoft Windows 2000 host could have undesirable side effects. For example, after swapping out one brand of an HBA for another (including driver installation), WinMgmt.exe might crash repeatedly and appear to be associated with an error in the Windows Event Log about being unable to retrieve data from the PerfLib subkey in the Registry. To solve this problem, reinstall the operating system.

"Fork Function Failed" Message on AIX Hosts

If a CIM extension running on AIX detects low physical or virtual memory when starting, a “Fork Function Failed” message appears.

A CIM extension on AIX uses additional memory and CPU resources at start time. If the resources on the AIX machine are already low, you might see the “Fork Function Failed” message. Depending on the AIX operating system or hardware, the host might crash after you see this message.

Known Driver Issues

Keep in mind the following:

- The software requires the driver to have a compliant SNIA HBA API. Emulex driver version 4.21e does not support the SNIA HBA API.
- If the driver has a compliant SNIA HBA API, make sure the driver is installed correctly.

Known Host Issues

The following table provides a description of the known device issues. You can find the latest information about device issues in the *Release Notes*.

- Support of recent changes to the Daylight Saving Time (DST) start and end dates is not included in all Java Runtime implementations. The OpenVMS/IA64 CIM Extension does not currently use the latest version of Java, so some features, such as log timestamps, are not be in synch with DST changes. A Daylight Saving Time patch is available from Sun Microsystems for most operating systems. After installing the agent on the host, download and unzip the JDK DST Timezone Update Tool from:

  http://www.oracle.com/technetwork/java/javase/downloads/index.html

  Run the tzupdater tool using the JRE included with the CIM extension (usually in directory /opt/APPQcime/jre). For example, on UNIX:

  ```
  setenv JAVA_HOME /opt/APPQcime/jre
  ```

```
$JAVA_HOME/bin/java -Djava.home=$JAVA_HOME -Djava.vendor="Sun
Microsystems Inc." -jar tzupdater.jar -u
```

- The Unmounted Volume box under Capacity Summary automatically displays 0 MB if you discovered the host but not the storage system connected to it. This might occur if you did not enter the IP address of the storage system when performing discovery, or if your license does not allow you to discover a particular storage system. See the support matrix for your edition to determine which storage systems you can discover. The List of Features is accessible from the Documentation Center (**Help** > **Documentation Center**).
- (Windows host VxVM) The SCSI bus number is always reported to be 1 in the SCSI bus column of the Disk Drives page.
- (Solaris hosts on VxVM) If you discover a host with any typical SAN disk groups off line, the storage volume page shows SAN mount points as local instead of external. These disks, however, are not accessible. When you perform Get Details with all disk groups online, disks on the SAN are shown as external. Hosts connected directly to a storage system are shown as local, except for hosts connected by fibre. Hosts connected directly to a storage system through fiber are shown as external.
- The following issues exists for Solaris hosts using HDLM:
  - If you discover the Solaris host by itself without the switches and storage, the storage volume page reports all drive types as local. Once you discover the host with the switches and storage, it reports its drives as being external. It reports the same result with Active-Active and Active-Standby.
  - Solaris HDLM disks cannot be monitored. If you try monitoring them, the management server displays a message saying “data is late or an error occurred.”
  - If you do a Get Details for the host by itself, on the bindings page, the controller number begins with c-1; for example, c-1t0d58. Perform Get Details on the host with storage and switches. The controller numbers are displayed correctly.
- (Solaris hosts using Sun SAN Foundation Suite driver (Leadville driver)) The bindings page reports a SCSI number that comes from the HBAAPI. This number cannot be seen by the user. For example SCSI target 267008 does not correlate to anything
- (AIX hosts)If you are receiving replication errors for an AIX host, the provider might be trying to connect to the host using the 0.0.0.0 IP address instead of the real host IP address. If this situation occurs, you see a message containing the following when you start the CIM extension:

  ```
  CXWS 3.1.0.144 on 0.0.0.0/0.0.0.0 now accepting connections
  ```

  To fix this, add the following line to the /opt/APPQcime/tools/start file on the AIX host:

  ```
  export NSORDER=local,bind
  ```

- (AIX hosts using an IBM Storage System) If you have an AIX host using an IBM storage system, not all bindings might be displayed on the bindings page on the Navigation tab. For example, assume diskA on host123 has six paths. All six bindings might not be displayed.

- (Hosts running SGI IRIX version 6.5.22 or 6.5.24) If a host is running SGI IRIX version 6.5.22 or 6.5.24, the HBA port page on the Navigation tab in System Manager displays 0 GB/s for HBA ports.
- (SGI IRIX hosts on CXFS file systems) The management server can only monitor CXFS file systems from the host generating the input/output. For example, assume the elements are part of a CXFS file system. When you generate input/output into the metadata server into /folder, only the metadata server is able to monitor the file system. For example, assume the metadata server generates 100 KB write, the management server displays 0 KB write for /folder on the metadata client. For more information, see About the Data from CXFS File Systems on page 518.
- When a host running IBM Director software is discovered by the management server, the discovery process will be delayed for up to 10 minutes waiting for a time-out period to expire.
- Internal volume information (mimage_0, mimage_1, log) is shown on the Volume Management page for a mirrored LVM volume on a RedHat host.
- The fdisk partitions p1-p4 on Solaris x86 hosts are not reported back to the management server. In addition to the disks not being displayed in the management server user interface, capacity numbers for the Solaris x86 hosts are also affected.
- Real time Processor Utilization is not shown for Tru64 hosts. The message "Data is late or an error occurred" appears.
- The Host Unused Capacity for HP-UX hosts include the capacity from any DVD device on the host.
- The HP-UX CIM Extension does not report capacity for a VxFS file system on HP-UX where the file system's size exceeds 2TB.
- The presence of special agile devices on HP-UX causes the local disk to appear on the Multipathing Tab for the host.
- Running Update Element Data (Single Element Refresh) for a host duplicates multipathing and capacity information for that host. Run Step 3 Discovery to clear the duplication.
- The model number for the AH403A HBA is not shown when installed on HP-UX 11.31 hosts due to an issue in the SNIA HBAAPI library.
- NonStop (NSK) Itanium Model Reported as “NSE-A”, J-Series reported as H-Series. The Model reported on the Host Navigation page for NonStop Itanium machines is shown as NSE-A. The Model reported on the Host Navigation page is H-Series NonStop J-Series machines.
- The Switch Ports displayed on the Host's Dependent Switches page reflect the switch port directly connected to the host for each fabric.
- Veritas with the Japanese Language Pack. Veritas software running with the Japanese Language Pack does not return the appropriate information to the Windows CIM extension. Hewlett-Packard is in contact with Veritas about this issue.
- Tru64 AdvFS and NonStop-specific file systems are not modeled properly in the management server software at this time. The capacity calculations for the NonStop host erroneously

combine the capacities of the OSS filesets on the host along with that of the Guardian volumes while computing the aggregated capacity of the host. Capacity for Tru64 AdvFS does not take into account that disk blocks can be shared in an AdvFS Domain.

- Unmounted Volume Capacity not Reported for Windows Hosts. Capacity Manager will not report unmounted volume capacity for Windows hosts. Unmounted volumes appear correctly on the Volume Management page.
- The LP9002 HBA Reports 0 Gbps Port Speed on OpenVMS. The port speed for the Emulex LP9002 Host Bus Adapter is reported as 0 Gbps on OpenVMS hosts.
- HBAs and HP-UX. The Link Failure counter does not report data for most HBAs supported on HP-UX. The A5158A HBA does report values correctly.
- Volume Management Does Not Display all Mount Points for Veritas Volumes with Multiple Mount Points. In a configuration where Veritas Volume Manager volumes have multiple mount points (mounted as a drive letter and then as a reparse point, for example), the Volume Management page shows only one mount point for the volume.
- Compaq RAID Arrays Are Incorrectly Reported Under Linux. Local disks on Compaq RAID arrays are not reported correctly under Linux.
- Drives Are Not Shown Attached. Drives are not shown attached to the Compaq Smart Array Controller on the Properties page. From the host Properties page, click the link for the array controller. No drives appear on this page.
- RAID Volume Capacity for Windows Dynamic or Veritas RAID5/Mirror Reparse Points Reported Incorrectly on the Storage Volumes Page. Capacity reported for RAID5 or Mirror Windows dynamic volumes or Veritas Volume Manager volumes that are mounted to a Windows directory (reparse point) without a drive letter do not display the correct capacity on the Storage Volumes page. Total Capacity, Total Used, Available, and Percentage Used consider the whole volume as usable storage when it should only be a portion of the volume when considering the RAID configuration.
- Issues with Solstice DiskSuite. Only Solstice DiskSuite slices that are in use will be reported on. The metadatabase slices are not reported on. Currently the descriptions given for DiskSuite slices are inconsistent.
- NonStop and Disk I/O Information. The management server is unable to gather or present disk I/O information for NonStop in this release.
- LUN Information Missing in QLogic Failover Configuration. In a QLogic Failover configuration, LUN information is missing from the Multipathing page unless the LUN is visible through all HBAs.

"Mailbox command 17 failure status FFF7" Message

If one or more of your Microsoft Windows hosts are using an Emulex HBA driver, you might see the following message in Windows Event Viewer:

```
mailbox command 17 failure status FFF7
```

This message can be safely ignored. The HBAAPI is being used to access data in the flash memory of the adapter that does not exist, and this is causing the event to be logged. This issue has been seen with version 5.2.2 of the driver.

"Process Has an Exclusive Lock" Message

You will receive a message like the following if a process locked the EMC Symmetrix storage system and you attempt a process that requires a lock on the Symmetrix storage system.

```
SYMAPI routine SymDevMaskSessionStart failed with error code 188: The operation failed because another process has an exclusive lock on the local Symmetrix.
```

The Symmetrix storage system can become locked for many reasons. For example, the storage system becomes locked when it performs LUN mapping, LUN masking, or Get Details. The Symmetrix storage system can also remain locked after a provisioning operation has failed.

After the management server detects the lock on the Symmetrix storage system, it tries to access the storage system for 15 minutes and then logs the errors.

If you receive the error message, determine if someone is performing an operation that requires a lock, such as LUN mapping, LUN masking, or Get Details. This also applies even if one of the processes is being used by a third-party product, such as for LUN masking. Wait until the process is complete before you remove the lock manually. Make sure that no other processes are occurring on the storage system. To learn how to remove the lock, see the documentation for the Symmetrix storage system.

If a provisioning failure caused the Symmetrix storage system to remain locked, you are alerted to this situation in Event Manager and on the Properties tab. You could receive a message like the following:

```
Unable to end device masking session. Symmetrix '000001835005700' may be locked.
```

Known Issues with Switches

- The Management URL for Cisco Virtual Switches is Incorrect. The Management URL presented on the properties page for Cisco Virtual Switches incorrect. It does not point to the expected target.
- Physical Switch Properties for Cisco Switches discovered through SNMP. The following physical switch properties are missing or incorrect for Cisco switches discovered through SNMP: Switch Sate; Status; Domain ID; Role; Hardware Zoning Capabilities; Software Zoning Capabilities; Current Zoning Enforcement; Max. Number of Modules; Zoning Installed.
- Aggregated Real Time Port Statistics are available for Cisco Switches. Real Time port statistics are available aggregated per physical switch, not on a per port basis, for Cisco switches discovered through SNMP or SMI-S. Cisco virtual switches do not show either aggregated or per port Real Time performance statistics.
- Cisco Switch Provisioning with switches discovered by both SMI-S and SNMP. In environments where some Cisco switches are discovered through SMI-S and some through

SNMP, adding a Zone Alias or WWN to a Zone shows the message "Provisioning not supported", and the provisioning operation does not proceed. To work around this issue, create a new Zone with the desired members, and replace the Zone you wanted to change with the new Zone. Activate the changes.

- Cannot remove a Zone Alias from a Zone with Cisco SMI-S Provisioning. An issue in the Cisco SMI-S provider prevents the removal of a Zone Alias from a Zone during a provisioning operation. Use the tools provided by Cisco to remove a Zone Alias from a Zone.
- Brocade Virtual Switches shown in Step 3 Discovery. When a Brocade proxy server is discovered, it returns a list of virtual switches to the management server. These are used in Step 3 Discovery. This differs from other switch vendor proxies which return physical switches. Brocade has been notified of this issue.
- Duplicate E Ports shown for CISCO Multi-VSAN ISL. Duplicate E ports are shown in the port list for all fabrics for multi-VSAN ISLs on CISCO switches. Logical ports are shown instead of physical ports. A correction will need to be made to the CISCO provider to resolve this issue.
- CISCO 8Gb/s switch ports reported as 0Gb/s. CISCO switches with ports higher than 4Gb/s report the switch port speed to HP Storage Essentials as 0Gb/s, which is what HP Storage Essentials then reports. CISCO has been made aware of this issue.
- Cisco SMI-S DomainID:PortNumber Format. The Cisco SMI-S provider does not always handle DomainID:PortNumber format properly when creating zones or reporting on zone content. Existing zones on Cisco switches with DomainID:PortNumber members might not be displayed correctly. The management server UI has been changed to prevent the creation of zones with DomainID:PortNumber style members.
- Some Inactive Zone Aliases do not Appear in the Associated Zones on Cisco SNMP Switches. On Cisco switches managed through SNMP, some inactive zone aliases are not shown in the zones to which they belong.
- Cisco Provisioning Fails if Enhanced Zoneset Feature is Enabled. If the Enhanced Zoneset feature is enabled on a Cisco switch, provisioning fails with a message similar to the following: “ZoneAlias creation failed with error code 1234.” Hewlett-Packard is working with Cisco to resolve this issue.
- Cisco Port Type. TE ports on Cisco switches discovered through SMI-S are shown in the management server as E ports. This is because the Cisco SMI-S provider returns the E port type for TE ports.
- Fabric Shown as the Source for some McDATA Switch Events. The McDATA SMI-S provider does not return source information for some events (indications), such as those events dealing with fan or power supply issues. Such events are reported in the management server user interface as coming from the fabric, not the individual switch.
- Brocade SMI-S Active Zone/Zone Set Does Not Contain Aliases. When you use Path Provisioning to provision a zone using aliases on a Brocade switch managed through SMI-S, the management server user interface will show you the results of the zoning operation from the active zone set. Although you might have used aliases to create the zones, the active zone set uses actual WWNs, not aliases. This can be misleading, but it is expected in this situation because of the way the Brocade SMI-S provider operates.

- Unmanaged CISCO Switches Display as Unmanaged Hosts. CISCO Switches discovered through SMI-S will show ISL'ed switches as unmanaged hosts if the ISL'ed switches are not managed by the management server. If you discover and manage the ISL'ed switches, they will be shown as switches.
- Changes to Inactive Zone Sets on Cisco Switches. On Cisco switches discovered through SMI-S, provisioning changes for inactive zones are reverted if the switch reboots or if the switch loses power. Changes made to active zones and zone sets are saved across switch reboots. To make changes to inactive zones and zone sets permanent, use the Cisco native tools. Follow these steps:
  - Start the Cisco Fabric Manager. See your Cisco documentation for more information if needed.
  - Under Physical Attributes, expand the Switches folder and select **Copy Configuration**.
  - In the upper left of the screen, select the desired switches for which you want to save the running configuration to the startup configuration.
- Cisco Fabrics Connected through FC-IP with IVR Returns Zonesets from only the First VSAN. In Cisco fabrics where the connections are made through FC-IP with IVR, the management server will report on zonesets that are in the first VSAN encountered in the configuration. Other zonesets will not be reported.

Known Issues with Arrays

- IBM DS5300 Capacity Numbers May Differ from the Native Point Tool. The "Unused Raw" capacity reported by the management server for IBM DS5300 arrays might differ from the information returned by the IBM tools for managing the array.
- "Primary key violation" while Collecting Performance Metrics for NetApp. When performance metrics are collected for NetApp devices, you may note "Primary key violation / unique constraint (APPIQ_SYSTEM.NAS_LUN_STATS_PK) violated" messages in the management server logs. Data collection is not affected, and these messages may be ignored.
- Space Efficient (Thin Provisioned) Volumes not shown in HSGs for IBM DS8000 Arrays. Due to an issue in the provider, Space Efficient (Thin Provisioned) volumes are not shown in HSGs for IBM DS8000 series arrays. The management server logs indicate an error similar to the following: "ERROR replicating ":IBMTSDS_ProtocolControllerForUnit.Antecedent ... Cannot find reference in database for: IBMTSDS_Volume".
- IBM DS5300 Pool Names Differ from the Native Tool. For the IBM DS5300, the management server displays Pool Names similar to "Volume Group 1", "Volume Group 2", and so on. The native tool to manage the array shows "Disk_Group_1", "Disk_Group_2", and so on.
- Capacity Manager "% Over-Allocated" Value for Thin-Provisioned Arrays). The "% Over-Allocated" value in Capacity Manager is calculated by dividing the "Virtual Allocated" by the "Total Capacity". The heading in Capacity Manager implies that the value given is the percentage exceeding the Total Capacity.

- EVA Post-RAID Capacity does not match the value from Command View EVA. Post-RAID capacity reported by the management server does not match the capacity reported by Command View EVA. Command View EVA includes items such as replication log file growth and meta-data overhead which are unavailable through the EVA SMI-S provider the management server uses to discover the EVA arrays.
- Celerra Processor Count is 1. The processor count reported for Celerra configurations is always 1 (one).
- Quota for "/" shown for EMC Celerra). The management server user interface shows a tree quota for "/" on the Volume Composition Page for the EMC Celerra.
- EMC Volume Property not populated until Detailed Discovery. The EMC Volume Property does not contain disk drive and member information for a newly created Meta volume until a Step 3 Detailed Discovery is complete.
- Port Speed 0 is shown for some Back End Ports on SVSP. When the management server requests information from an SVSP, some back end switch ports return a port speed of 0.
- Outdated Data from the SMI-S Provider for MSA Arrays. The management server requests information from MSA Arrays using the MSA SMI-S provider. The MSA SMI-S provider does not always report the latest configuration information from the MSA Array. There may be a delay between the time the array is changed and the time the MSA SMI-S provider updates the information it has for the array. For example, you may create a volume on an MSA Array using the tools provided with the MSA Array, but the management server will not see the new volume until the MSA's SMI-S provider contains the updated information and a Detailed Discovery or a Single Element Refresh is performed.
- HDS 9970V Returns 0 for Cache Statistics. The HDS 9970V (RAID-450) arrays return 0 for the Percent Write Pending Data and Percent Cache Used cache statistics, so this information will be displayed as 0 by the management server for this model of array. The vendor has been notified of this issue.
- Physical Disk Status, Associated Disk Group, and Ungrouped Disk Capacity not Updated after Grouping a Disk Drive Using “Add Disks” on the Disk Group Properties Page in Command View EVA . Physical disk status, associated disk group, and ungrouped disk capacity are not updated after grouping a disk drive using the “Add Disks” functionality on the disk group properties page in Command View EVA. As a workaround, use the “Group” or “Ungroup” functionality on the disk drive properties page in Command View EVA to affect the desired changes on the EVA. When done this way, HP Storage Essentials will be able to obtain the correct disk group property information from the array.
- XP Array Port Traffic and Continuous Access and External LUNs. When an XP array is enabled for Continuous Access or External LUNs, the statistics initiator does not return the Total I/O Rate and the Total Data Rate. As a workaround, use the associated switch or host port statistics.
- Anomalous Historical Performance Data Points from the Legacy Built-In EVA Provider. The legacy built-in EVA provider might return anomalous data points (spikes) for historical performance statistics. This is caused by an issue in the firmware on EVA arrays. A solution is to move to a newer version of Command View EVA (9.1 and later), and the management

server will begin to use the new external EVA SMI-S provider, which is unaffected by this issue.

- The External EVA SMI-S Provider Reports Incorrect Used and Available Space when creating certain Volumes. The external SMI-S provider for the EVA arrays reports incorrect used and available space after a storage pool with more than eight disks has been created. Eventually the SMI-S provider returns the correct information, although this is not immediate. The correct storage is reported if the same volume is created using the EVA point tools.
- Virtual Volume (V-VOL) and Virtual Array Group Utilization % is shown as 0. In Performance Manager the Utilization % for THP/Snapshot Virtual Volumes (V-VOLs) and Virtual Array Groups is shown incorrectly as 0%.
- Some NPIV Ports show as Connected to Unmanaged Storage Arrays. When an N_Port ID Virtualization (NPIV)-enabled switch port is physically connected to a Brocade Access Gateway (AG) switch port or Cisco N_Port Virtualizer (NPV) switch port, the NPIV switch port shows the remote AG or NPV port as an unmanaged storage array. In Topology the NPIV switch will display an unmanaged storage array for each port connected to an AG or NPV switch. Switches running in Access Gateway or NPV mode cannot be managed by the management server at this time. It is not possible to remove the unmanaged storage arrays that appear in the topology.
- The MSA SMI-S provider is unstable with 4 or more arrays (QCCR1G32013, QCCR1G32014 (formerly IEV-27310)). The MSA SMI-S provider has been shown to be unstable when managing four or more arrays. HP Storage Essentials is unable to connect to the provider, indicating “Connection refused.”
- Cannot Always Delete Selected Volume on MSA. MSA volumes must be deleted in the reverse order of their creation. For example, if you have six volumes, and you want to delete the second one you created, you must delete the volumes one at a time, starting with the volume created sixth and continuing with the fifth, fourth, third, and then the second. Attempting to delete other volumes will return a generic error code 4. For this operation you might wish to consider using the MSA native tool.
- Use Edit, not New Host Security Group, for EVA Provisioning if Host Storage Group (HSG) Exists. When editing an EVA Host Storage Group (HSG) in Wizard Provisioning, select the initiator if available in the table on the Step 3 Host Security Groups pane and use the edit function. Do not use New Host Security Group to modify an HSG. If you select New Host Security Group and then select an initiator that is already associated with an HSG, you will see a CIM_ERR_FAILED java.lang.reflect.InvocationTargetException error. New Host Security Group is for creating, not modifying, an HSG.
- Incorrect Size Range Displayed while Creating an EVA Storage Pool. Selecting advanced options when creating an EVA storage pool could show an incorrect size range for available space. Advanced options can include reserving a drive for parity, yet this drive space is not included in the size calculation of available space. Please keep in mind that when using advanced options for storage pool creation on EVAs that some space might not be available based on the options you choose.
- Replication Errors for HDS or HP XP Arrays. Customers with HDS or HP XP series arrays might see replication errors for the HdsFCPort and HdsProtocolControllerForPort.Antecedent

classes during the first Get All Details operation after upgrading to SP4. To resolve this issue, run Get All Details again. The errors will no longer appear, and the management server information about the arrays will be updated properly.

- IBM DS6000 Array Port Speed Shown as 0 Gbps. The array port speed reported for the IBM D6000 array is 0 Gbps. The IBM 5.2.1 CIMOM used to manage the array does not report the port speed for this array (it is reported as 0).
- Unused Raw Capacity on IBM DS Arrays. For IBM DS series arrays using IBM CIM Agent version 5.2.1 and possibly higher, raw capacity of ArraySites is used to compute capacity for IBM DS series arrays if any ArraySites are not formatted into Arrays. If the ArraySites are not formatted into Arrays, the raw capacity of the disks comprising the ArraySites is not used to represent Unused Raw Capacity.
- IBM SAN Volume Controller Mirrored Virtual Disks Show only in the Primary Storage Pool. Mirrored virtual disks on IBM SAN Volume Controllers are shown in the primary storage pool only.
- Provisioning is not Supported for 3PAR Arrays. Although permitted in the management server user interface, provisioning does not work correctly for 3PAR arrays and is not supported at this time.
- EVA Host Security Group Information Unavailable if iSCSI Link Down. Host Security Group information is not gathered for EVA arrays that have an iSCSI module where the network link is down. If you have several EVAs discovered through Command View, all EVAs will be affected if one EVA has a link down on an iSCSI module.
- Command View EVA HSG Folders are not Supported. Folders associated with Host Security Groups on EVAs managed through Command View are not shown in the management server user interface. Attempting to provision volumes using folders that differ in name by case can create a volume in the incorrect folder.
- MSA Capacity and Variable RAID Overhead. The RAID overhead for MSA disk groups varies depending on the RAID level and the number of disks. MSA capacity calculations done by the management server do not take into account the variability of the RAID overhead at this time.
- Problems Performing LUN Security Operations and Zoning on a CLARiiON Array. If there is a new host bus adapter (HBA) attached to a CLARiiON system, performing LUN Security operations and zoning on the array might fail due to the registering process of the new HBA to the CLARiiON port. The work around is to perform zone operation first (create zone) and then perform LUN Security.
- IBM Subsystem Device Driver (SDD) or MPIO (multipath I/O).If you discover an IBM storage system without SDD, incorrect stitching is displayed in System Manager for the storage system. You are shown only one path if the storage system is using MPIO instead of SDD.
- Disks from an XP array will remain on the Bindings page even if you remove the disks from the hosts an run a full Get Details. This is caused by an issue in the SNIA HBAAPI library in the way it reports disks no longer presented to a particular host.

Glossary

A

access point

The intersection of the IP address and the provider that discovered the IP address. It is displayed on the screens for discovery. A provider is software that is used to gather information about an element.

active zone set

The zone set in use. You can have only one zone set active at a time; however, you can have a zone in more than one zone set. Zones sets are usually created for a particular task.

Zones work by exclusion. Members of a zone can communicate with other members in the zone. If two devices are not within the same zone, they cannot communicate. Only elements in active zones can communicate with each other. When a zone is not active, it does not have any effect.

C

Common Information Model (CIM)

A common data model of an implementation-neutral schema for describing overall management information in a network/enterprise environment. CIM is comprised of a specification and a schema. The specification defines the details for integration with other management models (i.e. SNMP's MIBs or the DMTF's MIFs) while the schema provides the actual model descriptions.

Common Information Model Object Manager (CIM Object Manager)

A component in the CIM management infrastructure that handles the interaction between management applications and providers. The CIM Object Manager supports services such as event notification, remote access, and query processing. The CIM Object Manager also grants access to the CIM Object Manager repository.

D

device

This documentation set defines a device as a piece of hardware in the storage network.

EMC uses the term device to refer to a volume on one of its storage systems.

E

element

Anything on the network that can be detected by the management server, such as hosts and switches.

An element created in Chargeback Manager is a type of generic element. When you create a record for an element in Chargeback Manager, the element appears as a generic element in the topology.

F

frozen

When media is referred to as frozen in Backup Manager, the media will never become available again, but it is still available for restores.

File Server Storage Resource Management (SRM)

File Server Storage Resource Management (SRM) does a recursive lookup on the file system and stores the information in an embedded database. File System Viewer can scan files very quickly because of its structure in the database and because it uses a multi-threaded process. More than one process can be used at a time to scan the files.

G

generic element

An element is considered to be generic if the management server can detect the element but it cannot obtain additional information about the element during Get the Topology or Get Details.

global reporting view

A global reporting view contains information in the database that can be used for global reports.

Global Reporter server

A management server that has global reporting enabled.

H

hard zone

A hard zone is created by assigning a domain/port to a zone. Any device attached to the port is automatically in the zone.

host persistent binding

A system SCSI target ID assigned permanently to an element. The host binding is implemented on the host bus adapter (HBA), resulting in the HBA being tied to a certain LUN.

I

initiator WWN

The Worldwide Name (WWN) of a host bus adapter's port. The WWN differentiates the port from others.

inode file

An inode file stores information about a file, excluding the file’s data.

M

Managed Application Licenses (MALs)

Managed application licenses (MALs) are the number of detected instances of Microsoft Exchange, Oracle, SQL Server, Caché, and Sybase Adaptive Server Enterprise.

managed object

A hardware or software system component that is represented as an instance of the CIM class. Information about managed objects is supplied by data and event providers, as well as by the CIM Object Manager.

Managed Access Points (MAPs)

The sum of all storage access ports of all hardware elements that the management server manages.

materialized view

A materialized view is a snapshot of data, from the database, created from a query.

mapped

Capacity that is accessible by one or more hosts external to the array (aggregated capacity of volumes that are accessible from hosts external to the subsystem).

meta device

A term is used by EMC. A meta device is a device that is a concatenation of several devices.

metavolume

Metavolumes are created from a disk, slice, stripe, or other metavolumes. Metavolumes are extremely useful because they can expand their storage capacity, such as to mainframe volume sizes. Also referred to as “LDEVs” for HDS storage systems.

missing element

The management server was able to discover the element, but it lost contact with the element before more information could be gathered during “Get the Topology” or Get Details. A missing element can be managed if the management server lost contact with the element after Get Details was performed.

multipathing

The process of providing a server more than one path to a storage system. So that in the case of an emergency, the server will have continuous access to the storage system. Multipathing can be done many ways. For example, you can provide redundant switches for a server to access a storage system. Another example of multipathing is providing redundant paths from the server to the switch.

P

provider

Software used to gather information about an element.

proxy server

A device, such as a host, connected to a storage system. It is sometimes referred to as a storage system proxy or an API proxy. An example of a proxy server is the EMC Solutions Enabler or Hitachi HiCommand Device Manager. LSI storage systems do not require a proxy, as they can be accessed directly.

S

SAN

A Storage Area Network (SAN) is a high speed network configuration that is dedicated to transporting storage data among network devices, such as storage systems, hosts (servers), switches, and tape libraries to end users. In addition to connecting local elements to storage arrays, it might also be extended to off site or remote locations for the purposes of backup, archival or acting as a hot site in the event of a disaster.

A SAN can communicate via current technologies such as ESCON (mainframe), fibre channel, or newer technology such as iSCSI. SAN's can support several configurations such as disk mirroring, RAID 5, backup/restore, and data migration, as well as being able to incorporate Network Attached Storage (NAS).

SMTP

A TCP/IP protocol used in sending and receiving e-mail.

soft zone

A soft zone is created by assigning a world wide name (WWN) of a device port to a zone.

Storage Management Infrastructure Specification

A Storage Networking Industry Association (SNIA) standard for implementing data storage management using the Common Information Model (CIM).

storage pool

A group of volumes. Also known as volume group.

switch port zoning

A type of zoning in which the port of the switch is physically in the zone. Any device attached to the port is automatically in the zone.

V

Virtual Storage Area Networks (VSAN)

See VSAN.

Virtual Application

A placeholder you create for an unsupported application. For example, assume your company has created an internal application, and you want to be able to use the software to keep track of that application. You can create a virtual application for that product.

VSAN (Virtual Storage Area Networks)

Logical fabrics formed as subsets of physical FC (Fibre Channel) switch networks. A VSAN is defined as a set of Fx, E and TE ports-entry/exit, traditional ISL (Inter-switch Link) ports, and trunked ISL ports. TE ports are ISLs that might be shared between a named set of VSANs. All ports other than TE ports must be members of exactly one VSAN. On Cisco switches, there is a default VSAN (VSAN 1) that initially includes all ports, and an isolated VSAN (VSAN 4094) where ports end up if their owning VSAN is deleted.

W

Web-Based Enterprise Management (WBEM)

An initiative based on a set of management and Internet standard technologies developed to unify the management of enterprise computing environments. WBEM provides the ability for the industry to deliver a well-integrated set of standard-based management tools leveraging the emerging technologies such as CIM and XML.

Windows Management Instrumentation (WMI)

Microsoft created WMI as its implementation of Web-based Enterprise Management (WBEM). For more information about WMI, refer to the Microsoft Web site at http://www.microsoft.com.

The Windows CIM Extension enhances Windows Management Instrumentation (WMI) so that it can gather information from host bus adapters and make the information available to the management server.

Worldwide Name (WWN) zoning

A type of zoning in which the port is assigned to a Worldwide Name of a host or a storage system. It is not dependent on the switch.

WWN (Worldwide Name)

Fibre-channel Worldwide Name. Usually formatted as 16-hexadecimal digits. This name is globally unique, and it identifies the connection or set of connections to the network.

zone

A collection of zone aliases and ports.

zone alias

To avoid remembering a port's Worldwide Name (WWN), assign the port to a zone alias.

zone member

A port attached to a switch, a Worldwide Name (WWN) or a zone alias.

As a best practice, a zone should contain either zone aliases or ports, but not both.

zone set

A collection of zones. You can have only one zone set active in a fabric; however, you can have a zone in more than one zone set. Zones sets are usually created for a particular task.

Materialized Views

If you created reports in previous releases, you are most likely using old views. However, some of the views have changed and might no longer work for you. You should verify that your existing reports will work correctly against these new views. The following is an alphabetized list of the materialized views.

To see the mapping between the older and new views, see Views from Previous Releases on page 1008.

Materialized View Name	See
MVA	MVA on page 925
MVC_APPLICATIONRELATIONVW	MVC_APPLICATIONRELATIONVW on page 926
MVC_APPLICATIONSUMMARYVW	MISAPPLICATIONS on page 926
MVC_ASSETSUMMARYVW	MVC_ASSETSUMMARYVW on page 927
MVC_CARDSUMMARYVW	MVC_CARDSUMMARYVW on page 929
MVC_ CHANGERDEVICESUMMARYVW	MVC_CHANGERDEVICESUMMARYVW on page 930
MVC_DISCOVERYDETAILSVW	MVC_DISCOVERYDETAILSVW on page 931
MVC_DISKDRIVESUMMARYVW	MVC_DISKDRIVESUMMARYVW on page 932
MVC_DISK_EXTENTVW	MVC_DISK_EXTENTVW on page 933
MVC_DISKEXTENTSUMMARYVW	MVC_DISKEXTENTSUMMARYVW on page 933
MVC_EVENTVW	MVC_EVENTVW on page 933
MVC_EXTENTEXTENTMAPVW	MVC_EXTENTEXTENTMAPVW on page 935
MVC_FABRICSUMMARYVW	MVC_FABRICSUMMARYVW on page 935
MVC_HBAPORTTARGETS	MVC_HBAPORTTARGETS on page 935
MVC_HOSTCAPACITYVW	MVC_HOSTCAPACITYVW on page 936
MVC_HOSTDISKDRIVEVW	MVC_HOSTDISKDRIVEVW on page 936

Materialized View Name	See
MVC_HOSTRELATIONVW	MVC_HOSTRELATIONVW on page 937
MVC_HOSTSUMMARYVW	MVC_HOSTSUMMARYVW on page 937
MVC_HOSTVOLUMESUMMARYVW	MVC_HOSTVOLUMESUMMARYVW (Logical Volumes) on page 938
MVC_LOCALIZEDSTRINGVW	MVC_LOCALIZEDSTRINGVW on page 939
MVC_LUNMAPPINGVW	MVC_LUNMAPPINGVW on page 939
MVC_MULTIPATHVW	MVC_MULTIPATHVW on page 940
MVC_NASEXTENTSUMMARYVW	MVC_NASEXTENTSUMMARYVW on page 941
MVC_NASQUOTASUMMARYVW	MVC_NASQUOTASUMMARYVW on page 941
MVC_NASSNAPSHOTSUMMARYVW	MVC_NASSNAPSHOTSUMMARYVW on page 942
MVC_NAS_ FSDIRECTORYSUMMARYVW	MVC_NAS_FSDIRECTORYSUMMARYVW on page 943
MVC_NAS_FSSHARESUMMARYVW	MVC_NAS_FSSHARESUMMARYVW on page 943
MVC_OPTIONALTABLEVW	MVC_OPTIONALTABLEVW on page 944
MVC_ORGANIZATIONVW	MVC_ORGANIZATIONVW on page 944
MVC_ORGFILTERVW	MVC_ORGFILTERVW on page 945
MVC_PATHVW	MVC_PATHVW on page 945
MVC_PORTCONTROLLERMAPVW	MVC_PORTCONTROLLERMAPVW on page 946
MVC_PROTOCOLCONTROLLERVW	MVC_PROTOCOLCONTROLLERVW on page 946
MVC_PORTSUMMARYVW	MVC_PORTSUMMARYVW on page 947
MVC_ROLEVW	MVC_ROLEVW on page 948
MVC_STORAGEPOOLCONFIGVW	MVC_STORAGEPOOLCONFIGVW on page 949
MVC_ STORAGEPROCESSORSUMMARYVW	MVC_ STORAGEPROCESSORSUMMARYVW on page 952

Materialized View Name	See
MVC_STORAGESYNCHRONIZEDVW	MVC_STORAGESYNCHRONIZEDVW on page 952
MVC_STORAGESYSTEMCONFIGVW	MVC_STORAGESYSTEMCONFIGVW on page 953
MVC_ STORAGESYSTEMSUMMARYVW	MVC_STORAGESYSTEMSUMMARYVW on page 954
MVC_STORAGETIERDETAILVW	MVC_STORAGETIERDETAILVW on page 955
MVC_STORAGEVOLUMEPORTS	MVC_STORAGEVOLUMEPORTS on page 956
MVC_ STORAGEVOLUMESUMMARYVW	MVC_STORAGEVOLUMESUMMARYVW on page 956
MVC_SUBPATHVW	MVC_SUBPATHVW on page 958
MVC_SWITCHCONFIGVW	MVC_SWITCHCONFIGVW on page 959
MVC_SWITCHSUMMARYVW	MVC_SWITCHSUMMARYVW on page 959
MVC_TAPELIBRARYSUMMARYVW	MVC_TAPELIBRARYSUMMARYVW on page 961
MVC_USERROLEMAPVW	MVC_USERROLEMAPVWV on page 961
MVC_USERVW	MVC_USERVW on page 961
MVC_VOLUMEDISKDRIVEVW	MVC_VOLUMEDISKDRIVEVW on page 962
MVC_ZONEPORTVW	MVC_ZONEPORTVW on page 962
MVC_ZONESUMMARYVW	on page 954
MVC_ZONEVW	MVC_ZONEVW on page 963
MVCA_BU_CLIENTDETAILVW	MVCA_BU_CLIENTDETAILVW on page 964
MVCA_BU_DRIVESTATISTICSVW	MVCA_BU_DRIVESTATISTICSVW on page 964
MVCA_BU_DRIVESTATVW	MVCA_BU_DRIVESTATVW on page 965
MVCA_BU_IMAGEDETAILVW	MVCA_BU_IMAGEDETAILVW on page 965
MVCA_BU_JOBDETAILVW	MVCA_BU_JOBDETAILVW on page 966
MVCA_BU_JOB_SUMMARYVW	MVCA_BU_JOB_SUMMARYVW on page 967

Materialized View Name	See
MVCA_BU_LIBRARYDETAILVW	MVCA_BU_LIBRARYDETAILVW on page 968
MVCA_BU_ MASTERSERVERDETAILVW	MVCA_BU_MASTERSERVERDETAILVW on page 969
MVCA_BU_MEDIADETAILVW	on page 967
MVCA_BU_MEDIASERVERDETAILVW	MVCA_BU_MEDIASERVERDETAILVW on page 970
MVCA_BU_OPTIONALTABLEVW	MVCA_BU_OPTIONALTABLEVW on page 971
MVCA_BU_RESOURCE_ SUMMARYVW	MVCA_BU_RESOURCE_SUMMARYVW on page 971
MVCA_BU_TEMPLATECLIENTVW	MVCA_BU_TEMPLATECLIENTVW on page 972
MVCA_BU_TEMPLATEDETAILVW	MVCA_BU_TEMPLATEDETAILVW on page 972
MVCA_DBAPPINSTCAPACITYVW	MVCA_DBAPPINSTCAPACITYVW on page 973
MVCA_DBAPPPHYCAPACITYVW	MVCA_DBAPPPHYCAPACITYVW on page 973
MVCA_ EXCHANGESTORESUMMARYVW	MVCA_EXCHANGESTORESUMMARYVW on page 974
MVCA_EXCHANGEAPPCAPACITYVW	MVCA_EXCHAPPCAPACITYVW on page 974
MVCA_EXCHMAILBOXDETAILVW	MVCA_EXCHMAILBOXDETAILVW on page 975
MVCA_ EXCHPUBLICFOLDERDETAILVW	MVCA_EXCHPUBLICFOLDERDETAILVW on page 975
MVCA_ EXCHSTORGROUPSUMMARYVW	MVCA_EXCHSTORGROUPSUMMARYVW on page 976
MVCA_FSRM_ACL_DETAILSVW	MVCA_FSRM_ACL_DETAILSVW on page 976
MVCA_FSRM_AGEDFILEDETAILS	MVCA_FSRM_AGEDFILEDETAILS on page 976

Materialized View Name	See
MVCA_FSRM_AGEID_DETAILVW	MVCA_FSRM_AGEID_DETAILVW on page 977
MVCA_FSRM_AGESUMMARYVW	MVCA_FSRM_AGESUMMARYVW on page 977
MVCA_FSRM_ DIRDETAILSUMMARYVW	MVCA_FSRM_DIRDETAILSUMMARYVW on page 978
MVCA_FSRM_DIRREPORTDATAVW	MVCA_FSRM_DIRREPORTDATAVW on page 978
MVCA_FSRM_ EXTDETAILSUMMARYVW	MVCA_FSRM_EXTDETAILSUMMARYVW on page 979
MVCA_FSRM_FILEREPORTDATAVW	MVCA_FSRM_FILEREPORTDATAVW on page 979
MVCA_FSRM_LARGEDIRINFO	MVCA_FSRM_LARGEDIRINFO on page 980
MVCA_FSRM_REPORTRULEVW	MVCA_FSRM_REPORTRULEVW on page 980
MVCA_FSRM_RULE_VOLUME_ MAPVW	MVCA_FSRM_RULE_VOLUME_MAPVW on page 980
MVCA_FSRM_SIDDETAILSVW	MVCA_FSRM_SIDDETAILSVW on page 980
MVCA_FSRM_SIZEID_DETAILVW	MVCA_FSRM_SIZEID_DETAILVW on page 981
MVCA_FSRM_SIZE_DETAILSVW	MVCA_FSRM_SIZE_DETAILSVW on page 981
MVCA_FSRM_TEMPFILEDETAILS	MVCA_FSRM_TEMPFILEDETAILS on page 981
MVCA_FSRM_TOPNFILES	MVCA_FSRM_TOPNFILES on page 982
MVCA_FSRM_USER_DIR_MAPVW	MVCA_FSRM_USER_DIR_MAPVW on page 982
MVCA_FSRM_USERDETAILVW	MVCA_FSRM_USERDETAILVW on page 983
MVCA_FSRM_USERSUMMARYVW	MVCA_FSRM_USERSUMMARYVW on page 983
MVCA_FSRM_VOLUMESUMMARYVW	MVCA_FSRM_VOLUMESUMMARYVW on page 983

Materialized View Name	See
MVCA_NAMESPACESUMMARYVW	MVCA_NAMESPACESUMMARYVW on page 984
MVCA_VIRTUALAPPCAPACITYVW	MVCA_VIRTUALAPPCAPACITYVW on page 984
MVCS_EVACTRLSTATSVW	MVCS_EVACTRLSTATSVW on page 985
MVCS_EVADISKSTATSVW	MVCS_EVADISKSTATSVW on page 986
MVCS_EVAHOSTFCPORTSTATSVW	MVCS_EVAHOSTFCPORTSTATSVW on page 987
MVCS_EVASPAGVOLUMESTATSVW	MVCS_EVASPAGVOLUMESTATSVW on page 988
MVCS_ EVASTORAGESYSTEMSTATSVW	MVCS_EVASTORAGESYSTEMSTATSVW on page 989
MVCS_EVAVOLUMESTATSVW	MVCS_EVAVOLUMESTATSVW on page 989
MVCS_HOSTCPUSTATSVW	MVCS_HOSTCPUSTATSVW on page 991
MVCS_HOSTMEMORYSTATSVW	MVCS_HOSTMEMORYSTATSVW on page 991
MVCS_LSICTRLAGDISKSTATSVW	MVCS_LSICTRLAGDISKSTATSVW on page 991
MVCS_LSICTRLAGVOLUMESTATSVW	MVCS_LSICTRLAGVOLUMESTATSVW on page 992
MVCS_LSIDISKIOSTATSVW	MVCS_LSIDISKIOSTATSVW on page 993
MVCS_ LSINOPERFSTORVOLSTATSVW	MVCS_LSINOPERFSTORVOLSTATSVW on page 994
MVCS_LSISPAGDISKSTATSVW	MVCS_LSISPAGDISKSTATSVW on page 995
MVCS_LSISPAGVOLUMESTATSVW	MVCS_LSISPAGVOLUMESTATSVW on page 995
MVCS_LSISSAGDISKSTATSVW	MVCS_LSISSAGDISKSTATSVW on page 997
MVCS_LSISSAGVOLUMESTATSVW	MVCS_LSISSAGVOLUMESTATSVW on page 997

Materialized View Name	See
MVCS_ LSISTORAGEVOLUMESTATSVW	MVCS_LSISTORAGEVOLUMESTATSVW on page 999
MVCS_PORTSTATSVW	MVCS_PORTSTATSVW on page 1000
MVCS_XPCONTROLLERSTATSVW	MVCS_XPCONTROLLERSTATSVW on page 1000
MVCS_XPHOSTFCPORTSTATSVW	MVCS_XPHOSTFCPORTSTATSVW on page 1001
MVCS_XPSTORAGEPOOLSTATSVW	MVCS_XPSTORAGEPOOLSTATSVW on page 1001
MVCS_XPVOLUMESTATSVW	MVCS_XPVOLUMESTATSVW on page 1002
MVIEW_STATUS	MVIEW_STATUS on page 1003
MVIEWCORE_STATUS	MVIEWCORE_STATUS on page 1004
MVCS_CENTERADEVICEUSAGEVW	MVCS_CENTERADEVICEUSAGEVW on page 1004
MVCS_CENTERANODEUSAGEVW	MVCS_CENTERANODEUSAGEVW on page 1005
MVCS_CENTERAPOOLUSAGEVW	MVCS_CENTERAPOOLUSAGEVW on page 1006
MVCS_CENTERADISKUSAGEVW	MVCS_CENTERADISKUSAGEVW on page 1007
MVCS_CENTERAPOOLPRIVILEGE	MVCS_CENTERAPOOLPRIVILEGE on page 1008
MVC_CELERRAVOLUMECPOINTMAP	MVC_CELERRAVOLUMECPOINTMAP on page 1008

MVA

Column Name	Type	Description
HOSTID	NUMBER(38)	HostID
HOSTNAME	VARCHAR2(256)	Host Name
DOMAINID	NUMBER(38)	DomainID
VENDOR	VARCHAR2(256)	Host Vendor

Column Name	Type	Description
DESCRIPTION	VARCHAR2(1024)	Host Description
STATUS	NUMBER(38)	Operation status (provide map here)

MVC_APPLICATIONRELATIONVW

Column Name	Description
ApplicationClusterID	ID of the cluster application
ApplicationClusterName	Name of the cluster application
AppClusterMemberID	ID of the application cluster member
AppClusterMemberName	Name of the application cluster member

MISAPPLICATIONS

Column Name	Type	Description
APPLICATIONID	NUMBER(38)	Application Identifier
APPLICATIONNAME	VARCHAR2(256)	Application Name
DOMAINID	NUMBER(38)	Domain Identifier
APPLICATIONTYPE	NUMBER(38)	Application type
VENDOR	VARCHAR2(254)	Vendor
DESCRIPTION	VARCHAR2(1024)	Description of the application
STATUS	NUMBER(38)	Operational status
BUILDNUMBER	VARCHAR2(254)	Software build number
SERIALNUMBER	VARCHAR2(254)	Software serial number
TARGETOS	VARCHAR2(254)	Target operating system
PRODUCTNAME	VARCHAR2(254)	Product name
VERSION	VARCHAR2(254)	Software version
IDENTIFICATIONCODE	VARCHAR2(254)	Software Identifierentification code
SOFTWAREELEMENTID	VARCHAR2(254)	Software element Identifier
HOSTID	NUMBER(38)	ID of host where the application is running from

Column Name	Type	Description
HOSTNAME	VARCHAR2(256 CHAR)	Name of the Host
ISCLUSTER	NUMBER(1)	Flag to know whether it a cluster or not

MVC_ASSETSUMMARYVW

Column Name	Type	Description
DOMAINID	NUMBER(38)	Domain identifier
ASSETID	NUMBER(38)	ID of the asset depending on Assetclass; for example, when assetclass is Host, this is host.id.
ASSETCLASS	VARCHAR2(13)	Asset class: HOST, APPLICATION, STORAGESYSTEM, or SWITCH
NAME	VARCHAR2(256)	Name of the host
DATECREATED	DATE	Creation date
DATELASTMODIFIED	DATE	Date last modified
DESCRIPTION	VARCHAR2(255)	Asset description
STATUS	VARCHAR2(8)	Asset status: NEW, MISSING, REPAIRED, or IN USE
VENDOR	VARCHAR2(254)	Asset vendor
MODEL	VARCHAR2(254)	Asset model
SERIALNUMBER	VARCHAR2(254)	Asset serial number
BARCODE	VARCHAR2(255)	Asset bar code
ASSETCODE	VARCHAR2(255)	Asset code
ASSETTYPE	VARCHAR2(255)	Asset type
ASSETTAG	VARCHAR2(255)	Asset tag
ASSETCATEGORY	VARCHAR2(255)	Asset category

Column Name	Type	Description
GEOGRAPHICLOCATION	VARCHAR2(255)	Asset location
STORAGETIERCLASSIFICATION	VARCHAR2(255)	Asset storage tier name; for example, Ultra High Availability
STORAGETIERCOSTPERGB	NUMBER(36,2)	Asset storage tier cost
DEPARTMENTNO	VARCHAR2(255)	Asset department number
DEPARTMENTNAME	VARCHAR2(255)	Asset department name
PERCENTAGEOWNED	NUMBER(5, 2)	Percentage owned by department
ADMINISTRATOR	VARCHAR2(255)	Asset maintained by
STAFFNAME	VARCHAR2(255)	Staff Name
STAFFPHONENUMBER	VARCHAR2(255)	Staff PH#
STAFFDEPARTMENT	VARCHAR2(255)	Staff Department
STAFFEMAIL	VARCHAR2(255)	Staff e-mail
RACKNUMBER	VARCHAR2(100)	
FLOOR	VARCHAR2(100)	
DATACENTER	VARCHAR2(100)	
CITY	VARCHAR2(100)	
REGION	VARCHAR2(100)	
COUNTRY	VARCHAR2(50)	
CONTINENT	VARCHAR2(20)	
ADDRESS	VARCHAR2(1024)	
ZIPCODE	VARCHAR2(16)	
LICENSE	VARCHAR2(4000)	
PURCHASEORDERNUMBER	VARCHAR2(255)	Asset Purchase Order number
DATEPURCHASED	DATE	Asset date purchased
COST	NUMBER(36,2)	Asset cost

Column Name	Type	Description
VALUE	NUMBER(36,2)	
SALVAGECOST	NUMBER(36,2)	Asset Deprecated Salvage Cost
DEPRECITIONPERIOD	NUMBER(18)	Asset Deprecation Period
DEPRECITIONMETHOD	NUMBER(16)	Asset Deprecation Method (Straight line, Fixed declining, Double declining)
DEPRECITEDVALUE	NUMBER(36,2)	Asset Depreciated Value
RESELLER	VARCHAR2(255)	Asset Reseller
COMMENTS	VARCHAR2(4000)	Comments
ASSETFIXCOSTTAXPERDEPTPERYEAR	NUMBER	Asset Fixed cost tax per department per Year
STORAGFIXCOSTTAXPERDEPTPERYEAR	NUMBER	Storage Fixed cost tax per department per Year
SECONDARYSTAFFNAME	VARCHAR2(255 CHAR)	Secondary Staff Name
SECONDARYSTAFFPHONENUMBER	VARCHAR2(255 CHAR)	Secondary Staff Phone Number
SECONDARYSTAFFDEPARTMENT	VARCHAR2(255 CHAR)	Secondary Staff Department
SECONDARYSTAFFEMAIL	VARCHAR2(255 CHAR)	Secondary Staff Mail
BASE_TABLE_NAME	CHAR(10 CHAR)	Base Table Name

MVC_CARDSUMMARYVW

Column Name	Type	Description
CardID	NUMBER(38)	Card Identifier
CardName	VARCHAR2(256)	Card Name
ContainerID	NUMBER(38)	Container Identifier

Column Name	Type	Description
CardType	VARCHAR2(7)	Card Type (HBA, SCSI)
DomainID	NUMBER(38)	Domain Identifier
Vendor	VARCHAR2(256)	Card Vendor
Description	VARCHAR2(1024)	Card Description
Status	NUMBER(38)	Operational status
WWN	VARCHAR2(256)	Node WWN
Model	VARCHAR2(256)	Card model
SerialNumber	VARCHAR2(256)	Card Serial Number
Version	VARCHAR2(256)	Card Version
Firmware	VARCHAR2(256)	Firmware version
DriverVersion		Driver version
BASETABLENAME	CHAR(4)	Name of the base table for optional values
ELEMENT_TYPE	VARCHAR2(254 CHAR)	Element Type
CARDIP	VARCHAR2(1024 CHAR)	Card IP Address

MVC_CHANGERDEVICESUMMARYVW

Column Name	Type	Description
CHANGERDEVICEID	NUMBER(38)	Changer Device Identifier
BUSINESSCOST	NUMBER(36,2)	Business Cost for Changer Device
INSTALLDATE	TIMESTAMP(6)	Changer Device Installation Date
VENDOR	VARCHAR2(255 CHAR)	Name of Vendor
SERIALNUMBER	VARCHAR2(255 CHAR)	Serial Number
MODEL	VARCHAR2(255 CHAR)	Model Name

Column Name	Type	Description
VERSION	VARCHAR2(255 CHAR)	Version Number
APPIQNAME	VARCHAR2(256 CHAR)	Customized Name for Changer Device
APPIQCONTACTEDTIME	TIMESTAMP(6)	Contacted Time
PROVIDERNAME	VARCHAR2(256 CHAR)	Changer Device Name
APPIQSTATUS	NUMBER(38)	Changer Device Status
APPIQCREATIONTIME	TIMESTAMP(6)	Creation Time
DESCRIPTION	VARCHAR2(1024 CHAR)	Description
SYSTEMID	NUMBER(38)	Storage System Identifier
CARDID	NUMBER(38)	Card Identifier
MODELTYPE	NUMBER(38)	Model Type
OID	VARCHAR2(254 CHAR)	O Identifier
DOMAINID	NUMBER(38)	Domain Identifier

MVC_DISCOVERYDETAILSVW

Name	Description
ElementID	Quarantined element identifier
ElementName	Name of the quarantined element
Address	IP address of the element
ElementType	Type of element
DiscoveryGroup	Name of the discovery group
Enabled	Enabled quarantined or not
Status	Status of the element
ParentAddressID	Parent address
Mapping_version	

MVC_DISKDRIVESUMMARYVW

Column Name	Type	Description
DiskDriveID	NUMBER(38)	Disk Drive Identifier
DiskDriveName	VARCHAR2(256)	Name of the Disk Drive
DomainID	NUMBER(38)	Domain Identifier
OID	VARCHAR2(254)	Object Identifier of the Disk Drive
Vendor	VARCHAR2(254)	Vendor
Description	VARCHAR2(1024)	Description
Status	NUMBER(38)	Operational status
Model	VARCHAR2(254)	Disk Drive Model
Name	VARCHAR2(254)	Name coming from disk drive
CardID	NUMBER(38)	Card Identifier
DiskDriveStatus	VARCHAR2(254)	Disk Drive Status
SCSIBus	NUMBER(38)	SCSI Bus
SCSILUN	NUMBER(38)	SCSI LUN
SCSITargetID	NUMBER(38)	SCSI target Identifier
SCSIPort	NUMBER(38)	SCSI PORT Identifier
SystemID	NUMBER(38)	Container System Identifier
MaxMediaSize	NUMBER(38)	Maximum media size
MaxBlockSize	NUMBER(38)	Maximum block size
MinBlockSize	NUMBER(38)	Minimum block size
EnableStatus	VARCHAR2(254)	
Availability	VARCHAR2(254)	
BASETABLENAME	CHAR(11)	Name of the base table
SHAREID	VARCHAR2(254 CHAR)	Disk Drive Share Identifier
PATHSTATUS	VARCHAR2(128 CHAR)	Status of the Path

MVC_DISK_EXTENTVW

Column Name	Type	Description
ExtentID	NUMBER	Disk Extent Identifier
ContainerExtentID	NUMBER	Container Extent Identifier
DiskID	NUMBER	Disk Drive Identifier

MVC_DISKEXTENTSUMMARYVW

Column Name	Type	Description
DiskExtentID	NUMBER(38)	Disk Extent Identifier
DiskEnxtentName	VARCHAR2(256)	Name of the extent
DomainID	NUMBER(38)	Domain Identifier
DiskExtentDescription	VARCHAR2(1024)	Description of the extent
Status	NUMBER(38)	Operational status
Access_Type	VARCHAR2(254)	Access type
BlockSize	NUMBER(38)	Block Size
Number_Of_Blocks	NUMBER(38)	Number of total blocks
Consumable_Blocks	NUMBER(38)	Number of usable blocks
StoragePoolID	NUMBER(38)	Storage Pool Identifier
SystemID	NUMBER(38)	Container System Identifier
BASETABLENAME	CHAR(11)	Name of the base table
SHAREDID	VARCHAR2(254 CHAR)	Share Disk Identifier
CLUSTERSHARED	NUMBER(1)	Flag to know whether cluster shared or not

MVC_EVENTVW

Column Name	Type	Description
DOMAINID	NUMBER(38)	Domain Identifier
ID	NUMBER(38)	Event Identifier

Column Name	Type	Description
ELEMENT_ID	NUMBER(38)	Element Identifier on which event occurred
DEVICE_ID	NUMBER(38)	Device Identifer
SUMMARY_TEXT	VARCHAR2(4000 CHAR)	Summary of Event
TIME_REPORTED	DATE	Event Reported Time
SEVERITY	NUMBER(38)	Event Severity
CLEARED	NUMBER(1)	Flag to find whether or not event cleared
SOURCE	VARCHAR2(254 CHAR)	Event Source
TYPE	NUMBER(38)	Type
SUBTYPE	NUMBER(38)	Subtype
PROBABLE_CAUSE_ DESCRIPTION	VARCHAR2(4000 CHAR)	Probable Cause Description
RECOMMENDED_ACTIONS	VARCHAR2(4000 CHAR)	Recommended Actions
CORRELATED_ INDICATIONS	VARCHAR2(254 CHAR)	Correlated Indications
TYPE_DESCRIPTION	VARCHAR2(254 CHAR)	Event Type Description
OBJECT_TYPE	VARCHAR2(254 CHAR)	Object Type
NAME	VARCHAR2(256 CHAR)	Name of the Event
VENDOR	VARCHAR2(254 CHAR)	Name of the Vendor
MODEL	VARCHAR2(254 CHAR)	Model Name
EVENTTYPE	VARCHAR2(4000 CHAR)	Event Type
COMPONENT	VARCHAR2(4000 CHAR)	Component

MVC_EXTENTEXTENTMAPVW

Column Name	Type	Description
ID	NUMBER(38)	Extent Map Identifier
TARGETEXTENTID	NUMBER(38)	Target Extent Identifier
SOURCEEXTENTID	NUMBER(38)	Source Extent Identifier

MVC_FABRICSUMMARYVW

Column Name	Type	Description
FABRICID	NUMBER(38)	Fabric Identifier
FABRICNAME	VARCHAR2(256 CHAR)	Name of Fabric
FABRICWWN	VARCHAR2(254 CHAR)	Fabric WWN
SANID	NUMBER(38)	SAN Identifier
REPLICATIONSANID	NUMBER(38)	Replication SAN Identifier
VIRTUALFABRICID	NUMBER(38)	Virtual Fabric Identifier

MVC_HBAPORTTARGETS

Column Name	Type	Description
ID	NUMBER(38)	
PORT_ID	NUMBER(38)	Port Identifier
SCSI_BUS	NUMBER(38)	SCSI Bus Identifier
SCSI_TARGET_ID	NUMBER(38)	SCSI Target Identifier
LUN_ID	NUMBER(38)	LUN Identifier
PERSISTENT	NUMBER(1)	Persistent
OS_DEVICE_NAME	VARCHAR2(254 CHAR)	OS Device Name
TARGET_LUN_ID	NUMBER(38)	Target LUN Identifier
TARGET_PORT_WWN	VARCHAR2(32 CHAR)	Target Port of WWN
APPIQ_STATUS	NUMBER(38)	Status
CREATION_TIME	TIMESTAMP(6)	Creation Time
CONTACTED_TIME	TIMESTAMP(6)	Contacted Time

MVC_HOSTCAPACITYVW

Column Name	Type	Description
HostID	NUMBER(38)	Host Identifier
VolumeID	NUMBER(38)	Logical Volume Identifier
TimeStamp	TIMESTAMP	Time of data collection
Total	NUMBER(38)	Total capacity in megabytes
Used	NUMBER(38)	Used capacity in megabytes
Free	NUMBER(38)	Free capacity in megabytes
CAPACITYTYPE	VARCHAR2(256 CHAR)	Type of the Capacity
CAPACITYNUM	NUMBER	Capacity number
FileSystemType	VARCHAR2(254 CHAR)	File System Type
HostVolume	VARCHAR2(254 CHAR)	Volume of the Host
shareid		Share Identifier

MVC_HOSTDISKDRIVEVW

Column Name	Type	Description
HostID	NUMBER(38)	Host Identifier
DiskDriveID	NUMBER(38)	Disk Drive Identifier
ExtentID	NUMBER(38)	Disk Partition Identifier
DiskDrive	VARCHAR2(256)	Disk Drive Name
DiskPartition	VARCHAR2(256)	Disk Partition Name
DiskPartitionDescription	VARCHAR2(1024)	Description of the partition
DiskPartitionSPace	Number	Capacity of the disk partition, in megabytes
DPVENDOR	VARCHAR2(254 CHAR)	Vendor of the DP
DISKMODEL	NUMBER(38)	Model of the Disk

MVC_HOSTRELATIONVW

Name	Description
DomainID	Domain Identifier
ID	Cluster Identifier
ClusterHostID	Cluster Host Identifier
ClusterHostname	Cluster Host name
MemberHostID	Member host Identifier
MemberRelation	Member relation
State	State of the cluster
ClusterHostModelType	Model type
ClusterHostVendor	Vendor of the cluster
Description	Description
Status	Status of the cluster
ObjectType	Object type
MEMBERHOSTNAME	Name of the Member Host

MVC_HOSTSUMMARYVW

Column Name	Type	Description
HOSTID	NUMBER(38)	Host Identifier
HOSTNAME	VARCHAR2(256)	Host Name
DOMAINID	NUMBER(38)	Domain Identifier
VENDOR	VARCHAR2(256)	Host Vendor
DESCRIPTION	VARCHAR2(1024)	Host Description
STATUS	NUMBER(38)	Operation status (provide map here)
IP	VARCHAR2(16)	Host IP
DNS	VARCHAR2(50)	Host DNS Name
Model	VARCHAR2(256)	Host Model
Version	VARCHAR2(256)	Host Version number
OS	VARCHAR2(24)	Host Operating System

Column Name	Type	Description
TOTALPHYSICALMEM	NUMBER(38)	Total physical memory
NUMBERPROCESSOR	Number	Number of processors
SUPPORTFLAG	NUMBER(38)	Support flag (unused now)
BASETABLENAME	CHAR(4)	Name of the base table for optional values
model_type	NUMBER(38)	Host Model Type
iscluster	NUMBER(1)	Flag to know whether the host is a cluster or not
isvirtualServer	NUMBER(1)	Flag to know whether the host is a virtual server or not
isvirtualMachine	NUMBER(1)	Flag to know whether the host is a virtual machine or not
ISBACKUPCLIENT	NUMBER(1)	Boolean value to know whether or not host is backup client

MVC_HOSTVOLUMESUMMARYVW (Logical Volumes)

Column Name	Type	Description
LogicalVolumeID	NUMBER(38)	Storage Volume Identifier
LogicalVolumeName	VARCHAR2(256)	Name of the logical volume
DomainID	NUMBER(38)	Domain Identifier
Description	VARCHAR2(1204)	Description
HostID	NUMBER(38)	Container Host Identifier
DeviceID	VARCHAR2(254)	Logical Device Identifier
FileSystemType	VARCHAR2(254)	File System Type
Blocksize	NUMBER(38)	These 3 fields might not be needed for Host
NumberOfBlocks	NUMBER(38)	Logical Volumes
ConsumableBlocks	NUMBER(38)	
BASETABLENAME	CHAR(14)	Name of the base table for optional values

Column Name	Type	Description
VOLUMETYPE	VARCHAR2(13 CHAR)	Type of Volume
STATUS	NUMBER(38)	Status
HOSTNAME	VARCHAR2(256 CHAR)	Name of Host
SHAREID	VARCHAR2(254 CHAR)	Share Disk Identifier
MODELTYPE	NUMBER(38)	Model Type
SHARE_NAME	VARCHAR2(254 CHAR)	Name of the Share Disk
SHARE_SYSTEM_ NAME	VARCHAR2(254 CHAR)	Name of the Share System

MVC_LOCALIZEDSTRINGVW

Column Name	Type	Description
ID	NUMBER(38)	Localized String Identifier
VALUE	VARCHAR2(256 CHAR)	Value of Localized String

MVC_LUNMAPPINGVW

Column Name	Type	Description
ID	NUMBER(38)	LUN Mapping Identifier
INITIATOR	VARCHAR2(255 CHAR)	Initiator WWN
INITIATOR_FORMAT	VARCHAR2(255 CHAR)	Initiator Format
PROVIDERNAME	VARCHAR2(256 CHAR)	Provider Name
APPIQNAME	VARCHAR2(256 CHAR)	Customized Name
DESCRIPTION	VARCHAR2(1024 CHAR)	Description

Column Name	Type	Description
MODELTYPE	NUMBER(38)	Model Type
PRIVILEGEID	NUMBER(38)	Privilege Identifier
OID	VARCHAR2(254 CHAR)	O Identifier
CONTROLLERID	NUMBER(38)	Controller Identifier
STORAGE_VOLUME_ID	NUMBER(38)	Storage Volume Identifier
STORAGE_SYSTEM_PORT_ID	NUMBER(38)	Storage System Port Identifier
STORAGE_VOLUME_PORT_LUNID	NUMBER(38)	Storage Volume Port LUN Identifier
ACCESS_STATE	VARCHAR2(254 CHAR)	Access State
HOST_GROUP	VARCHAR2(254 CHAR)	Host Group
HOST_GROUP_NAME	VARCHAR2(254 CHAR)	Host Group Name
HOST_GROUP_MODES	VARCHAR2(1024 CHAR)	Host Group Modes
HBA_PORT_ID	NUMBER(38)	HBA Port Identifier
LUN_ID	NUMBER(38)	LUN Identifier
PERSISTENT	NUMBER(1)	Persistent
TARGET_LUN_ID	NUMBER	Target LUN Identifier

MVC_MULTIPATHVW

Column Name	Type	Description
PathID	NUMBER(38)	Path Identifier
MultiPathDiskExtentID	NUMBER(38)	Multipath disk extent identifier
VolumeManagerVolumeID	NUMBER(38)	Volume Manager Volume Identifier
VxvmDiskPartitionID	NUMBER(38)	Volume manager volume disk partition identifier
DISKDRIVEID	NUMBER(38)	Disk Drive Identifier

MVC_NASEXTENTSUMMARYVW

Column Name	Type	Description
DOMAINID	NUMBER(38)	Domain Identifier
HOSTID	NUMBER(38)	Host Identifier
HOSTNAME	VARCHAR2(256 CHAR)	Host Name
VOLUMENAME	VARCHAR2(256 CHAR)	Volume Name
FILESYSTEMTYPE	VARCHAR2(254 CHAR)	Type of File System
LOGICALVOLUMEID	NUMBER(38)	Logical Volume Identifier
TOTAL	NUMBER(38)	Total
USED	NUMBER(38)	Used
FREE	NUMBER(38)	Free
DISKEXTENTID	NUMBER(38)	Disk Extent Identifier
AGGREGATE	VARCHAR2(256 CHAR)	Aggregate
AGGREGATEID	NUMBER(38)	Aggregate Identifier
PLEX	VARCHAR2(256 CHAR)	Plex
PLEXID	NUMBER(38)	Plex Identifier
RAIDGROUPNAME	VARCHAR2(256 CHAR)	RAID Group Name
RAIDGROUPID	NUMBER(38)	RAID Group Identifier

MVC_NASQUOTASUMMARYVW

Column Name	Type	Description
QUOTAID	NUMBER(38)	Quota Identifier
QUOTANAME	VARCHAR2(256 CHAR)	Name of Quota
STATUS	NUMBER(38)	Status
RESOURCETYPE	VARCHAR2(32 CHAR)	Type of Resource
QUOTATYPE	VARCHAR2(32 CHAR)	Type of Quota
QUOTATARGET	VARCHAR2(128 CHAR)	Quota Target
SOFTLIMIT	NUMBER(38)	Soft Limit

Column Name	Type	Description
HARDLIMIT	NUMBER(38)	Hard Limit
THRESHOLD	NUMBER(38)	Threshold
AMOUNTUSED	NUMBER(38)	Amount Used
FILESYSTEMID	NUMBER(38)	File System Identifier
FILESYSTEMNAME	VARCHAR2(256 CHAR)	Name of File System
DIRECTORYID	NUMBER(38)	Directory Identifier
DIRECTORYNAME	VARCHAR2(256 CHAR)	Name of Directory
MODELTYPE	NUMBER(38)	Model Type
DOMAINID	NUMBER(38)	Domain Identifier

MVC_NASSNAPSHOTSUMMARYVW

Column Name	Type	Description
SNAPSHOTID	NUMBER(38)	Snapshot Identifier
SNAPSHOTNAMENAME	VARCHAR2(256 CHAR)	Name of Snapshot
STATUS	NUMBER(38)	Status
ACCESSTIME	NUMBER(38)	Access Time
BLOCKSIZE	NUMBER(38)	Block Size
TOTALBLOCKSKB	NUMBER(38)	Total Number of Blocks in KBs
CONSUMABLEBLOCKS	NUMBER(38)	Consumable Blocks
PCTTOTALBLOCKS	NUMBER(38)	Total Blocks Percentage
PCTUSEDBLOCKS	NUMBER(38)	Used Blocks Percentage
CUMPERCENTOFTOTALBLOCKS	NUMBER(38)	Total Blocks Cumulative Percentage
CUMPERCENTOFUSEDBLOCKS	NUMBER(38)	Used Blocks Cumulative Percentage
CUMULATIVETOTALKB	NUMBER(38)	Total Cumulative Blocks in KB

Column Name	Type	Description
DEPENDENCY	VARCHAR2(32 CHAR)	Dependency
HOSTID	NUMBER(38)	Host Identifier
VOLUMEID	NUMBER(38)	Volume Identifier
DOMAINID	NUMBER(38)	Domain Identifier

MVC_NAS_FSDIRECTORYSUMMARYVW

Column Name	Type	Description
DIRECTORYID	NUMBER(38)	Directory Identifier
BUSINESSCOST	NUMBER(36,2)	Business Cost
INSTALLDATE	TIMESTAMP(6)	Date of Installation
VENDOR	VARCHAR2(254 CHAR)	Name of Vendor
DIRECTORYNAME	VARCHAR2(256 CHAR)	Name of Directory
APPIQCONTACTEDTIME	TIMESTAMP(6)	Contacted Time
FILESYSTEMNAME	VARCHAR2(255 CHAR)	Name of File System
STATUS	NUMBER(38)	Status
APPIQCREATIONTIME	TIMESTAMP(6)	Creation Time
DESCRIPTION	VARCHAR2(1024 CHAR)	Description
FILESYSTEMID	NUMBER(38)	ID of File System
MODELTYPE	NUMBER(38)	Model Type
OID	VARCHAR2(254 CHAR)	OID
DOMAINID	NUMBER(38)	Domain Identifier

MVC_NAS_FSSHARESUMMARYVW

Column Name	Type	Description
SHAREID	NUMBER(38)	ID of the File System Share
BUSINESSCOST	NUMBER(36,2)	Business Cost

Column Name	Type	Description
INSTALLDATE	TIMESTAMP(6)	Date of Installation
VENDOR	VARCHAR2(254 CHAR)	Name of Vendor
SHARENAME	VARCHAR2(256 CHAR)	Name of File System Share
DIRECTORYID	NUMBER(38)	ID of the Directory
APPIQCONTACTEDTIME	TIMESTAMP(6)	Contacted Time
STATUS	NUMBER(38)	Status
MOUNTPOINT	VARCHAR2(255 CHAR)	Mount Point of Share
APPIQCREATIONTIME	TIMESTAMP(6)	Creation Time
DESCRIPTION	VARCHAR2(1024 CHAR)	Description
MODELTYPE	NUMBER(38)	Model Type
OID	VARCHAR2(254 CHAR)	OID
DOMAINID	NUMBER(38)	Domain Identifier

MVC_OPTIONALTABLEVW

Column Name	Type	Description
BaseTableName	VARCHAR2(30)	Name of the base table
BaseTableID	NUMBER(38)	ID of the base table record
OptionalName	VARCHAR2(50)	Name of the optional value
OptionalValue	VARCHAR2(4000)	Value of the optional value

MVC_ORGANIZATIONVW

Column Name	Type	Description
ORGID	NUMBER(38)	ID from Organization table
ORGNAME	VARCHAR2(256)	Organization name
ORGDESCRIPTION	VARCHAR2(4000)	Description of this Organization
ELEMENTID	NUMBER(38)	Element Identifier mapped to this Organization
DOMAINID	NUMBER(38)	Domain Identifier

Column Name	Type	Description
EVERYTHING	NUMBER(1)	Flag to know whether it belongs to default organization

MVC_ORGFILTERVW

Column Name	Type	Description
ORGID	NUMBER(38)	HostID
ELEMENTID	NUMBER(38)	Host Name
INCLUDEFLAG	NUMBER(1)	DomainID
EXCLUDETABLE	VARCHAR2(128 CHAR)	Host Vendor
SUBELEMENTID	NUMBER(38))	Host Description

MVC_ORGRELATIONVW

Column Name	Type	Description
ParentOrgID	NUMBER(38)	Parent Organization Identifier
ChildOrgID	NUMBER(38)	Child Organization Identifier
DOMAINID	NUMBER(38)	Domain Identifier

MVC_PATHVW

Column Name	Type	Description
PathID	NUMBER(38)	Path Identifier
HostID	NUMBER(38)	Host Identifier on this path
LogicalVolumeID	NUMBER(38)	Logical Volume Identifier
DiskPartitionID	NUMBER(38)	Disk Partition ID if raw partition is present
IsMounted	VARCHAR2(5)	(TRUE, FALSE)
ContainerPathID	NUMBER(38)	Currently not in use
ApplicationID	NUMBER(38)	Application ID for the application file
ApplicationFileID	NUMBER(38)	Application File Identifier
FilePath	VARCHAR2(256)	Application file path

Column Name	Type	Description
FileName	VARCHAR2(128)	Application file name
ISMOUNTEDNUM	NUMBER	Flag to know whether Mounted or Not

MVC_PORTCONTROLLERMAPVW

Column Name	Type	Description
ID	NUMBER(38)	
CONTROLLERID	NUMBER(38)	
APPIQSTATUS	NUMBER(38)	Status
PORTID	NUMBER(38)	ID of Port

MVC_PROTOCOLCONTROLLERVW

Column Name	Type	Description
ID	NUMBER(38)	ID of Protocol Controller
VENDOR	VARCHAR2(254 CHAR)	Name of Vendor
NAME	VARCHAR2(256 CHAR)	Name of Protocol Controller
APPIQCONTACTEDTIME	TIMESTAMP(6)	Contacted Time
APPIQSTATUS	NUMBER(38)	Status
APPIQCREATIONTIME	TIMESTAMP(6)	Creation Time
DESCRIPTION	VARCHAR2(1024 CHAR)	Description
SYSTEMID	NUMBER(38)	ID of Storage System
PARENTID	NUMBER(38)	ID of Parent Controller
MASKINGCAPABILITIESID	NUMBER(38)	ID of Masking Capabilities
MODELTYPE	NUMBER(38)	Type of Model
OID	VARCHAR2(254 CHAR)	OID
DOMAIN	NUMBER(38)	Domain Identifier

MVC_PORTSUMMARYVW

Column Name	Type	Description
PortID	NUMBER(38)	Port Identifier
PortName	VARCHAR2(256)	Port Name
DomainID	NUMBER(38)	Domain Identifier
Description	VARCHAR2(1024)	Description
Status	NUMBER(38)	Operational status (provide map)
WWN	VARCHAR2(32)	Port WWN
ContainerID	NUMBER(38)	Container Identifier of this port
TargetPort	NUMBER(38)	Target Port Identifier (never populated)
Connected_To_WWN	VARCHAR2(32)	WWN of connected port
Device_ID	VARCHAR2(254)	Port Device Identifier
PortState	VARCHAR2(254)	
PortStatus	VARCHAR2(254)	
Physical_State	VARCHAR2(254)	
Port_ID	NUMBER(38)	
Port_Speed	NUMBER(38)	
Max_Speed	NUMBER(38)	Port Max Speed, bit/second
PortNumber	NUMBER(18)	
SCSIPort	NUMBER(18)	
ConnectedToNodeWWN	VARCHAR2(32)	
PortType	VARCHAR2(254)	
BASETABLENAME	CHAR(4)	Name of the base table for optional values
CONTAINERNAME	VARCHAR2(256 CHAR)	Container Name
PORT_SYMBOLIC_ NAME	VARCHAR2(4000 CHAR)	Port Symbolic Name
LINK_TECHNOLOGY	VARCHAR2(254 CHAR)	Link Technology

Column Name	Type	Description
TRUNKEDSTATE	VARCHAR2(22 CHAR)	Port Trunked State
ParentPortID	NUMBER(38)	Parent port identifier
PortFlavor	NUMBER(38)	Port flavor

MVC_ROLEVW

Column Name	Type	Description
ROLEID	NUMBER(38)	
ROLENAME	VARCHAR2(256 CHAR)	Name of Role
DESCRIPTION	VARCHAR2(1024 CHAR)	Role Description
DOMAINADMIN	NUMBER(1)	Flag to Know Domain Admin or Not
APPLICATION	NUMBER(38)	Role Application (Permission Types)
SERVER	NUMBER(38)	Role Server (Permission Types)
SUBSYSTEM	NUMBER(38)	Role Sub System (Permission Types)
SWITCH	NUMBER(38)	Role Switch (Permission Types)
TAPELIBRARY	NUMBER(38)	Role Tape Library (Permission Types)
OTHER	NUMBER(38)	Role Others (Permission Types)
APPLICATIONEXPLORER	NUMBER(1)	Flag to Know Role on Application Viewer
SYSTEMEXPLORER	NUMBER(1)	Flag to Know Role on System Manager
EVENTS	NUMBER(1)	Flag to Know Role on Events
PROVISIONING	NUMBER(1)	Flag to Know Role on Provisioning
ASSETMANAGEMENT	NUMBER(1)	Flag to Know Role on Asset Management

Column Name	Type	Description
POLICYMANAGER	NUMBER(1)	Flag to Know Role on Policy Manager
REPORTING	NUMBER(1)	Flag to Know Role on Reporting
SYSTEMCONFIG	NUMBER(1)	Flag to Know Role on System Configuration
ACCESSCLI	NUMBER(1)	Flag to Know Role on CLI Access
ACCESSFSRM	NUMBER(1)	Flag to Know Role on FSRM Access
BUSINESSTOOLS	NUMBER(1)	Flag to Know Role on Business Tools
ENTERPRISEREPORTING	NUMBER(1)	Flag to Know Role on Enterprise Reporting
CAPACITYEXPLORER	NUMBER(1)	Flag to Know Role on Capacity Manager
PERFORMANCEEXPLORER	NUMBER(1)	Flag to Know Role on Performance Manager
PROTECTIONEXPLORER	NUMBER(1)	Flag to Know Role on Backup Manager
PROVISIONADMIN	NUMBER(1)	Flag to Know Role Provision Admin or Not

MVC_STORAGEPOOLCONFIGVW

Column Name	Type	Description
StoragePoolID	NUMBER(38)	Storage Pool Identifier
CollectionTime	TIMESTAMP (6)	Configuration stats collection time
ExportedMB	NUMBER	Storage exposed in megabytes
UnexportedMB	NUMBER	Volume unexposed in megabytes
AvailableMB	NUMBER	Available storage left on storage pool
ProvisionedMB	NUMBER	Provisioned storage in megabytes
TotalMB	NUMBER	Total storage in megabytes
StorageCapability	VARCHAR2(255)	Storage pool capability

Column Name	Type	Description
MainframePool	VARCHAR2(13)	Indicate whether reserved for mainframe: MAINFRAMEPOOL
CAPACITYTYPE	VARCHAR2(256 CHAR)	Type of the Capacity
CAPACITYNUM	NUMBER	Capacity Number
EXTERNALPOOL	NUMBER	External Pool
MAINFRAMEPOOLNUM	VARCHAR2(256 CHAR)	Mainframe Pool Number
ConsumableMB	NUMBER	Consumable in megabytes
ConsumedMB	NUMBER	Consumed in megabytes
ConsumedExportedMB	NUMBER	Consumed exported in megabytes

MVC_STORAGEPOOLSUMMARYVW

Column Name	Type	Description
STORAGEPOOLID	NUMBER(38)	
DOMAINID	NUMBER(38)	Domain Identifier
STORAGEPOOLNAME	VARCHAR2(256 CHAR)	Name of Storage Pool
STORAGEPOOLDESCRIPTION	VARCHAR2(1024 CHAR)	Storage Pool Description
STATUS	NUMBER(38)	Status
STORAGESYSTEMID	NUMBER(38)	
STORAGESYSTEMNAME	VARCHAR2(256 CHAR)	Name of Storage System
STORAGESYSTEMVENDOR	VARCHAR2(254 CHAR)	Name of Storage System Vendor
STORAGEDESCRIPTION	VARCHAR2(1024 CHAR)	Description of Storage Pool
POOLSETTINGID	NUMBER(38)	
PARENTPOOLID	NUMBER(38)	Identifier of Parent Storage Pool

Column Name	Type	Description
TOTALAVAILABLESPACE	NUMBER(38)	Total Available Space in the Storage Pool
CIMPOOLID	VARCHAR2(254 CHAR)	
POOLTYPE	NUMBER(38)	Type of Storage Pool
STORAGECAPABILITYNAME	VARCHAR2(254 CHAR)	Name of Storage Capabilities
NOSINGLEPTOFFAILURE	NUMBER(1)	Flag to Know No Single Point of Failure
DEFAULTNOSINGLEPTOFFAILURE	NUMBER(1)	Flag to Know Default No Single Point of Failure
MINDATAREDUNDANCY	NUMBER(18)	Minimum Data Redundancy
MAXDATAREDUNDANCY	NUMBER(18)	Maximum Data Redundancy
DEFAULTDATAREDUNDANCY	NUMBER(18)	Default Data Redundancy
MINSPINDLEREDUNDANCY	NUMBER(18)	Minimum Spindle Redundancy
MAXSPINDLEREDUNDANCY	NUMBER(18)	Maximum Spindle Redundancy
DEFAULT_SPINDLE_REDUNDANCY	NUMBER(18)	Default Spindle Redundancy
MINDELTARESERVATION	NUMBER(18)	Minimum Delta Reservation
MAXDELTARESERVATION	NUMBER(18)	Maximum Delta Reservation
DEFAULTDELTARESERVATION	NUMBER(18)	Default Delta Reservation
STORAGECAPABILITYCOMMONNAME	VARCHAR2(256 CHAR)	Common Name for Storage Capability
STORAGECAPABILITYDESCRIPTION	VARCHAR2(1024 CHAR)	Description or Storage Capability

MVC_STORAGEPROCESSORSUMMARYVW

Column Name	Type	Description
SystemProcessorID	NUMBER(38)	Storage System Processor Identifier
SystemProcessorName	VARCHAR2(256)	Name of the system processor
DomainID	NUMBER(38)	Domain Identifier
Vendor	VARCHAR2(254)	Vendor
Description	VARCHAR2(1024)	
Status	NUMBER(38)	Operational status
IP	VARCHAR2(16)	
DNS	VARCHAR2	
WWN	VARCHAR(16)	
Model	VARCHAR2(254)	
PowerManagement	VARCHAR2(254)	
SerialNumber	VARCHAR2(254)	
Version	VARCHAR2(254)	
ContainerID	NUMBER(38)	Container system Identifier
ProcessorStatus	NUMBER	Status of the processor
ResetCapability	VARCHAR2(254)	Reset Capability
Roles	VARCHAR2(254)	Roles
ProviderTag	VARCHAR2(254)	Provider name tag
SupportFlags	NUMBER(38)	Support flags
BASETABLENAME	CHAR(14)	Name of the base table

MVC_STORAGESYNCHRONIZEDVW

Column Name	Type	Description
SYNCHRONIZED	NUMBER(38)	Source Synchronizer Identifier
SOURCEID	NUMBER(38)	Source Identifier
SOURCENAME	VARCHAR2(256 CHAR)	Name of Source
SOURCETYPE	VARCHAR2(14 CHAR)	Type of Source

Column Name	Type	Description
SOURCESYSTEMID	NUMBER(38)	
SOURCESYSTEMNAME	VARCHAR2(256 CHAR)	Name of Source Ssytem
SOURCEVENDOR	VARCHAR2(254 CHAR)	Vendor of Source
SOURCEMODEL	VARCHAR2(254 CHAR)	Model of Source
SOURCEPROVIDERTAG	VARCHAR2(254 CHAR)	Source Provider Tag
TARGETID	NUMBER(38)	
TARGETNAME	VARCHAR2(256 CHAR)	Name of Target
TARGETTYPE	VARCHAR2(14 CHAR)	Type of Target
TARGETSYSTEMID	NUMBER(38)	Target System Identifier
TARGETSYSTEMNAME	VARCHAR2(256 CHAR)	Target System Name
TARGETVENDOR	VARCHAR2(254 CHAR)	Target Vendor
TARGETMODEL	VARCHAR2(254 CHAR)	Target Model
TARGETPROVIDERTAG	VARCHAR2(254 CHAR)	Target Provider Tag
REMOTEELEMENTID	VARCHAR2(1024 CHAR)	Remote Element Identifier
REMOTESYSTEMID	VARCHAR2(1024 CHAR)	Remote System Identifier
SYNCSTATE	VARCHAR2(128 CHAR)	Sync State
SYNCTIME	TIMESTAMP(6)	Sync Time
COPYTYPE	VARCHAR2(64 CHAR)	Copy Type
REPLICATYPE	VARCHAR2(64 CHAR)	Replica Type
SYNCMAINTAINED	NUMBER(1)	Sync Maintained
SYNCEDTIMESTAMP	TIMESTAMP(6)	Synced Timestamp
LOCALITY	VARCHAR2(254 CHAR)	Locality

MVC_STORAGESYSTEMCONFIGVW

Column Name	Type	Description
StorageSystemID	NUMBER(38)	Storage System Identifier
CollectionTime	TIMESTAMP(6)	Configuration statistics collection time

Column Name	Type	Description
ExportedMB	NUMBER	Storage exposed in megabytes
UnexportedMB	NUMBER	Volume unexposed in megabytes
AvailableMB	NUMBER	Available storage left on storage pool
ProvisionedMB	NUMBER	Provisioned storage in megabytes
RawStorageMB	NUMBER	Unconfigured storage in megabytes
TotalMB	NUMBER	Total storage in megabytes
CAPACITYTYPE	VARCHAR2(256 CHAR)	Type of Capacity
CAPACITYNUM	NUMBER	Capacity Number
RAWTOTALMB	NUMBER	Total amount of RAW MB
TOTALINTERNALMB	NUMBER	Total amount of Internal MB
AvailablePorts	NUMBER(38)	Number of available ports
ConnectedPorts	NUMBER(38)	Number of used ports
TotalPorts	NUMBER(38)	Total system ports
ActualProvisionedMB	NUMBER	Actual provisioned in megabytes
ActualExportedMB	NUMBER	Actual exported in megabytes
ActualUnexportedMB	NUMBER	Actual unexported in megabytes
ReservedMB	NUMBER	Reserved in megabytes
ReservedConsumableMB	NUMBER	Reserved consumable in megabytes
ReservedConsumedMB	NUMBER	Reserved consumed in megabytes

MVC_STORAGESYSTEMSUMMARYVW

Column Name	Type	Description
StorageSystemID	NUMBER(38)	Storage system Identifier
StorageSystemName	VARCHAR2(256)	Storage system Name

Column Name	Type	Description
DomainID	NUMBER(38)	Domain Identifier
Vendor	VARCHAR2(254)	Vendor
Description	VARCHAR2(1024)	Description of the Storage System
Status	NUMBER(38)	Operational Status (provide map here)
IP	VARCHAR2(16)	Not used
Model	VARCHAR2(254)	Model
SerialNumber	VARCHAR2(254)	Serial Number
Version	VARCHAR2(254)	Version
StorageSystemStatus	VARCHAR2(254)	Intrinsic status of the system
ResetCapability	VARCHAR2(254)	Indicating reset capability
ProvisionCapabilities	NUMBER(38)	Provide map here
SupportFlag	NUMBER(38)	Provide map here
BASETABLENAME	Varchar	Name of the base table for optional values
PRIMARY_OWNER_NAME	VARCHAR2(254 CHAR)	Primary Owner Name
PROVIDER_TAG	VARCHAR2(254 CHAR)	Provider Tag

MVC_STORAGETIERDETAILVW

Name	Description
Memberid	
Domainid	Domain of the member
Member	Name of the member
SSID	Storage system Identifier of the member
Storagesystem	Name of the storage system
TierID	Storage tier Identifier
TierName	Name of the storage tier

Name	Description
CostPerGB	cost per GB of the tier
TotalCapacity	Total capacity of the member
Storagetype	Type of storage

MVC_STORAGEVOLUMEPORTS

Column Name	Type	Description
ID	NUMBER(38)	
Storage_Volume_ID	NUMBER(38)	Storage Volume Identifier
Port_ID	NUMBER(38)	
LUN_ID	NUMBER(38)	
Access_Mode	Varchar(254)	
Access_State	Varchar(254)	
Host_Group	Varchar(254)	
Host_Group_Name	Varchar(254)	
Host_Group_Modes	Varchar(1024)	

MVC_STORAGEVOLUMESUMMARYVW

Column Name	Type	Description
StorageVolumeID	NUMBER(38)	StorageVolume Identifier
StorageVolumeName	VARCHAR2(256)	StorageVolume Name
DomainId	NUMBER(38)	Domain Identifier
OID	VARCHAR2(254)	Reserved
Status	NUMBER(38)	Operational status (provide map here)
StorageSystemID	NUMBER(38)	ID of the storage system that contains this volume
StorageCapabilityID	NUMBER(38)	Storage Capability Identifier
VolumeDeviceID	VARCHAR2(254)	Device Identifier

Column Name	Type	Description
AccessType	VARCHAR2(254)	Volume access type
Blocksize	NUMBER(38)	Size per block in bytes
NumberOfBlocks	NUMBER(38)	Total number of blocks in the volume
ConsumableBlocks	NUMBER(38)	Total number of consumable blocks
SeqAccess	NUMBER(1)	Sequential access
Availability	VARCHAR2(254)	Availability indication
StatusInfo	VARCHAR2(254)	Status of the volume
PoolID	NUMBER(38)	ID of the Storage Pool that contains this volume
VolumeType	NUMBER(38)	Type of volume
BASETABLENAME	CHAR(14)	Name of the base table for optional values
STORAGESYSTEMNAME	VARCHAR2(256 CHAR)	Storage System Name
STORAGECAPABILITIES	VARCHAR2(256 CHAR)	Storage Capabilities
PURPOSE	VARCHAR2(254 CHAR)	Purpose
VOLDFLTSNGLPTOFFAILURE	NUMBER(1)	Flag to know Volume Default single of failure
VOLNOSINGLEPOINTOFFAILURE	NUMBER(1)	Flag to know Volume No single Point of Failure
VOLMINDATAREDUNDANCY	NUMBER(18)	Volume Minimum Data Redundancy
VOLMAXDATAREDUNDANCY	NUMBER(18)	Volume Maximum Data Redundancy
VOLDEFAULTDATAREDUNDANCY	NUMBER(18)	Volume Default Data Redundancy
VOLMINSPINDLEREDUNDANCY	NUMBER(18)	Volume Minimum Spindle Redundancy

Column Name	Type	Description
VOLMAXSPINDLEREDUNDANCY	NUMBER(18)	Volume Maximum Spindle Redundancy
VOLDEFAULTSPINDLEREDUNDANCY	NUMBER(18)	Volume Default Spindle Redundancy
VOLMINDELTARESERVATION	NUMBER(18)	Volume Minimum Delta Reservation
VOLMAXDELTARESERVATION	NUMBER(18)	Volume Maximum Delta Reservation
DEFAULTDELTARESERVATION	NUMBER(18)	Volume Default Delta Reservation
ConsumedBlocks	NUMBER(38)	Consumed blocks
ThinlyProvisioned	NUMBER(1)	To find out whether or thinly provisioned

MVC_SUBPATHVW

Column Name	Type	Description
SubPathID	NUMBER(38)	SubPath Identifier
PathID	NUMBER(38)	Parent Path Identifier
HostID	NUMBER(38)	Host Identifier
DiskDriveID	NUMBER(38)	Disk Drive Identifier
HBACardID	NUMBER(38)	HBA Card Identifier
HBAPortID	NUMBER(38)	HBA Port Identifier
HostSwitchPortID	NUMBER(38)	ID of Host Switch Port
SystemSwitchPortID	NUMBER(38)	ID of the system switch port
StorageSystemPortID	NUMBER(38)	
StorageVolumeID	NUMBER(38)	
StorageSystemID	NUMBER(38)	
LUN	NUMBER(38)	
IsLocal	VARCHAR2(6)	
FabricID	NUMBER(38)	

Column Name	Type	Description
MultipathDeviceID	NUMBER(38)	
PathSwitchID	NUMBER(38)	

MVC_SWITCHCONFIGVW

Column Name	Type	Description
SwitchID	NUMBER(38)	Switch Identifier
AvailablePorts	NUMBER(38)	Number of available ports
ConnectedPorts	NUMBER(38)	Number of used ports
TotalPorts	NUMBER(38)	Number of total ports of this switch
COLLECTIONTIME	DATE	Timestamp of the Collection
CAPACITYTYPE	VARCHAR2(256 CHAR)	Capacity Type
CAPACITYNUM	NUMBER	Capacity Number

MVC_SWITCHSUMMARYVW

Column Name	Type	Description
SwitchID	NUMBER(38)	Switch Identifier
SwitchName	VARCHAR2(256)	Switch Name
DomainID	NUMBER(38)	Domain Identifier
Vendor	VARCHAR2(254)	Switch Vendor
Description	VARCHAR2(1024)	Description of the Switch
Status	NUMBER(38)	Operational status (provide map here)
IP	VARCHAR2(16)	Switch IP
DNS	VARCHAR2(50)	DNS of the Switch
WWN	VARCHAR2(254)	WWN of the Switch
Model	VARCHAR2(254)	
SerialNumber	VARCHAR2(254)	Serial Number of the Switch
Version	VARCHAR2(254)	Switch's hardware version

Column Name	Type	Description
LoginName	VARCHAR2(254)	Login name for this Switch
LoginPwd	VARCHAR2(254)	Login password for this Switch
HardZoningCapability	VARCHAR2(254)	
SoftZoningCapability	VARCHAR2(254)	
ZoningInstalled	NUMBER(1)	
MaxModuleNumber	NUMBER(38)	
CurrentZoningEnforcement	VARCHAR2(254)	
SwitchDomainID	NUMBER(38)	
SwitchStatus	VARCHAR2(254)	
SwitchState	VARCHAR2(254)	
IPGateway	VARCHAR2(254)	
IPNetwork	VARCHAR2(16)	
FCAddress	VARCHAR2(254)	
FCNetmask	VARCHAR2(16)	
SwitchRole	VARCHAR2(254)	
ProvisionSupportFlag	NUMBER(1)	
FabricWWN	Varchar	
FabricID	NUMBER(38)	Fabric Identifier
BASETABLENAME	CHAR(6)	Name of the base table for optional values
APPIQ_LAST_ CONTACTED	DATE	Appiq Last Contacted
CIMDOMAINID	NUMBER(38)	CIM Domain Identifier
FABRICNAME	VARCHAR2(256 CHAR)	Fabric Name
ParentID	NUMBER(38)	Parent Identifier
SanID	NUMBER(38)	SAN Identifier
ReplicationSanID	NUMBER(38)	Replication SAN Identifier
VirtualFabricID	NUMBER(38)	Virtual fabric Identifier

MVC_TAPELIBRARYSUMMARYVW

Column Name	Type	Description
TAPELIBRAYID	NUMBER(38)	ID of Tape Library
TAPELIBRARYNAME	VARCHAR2(256 CHAR)	Name of Tape Library
DOMAINID	NUMBER(38)	Domain Identifier
VENDOR	VARCHAR2(254 CHAR)	Name of Vendor
DESCRIPTION	VARCHAR2(1024 CHAR)	Description of Tape Library
STATUS	NUMBER(38)	Status
IP	VARCHAR2(1024 CHAR)	IP Address
MODEL	VARCHAR2(254 CHAR)	Model Name
PRIMARY_OWNER_NAME	VARCHAR2(254 CHAR)	Primary Owner Name
SERIALNUMBER	VARCHAR2(1024 CHAR)	Serial Number
VERSION	VARCHAR2(254 CHAR)	Version
PRODUCT	VARCHAR2(254 CHAR)	Product
BASETABLENAME	CHAR(12 CHAR)	Base Table Name

MVC_USERROLEMAPVWV

Column Name	Type	Description
DOMAINID	NUMBER(38)	Domain Identifier
USERID	NUMBER(38)	ID of User
ROLEID	NUMBER(38)	ID of Role
ORGID	NUMBER(38)	ID of Organization
IS_ACTIVE	NUMBER(1)	Flag to Know if it is Active or Not

MVC_USERVW

Column Name	Type	Description
USERID	NUMBER(38)	ID of User
NAME	VARCHAR2(256 CHAR)	Name of User

Column Name	Type	Description
FULLNAME	VARCHAR2(512 CHAR)	Full Name of User
PASSWORD	RAW(256)	Password of User
ADMIN	NUMBER(1)	Flag to Know if the User is an Admin
EMAILADDRESS	VARCHAR2(254 CHAR)	Email address of User
PHONENUMBER	VARCHAR2(32 CHAR)	Phone Number of User
CONTACTINFO	VARCHAR2(254 CHAR)	Contact Information for User

MVC_VOLUMEDISKDRIVEVW

Column Name	Type	Description
VolumeID	NUMBER(38)	Storage Volume Identifier
DiskDriveID	NUMBER(38)	Disk Drive Identifier
ExtentID	NUMBER(38)	Disk Extent Identifier

MVC_ZONEPORTVW

Column Name	Type	Description
ZONEID	NUMBER(38)	ID of Zone
ZONEMEMBERID	NUMBER(38)	ID of Zone Member
PORTWWN	VARCHAR2(32 CHAR))	WWN of Port

Column Name	Type	Description
ZoneID	NUMBER(38)	Zone Identifier
ZoneName	VARCHAR2(254)	Zone Name
DominaID	NUMBER(38)	Domain Identifier (currently only one domain)
CimClassName	VARCHAR2(28)	
Status	NUMBER(38)	AppIQ status
ActiveZone	VARCHAR2(3)	

Column Name	Type	Description
ZoneType	VARCHAR2(254)	
ProtocolType	VARCHAR2(254)	
ReadOnly	NUMBER(1)	
FabricID	NUMBER(38)	Fabric Identifier
FabricWWN	VARCHAR2(254)	Fabric WWN
FabricCName	VARCHAR2(256)	
ZoneCapID	NUMBER(38)	Zone Capability Identifier
ZoneCapabilitiesName	VARCHAR2(254)	Name of the zone capabilities
ZC_MaxName_length	NUMBER(18)	Name length limit
MaxZoneSets	NUMBER(18)	Number of maximum zone sets
MaxZones	NUMBER(18)	
MaxZoneMembers	NUMBER(18)	
MaxZonePerZoneSet	NUMBER(18)	
MaxZoneAliases	NUMBER(18)	
EnhencdZoning	NUMBER(1)	
OBJECTTYPE	VARCHAR2(254 CHAR)	Type of the Object
SanID	NUMBER(38)	SAN Identifier
ReplicationSanID	NUMBER(38)	Replication SAN Identifier
VirtualFabricID	NUMBER(38)	Virtual fabric Identifier

MVC_ZONEVW

Column Name	Type	Description
ZoneSetID	NUMBER(38)	Zone Set Identifier
ZoneSetName	VARCHAR2(256)	Zone Set Name
ZoneID	NUMBER(38)	Zone Identifier
FabricID	NUMBER(38)	ID of the fabric which the zone belongs
ZoneMemberID	NUMBER(38)	Zone Member Identifier

Column Name	Type	Description
ZoneMemberName	VARCHAR2(254)	Name of the zone member
ZoneMemberType	VARCHAR2(254)	Type of the zone member
ZoneMemberInFabric	NUMBER(1)	
ZonePortWWN	VARCHAR2(32)	WWN of the zone port
ZoneAlias	VARCHAR2(256)	
ZONEALIASID	NUMBER	Zone Alias Identifier

MVCA_BU_CLIENTDETAILVW

Column Name	Type	Description
CLIENTID	NUMBER	Client Identifier
CLIENTNAME	VARCHAR2(256)	Client Name
MASTERSERVERID	NUMBER	Master Server Identifier
HOSTID	NUMBER	ID of Host
VENDOR	VARCHAR2(254 CHAR)	Name of Vendor
DESCRIPTION	VARCHAR2(1024)	Description
STATUS	NUMBER(38)	Status
DOMAINID	NUMBER(38)	Domain Identifier

MVCA_BU_DRIVESTATISTICSVW

Column Name	Type	Description
STAT	VARCHAR2(800 CHAR)	Stat Name
DRIVE	NUMBER	Drive Identifier
COLLECTIONTIME	VARCHAR2(800 CHAR)	Time Stamp of the Collection
MASTERSERVERID	NUMBER	Master Server Identifier

MVCA_BU_DRIVESTATVW

Name	Description
Driveid	ID of the drive
Collectiontime	Timestamp of the collection
ID	
Robotnumber	Robot number
Serialnumber	Serial number
Status	Status of the drive
Devicetime	
APPLICATION	Application name
Mediainuse	Media usage
Duration	Duration

MVCA_BU_IMAGEDETAILVW

Column Name	Type	Description
IMAGEID	NUMBER(38)	HostID
CLIENTNAME	VARCHAR2(1024 CHAR)	Client Name
BACKUPID	VARCHAR2(512 CHAR)	Backup Identifier
POLICYNAME	VARCHAR2(1024 CHAR)	Policy Name
POLICYTYPE	VARCHAR2(512 CHAR)	Policy Type
SCHEDULENAME	VARCHAR2(1024 CHAR)	Schedule Name
SCHEDULETYPE	VARCHAR2(512 CHAR)	Schedule Type
RETENTIONLEVEL	VARCHAR2(512 CHAR)	Retention Level
FILES	NUMBER	Number of Files
EXPIRY	DATE	Expiry Date
COMPRESSION	VARCHAR2(256 CHAR)	Compression
ENCRYPT	VARCHAR2(256)	Encrypt
FRAGMENTS	NUMBER	Fragments
FRAGMENTSSIZE	NUMBER	Fragments Size

Column Name	Type	Description
MEDIAID	VARCHAR2(512 CHAR)	Media Identifier
MASTERSERVERNAME	VARCHAR2(128 CHAR)	Master Server Name
MEDIASERVERNAME	VARCHAR2(1024 CHAR)	Media Server Name

MVCA_BU_JOBDETAILVW

Column Name	Type	Description
JOBID	NUMBER	ID of Job
CREATIONDATE	DATE	Creation Date
TEMPLATEID	NUMBER	Template Identifier
TEMPLATENAME	VARCHAR2(1024 CHAR)	Template Name
MASTERSERVERID	NUMBER	Master Server Identifier
CLIENTID	NUMBER	Client Identifier
CLIENTNAME	VARCHAR2(1024 CHAR)	Client Name
MEDIAUSED	VARCHAR2(512 CHAR)	Media Used
BUJOBID	VARCHAR2(512 CHAR)	Backup Job Identifier
BUJOBTYPE	VARCHAR2(512 CHAR)	Backup Job Type
JOBSTATE	VARCHAR2(512 CHAR)	Backup Job State
JOBSTATUS	VARCHAR2(256 CHAR)	Backup Job Status
SCHEDULENAME	VARCHAR2(1024 CHAR)	Schedule Name
STORAGEUNIT	VARCHAR2(1024 CHAR)	Storage Unit
BUTARGETSERVER	VARCHAR2(1024 CHAR)	Backup Target Server
FILESLASTWRITTEN	NUMBER	Number of Files Last Written
STARTTIME	DATE	Start Time
ENDTIME	DATE	End Time
DESCRIPTION	VARCHAR2(2048 CHAR)	Job Description
TIME	NUMBER	Time
RETENTIONPERIOD	VARCHAR2(256 CHAR)	Retention Period

Column Name	Type	Description
COMPRESSION	VARCHAR2(256 CHAR)	Compression
PRIORITY	VARCHAR2(256 CHAR)	Priority
KBLASTWRITTEN	NUMBER	Last Written in KB
FILELISTCOUNT	NUMBER	Number of Files Listed
MEDIASERVER	VARCHAR2(1024 CHAR)	Media Server
DRIVEUSED	VARCHAR2(1024 CHAR)	Drive Used
JOBSTARTTIME	DATE	Job Start Time
JOBFINISHTIME	DATE	Job Finish Time
SEQUENCE	NUMBER(38)	Sequence
JOBMETATYPE	NUMBER(38)	Job Meta Type

MVCA_BU_JOB_SUMMARYVW

Column Name	Type	Description
JOBSUMMARYID	NUMBER	Job Summary Identifier
MASTERSERVERNAME	VARCHAR2(256)	Master Server Name
MEDIASERVERNAME	VARCHAR2(1024 CHAR)	Media Server Name
VENDOR	VARCHAR2(1024 CHAR)	Name of Vendor
CAPACITY	NUMBER(38)	Capacity
TOTALCLIENTS	NUMBER(38)	Total Number of Clients
TOTALJOBS	NUMBER(38)	Total Number of Jobs
TOTALSUCCESSFULJOBS	NUMBER(38)	Total Number of Successful Jobs
TOTALFAILEDJOBS	NUMBER(38)	Total Number of Failed Jobs
TOTALWARNINGJOBS	NUMBER(38)	Total Number of Warning Jobs

Column Name	Type	Description
TOTALPARTIALSUCCESSJOBS	NUMBER(38)	Total Number of Partial Success Jobs
KB	NUMBER(38)	KBytes
COLLECTIONTIME	TIMESTAMP(6)	Collection Time

MVCA_BU_LIBRARYDETAILVW

Column Name	Type	Description
TAPELIBRARYID	NUMBER	HostID
TAPELIBRARYNAME	VARCHAR2(256)	Host Name
PROVIDERNAME	VARCHAR2(256)	DomainID
VENDOR	VARCHAR2(254)	Host Vendor
DESCRIPTION	VARCHAR2(1024)	Host Description
MEDIASERVERID	NUMBER	Operation status (provide map here)
MEDIASERVERNAME	VARCHAR2(256 CHAR)	
MASTERSERVERID	NUMBER	
ROBOTYPE	VARCHAR2(1024 CHAR)	
ROBOTNUMBER	NUMBER	
TOTALNOOFSLOTS	NUMBER	
TOTALSLOTSINUSE	NUMBER	
TOTALNUMBEROFDRIVES	NUMBER	
ROBOTDEVICEPATH	VARCHAR2(2048 CHAR)	
DOMAINID	NUMBER(38)	

MVCA_BU_MASTERSERVERDETAILVW

Column Name	Type	Description
MASTERSERVERID	NUMBER	Master Server Identifier
MASTERSERVERNAME	VARCHAR2(256)	Master Server Name
HOSTID	NUMBER(38)	Host Identifier
VENDOR	VARCHAR2(254 CHAR)	Name of Vendor
DESCRIPTION	VARCHAR2(1024 CHAR)	Description
STATUS	NUMBER(38)	Master Server Status
PRODUCTNAME	VARCHAR2(254)	Product Name
LICENSEKEY	VARCHAR2(512)	License Key
LICENSEFEATURES	VARCHAR2(2048 CHAR)	License Features
DOMAINID	NUMBER(38)	Domain Identifier

Column Name	Type	Description
MEDIAID	NUMBER	Media Identifier
MEDIANAME	VARCHAR2(256 CHAR)	Media Name
TAPELIBRARYNAME	VARCHAR2(256 CHAR)	Tape Library Name
LOCATION	VARCHAR2(513 CHAR)	Location
MASTERSERVERID	NUMBER	Master Server Identifier
TLMEDIAID	VARCHAR2(1024 CHAR)	Tape Library Media
TYPE	VARCHAR2(1024 CHAR)	Media Type
BARCODE	VARCHAR2(1024 CHAR)	Bar Code
MEDIAPOOLNAME	VARCHAR2(256 CHAR)	Media Pool Name
ROBOTTYPE	VARCHAR2(1024 CHAR)	Robot Type
ROBOTNUMBER	NUMBER	Robot Number
ROBOTSLOT	NUMBER	Robot Slot
ROBOTHOST	VARCHAR2(1024 CHAR)	Robot Host
VOLUMEGROUP	VARCHAR2(1024 CHAR)	Volume Group

Column Name	Type	Description
CREATED	DATE	Created
ASSIGNED	DATE	Assigned
LASTMOUNTED	DATE	Last Mounted
FIRSTMOUNTED	DATE	First Mounted
EXPIRATIONDATE	DATE	Expiration Date
NUMBEROFMONTHS	NUMBER	Number of Months
MAXMOUNTSALLOCATED	NUMBER	Maximum Allocated Mounts
DENSITY	VARCHAR2(512 CHAR)	Media Density
TIMEALLOCATED	DATE	Time Allocated
LASTWRITTEN	DATE	Last Written Date
EXPIR	VARCHAR2(512 CHAR)	Media Expire
LASTREAD	DATE	Media Last Read
MBYTES	NUMBER	Media Used in Mega Bytes
NIMAGES	NUMBER	N Images
VIMAGES	NUMBER	V Images
RL	VARCHAR2(512 CHAR)	Media Retention Level
TOTALRESTORES	NUMBER	Total Number of Restores
MEDIASTATUS	VARCHAR2(512 CHAR)	Status of the Media
VENDOR	VARCHAR2(254 CHAR)	Name of the Vendor
DESCRIPTION	VARCHAR2(1024 CHAR)	Media Description
DOMAINID	NUMBER(38)	Domain Identifier

MVCA_BU_MEDIASERVERDETAILVW

Column Name	Type	Description
MEDIASERVERID	NUMBER	Media Server Identifier
MEDIASERVERNAME	VARCHAR2(256)	Media Server Name
MASTERSERVERID	NUMBER	Master Server Identifier

Column Name	Type	Description
HOSTID	NUMBER(38)	ID of Host
VENDOR	VARCHAR2(254)	Name of Vendor
DESCRIPTION	VARCHAR2(1024 CHAR)	Description
STATUS	NUMBER(38)	Status
PRODUCTNAME	VARCHAR2(254 CHAR)	Product Name
LICENSEKEY	VARCHAR2(512 CHAR)	License Key
LICENSEFEATURES	VARCHAR2(2048 CHAR)	License Features
DOMAINID	NUMBER(38)	Domain Identifier

MVCA_BU_OPTIONALTABLEVW

Name	Description
Basetableid	ID of the basetable
Basetablename	Name of the basetable
Optionalname	Optional name
Optionalvalue	Optional value

MVCA_BU_RESOURCE_SUMMARYVW

Column Name	Type	Description
RESOURCESUMMARYID	NUMBER	Resource Summary Identifier
MASTERSERVERNAME	VARCHAR2(256)	Master Server Name
MEDIASERVERNAME	VARCHAR2(256)	Media Server Name
POOLNAME	VARCHAR2(256)	Pool Name
LIBRARYNAME	VARCHAR2(256)	Library Name
TOTALAVAILABLEMEDIA	NUMBER(38)	Total Available Media
TOTALALLOCATEDMEDIA	NUMBER(38)	Total Allocated Media
TOTALFROZENMEDIA	NUMBER(38)	Total Frozen Media
TOTALSUSPENDEDMEDIA	NUMBER(38)	Total Suspended Media
COLLECTIONTIME	TIMESTAMP(6)	Timestamp of Collection

MVCA_BU_TEMPLATECLIENTVW

Column Name	Type	Description
TEMPLATEID	NUMBER	Template Identifier
TEMPLATENAME	VARCHAR2(1024 CHAR)	Template Name
CLIENTID	NUMBER	Client Identifier
CLIENTNAME	VARCHAR2(256)	Client Name
CLIENTHOSTID	NUMBER	Client Host Identifier
MASTERSERVERID	NUMBER	Master Server Identifier
MASTERSERVERNAME	VARCHAR2(256 CHAR)	Master Server Name

MVCA_BU_TEMPLATEDETAILVW

Column Name	Type	Description
TEMPLATEID	NUMBER	Template Identifier
MASTERSERVERID	NUMBER	Master Server Identifier
TEMPLATENAME	VARCHAR2(1024 CHAR)	Template Name
TEMPLATETYPE	VARCHAR2(512 CHAR)	Template Type
STATUS	VARCHAR2(512 CHAR)	Status
CLIENTCOMPRESS	VARCHAR2(16 CHAR)	Client Compress
FOLLOWNFSMOUNTS	VARCHAR2(512 CHAR)	Follow NFS Mounts
COLLECTTIRINFO	VARCHAR2(1024 CHAR)	Collect TIR Information
BLOCKINCREMENTAL	VARCHAR2(512 CHAR)	Block Incremental
MULTIDATASTREAM	VARCHAR2(16 CHAR)	Multi Data Stream
MAXJOBS	VARCHAR2(16 CHAR)	Maximum Jobs
CLIENTENCRIPT	VARCHAR2(16 CHAR)	Client Encrypt
FSINCLUDELIST	VARCHAR2(4000 CHAR)	File Server Include List
FSEXCLUDELIST	VARCHAR2(4000 CHAR)	File Server Exclude List
FILERESTORERAW	VARCHAR2(16 CHAR)	File Restore Raw
TEMPLATEPRIORITY	NUMBER	Template Priority Number

Column Name	Type	Description
DISASTERRECOVERY	NUMBER	Disaster Recovery
STORAGEUNIT	VARCHAR2(1024 CHAR)	Storage Unit
VOLUMEPOOL	VARCHAR2(1024 CHAR)	Volume Pool
SEQUENCE	NUMBER(38)	Sequence

MVCA_DBAPPINSTCAPACITYVW

Column Name	Type	Description
DBAPPLICATIONID	NUMBER(38)	ID of Database Application
HOSTID	NUMBER(38)	ID of Host
CAPACITYTYPE	VARCHAR2(256 CHAR)	Type of Capacity
CAPACITYNUM	NUMBER	Capacity Number
TIMESTAMP	DATE	Timestamp of the Collection
DBINSTANCENAME	VARCHAR2(255 CHAR)	Database Instance Name
TOTALMB	NUMBER	Total Size in MB
FREEMB	NUMBER	Free Size in MB
USEDMB	NUMBER	Used Size in MB

MVCA_DBAPPPHYCAPACITYVW

Column Name	Type	Description
DBAPPLICATIONID	NUMBER(38)	ID of the Database Application
HOSTID	NUMBER(38)	Host Identifier
CAPACITYTYPE	VARCHAR2(256)	Type of the Capacity
CAPACITYNUM	NUMBER	Capacity Number
TIMESTAMP	DATE	Timestamp of the Collection
DBINSTANCENAME	VARCHAR2(255 CHAR)	Database Instance Name
DBLOGICNAME	VARCHAR2(265 CHAR)	Database Logical Name
DBPHYSICALNAME	VARCHAR2(512 CHAR)	Database Physical Name

Column Name	Type	Description
TOTALMB	NUMBER	Total Size in MB
FREEMB	NUMBER	Free Size in MB
USEDMB	NUMBER	Used Size in MB

MVCA_EXCHANGESTORESUMMARYVW

Name	Description
StoreID	ID of the store
StorageGroupID	Storage group Identifier
ActiveDirectoryName	Directory name
StoreName	Name of the store
Filepath	File path of the store
Isonline	Online status
Private_store	

MVCA_EXCHAPPCAPACITYVW

Column Name	Type	Description
ExchangeAppID	NUMBER(38)	
CapacityType	Varchar2(7)	
Timestamp	Timestamp	
StorageGroupID	NUMBER(38)	
ExchangeFilePath	Varchar(512)	
TotalMB	NUMBER	
FreeMB	NUMBER	
UsedMB	NUMBER	
CAPACITYNUM	NUMBER	Capacity Number
STORAGEGROUPNAME	VARCHAR2(255 CHAR)	Name of the Storage Group

MVCA_EXCHMAILBOXDETAILVW

Name	Description
UserMailboxID	Mailbox Identifier
UserMailBoxName	Mailbox name
ServerName	Name of Exchange server
StoreID	Store Identifier of the mailbox
MailboxMessageSizeBytes	Messages size
UserMailBoxSizebytes	Mailbox size
Legacy_dn	
EmailAddress	Email address
NTUserName	Name of NT user
LastLogonTime	Timestamp of last logon
Storage_limit_info	Storage limit
CountofNormalMessages	Count of normal messages
CountofAssociatedMessages	Count of attachment messages
Applicationid	ID of the exchange application

MVCA_EXCHPUBLICFOLDERDETAILVW

Name	Description
PublicFolderID	Public folder Identifier
PublicFolderName	Public folder name
EmailAddress	Email address of the user
Legacy_dn	
Server_name	Name of the server
Applicationid	ID of the exchange application
FolderPath	Path of the folder
StoreID	ID of the store
CountofContacts	Count of contacts in the mailbox
CountofMessages	Count of messages

Name	Description
Associated_content_count	Associated messages count
MessageSizeinFolderbytes	Message size in bytes
FolderSizeBytes	Public folder size
CreationTime	Timestamp of creation
LastAccessTime	Timestamp of last access
LastModifiedTime	Timestamp of the last modification of the file
CountofOwners	Count of owner of folder

MVCA_EXCHSTORGROUPSUMMARYVW

Name	Description
StorageGroupID	ID of the storage group
StorageGroupName	Name of the storage group
ActiveDirectoryName	Name of active directory
ApplicationID	Application Identifier

MVCA_FSRM_ACL_DETAILSVW

Column Name	Type	Description
VOLUMEID	NUMBER(38)	ID of File Server Volume
OWNERID	NUMBER(38)	Owner Identifier
TRUSTEENAME	VARCHAR2(256)	Trustee Name
TOTALSIZE	NUMBER(38)	Total Size
TOTALDIRS	NUMBER(38)	Total Directories
TIMESTAMP	DATE	Timestamp of the Collection

MVCA_FSRM_AGEDFILEDETAILS

Column Name	Type	Description
VolumeID	NUMBER(38)	

Column Name	Type	Description
FileName	Varchar2(254)	
FileSize	NUMBER(38)	
FileAge	NUMBER(38)	
Timestamp	Timestamp(6)	
DomainID	NUMBER(38)	
FILECREATED	DATE	File Created Date
FILEMODIFY	DATE	File Modified Date
FILEOWNERID	NUMBER(38)	File Owner Identifier
FILEOWNERNAME	VARCHAR2(256 CHAR)	File Owner Name

MVCA_FSRM_AGEID_DETAILVW

Column Name	Type	Description
AGEID	NUMBER(38)	ID of the Age
MINAGE	NUMBER(38)	Minimum Age
MAXAGE	NUMBER(38)	Maximum Age
AGETYPE	VARCHAR2(10 CHAR)	Age Type

MVCA_FSRM_AGESUMMARYVW

Column Name	Type	Description
AgeID	NUMBER(38)	
VolumeID	NUMBER(38)	
TotalFiles	NUMBER(38)	
TotalSize	NUMBER(38)	
Timestamp	Timestamp(6)	

MVCA_FSRM_DIRDETAILSUMMARYVW

Column Name	Type	Description
DirKey	NUMBER(38)	
DirName	Varchar2(254)	
DirSize	NUMBER(38)	
TotalSubDirectories	NUMBER(38)	
TotalFiles	NUMBER(38)	
VolumeID	NUMBER(38)	
Timestamp	Timestamp(6)	
OWNERID	NUMBER(38)	ID of the Owner
OWNERNAME	VARCHAR2(256 CHAR)	Name of the Owner

MVCA_FSRM_DIRREPORTDATAVW

Name	Description
Volumeid	ID of FSRM volume
Volumename	Name of volume
Ruleid	Rule Identifier
Rulename	Name of the rule created
Fullpath	Path of the file
Filename	Name of the file
Totalsize	Total size of file
Accesstime	Timestamp of access time
Createtime	Timestamp of creation time
Modifytime	Timestamp of modified time
Attributes	
Owner	Owner name of file
TOTALSUBDIRS	Total Sub Directories
TOTALFILES	Total Number of Files
CURRENTDIRSIZE	Current Directory Size

Name	Description
CURRENTDIRSUBDIRS	Current directory Sub Directories
CURRENTDIRFILES	5. Current Directory Files

MVCA_FSRM_EXTDETAILSUMMARYVW

Column Name	Type	Description
ExtName	Varchar2(254)	
VolumeID	NUMBER(38)	
TotalFiles	NUMBER(38)	
TotalSize	NUMBER(38)	
Timestamp	Timestamp(6)	

MVCA_FSRM_FILEREPORTDATAVW

Name	Description
Volumeid	ID of FSRM volume
Volumename	Name of volume
Ruleid	Rule Identifier
Rulename	Name of the rule created
Fullpath	Path of the file
Filename	Name of the file
Totalsize	Total size of file
Accesstime	Timestamp of access time
Createtime	Timestamp of creation time
Modifytime	Timestamp of modified time
Attributes	
Owner	Owner name of file

MVCA_FSRM_LARGEDIRINFO

Column Name	Type	Description
VolumeID	NUMBER(38)	
DirName	Varchar2(256)	
DirSize	NUMBER(38)	
TotalFiles	NUMBER(38)	
TotalDirs	NUMBER(38)	
Timestamp	Timestamp(6)	
DomainID	NUMBER(38)	
OWNERID	NUMBER(38)	Directory Owner Identifier
OWNERNAME	VARCHAR2(256 CHAR)	Directory Owner Name

MVCA_FSRM_REPORTRULEVW

Name	Description
Ruleid	ID of the FSRM rule created
Rule_name	Name of the rule created
Description	Description
Collection_type	Type of collection
CONDITIONSQL	Report Rule SQL Condition

MVCA_FSRM_RULE_VOLUME_MAPVW

Column Name	Type	Description
RULEID	NUMBER(38)	Rule Identifier
VOLUMEID	NUMBER(38)	ID of the File Server Volume

MVCA_FSRM_SIDDETAILSVW

Column Name	Type	Description
VOLUMEID	NUMBER(38)	ID of File Server Volume
OWNERID	NUMBER(38)	ID of File Server Owner

Column Name	Type	Description
SIDSIZE	NUMBER(38)	Size of SID
NOOFFILES	NUMBER(38)	Number of Files
LARGEFILEPOSITIONS	VARCHAR2(2000 CHAR)	Large File Positions
AGEDFILEPOSITIONS	VARCHAR2(2000 CHAR)	Aged File Positions
FIRSTRECORDPOSITON	NUMBER(38)	First Record Position

MVCA_FSRM_SIZEID_DETAILVW

Column Name	Type	Description
SIZEID	NUMBER(38)	Size Identifier
MINSIZE	NUMBER(38)	Minimum Size
MAXSIZE	NUMBER(38)	Maximum Size

MVCA_FSRM_SIZE_DETAILSVW

Column Name	Type	Description
VOLUMEID	NUMBER(38)	ID of the File Server Volume
SIZEID	NUMBER(38)	ID of the Size
TOTALFILES	NUMBER(38)	Total Number of Files
TOTALSIZE	NUMBER(38)	Total Size
AGEDFILEPOSITIONS	VARCHAR2(2000)	Aged File Positions
FIRSTRECORDPOSITON	NUMBER(38)	First Record Position
TIMESTAMP	DATE	Timestamp of the Collection

MVCA_FSRM_TEMPFILEDETAILS

Column Name	Type	Description
VOLUMEID	NUMBER(38)	Volume Identifier
FILENAME	VARCHAR2(4000 CHAR)	Name of the File
FILESIZE	NUMBER(38)	Size of the File

Column Name	Type	Description
FILEAGE	DATE	Age of the File
FILECREATED	DATE	File Created Time
FILEMODIFY	DATE	File Modified Time
TIMESTAMP	DATE	Timestamp
DOMAINID	NUMBER(38)	Domain Identifier
FILEOWNERID	NUMBER(38)	File Owner Identifier
FILEOWNERNAME	VARCHAR2(256 CHAR)	File Owner Name

MVCA_FSRM_TOPNFILES

Column Name	Type	Description
FSID	NUMBER(38)	File System Identifier
DomainID	NUMBER(38)	Domain Identifier
VOLUMEID	NUMBER(38)	ID of the File Server Volume.
FILENAME	VARCHAR2(4000 CHAR)	Name of the File
FILESIZE	NUMBER(38)	Size of the File
FILEAGE	DATE	Age of the File
FILECREATED	DATE	File Creation Date
FILEMODIFY	DATE	File Modified Date
FILEOWNERID	NUMBER(38)	File Owner Identifier
FILEOWNERNAME	VARCHAR2(256 CHAR)	File Owner Name
FileAttributes	NUMBER(8)	File attributes

MVCA_FSRM_USER_DIR_MAPVW

Column Name	Type	Description
USERID	NUMBER(38)	ID of File Server User
DIRNAME	VARCHAR2(4000 CHAR)	Directory Name

MVCA_FSRM_USERDETAILVW

Column Name	Type	Description
USERID	NUMBER(38)	ID of the File Server User
FSID	NUMBER(38)	File Server Identifier
VOLUMEID	NUMBER(38)	ID of the File Server Volume
USEDSPACE	NUMBER(38)	Used Space
USESPACEIGNORECASE	NUMBER(38)	Used Space Ignore Case

MVCA_FSRM_USERSUMMARYVW

Column Name	Type	Description
UserID	NUMBER(38)	
FSID	NUMBER(38)	
UserProvidederID	Varchar2(254)	
UserName	Varchar2(254)	
Department	Varchar2(254)	
Email	Varchar2(254)	
Quota	NUMBER(38)	
DomainID	NUMBER(38)	
DEPARTMENTID	VARCHAR2(256 CHAR)	Id of the Department
OWNERNAME	VARCHAR2(256 CHAR)	Name of the Owner
HOMEDIRECTORY	VARCHAR2(4000 CHAR)	File Server User Home Directory
UserStatus	VARCHAR2(256 CHAR)	User status
Manager	VARCHAR2(256 CHAR)	Manager
ModifyTime	TIMESTAMP(6)	Modify time

MVCA_FSRM_VOLUMESUMMARYVW

Column Name	Type	Description
VolumeID	NUMBER(38)	
VolumeName	Varchar2(256)	

Column Name	Type	Description
FSID	NUMBER(38)	
TotalDirectories	NUMBER(38)	
TotalFiles	NUMBER(38)	
DomainID	NUMBER(38)	
Timestamp	Timestamp(6)	
OWNERBASEDUSER	NUMBER(1)	File Server User based on Owner

MVCA_NAMESPACESUMMARYVW

Column Name	Type	Description
NAMESPACEID	NUMBER(38)	Name Space Identifier
NAMESPACENAME	VARCHAR2(128 CHAR)	Name Space Name
INSTANCEID	NUMBER(38)	Instance Identifier
INSTANCENAME	VARCHAR2(128 CHAR)	Instance Name
DATABASEID	NUMBER(38)	Database Identifier
DBNAME	VARCHAR2(128 CHAR)	Database Name
TYPEOFDB	VARCHAR2(32 CHAR)	Type of Database
APPLICATIONID	NUMBER(38)	Application Identifier
APPLICATIONTYPE	VARCHAR2(32 CHAR)	Application Type

MVCA_VIRTUALAPPCAPACITYVW

Column Name	Type	Description
VirtualAppID	NUMBER(38)	
Timestamp	Date	
VirtualPath	Varchar2(512)	
TotalMB	NUMBER	
FreeMB	NUMBER	
UsedMB	NUMBER	

Column Name	Type	Description
CAPACITYTYPE	VARCHAR2(256 CHAR)	Capacity type
CAPACITYNUM	NUMBER	Capacity number

MVCS_EVACTRLSTATSVW

Name	Description
ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
AVGREADLATENCY	
AVGREADSIZE	
AVGWRITELATENCY	
AVGWRITESIZE	
CPUPERCENT	
DATAXFERPERCENT	
DELTAREADIOS	
DELTAREADLATENCY	
DELTAWRITEIOS	
DELTAWRITELATENCY	
PCTREADIOS	
PCTWRITEIOS	
READDATARATE	
READRATE	
TOTALDATARATE	
TOTALIORATE	
WRITEDATARATE	
WRITERATE	

MVCS_EVADISKSTATSVW

Name	Description
COLLECTIONTIME	
ID	
STATSTYPE	
DEVICETIME	
DURATION	
AVGDRIVELATENCY	
AVGQUEUEDEPTH	
AVGREADLATENCY	
AVGREADSIZE	
AVGWRITELATENCY	
AVGWRITESIZE	
DELTADRIVELATENCY	
DELTAREADIOS	
DELTAREADLATENCY	
DELTATOTALIOS	
DELTAWRITEIOS	
DELTAWRITELATENCY	
PCTREADIOS	
PCTWRITEIOS	
READDATARATE	
READRATE	
TOTALDATARATE	
TOTALIORATE	
WRITEDATARATE	
WRITERATE	

MVCS_EVAHOSTFCPORTSTATSVW

Name	Description
ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
AVGQUEUEDEPTH	
AVGREADLATENCY	
AVGWRITELATENCY	
BADCRCERR	
DELTAREADIOS	
DELTAREADLATENCY	
DELTAWRITEIOS	
DELTAWRITELATENCY	
DISCARDFRAMES	
LINKFAILURE	
LOSSOFSIGNAL	
LOSSOFSYNCH	
PCTREADIOS	
PCTWRITEIOS	
PROTOCOLERROR	
READDATARATE	
READRATE	
RECEIVEEOFA	
TOTALDATARATE	
TOTALIORATE	
WRITEDATARATE	
WRITERATE	

MVCS_EVASPAGVOLUMESTATSVW

Name	Description
ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
AVGREADHITLATENCY	
AVGREADMISSLATENCY	
AVGREADSIZE	
AVGWRITELATENCY	
AVGWRITESIZE	
DELTAREADHITIOS	
DELTAREADHITLATENCY	
DELTAREADMISSIOS	
DELTAREADMISSLATENCY	
DELTAWRITEIOS	
DELTAWRITELATENCY	
FLUSHDATARATE	
FLUSHRATE	
MIRRORDATARATE	
PCTREADIOS	
PCTWRITEIOS	
PREFETCHDATARATE	
READDATARATE	
READHITDATARATE	
READHITRATE	
READMISSDATARATE	

Name	Description
READMISSRATE	
READRATE	
TOTALDATARATE	
TOTALIORATE	
WRITEDATARATE	
WRITERATE	

MVCS_EVASTORAGESYSTEMSTATSVW

Name	Description
ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
TOTALDATARATE	
TOTALIORATE	

MVCS_EVAVOLUMESTATSVW

Name	Description
ID	
COLLECTIONTIME	
STATSTYPE	
DEVICETIME	
DURATION	
AVGREADHITLATENCY	
AVGREADMISSLATENCY	
AVGREADSIZE	

Name	Description
AVGWRITELATENCY	
AVGWRITESIZE	
DELTAREADHITIOS	
DELTAREADHITLATENCY	
DELTAREADMISSIOS	
DELTAREADMISSLATENCY	
DELTAWRITEIOS	
DELTAWRITELATENCY	
FLUSHDATARATE	
FLUSHRATE	
MIRRORDATARATE	
PCTREADIOS	
PCTWRITEIOS	
PREFETCHDATARATE	
READDATARATE	
READHITDATARATE	
READHITRATE	
READMISSDATARATE	
READMISSRATE	
READRATE	
TOTALDATARATE	
TOTALIORATE	
WRITEDATARATE	
WRITERATE	

MVCS_HOSTCPUSTATSVW

Name	Description
Hostid	Host Identifier
CAPACITYTYPE	Type of capacity (Raw,Daily,Weekly,Monthly)
Processorid	ID of the processor
Name	Name of the processor
Timestamp	Collection time
Pctprocessortime	

MVCS_HOSTMEMORYSTATSVW

Name	Description
Hostid	Host Identifier
CAPACITYTYPE	Type of capacity (Raw, Daily, Weekly, Monthly)
Timestamp	Collection time
Percentphysicalused	Percentage of physical memory used
Freephysicalmemory	Percentage of physical memory free
Percentvirtualused	Percentage of virtual memory used
Freevirtualmemory	Percentage of virtual memory free

MVCS_LSICTRLAGDISKSTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTIONTIME	DATE	Timestamp of the Collection
STATSTYPE	VARCHAR2(256)	Stats Type
ADDRECOVEREDERRORS	NUMBER	Controller Recovered Errors
ADDTOTALBANDWIDTH	NUMBER	Controller Total Bandwidth
ADDREADIORATE	NUMBER	Controller Read IO Rate
ADDTOTALIOS	NUMBER	Controller Total IOs
ADDTOTALIORATE	NUMBER	Controller Total IO Rate

Column Name	Type	Description
ADDWRITEIORATE	NUMBER	Controller Write IO Rate
ADDUNRECOVEREDERRORS	NUMBER	Controller Unrecovered Errors
ADDTIMEOUTS	NUMBER	Controller Timeouts
ADDBYTESTRANSFERED	NUMBER	Controller Bytes Transferred
ADDWRITEIOS	NUMBER	Controller Write IOs
ADDREADIOS	NUMBER	Controller Read IOs
ADDRETRIEDREQUESTS	NUMBER	Controller Retried Requests
ADDAVERAGETOTALSIZE	NUMBER	Controller Average Total Size
ADDPRECENTREAD	NUMBER	Controller Percent Read

MVCS_LSICTRLAGVOLUMESTATSVW

Column Name	Type	Description
ID	NUMBER	ID
COLLECTIONTIME	NUMBER	Timestamp of the Collection
STATSTYPE	VARCHAR2(256 CHAR)	Stats Type
AVREADIOS	NUMBER	Average Read IO
AVBYTESWRITTENLARGE	NUMBER	Average Bytes Written Large
AVTOTALIORATE	NUMBER	Average Total IO Rate
AVWRITEIOSLARGE	NUMBER	Average Write IO Large
AVREADIORATE	NUMBER	Average Read IO Rate
AVREADBANDWIDTH	NUMBER	Average Read Bandwidth
AVWRITEBANDWIDTH	NUMBER	Average Write Bandwidth
AVAVERAGETOTALSIZE	NUMBER	Average Total Size
AVAVERAGEREADSIZE	NUMBER	Average Read Size
AVWRITEIOS	NUMBER	Average Write IOs

Column Name	Type	Description
AVBYTESWRITTEN	NUMBER	Average Bytes Written
AVWRITEIORATE	NUMBER	Average Write IO Rate
AVREADIOSLARGE	NUMBER	Average Read IO Large
AVBYTESTRANSFERED	NUMBER	Average Bytes Transferred
AVPERCENTREAD	NUMBER	Average Percent Read
AVAVERAGEREADSIZELARGE	NUMBER	Average Read Size Large
AVAVERAGEWRITESIZE	NUMBER	Average Write Size
AVPERCENTREADHITS	NUMBER	Average Percent of Read Hits
AVTOTALIOS	NUMBER	Average Total IOs
AVBYTESREAD	NUMBER	Average Bytes Read
AVAVERAGEWRITESIZELARGE	NUMBER	Average Write Size Large
AVTOTALBANDWIDTH	NUMBER	Average Total Bandwidth
AVBYTESREADLARGE	NUMBER	Average Bytes Read Large
AVREADHITIOS	NUMBER	Average Read Hits IOs

MVCS_LSIDISKIOSTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTIONTIME	DATE	Timestamp of the Collection
STATSTYPE	VARCHAR2(256 CHAR)	Stats Type
AVREADIOS	NUMBER	Average Read IOs
READIOS	NUMBER	Read IOs
TOTALIORATE	NUMBER	Total IO Rate
UNRECOVEREDERRORS	NUMBER	Unrecovered Errors
RETRIEDREQUESTS	NUMBER	Retried Requests
BYTESTRANSFERED	NUMBER	Bytes Transferred
READIORATE	NUMBER	Read IO Rate

Column Name	Type	Description
TOTALBANDWIDTH	NUMBER	Total Bandwidth
TOTALIOS	NUMBER	Total IOs
TIMEOUTS	NUMBER	Timeouts
WRITEIORATE	NUMBER	Write IO Rate
AVERAGETOTALSIZE	NUMBER	Average Total Size
PERCENTREAD	NUMBER	Percent Read
WRITEIOS	NUMBER	Write IOs
RECOVEREDERRORS	NUMBER	Recovered Errors

MVCS_LSINOPERFSTORVOLSTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTIONTIME	DATE	Timestamp of the Collection
STATSTYPE	VARCHAR2(256 CHAR)	Stats Type
READIOS	NUMBER	Read IOs
TOTALIOS	NUMBER	Total IOs
WRITEIOS	NUMBER	Write IOs
BYTESTRANSFERED	NUMBER	Bytes Transferred
WRITEOPERATIONS	NUMBER	Write Operations
REQUESTSSERVICES	NUMBER	Request Services
READREQUESTS	NUMBER	Read Requests
AVERAGEBLOCKSREQUESTED	NUMBER	Average Blocks Requested

MVCS_LSISPAGDISKSTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTIONTIME	DATE	Timestamp of the Collection
STATSTYPE	VARCHAR2(256 CHAR)	Stats Type
ADDRECOVEREDERRORS	NUMBER	Recovered Errors
ADDTOTALBANDWIDTH	NUMBER	Total Bandwidth
ADDREADIORATE	NUMBER	Read IO Rate
ADDTOTALIOS	NUMBER	Total IOs
ADDTOTALIORATE	NUMBER	Total IO Rate
ADDWRITEIORATE	NUMBER	Write IO Rate
ADDUNRECOVEREDERRORS	NUMBER	Unrecovered Errors
ADDTIMEOUTS	NUMBER	Timeouts
ADDBYTESTRANSFERED	NUMBER	Bytes Transferred
ADDWRITEIOS	NUMBER	Write IOs
ADDREADIOS	NUMBER	Read IOs
ADDRETRIEDREQUESTS	NUMBER	Retried Requests
ADDAVERAGETOTALSIZE	NUMBER	Average Total Size
ADDPERCENTREAD	NUMBER	Percent Read

MVCS_LSISPAGVOLUMESTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTIONTIME	DATE	Timestamp of the Collection
STATSTYPE	VARCHAR2(256 CHAR)	Stats Type
AVREADIOS	NUMBER	Average Read IOs

Column Name	Type	Description
AVBYTESWRITTENLARGE	NUMBER	Average Bytes Written Large
AVTOTALIORATE	NUMBER	Average Total IO Rate
AVWRITEIOSLARGE	NUMBER	Average Write IOs Large
AVREADIORATE	NUMBER	Average Read IO Rate
AVREADBANDWIDTH	NUMBER	Average Read Bandwidth
AVWRITEBANDWIDTH	NUMBER	Average Write Bandwidth
AVAVERAGETOTALSIZE	NUMBER	Average Total Size
AVAVERAGEREADSIZE	NUMBER	Average Read Size
AVWRITEIOS	NUMBER	Average Write IOs
AVBYTESWRITTEN	NUMBER	Average Bytes Written
AVWRITEIORATE	NUMBER	Average Write IO Rate
AVREADIOSLARGE	NUMBER	Average Read IOs Large
AVBYTESTRANSFERED	NUMBER	Average Bytes Transferred
AVPERCENTREAD	NUMBER	Average Percent Read
AVAVERAGEREADSIZELARGE	NUMBER	Average Read Size Large
AVAVERAGEWRITESIZE	NUMBER	Average Write Size
AVPERCENTREADHITS	NUMBER	Average Percent of Read Hits
AVTOTALIOS	NUMBER	Average Total IOs
AVBYTESREAD	NUMBER	Average Bytes Read
AVAVERAGEWRITESIZELARGE	NUMBER	Average Write Size Large
AVTOTALBANDWIDTH	NUMBER	Average Total Bandwidth
AVBYTESREADLARGE	NUMBER	Average Bytes Read Large
AVREADHITSIOS	NUMBER	Average Read Hit IOs

MVCS_LSISSAGDISKSTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTIONTIME	DATE	Timestamp of the Collection
STATSTYPE	VARCHAR2(256 CHAR)	Stats Type
ADDRECOVEREDERRORS	NUMBER	Recovered Errors
ADDTOTALBANDWIDTH	NUMBER	Total Bandwidth
ADDREADIORATE	NUMBER	Read IO Rate
ADDTOTALIOS	NUMBER	Total IOs
ADDTOTALIORATE	NUMBER	Total IO Rate
ADDWRITEIORATE	NUMBER	Write IO Rate
ADDUNRECOVEREDERRORS	NUMBER	Unrecovered Errors
ADDTIMEOUTS	NUMBER	Timeouts
ADDBYTESTRANSFERED	NUMBER	Bytes Transferred
ADDWRITEIOS	NUMBER	Write IOs
ADDREADIOS	NUMBER	Read IOs
ADDRETRIEDREQUESTS	NUMBER	Retried Requests
ADDAVERAGETOTALSIZE	NUMBER	Average Total Size
ADDPERCENTREAD	NUMBER	Percentage Read

MVCS_LSISSAGVOLUMESTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTIONTIME	DATE	Timestamp of the Collection
STATSTYPE	VARCHAR2(256 CHAR)	Stats Type
AVREADIOS	NUMBER	Average Read IOs

Column Name	Type	Description
AVBYTESWRITTENLARGE	NUMBER	Average Bytes Written Large
AVTOTALIORATE	NUMBER	Average Total IO Rate
AVWRITEIOSLARGE	NUMBER	Average Write IOs Large
AVREADIORATE	NUMBER	Average Read IO Rate
AVREADBANDWIDTH	NUMBER	Average Read Bandwidth
AVWRITEBANDWIDTH	NUMBER	Average Write Bandwidth
AVAVERAGETOTALSIZE	NUMBER	Average Total Size
AVAVERAGEREADSIZE	NUMBER	Average Read Size
AVWRITEIOS	NUMBER	Average Write IOs
AVBYTESWRITEN	NUMBER	Average Bytes Written
AVWRITEIORATE	NUMBER	Average Write IO Rate
AVREADIOSLARGE	NUMBER	Average Read IOs Large
AVBYTESTRANSFERED	NUMBER	Average Bytes Transferred
AVPERCENTREAD	NUMBER	Average Percent Read
AVAVERAGEREADSIZELARGE	NUMBER	Average Read Size Large
AVAVERAGEWRITESIZE	NUMBER	Average Write Size
AVPERCENTREADHITS	NUMBER	Average Percent of Read Hits
AVTOTALIOS	NUMBER	Average Total IOs
AVBYTESREAD	NUMBER	Average Bytes Read
AVAVERAGEWRITESIZELARGE	NUMBER	Average Write Size Large
AVTOTALBANDWIDTH	NUMBER	Average Total Bandwidth
AVBYTESREADLARGE	NUMBER	Average Bytes Read Large
AVREADHITIOS	NUMBER	Average Read Hit IOs

MVCS_LSISTORAGEVOLUMESTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTIONTIME	DATE	Timestamp of the Collection
STATSTYPE	VARCHAR2(256 CHAR)	Stats Type
BYTESREAD	NUMBER	Bytes Read
READIOS	NUMBER	Read IOs
READBANDWIDTH	NUMBER	Read Bandwidth
BYTESWRITTENLARGE	NUMBER	Bytes Written Large
CACHEREADCHECKS	NUMBER	Cache Read Checks
READHITIOS	NUMBER	Read Hit IOs
TOTALBANDWIDTH	NUMBER	Total Bandwidth
TOTALIOS	NUMBER	Total IOs
AVERAGEWRITESIZE	NUMBER	Average Write Size
BYTESREADLARGE	NUMBER	Bytes Read Large
AVERAGEREADSIZE	NUMBER	Average Read Size
BYTESWRITTEN	NUMBER	Bytes Written
AVERAGEWRITESIZELARGE	NUMBER	Average Write Size Large
READIOSLARGE	NUMBER	Read IOs Large
WRITEIOS	NUMBER	Write IOs
TOTALIORATE	NUMBER	Total IO Rate
PERCENTREADHITS	NUMBER	Percent of Read Hits
WRITEIOSLARGE	NUMBER	Write IOs Large
BYTESTRANSFERED	NUMBER	Bytes Transferred
READIORATE	NUMBER	Read IO Rate
WRITEIORATE	NUMBER	Write IO Rate
AVERAGEREADSIZELARGE	NUMBER	Average Read Size Large

Column Name	Type	Description
AVERAGETOTALSIZE	NUMBER	Average Total Size
WRITEBANDWIDTH	NUMBER	Write Bandwidth
PERCENTREAD	NUMBER	Percent Read

MVCS_PORTSTATSVW

Column Name	Type	Description
PORTID	NUMBER(38)	ID of the Port
TIMESTAMP	DATE	Timestamp of the Collection
STATSTYPE	VARCHAR2(256 CHAR)	Stats Type
PORTTYPE	VARCHAR2(17 CHAR)	Type of the Port
BYTESRECEIVED	NUMBER	Bytes Received
BYTESTRANSMITTED	NUMBER	Bytes Transmitted
LINKFAILURES	NUMBER	Link Failures
CRCERRORS	NUMBER	CRC Errors

MVCS_XPCONTROLLERSTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTION TIME	DATE	Timestamp of the Collection
STATTYPE	VARCHAR2(256 CHAR)	Stats Type
DURATION	VARCHAR2(1024)	Duration
UTILIZATION0	NUMBER	Percentage Utilization for Controller's Process MP0
UTILIZATION1	NUMBER	Percentage Utilization for Controller's Process MP1
UTILIZATION2	NUMBER	Percentage Utilization for Controller's Process MP2

Column Name	Type	Description
UTILIZATION3	NUMBER	Percentage Utilization for Controller's Process MP3
UTILIZATION4	NUMBER	Percentage Utilization for Controller's Process MP4
UTILIZATION5	NUMBER	Percentage Utilization for Controller's Process MP5
UTILIZATION6	NUMBER	Percentage Utilization for Controller's Process MP6
UTILIZATION7	NUMBER	Percentage Utilization for Controller's Process MP7

MVCS_XPHOSTFCPORTSTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTION TIME	DATE	Timestamp of the Collection
STATTYPE	VARCHAR2(256 CHAR)	Stats Type
DURATION	NUMBER	Duration
TOTALDATARATE	NUMBER	Total Data Rate
TOTALIORATE	NUMBER	Total IO Rate

MVCS_XPSTORAGEPOOLSTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTION TIME	DATE	Timestamp of the Collection
STATTYPE	VARCHAR2(256 CHAR)	Stats Type
DURATION	NUMBER	Duration
AVGREADSIZE	NUMBER	Average Read Size
AVGWRITESIZE	NUMBER	Average Write Size
PCTREADHITS	NUMBER	Read Hits Percentage

Column Name	Type	Description
PCTREADHITSRANDOM	NUMBER	Random Read Hits Percentage
PCTREADHITSSEQ	NUMBER	Sequential Read Hits Percentage
PCTREADIOS	NUMBER	Read IOs Percentage
PCTWRITEIOS	NUMBER	Write IOs Percentage
READDATARATE	NUMBER	Read Data Rate
READDATARATERANDOM	NUMBER	Random Read Data Rate
READDATARATESEQ	NUMBER	Sequential Read Data Rate
READHITRATE	NUMBER	Read Hit Rate
READRATE	NUMBER	Read Rate
READRESPONSETIME	NUMBER	Read Response Time
TOTALDATARATE	NUMBER	Total Data Rate
TOTALIORATE	NUMBER	Total IO Rate
UTILIZATION	NUMBER	Utilization
WRITEDATARATE	NUMBER	Write Data Rate
WRITEDATARATERANDOM	NUMBER	Random Write Data Rate
WRITEDATARATESEQ	NUMBER	Sequential Write Data Rate
WRITERATE	NUMBER	Write Rate
WRITERESPONSETIME	NUMBER	Write Response Time

MVCS_XPVOLUMESTATSVW

Column Name	Type	Description
ID	NUMBER(38)	ID
COLLECTION TIME	DATE	Timestamp of the Collection
STATTYPE	VARCHAR2(256 CHAR)	Stats Type
DURATION	NUMBER	Duration

Column Name	Type	Description
AVGREADSIZE	NUMBER	Average Read Size
AVGWRITESIZE	NUMBER	Average Write Size
PCTREADHITS	NUMBER	Read Hits Percentage
PCTREADHITSRANDOM	NUMBER	Random Read Hits Percentage
PCTREADHITSSEQ	NUMBER	Sequential Read Hits Percentage
PCTREADIOS	NUMBER	Read IOs Percentage
PCTWRITEIOS	NUMBER	Write IOs Percentage
READDATARATE	NUMBER	Read Data Rate
READDATARATERANDOM	NUMBER	Random Read Data Rate
READDATARATESEQ	NUMBER	Sequential Read Data Rate
READHITRATE	NUMBER	Read Hit Rate
READRATE	NUMBER	Read Rate
READRESPONSETIME	NUMBER	Read Response Time
TOTALDATARATE	NUMBER	Total Data Rate
TOTALIORATE	NUMBER	Total IO Rate
UTILIZATION	NUMBER	Utilization
WRITEDATARATE	NUMBER	Write Data Rate
WRITEDATARATERANDOM	NUMBER	Random Write Data Rate
WRITEDATARATESEQ	NUMBER	Sequential Write Data Rate
WRITERATE	NUMBER	Write Rate
WRITERESPONSETIME	NUMBER	Write Response Time

MVIEW_STATUS

Name	Type
MVIEWNAME	NOT NULL VARCHAR2(30)
LAST_REFRESH_TIME	DATE

Name	Type
TOTALREFRESHTIME	
STATUS	VARCHAR2(10)

MVIEWCORE_STATUS

Name	Type
MVIEWNAME	NOT NULL VARCHAR2(30)
LAST_REFRESH_TIME	DATE
TOTALREFRESHTIME	VARCHAR2(32)
STATUS	VARCHAR2(10)

MVCS_CENTERADEVICEUSAGEVW

Name	Type	Description
ID	NUMBER(38)	Centera device identifier
CAPACITYTYPE	VARCHAR2(256 CHAR)	Type of the capacity collected
TIMESTAMP	DATE	Timestamp of the Centera Device statistics collection
TOTALCAPACITY	NUMBER	Total device capacity
SPARECAPACITY	NUMBER	The capacity that is available on nodes that do not have the storage role assigned.
USEDCAPACITY	NUMBER	The capacity that is used or otherwise not available to store data.
FREECAPACITY	NUMBER	The amount of capacity available to store data.
SYSTEMRESOURCE	NUMBER	The capacity that is used by the Centra Star software and is never available for storing data.
OFFLINE CAPACITY	NUMBER	The capacity that is temporarily unavailable.

Name	Type	Description
AUDITMETADATA	NUMBER	This includes indexes, databases, and internal queues.
PROTECTEDUSERDATA	NUMBER	The capacity taken by user data, including CDFs, reflections, and protected copies of user files.
USEDOBJECTCOUNT	NUMBER(38)	The total number of stored objects
TOTALOBJECTCOUNT	NUMBER(38)	The maximum number of objects that can be stored.
SUPPORTEDOBJECTCOUT	NUMBER(38)	The object count supported on the node.
SYSTEMBUFFER	NUMBER	Allocated space that allows internal databases and indexes to safely grow and failover.
PERCENTFREE	NUMBER	Percentage of device free capacity
DEVICETIME	TIMESTAMP(6)	Time of th Device
DURATION	NUMBER(38)	Duration
CAPACITYNUM	NUMBER	Number given to the capacity type

MVCS_CENTERANODEUSAGEVW

Name	Type	Description
ID	NUMBER(38)	Centera Node Identifier
CAPACITYTYPE	VARCHAR2(256 CHAR)	Type of the capacity collected
TIMESTAMP	DATE	Timestamp of the Centera Node statistics collection
TOTALCAPACITY	NUMBER	Total device capacity
SPARECAPACITY	NUMBER	The capacity that is available on nodes that do no have the storage role assigned.
USEDCAPACITY	NUMBER	The capacity that is used or otherwise not available to store data.

Name	Type	Description
FREECAPACITY	NUMBER	The amount of capacity available to store data.
SYSTEMRESOURCE	NUMBER	The capacity that is used by the Centra Star software and is never available for storing data.
OFFLINECAPACITY	NUMBER	The capacity that is temporarily unavailable.
AUDITMETADATA	NUMBER	This includes indexes, databases, and internal queues.
PROTECTEDUSERDATA	NUMBER	The capacity taken by user data, including CDF's reflections, and protected copies of user files.
USEDOBJECTCOUNT	NUMBER(38)	The total number of stored objects.
TOTALOBJECTCOUNT	NUMBER(38)	The maximum number of objects that can be stored.
SUPPORTEDOBJECTCOUNT	NUMBER(38)	The object count supported on the node.
SYSTEMBUFFER	NUMBER	Allocated space that allows internal databases and indexes to safely grow and failover
PERCENTFREE	NUMBER	Percentage of device free capacity
DEVICETIME	TIMESTAMP(6)	Time of the Device
DURATION	NUMBER(38)	Duration
CAPACITYNUM	NUMBER	Number given to the capacity type

MVCS_CENTERAPOOLUSAGEVW

Name	Type	Description
ID	NUMBER(38)	Centera pool Identifier
CAPACITYTYPE	VARCHAR2(256 CHAR)	Type of the capacity collected
TIMESTAMP	DATE	Timestamp of the Centera Pool statistics collection

Name	Type	Description
USEDCAPACITY	NUMBER	The capacity that is used or otherwise not available to store data.
FREECAPACITY	NUMBER	The amount of capacity available to store data.
USEDOBJECTCOUNT	NUMBER(38)	The total number of stored objects.
USEDCLIPCOUNT	NUMBER(38)	Number of C-Clips stored in the pool
PERCENTFREE	NUMBER	Percentage of Pool free capacity
QUOTA	NUMBER	Total capacity of the pool
DURATION	NUMBER(38)	Duration
CAPACITYNUM	NUMBER	Number given to the capacity type

MVCS_CENTERADISKUSAGEVW

Name	Type	Description
ID	NUMBER(38)	Centera disk Identifier
CAPACITYTYPE	VARCHAR2(256 CHAR)	Type of the capacity collected
TIMESTAMP	DATE	Timestamp of the Centera Disk statistics collection
TOTALCAPACITY	NUMBER	Total disk capacity
USEDCAPACITY	NUMBER	The capacity that is used or otherwise not available to store data.
FREECAPACITY	NUMBER	The amount of capacity available to store data.
SYSTEMRESOURCE	NUMBER	The capacity that is used by the Centra Star software and is never available for storing data.
AUDITMETADATA	NUMBER	This includes indexes, databases, and internal queues.
PROTECTEDUSERDATA	NUMBER	The capacity taken by user data, including CDF's, reflections, and protected copies of user files.
USEDOBJECTCOUNT	NUMBER(38)	The total number of stored objects.

Name	Type	Description
SYSTEMBUFFER		Allocated space that allows internal databases and indexes to safely grow and failover.
DURATION	NUMBER(38)	Duration
CAPACITYNUM	NUMBER	Number given to the capacity type

MVCS_CENTERAPOOLPRIVILEGE

Name	Type	Description
ID	NUMBER(38)	Centera privilege data identifier
PRIVILEGEID	NUMBER(38)	Identifies the privilege in Centera
POOLID	NUMBER(38)	Identifies the virtual pool in Centera

MVC_CELERRAVOLUMECPOINTMAP

Name	Type	Description
ID	NUMBER(38)	Celerra device identifier
STORAGEVOLUMEID	NUMBER(38)	Celerra volume identifier
CHECKPOINTID	NUMBER(38)	Celerra checkpoint volume identifier

Views from Previous Releases

The materialized views in this release were renamed and revised. The following views were dropped:

- MV_STORAGESYSTEMCAPSUMMARYVW
- MV_HOSTDETAILVW
- MV_UNITACCESSVW

The following are the views from earlier releases and the corresponding new views. Use the new views for any new report development.

Legacy View	Alternate Core Views
MV_ STORAGESYSTEMPORTUTILVW	MVC_STORAGESYSTEMCONFIGVW MVC_STORAGESYSTEMSUMMARYVW

Legacy View	Alternate Core Views
MV_SSFRONTENDVW	MVC_LUNMAPPINGVW MVC_PORTSUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_STORAGEVOLUMESUMMARYVW MVC_PORTSUMMARYVW MVC_STORGAEPOOLSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW
MV_SSBACKENDDETAILVW	MVC_CARDSUMMARYVW MVC_DISKEXTENTSUMMARYVW MVC_DISK_EXTENTVW MVC_DISKDRIVESUMMARYVW MVC_STORAGESYSTEMSUMMARYVW
MV_SSLOGICALDETAILVW	MVC_STORGAEPOOLSUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_VOLUMEDISKDRIVEVW MVC_DISKEXTENTSUMMARYVW
MV_SSAVAILABLEVOLUMEVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_STORAGEPROCESSORSUMMARYVW
MV_ TEMPMAPPEDVOLSUMMARYVW	MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_PROTOCOLCONTROLLERVW MVC_LUNMAPPINGVW
MV_LUNSPERFASUMMARYVW	MVC_STORAGESYSTEMCONFIGVW MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_PORTSUMMARYVW MVC_PORTCONTROLLERMAPVW MVC_PROTOCOLCONTROLLERVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_LUNMAPPINGVW MVC_CARDSUMMARYVW, MV_HOSTSUM

Legacy View	Alternate Core Views
MV_FABRICADAPCAPLUNVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_PORTSUMMARYVW MVC_PORTCONTROLLERMAPVW MVC_PROTOCOLCONTROLLERVW MVC_STORAGEVOLUMEPORTS MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW
MV_ TEMPSTORSYSTEMSUMMARYVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGESYSTEMCONFIGVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGEPOOLCONFIGVW MVC_STORGAEPOOLSUMMARYVW
MV_ TEMPSTORPOOLSUMMARYVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGEPOOLCONFIGVW MVC_STORGAEPOOLSUMMARYVW
MV_TEMPFRONTENDVW	MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGEPOOLCONFIGVW MVC_POOLSUMMARYVW
MV_LUNMAPPINGVW	MVC_LUNMAPPINGVW
MV_HOSTSUMMARYVW	MVC_HOSTSUMMARYVW MVC_HBASUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW
MV_TEMPHOSTLOGICALVW	MVC_HOSTDISKDRIVEVW MVC_SUBPATHVW MVC_PATHVW MVC_HOSTVOLUMESUMMARYVW MVC_HOSTSUMMARYVW MVC_HBASUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW
MV_TEMPHOSTCARDVW	MVC_PORTSUMMARYVW MVC_HBAPORTTARGETS MVC_PORTSUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW

Legacy View	Alternate Core Views
MV_HOSTSTORAGESUMMARYVW	MVC_HOSTSUMMARYVW MVC_HOSTVOLUMESUMMARYVW MVC_HOSTCAPACITYVW MVC_OPTIONALTABLEVW
MV_HOSTSTORAGEBYOSVW	MVC_HOSTSUMMARYVW MVC_HOSTCAPACITYVW
MV_ HOSTSTORAGEALLOCATIONVW	MVC_HOSTSUMMARYVW MVC_HOSTCAPACITYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGEVOLUMEPORTS MVC_STORGAEPOOLSUMMARYVW MVC_APPLICATIONSUMMARYVW MVC_PORTSUMMARYVW MVC_SWITCHSUMMARYVW
MV_HOSTCONNECTIVITYVW	MVC_PATHVW MVC_HOSTSUMMARYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW
MV_HOSTVXVMVW	MVC_DISKEXTENTSUMMARYVW MVC_HOSTSUMMARYVW MVC_DISK_EXTENTVW MVC_DISKDRIVESUMMARYVW MVC_OPTIONALTABLEVW, MVC_PATHVW MVC_SUBPATHVW MVC_HOSTDISKDRIVEVW MVC_HOSTVOLUMESUMMARYVW
MV_HOSTVMVW	MVC_HOSTDISKDRIVEVW MVC_PATHVW MVC_HOSTVOLUMESUMMARYVW MVC_HOSTSUMMARYVW MVC_HOSTDISKDRIVEVW
MV_HOSTLOGICALVOLUMEVW	MVC_HOSTCAPACITYVW

Legacy View	Alternate Core Views
MV_HOSTFSVOLUMEVW	MVX_HOSTSUMMARYVW MVC_HOSTCAPACITYV MVC_PATHVW MVC_SUBPATHVW MVC_HOSTVOLUMESUMMARYVW MVC_STORGAEPOOLISUMMARYVW MVC_STORAGEVOLUMESUMMARYVW
MV_HOSTRAWVOLUMEVW	MVC_PATHVW MVC_HOSTSUMMARYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW MVC_DISKEXTENTSUMMARYVW
MV_HOSTUNUSEDVOLUMEVW	MVC_PATHVW MVC_HOSTSUMMARYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORGAEPOOLSUMMARYVW
MV_HOSTDISKDRIVEVW	MVC_HOSTSUMMARYVW MVC_DISKDRIVESUMMARYVW MVC_DISK_EXTENTVW MVC_DISKEXTENTSUMMARYVW
MV_HOSTSSDEPENDECYVW	MVC_HOSTSUMMARYVW MVC_SUBPATHVW MVC_STORAGEVOLUMESUMMARYVW MVC_STORAGESYSTEMSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_PORTSUMMARYVW MVC_PATHVW MVC_HOSTVOLUMESUMMARYVW
MV_HOSTAPPDEPENDECYVW	MVC_APPLICATIONSUMMARYVW MVC_HOSTVOLUMESUMMARYVW MVC_STORAGEVOLUMESUMMARYVW MVC_SUBPATHVW MVC_PATHVW MVC_PORTSUMMARYVW

Legacy View	Alternate Core Views
MV_HOSTSPERFAVW	MVC_HOSTSUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW MVC_LUNMAPPINGVW MVC_STORAGEVOLUMESUMMARYVW MVC_PROTOCOLCONTROLLERVW MVC_PORTCONTROLLERMAPVW MVC_PORTSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW
MV_HOSTDISKPARTITIONVW	MVC_SUBPATHVW MVC_VOLUMEDISKDRIVEVW MVC_HOSTDISKDRIVEVW MVC_HOSTVOLUMESUMMARYVW
MV_ TEMPSWITCHBYTESINTERVALVW	Switch Port statistics such as Bytes sent and Bytes received. This Materialized View does not have any alternative Core Views.
MV_ TEMPSWITCHCONNECTEDVW	MVC_PORTSUMMARYVW MVC_SWITCHSUMMARYVW
MV_TEMPCONNECTEDHOSTVW	MVC_PORTSUMMARYVW MVC_CARDSUMMARYVW MVC_HOSTSUMMARYVW MVC_SWITCHSUMMARYVW
MV_ TEMPCONNECTEDSTORAGEVW	MVC_PORTSUMMARYVW MVC_STORAGEPROCESSORSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW MVC_SWITCHSUMMARYVW MVC_HOSTSUMMARYVW
MV_TEMPZONEVW	MVC_ZONEVW MVC_ZONESUMMARY
MV_SWITCHDETAILVW	MVC_SWITCHSUMMARYVW MVC_ZONESUMMARY MVC_ZONEVW MVC_PORTSUMMARYVW MVC_CARDSUMMARYVW MVC_HOSTSUMMARYVW
MV_AVAILABLEPORTVW	MVC_SWITCHSUMMARYVW
MV_TOTALPORTSVW	MVC_PORTSUMMARYVW

Legacy View	Alternate Core Views
MV_SANZONEPORTWWNVW	MVC_PORTSUMMARYVW MVC_SWITCHSUMMARYVW MVC_CARDSUMMARYVW MVC_PORTSUMMARYVW MVC_ZONEVW
MV_SANCOMNOTLOGINVW	MVC_PORTSUMMARYVW MVC_ZONESUMMARY MVC_ZONEVW MVC_SWITCHSUMMARYVW MVC_CARDSUMMARYVW
MV_ZONEDETAILSVW	MVC_SWITCHSUMMARYVW MVC_ZONESUMMARY MVC_ZONEVW
MV_EVENTVW	MVC_EVENTVW
MV_ORGANIZATIONVW	MVC_ORGANIZATIONVW
MV_ORGRELATIONVW	MVC_ORGRELATIONVW
MV_APPLICATIONVW	MVC_APPLICATIONSUMMARYVW
MV_DBAPPCHARGEBACKVW	MVC_APPLICATIONSUMMARYVW MVCA_DBAPPINSTCAPACITYVW MVCA_DBAPPPHYCAPACITYVW MVCA_EXCHAPPCAPACITYVW MVCA_VIRTUALAPPCAPACITYVW
MV_APPDEPENDENCYVW	MVC_APPLICATIONSUMMARYVW MVC_HOSTVOLUMESUMMARYVW MVC_SUBPATHVW MVC_PATHVW MVC_PORTSUMMARYVW
MV_ASSETSUMMARYVW	MVC_ASSETSUMMARYVW MVC_OPTIONALTABLEVW
MV_ASSETCOUNTVW	MVC_APPLICATIONSUMMARYVW MVC_HOSTSUMMARYVW MVC_STORAGESYSTEMSUMMARYVW MVC_SWITCHSUMMARYVW MVC_TAPELIBRARYSUMMARYVW
MV_ FILESERVERVOLUMEDETAILVW	MV_TEMPFILESERVERVOLUMEVW MVC_APPLICATIONSUMMARYVW

Legacy View	Alternate Core Views
MV_FILESERVERVW	MVC_APPLICATIONSUMMARYVW MVC_HOSTSUMMARYVW MV_TEMPFILESERVERVOLUMEVW MVCA_FSRM_USERSUMMARYVW
MV_ TEMPFILESERVERVOLUMEVW	MVC_HOSTVOLUMESUMMARYVW MVC_HOSTCAPACITYVW MVCA_FSRM_USERDETAILVW MVCA_FSRM_AGESUMMARYVW