IBM® Storage Networking

IBM Network Advisor SAN User Manual

Supporting IBM Network Advisor version 14.2.1

NOTE

This product contains software that is licensed under written license agreements. Your use of such software is subject to the license agreements under which they are provided.

SC27-6619-03

IBM® Storage Networking

IBM Network Advisor SAN User Manual

Supporting IBM Network Advisor version 14.2.1

SC27-6619-03

SC27-6619-03

Contents

User Account Management

Fabric Insight Portal

About This Document

Supported hardware and software

In those instances in which procedures or parts of procedures documented here apply to some devices but not to others, this guide identifies exactly which devices are supported and which are not.

Although many different software and hardware configurations are tested and supported by IBM Network Advisor 14.2.1, documenting all possible configurations and scenarios is beyond the scope of this document.

Fabric OS hardware and software support

The following firmware platforms are supported by this release of IBM Network Advisor 14.2.1:

- Fabric OS 7.0 or later
- Fabric OS 8.0 or later
- Fabric OS 8.1 or later

NOTE
For platform specific Fabric OS requirements, refer to the Firmware level required column in Table 1.

NOTE
Discovery of a Secure Fabric OS fabric in strict mode is not supported.

Table 1 provides a list of the hardware platforms supported by this release of IBM Network Advisor 14.2.1 as well as any platform specific Fabric OS requirements.

TABLE 1 Supported Hardware

IBM Name	Terminology used in documentation	Firmware level required
SAN24B-4	24-port, 8 Gbps FC Switch	Fabric OS v7.0.0 or later
SAN40B-4	40-port, 8 Gbps FC Switch	Fabric OS v7.0.0 or later
SAN80B-4	80-port, 8 Gbps FC Switch	Fabric OS v7.0.0 or later
SAN24B-5	24-port, 16 Gbps Edge switch	Fabric OS v7.0.1 or later
SAN48B-5	48-port, 16 Gbps switch	Fabric OS v7.0.0 or later
SAN96B-5	96-port, 16 Gbps switch	Fabric OS v7.1.0 or later
IBM Flex System FC5022 16Gb SAN Scalable Switches (ScSM)	48-port, 16 Gbps embedded switch	Fabric OS v7.2.0 or later
SAN04B-R	4 Gbps Extension Switch	Fabric OS v7.0.0 or later

TABLE 1 Supported Hardware

IBM Name	Terminology used in documentation	Firmware level required
SAN06B-R	8 Gbps Extension Switch	Fabric OS v7.0.0 or later
SAN42B-R	16 Gbps 24-FC port, 18 GbE port Switch	Fabric OS v7.3.0 or later
IBM Converged Switch B32	8 Gbps 8-FC-port, 10 GbE 24-CEE port Switch	Fabric OS v6.1.2_CEE
SAN32B-E4 Encryption Switch	8 Gbps Encryption Switch	Fabric OS v6.1.1_enc or later
SAN24B-6	**24-port, 32 Gbps switch**	**Fabric OS v8.1.0 or later**
SAN64B-6	64-port, 32 Gbps switch	Fabric OS v8.0.0 or later
SAN768B-2[1, 2]	16 Gbps 8-slot Backbone Chassis	Fabric OS v7.0.0 or later
SAN768B-2[1, 2] with FC8-64 and FX8-24 Blades	16 Gbps 8-slot Backbone Chassis with 8 Gbps 64-port and 8 Gbps 12-FC port, 10 GbE ports, 2-10 GbE ports Extension blades	Fabric OS v7.0.0 or later
SAN768B-2[1, 2] with FC16-32 and FC16-48 Blades	16 Gbps 8-slot Backbone Chassis with 16 Gbps 32-port and 16 Gbps 48-port blades	Fabric OS v7.0.0 or later
SAN768B-2[1, 2] with FCoE 10-24 Blades	16 Gbps 8-slot Backbone Chassis with 8 Gbps 24-port FCoE Blade	Fabric OS v7.0.0 or later
SAN768B-2[1, 2] with FC16-64 Blade	16 Gbps 8-slot Backbone Chassis with 16 Gbps 64-FC port blade	Fabric OS v7.3.0 or later
SAN256B-6[1, 2]	32 Gbps, 4-slot Backbone Chassis	Fabric OS v8.0.1 or later
SAN512B-6[1, 2]	32 Gbps, 8-slot Backbone Chassis	Fabric OS v8.0.1 or later
FC8-16 Blade	FC 8 GB 16-port Blade	Fabric OS v7.0.0 or later
FC8-32 Blade	FC 8 GB 32-port Blade	Fabric OS v7.0.0 or later
FC8-32E Blade[1]	FC 8 GB 32-port Blade	Fabric OS v7.0.1 or later
FC8-48 Blade	FC 8 GB 48-port Blade	Fabric OS v7.0.0 or later
FC8-48E Blade[1]	FC 8 GB 48-port Blade	Fabric OS v7.0.1 or later
FC8-64 Blade	FC 8 GB 64-port Blade	Fabric OS v7.0.0 or later
FC10-6 Blade	FC 10 - 6 ISL Blade	Fabric OS v7.0.0 or later
FC16-32 Blade[1]	16 Gbps 32-port blade	Fabric OS v7.0.0 or later
FC16-48 Blade[1]	16 Gbps 48-port blade	Fabric OS v7.0.0 or later
FC16-64 Blade[1]	16 Gbps 64-FC port blade	Fabric OS v7.3.0 or later
FCoE10-24 Blade[3]	10 Gig FCoE Port Router Blade	Fabric OS v7.0.0 or later
FR4-18i Extension Blade	4 Gbps Router, Extension Blade	Fabric OS v7.0.0 or later
FR8-24 Extension Blade	8 Gbps Router, Extension Blade	Fabric OS 7.0.0 or later
FS8-18 Encryption Blade	Encryption Blade	Fabric OS v6.1.1_enc or later
FX8-24 Extension Blade [1, 2]	8 Gbps Extension Blade	Fabric OS v6.3.1_CEE
FC32-48 Port Blade[1, 2]	32 Gbps 48-port blade	Fabric OS v8.0.1 or later
SX6 Extension Blade[1, 2]	32 Gbps, Router Extension blade	Fabric OS v8.0.1 or later

1. Only supported on the SAN384B-2 and SAN768B-2 chassis.

What's new in this document

The following table describes information added to this guide for Network Advisor release 14.2.1.

TABLE 2 Summary of enhancements

Feature	Description	Location
FC troubleshooting	• Added electrical and optical loop back test support for 32G QSFPS.	Refer to "Port diagnostics requirements".
MAPS	• Added Rule on Rule (RoR) support. • Removed Uninstall vTap support in 14.2.1.	• Refer to "Configuring a MAPS policy" and "User-defined policies".
Call Home	• Updated "Sent Test" button support.	• Refer to "Editing the EMC Call Home center".
COMPASS	• Added note on AAA server support.	• Refer to "Configuration blocks".
FCIP	• Added note to update the allowed circuits, information message.	• Refer to "Adding an FCIP circuit".
Fabric Insight Portal	• Added new "Collections" feature. • Added Filter Management conventions, Filter and LUN range, and Collections filter type support. • Added Ignore Default Rules and updated New Custom Rule Sets dialog box.	• Refer to "Collections" section. • Refer to "Filter Management" section. • Refer to "Collections Management" section.

Document conventions

This section describes text formatting conventions and important notice formats used in this document.

Text formatting

The narrative-text formatting conventions that are used are as follows:

bold text
- Identifies command names
- Identifies the names of user-manipulated GUI elements
- Identifies keywords and operands
- Identifies text to enter at the GUI or CLI

italic text
- Provides emphasis
- Identifies variables
- Identifies paths and Internet addresses
- Identifies document titles

`code` text
- Identifies CLI output
- Identifies command syntax examples

For readability, command names in the narrative portions of this guide are presented in mixed lettercase: for example, switchShow. In actual examples, command lettercase is often all lowercase. Otherwise, this manual specifically notes those cases in which a command is case sensitive.

Notes, cautions, and warnings

The following notices and statements are used in this manual. They are listed below in order of increasing severity of potential hazards.

NOTE
A note provides a tip, guidance or advice, emphasizes important information, or provides a reference to related information.

ATTENTION
An Attention statement indicates potential damage to hardware or data.

Key terms

For definitions of SAN-specific terms, visit the Storage Networking Industry Association online dictionary at:

http://www.snia.org/education/dictionary

Additional information

This section lists additional IBM-specific documentation that you might find helpful.

For support information for this product and other SAN products, see the following Web site: *www.ibm.com/supportportal/*

Visit *www.ibm.com/contact/* for the contact information for your country or region. You can also contact IBM within the United States at 1-800-IBMSERV (1-800-426-7378). For support outside the United States, you can find the service number at: *www.ibm.com/planetwide/*.

Getting technical help

Contact IBM support for hardware, firmware, and software support, including product repairs and part ordering. To expedite your call, have the following information available:

1. General Information
   - Switch model
   - Switch operating system version
   - Error numbers and messages received
   - **supportSave** command output
   - Detailed description of the problem, including the switch or fabric behavior immediately following the problem, and specific questions
   - Description of any troubleshooting steps already performed and the results
   - Serial console and Telnet session logs
   - syslog message logs
2. Switch Serial Number

   The switch serial number and corresponding bar code are provided on the serial number label, as illustrated below.

The serial number label is located as follows:

- SAN24B-4, SAN24B-5, SAN24B-6, SAN42B-R, SAN64B-6, SAN40B-4, SAN80B-4, SAN96B-5, SAN06B-R, and IBM Converged Switch B32—On the switch ID pull-out tab located inside the chassis on the port side on the left
- SAN48B-5—On the pull-out tab on the front of the switch
- SAN256B—Inside the chassis next to the power supply bays
- SAN768B and SAN768B-2—On the bottom right on the port side of the chassis
- SAN384B and SAN384B-2—On the bottom left on the port side of the chassis
- SAN256B-6 and SAN512B-6—On the upper portion of the chassis to the left of the fan assemblies

3. World Wide Name (WWN)

   Use the **wwn** command to display the switch WWN.

   You can also obtain the WWN from the same place as the serial number. For the SAN768B, SAN384B, SAN768B-2, SAN256B-6, and SAN512B-6, access the numbers on the WWN cards by removing the WWN bezel at the top of the nonport side of the chassis.

How to send your comments

Your feedback is important in helping us provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, send us your comments by e-mail to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title (paste into the e-mail subject line)
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

Getting Started

User interface components

NOTE
The Management application does not support I18N internationalization and localization.

The Management application provides easy, centralized management of the network, as well as quick access to all product configuration applications. Using this application, you can configure, manage, and monitor your networks with ease.

The Management application's main window contains a number of areas. The following figures illustrate the various areas, and descriptions of the areas follow the figures.

NOTE
Some widgets may be hidden. To display a widget to the **Dashboard** tab, click the Customize Dashboard icon ("Customizing the dashboard widgets and monitors" on page 207).

FIGURE 1 Main window

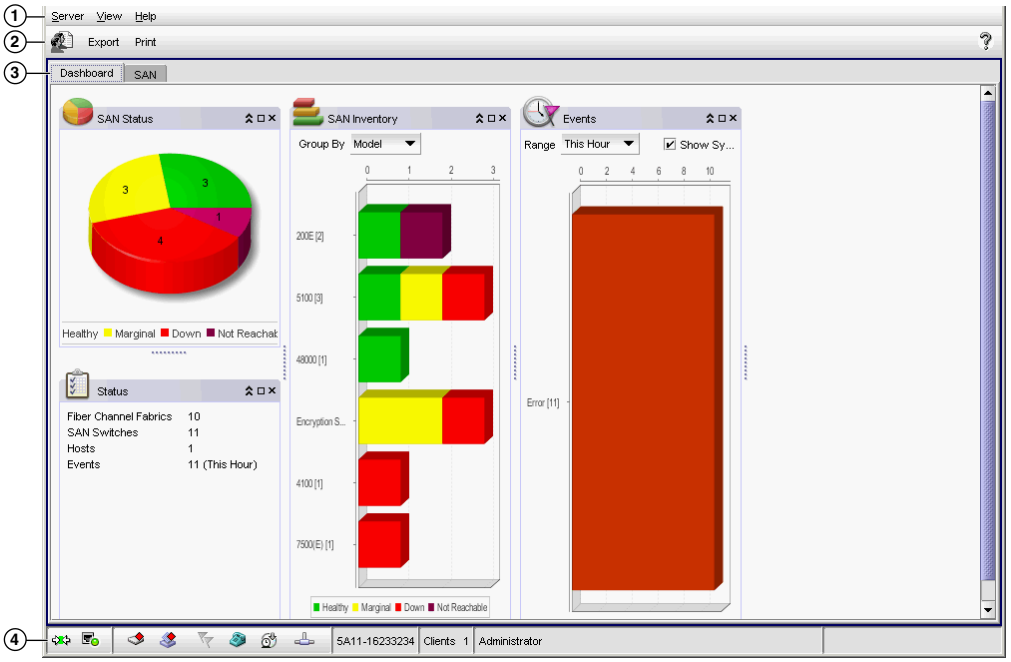

1. **Menu bar —** Lists commands you can perform on the Management application. The available commands vary depending on which tab (SAN or Dashboard) you select. For a list of available commands, refer to "Application Menus".
2. **Toolbar —** Provides buttons that enable quick access to dialog boxes and functions. The available buttons vary depending on which tab (SAN or Dashboard) you select. For a list of available commands, refer to "SAN main toolbar" on page 294, , or "Dashboard toolbar" on page 201.
3. **Tabs —** Provides quick access to the following views:
   - **Dashboard tab —** Provides a high-level overview of the network managed by Management application server. For more information, refer to "Dashboard Management" on page 199.
   - **SAN tab —** Displays the Master Log, Minimap, Connectivity Map (topology), and Product List. For more information, refer to the "SAN tab overview" on page 293.
4. **Master Log** — Displays the Master Log.
5. **Status bar —** Displays the connection, port, product, fabric, special event, Call Home, and backup status, as well as Server and User data.

Management server and client

The Management application has two parts: the server and the client. The server is installed on one machine and stores device-related information; it does not have a user interface. To view information through a user interface, you must log in to the server through a client. The server and clients may reside on the same machine, or on separate machines. If you are running Professional, the server and the client must be on the same machine.

Logging in to a server from the server machine

You must log in to a server to monitor your network.

NOTE
You must have an established user account on the server to log in.

To log in to a server, complete the following steps.

1. From the server machine, double-click the desktop icon or open the application from the **Start** menu.

   The **Log In** dialog box displays (Figure 2).

FIGURE 2 Log In dialog box

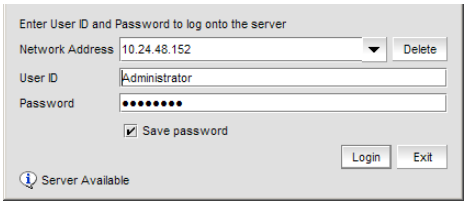

2.Remove a server from the **Network Address** list by selecting the IP address and clicking **Delete**.

3. Choose one of the following options:
   - If you configured authentication to CAC, enter your PIN in the CAC PIN field.
   - If you configured authentication to the local database, an external server (RADIUS, LDAP, or TACACS+), or a switch, complete the following steps.

a. Enter your user name and password.
The defaults are **Administrator** and **password**, respectively.

NOTE
Do not enter *Domain\User_Name* in the **User ID** field for LDAP server authentication.

b. Select or clear the **Save password** check box to choose whether you want the application to remember your password the next time you log in.
To change your password, refer to "Changing your password" on page 154.

4. Click **Login**.
5. Click **OK** on the **Login Banner** dialog box.

The Management application displays.

NOTE
When you launch the Management application or navigate to a new view, the **SAN** tab displays with a gray screen over the Product List and Topology Map while data is loading.

Launching a remote client

The remote client link in the **Start** menu does not automatically upgrade when you upgrade the Management application. You must clear the previous version from the Java cache. To clear the previous version, refer to "Clearing previous versions of the remote client" on page 4.

The remote client requires Oracle JRE. For the current supported JRE version for the Management application, refer to the Release Notes. For the website listing patch information, go to *http://www.oracle.com/technetwork/java/javase/downloads/index.html*.

NOTE
For higher performance, use a 64-bit JRE.

NOTE
If you are managing more than 9000 SAN ports, the client is not supported on 32-bit systems.

To launch a remote client, complete the following steps.

1. Choose one of the following options:
   - Open a web browser and enter the IP address of the Management application server in the **Address** bar.

     If the web server port number does not use the default (443 if is SSL Enabled; otherwise, the default is 80), you must enter the web server port number in addition to the IP address. For example, *IP_Address*:*Port_Number*.

     If this is the first time you are accessing this version of the Management application, this creates a start menu shortcut automatically in the Management application program directory.

     For Linux systems, remote client shortcuts are not created.
   - Select *Management_Application* (*Server_IP_Address*) in the Management application program directory from the start menu.

   The web client login page displays.
2. Click **Desktop Client**.

   The Management application web start page displays.

3. Click the **Web Start the Client** link.

   The **Log In** dialog box displays.

4. Log in to another server by entering the IP address to the other server in the **Network Address** field.

   NOTE
   The server must be the exact same version, edition, starting port number, and network size as the client.

5. Remove a server from the **Network Address** list by selected the IP address and clicking **Delete**.

6. Choose one of the following options:

   - If you configured authentication to CAC, enter your PIN in the CAC PIN field.
   - If you configured authentication to the local database, an external server (RADIUS, LDAP, or TACACS+), or a switch, complete the following steps.

     a. Enter your user name and password.
     The defaults are **Administrator** and **password**, respectively.

     NOTE
     Do not enter *Domain\User_Name* in the **User ID** field for LDAP server authentication.

     b. Select or clear the **Save password** check box to choose whether you want the application to remember your password the next time you log in.
     To change your password, refer to "Changing your password" on page 154.

7. Click **Login**.

8. Click **OK** on the **Login Banner** dialog box.

   The Management application displays.

   NOTE
   When you launch the Management application or navigate to a new view, the **SAN** tab displays with a gray screen over the Product List and Topology Map while data is loading.

Clearing previous versions of the remote client

The remote client link in the **Start** menu does not automatically upgrade when you upgrade the Management application. You must clear the previous version from the Java cache.

To clear the Java cache, complete the following steps.

1. Select **Start > Settings > Control Panel > Java**.

   The **Java Control Panel** dialog box displays.

2. Click **View** on the **General** tab.

   The **Java Cache Viewer** dialog box displays.

3. Right-click the application and select **Delete**.

4. Click **Close** on the **Java Cache Viewer** dialog box.

5. Click **OK** on the **Java Control Panel** dialog box.

   To create a remote client link in the **Start** menu, refer to "Launching a remote client" on page 3.

Logging in to the web client

You must log in to a Management application server to monitor the network. To launch a web client, complete the following steps.

1. Choose one of the following options:

   - Open a web browser and enter the IP address of the Management application server in the **Address** bar.

     If the web server port number does not use the default (443 if is SSL Enabled; otherwise, the default is 80), you must enter the web server port number in addition to the IP address. For example, *IP_Address*:*Port_Number*.

     If this is the first time you are accessing this version of the Management application, this creates a start menu shortcut automatically in the Management application program directory.

     For Linux systems, remote client shortcuts are not created.

   - Select *Management_Application* (*Server_IP_Address*) in the Management application program directory from the start menu.

   The web client login page displays.

   The web client login page displays with the server name and IP address in the upper left. You can launch the Java client from any page of the web client by clicking **Desktop Client**. You can download the client bundle (64-bit OS only), JRE, or MIB files by clicking **JRE and MIB files**.

   FIGURE 3 Management application web client log in page

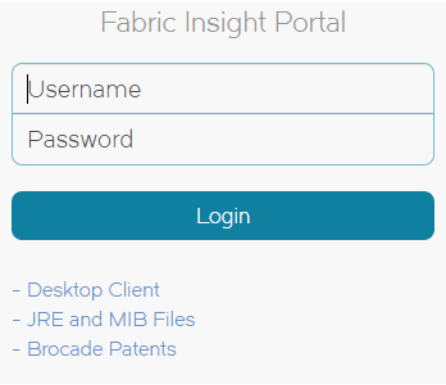

2. Enter your user name and password.

   NOTE
   Do not enter *Domain**User_Name* in the **User ID** field for LDAP server authentication.

3. Press **Enter** or click the login arrow icon.
4. Click **OK** on the **Login Banner**.

   The Management application web client displays.

   NOTE
   If the Administrator disconnects the web client using the **Active Sessions** dialog box (**Server > Active Sessions**), the web client redirects to the login page after three minutes or as soon as you make a selection.

Launching the Configuration Wizard

You can re-launch the Configuration wizard to change the following configurations:

- FTP server
- Server IP
- Server Ports
- SMI Agent

NOTE
Changes to these configurations require a server restart.

NOTE
You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **14.2.1 > Server Management Console**).

1. Choose one of the following options:
   - On Windows systems, select **Start > Programs >** *Management_Application_Name* **14.2.1>** *Management_Application_Name* **Configuration**.
   - On UNIX systems, execute `sh Install_Home/bin/configwizard` on the terminal.
2. Click **Next** on the **Welcome** screen.
3. Click **Yes** on the confirmation message.
4. Complete the following steps on the **FTP/SCP/SFTP Server** screen.
   a. Choose one of the following options:
      - Select **Built-in FTP/SCP/SFTP Server** to configure an internal FTP/SCP/SFTP server and select one of the following options:
      - Select **Built-in FTP Server** to configure an internal FTP server
        The internal FTP server uses a default account and port 21. You can configure your own account from the **Options** dialog box. For instructions, refer to "Configuring an internal FTP server" on page 123.
      - Select **Built-in SCP/SFTP Server** to configure an internal SCP/SFTP server
        The internal SCP/SFTP server uses a default account and port 22. You can configure your own account from the **Options** dialog box. For instructions, refer to "Configuring an internal SCP or SFTP server" on page 124.
      - Select **External FTP/SCP/SFTP Server** to configure an external FTP server.
        You can configure the external FTP server settings from the **Options** dialog box. For instructions, refer to "Configuring an external FTP, SCP, or SFTP server" on page 125.
   b. Click **Next**.

   If port 21 or 22 is busy, a message displays. Click **OK** to close the message and continue. Once the Management application is configured, make sure port 21 or 22 is free and restart the server to start the FTP/SCP/SFTP service.

   NOTE
   If you use an FTP/SCP/SFTP server which is not configured on the same machine as the Management application, the Firmware Repository feature will not be available.

5. Complete the following steps on the **Server IP Configuration** screen.

   NOTE
   If the Management server or client has multiple Network Interface Cards and if any of these interfaces are not plugged in, you must disable them; otherwise, the following features do not work properly:
   Server impact
   - Configuration wizard (does not display all IP addresses)
   - Trap and Syslog auto registration
   - Report content (Ipconfiguration element does not display all server IP addresses)
   - Network OS configuration backup through FTP
   - Trace dump through FTP
   Client impact
   - Options dialog box (does not display all IP addresses)
   - Firmware import and download dialog box
   - Firmware import for Fabric OS and Network OS products
   - FTP button in Technical Support Repository dialog box
   - Technical supportSave of Fabric OS, Network OS, and Host products through FTP

   a. Select an address from the **Server IP Configuration** list.

   b. Select an address from the **Switch - Server IP Configuration Preferred Address** list.

      NOTE
      The host name does not display in the list if it contains invalid characters. Valid characters include alphanumeric and dash (-) characters. The IP address is selected by default.

      If DNS is not configured for your network, do not select the "hostname" option from the **Server IP Configuration** list. Selecting the "hostname" option prevents clients and devices from communicating with the server.

   c. Select an IP address from the **Switch - Server IP Configuration Preferred Address** list. The preferred IP address is used for switch and server communication.

      or

      Select **Any** from the **Switch - Server IP Configuration Preferred Address** list to enable switch and server communication with one of the reachable IP address present in the server. By default, **Any** option is selected.

      If you select a specific IP address from the **Server IP Configuration** screen and the selected IP address changes, you will not be able to connect to the server. To change the IP address, refer to "Configuring an explicit server IP address" on page 114.

   d. Click **Next**.

6. Complete the following steps on the **Server Configuration** screen.

FIGURE 4 Server Configuration screen

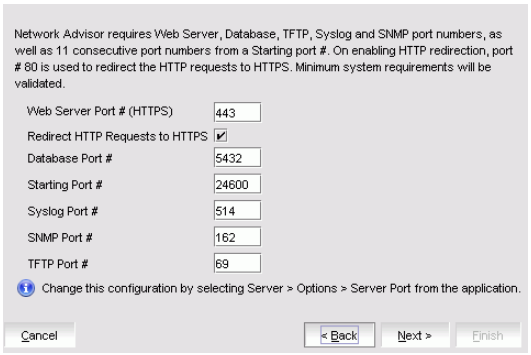

a. Enter a port number in the **Web Server Port # (HTTPS)** field (default is 443).

b. Enable HTTP redirection to HTTPS by selecting the **Redirect HTTP Requests to HTTPS** check box.

When you enable HTTP redirection, the server uses port 80 to redirect HTTP requests to HTTPS. You can configure the server port settings from the **Options** dialog box (**Server Port** pane). For instructions, refer to "Configuring the server port" on page 127.

c. Enter a port number in the **Database Port #** field (default is 5432).

d. Enter a port number in the **Starting Port Number** field (default is 24600).

NOTE
For Professional software, the server requires 11 consecutive free ports beginning with the starting port number.

NOTE
For Trial and Licensed software, the server requires 11 consecutive free ports beginning with the starting port number.

e. Enter a port number in the **Syslog Port Number** field (default is 514).

NOTE
If the default syslog port number is already in use, you will not receive any syslog messages from the device. To find and stop the process currently running on the default Syslog port number, refer to the *Installation and Migration Guide*.

f. Enter a port number in the **SNMP Port Number** field (default is 162).

g. Click **Next**.

If you enter a syslog port number already in use, a message displays. Click **No** on the message to remain on the **Server Configuration** screen and edit the syslog port number (return to step 6a). Click **Yes** to close the message and continue with step 7.

If you enter a port number already in use, a Warning displays next to the associated port number field. Edit that port number and click **Next**.

7. Complete the following steps on the **SMI Agent Configuration** screen.

a. Enable the SMI Agent by selecting the **Enable SMI Agent** check box.

b. Enable the SLP by selecting the **Enable SLP** check box.

c. Enable the SSL by selecting the **Enable SSL** check box.

d. Enter the SMI Agent port number in the **SMI Agent Port #** field (default is 5989 if SSL is enabled; otherwise, default is 5988).

e. Click **Next**.

8. Complete the following steps on the **Inventory Upload Configuration** screen.

a. Select **Enable** check box to enable inventory upload configuration.

b. Enter a valid E-mail address in **E-mail** field and click **Next**.

NOTE
E-mail server must be configured before enabling the Inventory upload configuration.

9. Verify your configuration information on the **Server Configuration Summary** screen and click **Next**.

10. Complete the following steps on the **Start Server** screen:

a. Select the **Start SMI Agent** check box, if necessary.

b. Select the **Start SLP** check box, if necessary.

c. Select the **Start Client** check box, if necessary.

d. Click **Finish**.

After all of the services (Server, SLP, SMI Agent, and Client) are started, the **Log In** dialog box displays.

11. Click **Yes** on the restart server confirmation message.

12. Choose one of the following options:

- If you configured authentication to CAC, enter your PIN in the **CAC PIN** field.
- If you configured authentication to the local database, an external server (RADIUS, LDAP, or TACACS+) or a switch, enter your user name and password.

The defaults are **Administrator** and **password**, respectively.

NOTE
Do not enter *Domain\User_Name* in the **User ID** field for LDAP server authentication.

13. Click **Login**.

14. Click **OK** on the Login Banner.

NOTE
When you launch the Management application or navigate to a new view, the **SAN** tab displays with a gray screen over the Product List and Topology Map while data is loading.

Viewing active sessions

To view the Management application active sessions, complete the following steps.

1. Select **Server > Active Sessions**.

   The **Active Sessions** dialog box displays (Figure 5).

   FIGURE 5 Active Sessions dialog box

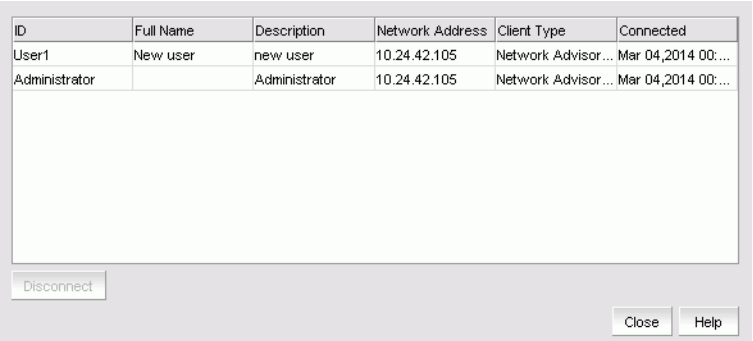

2. Review the active session information.

   The following information displays:

   - **ID** — Displays the name of the user (for example, Administrator).
   - **Full Name** — Displays the full name of the user.
   - **Description** — Displays the description of the user (for example, Operator).
   - **Network Address** — Displays the network address of the user.
   - **Client Type** — Displays the type of Management application client.
   - **Connected** — Displays the date and time the user connected to the server.

3. Click **Close**.

Disconnecting users

To disconnect a user, complete the following steps.

1. Select **Server > Active Sessions**.

   The **Active Sessions** dialog box displays.

2. Select the user you want to disconnect and click **Disconnect**.
3. Click **Yes** on the confirmation message.

   The user you disconnected receives the following message:

   The Client has been disconnected by *User_Name* from *IP_Address* at *Disconnected_Date_and_Time*.

4. Click **Close**.

   When you disconnect a client using the **Active Sessions** dialog box, the following event displays in the Master Log: Disconnect Client *User_Name @ IP_Address*.

Viewing server properties

To view the Management application server properties, complete the following steps.

1. Select **Server > Server Properties**.

   The **Server Properties** dialog box displays.

   FIGURE 6 Server Properties dialog box

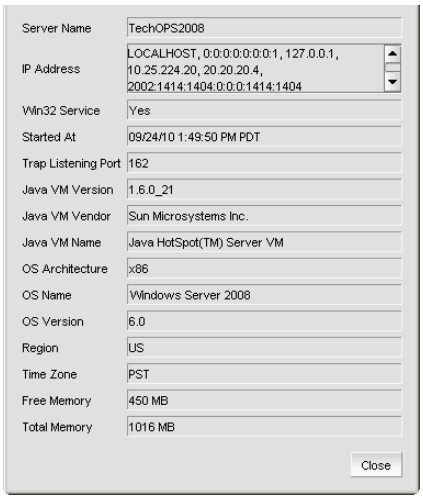

2. Review the information.

   TABLE 3 Server Properties

Field/Component	Description
Free Memory	The amount of free memory on the server.
IP Address	The IP address in IPv4 or IPv6 format.
Win32 Service	Specifies whether the Win32 service is available on the server. On UNIX servers, displays as “No”.
Java VM Name	The Java Virtual Machine name.
Java VM Vendor	The Java Virtual Machine vendor.
Java VM Version	The Java Virtual Machine version running on the server.
Server Name	The server’s name.
OS Architecture	The operating system architecture on the server.
OS Name	The name of the operating system running on the server.
OS Version	The operating system version running on the server.
Region	The server’s geographical region.
Started At	The time the server was started.
Time Zone	The server’s time zone.
Total Memory	The total amount of memory on the server.
Trap Listening Port	The number of the UDP port that listens for SNMP traps.

3. Click **Close** to close the **Server Properties** dialog box.

Viewing port status

The **Port Status** dialog box enables you to determine the availability of ports required for key Management application features. You can view the port status for the following ports:

- CIM Indication for Event Handling — Port 24618
- CIM Indication for HCM Proxy — Port 24619
- FTP — Port 21
- SCP/SFTP — Port 22
- SNMP Trap — Port 162
- Syslog — Port 514
- Web Server (HTTP) — Port 80
- Web Server (HTTPS) — Port 443

To view the port status, complete the following steps.

1. Click the port status icon ( ).

   The **Port Status** dialog box displays.

   FIGURE 7 Port Status dialog box

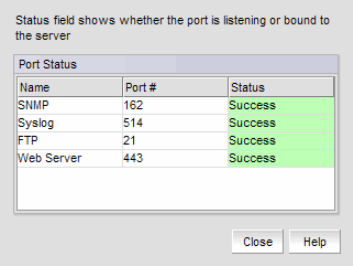

2. Review the port status details:
   - **Name** — The Port name. Options include CIM Indication for Event Handling, CIM Indication for HCM Proxy, FTP, SCP/SFTP, SNMP Trap, Syslog, Web Server (HTTP), and Web Server (HTTPS).
   - **Port #** — The required port number.
   - **Status** — The status of the port. The status options are as follows:
     - Success — The port is listening or bound to the server.
     - Failed — The port fails to listen or bind to the server. It is occuppied by another process.
     - Paritally Failed — The port is used by the server as well as other applications.
     - Disabled (external FTP port only) — This is considered a normal status.
   - **Running Process** – The name of the process using the port (not the Management applciation). Blank when the port is only used by the Management applciation server. If multiple processes occupy the same port, the process names display in a comma-separated list.
   - **Recommended Actions** — Suggested action to take to resolve the issues.
3. Click **Close**.

Server and client ports

In some cases, a network may utilize virtual private network (VPN) or firewall technology, which can prohibit communication between Products and the Servers or Clients. In other words, a Server or Client can find a Product, appear to log in, but is immediately logged out because the Product cannot reach the Server or Client. To resolve this issue, check to determine if the ports in the table below need to be opened up in the firewall.

NOTE
Professional edition does not support remote clients.

Table 4 lists the default port numbers and whether or not it needs to be opened up in the firewall and includes the following information:

- **Port Number** — The port at the destination end of the communication path.
- **Ports** — The name of the port.
- **Transport** — The transport type (TCP or UDP).
- **Description** — A brief description of the port.
- **Communication Path** — The "source" to "destination" vaules. Client and Server refer to the Management application client and server unless stated otherwise. Product refers to the Fabric OS, Network OS, or IronWare devices.
- **Open in Firewall** — Whether the port needs to be open in the firewall.

NOTE
For bi-directional protocols, you must open the firewall port bi-directionally.

TABLE 4 Port usage and firewall requirements

Port Number	Ports	Transport	Description	Communication Path	Open in Firewall
20[1]	FTP Port (Control)	TCP	FTP Control port for internal FTP server	Client–Server Product–Server	Yes
21[1]	FTP Port (Data)	TCP	FTP Data port for internal FTP server	Client–Server Product–Server	Yes
22[2]	SSH or SCP or SFTP	TCP	Secure telnet and secure upload and download to product	Server–Product Client –Product Product – Server	Yes
23	Telnet	TCP	Telnet port from server/client to product	Server–Product Client–Product	Yes
25[2]	SMTP Server port	TCP	SMTP Server port for e-mail communication if you use e-mail notifications without SSL	Server–SMTP Server	Yes
49[2]	TACACS+ Authentication port	TCP	TACACS+ server port for authentication if you use TACACS+ as an external authentication	Server–TACACS+ Server	Yes
69	TFTP	UDP	File upload/download to product	Product-Server	Yes
80[2]	Management application HTTP server	TCP	Non-SSL HTTP/1.1 connector port if you use secure client-server communication. You need this port for HTTP redirection	Client–Server	Yes
80[1]	Product HTTP server	TCP	Product non-SSL http port for http and CAL communication if you do not use secure communication to the product	Server–Product	Yes

TABLE 4 Port usage and firewall requirements (Continued)

Port Number	Ports	Transport	Description	Communication Path	Open in Firewall
			Product non-SSL http port for http and CAL communication if you do not use secure communication to the product and you do not use the Management application server proxy	Client–Product	Yes
161[2]	SNMP port	UDP	Default SNMP port	Server–Product	Yes
162[2]	SNMP Trap port	UDP	Default SNMP trap port	Product-Server	Yes
389[2]	LDAP Authentication Server Port	UDP TCP	LDAP server port for authentication if you use LDAP as an external authentication	Server–LDAP Server	Yes
443[1,2]	HTTPS server	TCP	HTTPS (HTTP over SSL) server port if you use secure client - server communication	Client-Server	Yes
443[2]			HTTPS (HTTP over SSL) server port if you use secure communication to the product	Server–Product	Yes
443			HTTPS (HTTP over SSL) server port if you use secure communication to the product and you do not use the Management application server proxy	Client–Product	Yes
443[2]			HTTPS (HTTP over SSL) server port if you use vCenter discovery	Server-vCenter Server	Yes
465[2]	SMTP Server port for SSL	TCP	SMTP Server port for e-mail communication if you use e-mail notifications with SSL	Server-SMTP Server	Yes
514[2]	Syslog Port	UDP	Default Syslog Port	Product–Server Managed Host - Server	Yes
636[2]	LDAP Authentication SSL port	TCP	LDAP server port for authentication if you use LDAP as an external authentication and SSL is enabled	Server–LDAP Server	Yes
1812[2]	RADIUS Authentication Server Port	UDP	RADIUS server port for authentication if you use RADIUS as an external authentication	Server–RADIUS Server	Yes
1813[2]	RADIUS Accounting Server Port	UDP	RADIUS server port for accounting if you use RADIUS as an external authentication	Server–RADIUS Server	Yes
5432	Database port	TCP	Port used by database if you access the database remotely from a third-party application	Remote ODBC–Database	Yes
5988	SMI Server port	TCP	SMI server port on the Management application and the CIM/SMI port on HBAs if you use SMI Agent without SSL	SMI Client- Server	Yes
				Server-Managed Host	Yes
5989[1,2]	SMI Server port with SSL enabled	TCP	SMI Agent port on the Management application and the CIM/SMI port on HBAs if you use SMI Agent with SSL	SMI Agent Server–Client	Yes
				Server-Managed Host	Yes
6343[2]	sFlow	UDP	Receives sFlow data from products if you are monitoring with sFlow	Product-Server	Yes
24600[2]	JBoss remoting connector port	TCP	Use for service location. Uses SSL for privacy.	Client–Server	Yes

TABLE 4 Port usage and firewall requirements (Continued)

Port Number	Ports	Transport	Description	Communication Path	Open in Firewall
24601[2]	JBoss Transaction Services Recovery Manager port	TCP	Not used remotely.	Server	Yes
24602[2]	JBoss Transaction Status Manager port	TCP	Not used remotely.	Server	Yes
24603[2]	HornetQ Netty port	TCP	Use for JMS (Java Message Service), async messages from server to client. Uses SSL for privacy.	Client–Server	Yes
24604[2]	JMX remoting connector port	TCP	Management console port for native connector (JMX)	Client–Server	Yes
24605[2]	JBoss https management port TCP	TCP	Management console port for HTTPS based management	Client–Server	Yes
24606[2]	Fault Management CIM Indication Listener Port	TCP	Used for HBA management	Managed Host - Server	Yes
24607[2]	HCM Proxy CIM Indication Listener port	TCP	Used for HBA management	Managed Host - Server	Yes
24608[2]	Reserved for future use	TCP	Not used	Client - Server	No
24609[2]	Reserved for future use	TCP	Not used	Client - Server	No
24610[2]	Reserved for future use	TCP	Not used	Client - Server	No
34568	HCM Agent discovery port	TCP	Used for HBA management via JSON	Server - Managed Host	Yes
55556	Launch in Context (LIC) client hand shaking port	TCP	Client port used to check if a Management application client opened using LIC is running on the same host **NOTE:** If this port is in use, the application uses the next available port.	Client	No

1. Port does not need to be open in the firewall for Professional edtion.

2. The default port number. You must use the same port number for all products or hosts managed by the Management server. This port is configurable in the Management server; however, some products and firmware versions do not allow you to configure a port.

Accessibility features for the Management application

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

The following list includes the major accessibility features in the Management application:

- Keyboard shortcuts
- Look and Feel

Keyboard shortcuts

You can use the keystrokes shown in the table below to perform common functions.

NOTE
To open a menu using keystrokes, press ALT plus the underlined letter. To open a submenu, open the menu, then press the key for the underlined letter (SHIFT plus letter for capitals) of the submenu option.

TABLE 5 Keyboard shortcuts

Menu Item or Function	Keyboard Shortcut
All Panels	F12
Collapse	CTRL + L
Command Tool	SHIFT + F4
Connectivity Map	F7
Copy	CTRL + C
Cut	CTRL + X
Delete	Delete
Delete All	CTRL + Delete
Help	F1
Internet Explorer	SHIFT + F2
Master Log	F5
FireFox	SHIFT + F1
Paste	CTRL + V
Product List	F9
Properties	Alt-Enter
Select All	CTRL + A
Show Ports	F4
SSH	Shift-F5
View Utilization	CTRL + U
Zoom In	CTRL + NumPad+
Zoom Out	CTRL + NumPad-

Look and feel customization

You can configure the Management application to mimic your system settings as well as define the size of the font.

"Look" refers to the appearance of graphical user interface widgets and "feel" refers to the way the widgets behave.

The Management application currently uses the "*Management_Application* Default Look and Feel" for some of the components (for example, Layout, Minimap, and so on) and the "Java Metal Look and Feel" for others.

Setting the look and feel

NOTE
Setting the look and feel is only supported on Windows systems.

The following table details the Management application components that change when you set the look and feel as well as those components that do not change.

TABLE 6 Look and feel changes

Components Affected	Components Not Affected
All Java native components with Metal Look And Feel are affected.	The Connectivity map does not change when devices are present. You must change the theme using the map display settings (**View > Map Display**).
The Menu bar, Tool bar, Status bar, as well as all tables and dialog boxes are affected.	All icons and images are not affected.
Layout is affected only when it is empty.	The Minimap is not affected.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Look and Feel** in the **Category** list.
3. Choose from one of the following options:
   - Select **Default** to configure the look and feel back to the Management application defaults.
   - Select **System** to configure the Management application to have the look and feel of your system.

     This changes the look and feel for the components that use "Java Metal Look and Feel". For example, if you have your system display color scheme set to 'High Contrast #1', then the Management application will be set to "High Contrast #1". Font size of the components is not affected by theme changes.
4. Click **Apply** or **OK** to save your work.
5. Click **OK** on the message.

   NOTE
   Changes do not take effect until after you restart the client.

Changing the font size

NOTE
Changing the font size is only supported on Windows systems.

Font size changes proportionately in relation to the system resolution. For example, if the system resolution is 1024 x 768, the default font size would be 8 and large font size would be 10.

To change the font size for all components including the Connectivity Map of the Management application interface, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Look and Feel** in the **Category** list.

3. Select one of the following options from the **Font Size** list:
    - Select **Default** to return to the default font size.
    - Select **Small** to change the font to a smaller font size.
    - Select **Large** to change the font to a larger font size.

        NOTE
        Changing the font size to **Large** may cause the interface components (for example, text and button labels) to display incorrectly.
4. Click **Apply** or **OK** to save your work.
5. Click **OK** on the message.

    NOTE
    Changes do not take effect until after you restart the client.

Highlighting events in the Master Log

To configure the Management application to highlight events in the Master Log based on event severity, complete the following steps.

1. Select **Server > Options**.

    The **Options** dialog box displays.
2. Select **Look and Feel** in the **Category** list.
3. Select the **Highlight Events Enable** check box.
4. Click **Apply** or **OK** to save your work.
5. Click **OK** on the message.

    NOTE
    Changes do not take effect until after you restart the client.

PostgreSQL database

You can connect to the database using one of the following options:

- pgAdmin III
- ODBC client
- Command line interface

Connecting to the database using pgAdmin III

To access the PostgreSQL database, complete the following steps.

1. Choose one of the following options:
    - On Windows systems, launch the dbadmin.bat script in the *Install_Home*\bin\ directory.
    - On UNIX systems, launch the dbadmin script in the *Install_Home*\bin\ directory.

2. Select **File > Add Server**.

   The **New Server Registration** dialog box displays.

3. Enter the *DB_server_IP_address* or "localhost" in the **Host** field.
4. Enter the port number (default is 5432) on which the PostgreSQL server is running in the **Port** field.
5. Enter your user name (default is dcmuser) in the **Username** field.
6. Enter your password (password) in the **Password** field.
7. Click **OK** on the **New Server Registration** dialog box.

   The **pgAdmin III** application displays.

8. To browse data in the database, complete the following steps.
   a. Expand the **Tables** tree in the **Object browser** pane.
   b. Right-click a table in the list and select **View Data > View All Rows**.
9. To execute a freestyle SQL query in the database, complete the following steps.
   a. Expand the **Tables** tree in the **Object browser** pane.
   b. Right-click a table in the list and select **Scripts > SELECT script**.

      The **Query** dialog box displays.

10. Select **File > Exit** to close the **pgAdmin III** application.

Connecting to the database using the ODBC client (Windows systems)

The Open Database Connectivity (ODBC) driver enables you to configure the data source name (DSN) for the database.

To install the ODBC driver and create a new data source, complete the following steps.

1. Double-click edb_psqlodbc.exe located on the DVD (*DVD_Drive*/*Management_Application*/odbc/Windows).
2. Install the file to the usual location for your system's application files (for example, C:\Program Files*Management_Application* ODBC Driver) on the **Select Install Folder** screen and click **Next**.

   NOTE
   If you select an invalid location, the ODBC driver is installed in a different location than where the ODBC executable drivers are located.

3. On the **Ready to Install** screen click **Next**.
4. Click **Finish** to complete the installation.
5. Choose one of the following opitions:
   - (32-bit OS) Select **Start > Settings > Control Panel > Administrative Tools > Data Sources (ODBC)**.

     NOTE
     Clients edition are supported on 32-bit operating systems.

   - (64-bit OS) (Windows only) Select **Start > Run**, type `%windir%\SysWOW64\odbcad32.exe` and press **Enter**.

   The **ODBC Data Source Administrator** dialog box displays.

6. Click the **System DSN** tab.
7. Click **Add**.

   The **Create a New Data Source** dialog box displays.
8. Select **PostgreSQL Unicode**.
9. Click **Finish**.

   The **PostgreSQL Unicode ODBC Driver (psqlODBC) Setup** dialog box displays.
10. Enter a name for the data source in the **Datasource** field.
11. Enter the description of the database in the **Description** field.
12. Enter the name of the database in the **Database** field.
13. Select **enable** or **disable** from the **SSL Mode** list to specify whether or not to use SSL when connecting to the database.
14. Enter the IP address or host name of the Management application server in the **Server** field.
15. Enter the database server port number (default is 5432) in the **Port Number** field.
16. Enter the database user name in the **User Name** field.
17. Enter the password in the **Password** field.
18. Click **Test** to test the connection.
19. Click **OK** on the **Connection Test** dialog box.
20. Click **Save**.
21. Click **OK** on the **ODBC Data Source Administrator** dialog box.
22. To export data, select **Data > Import External Data > New Database Query** and complete the steps in the **Data Connection Wizard**.

Connecting to the database using the ODBC client (Linux systems)

NOTE
The ODBC driver is not supported on 64-bit Linux systems.

You must have the Open Database Connectivity (ODBC) driver to allow remote clients to export data and generate reports. The ODBC driver enables you to configure the data source name (DSN) for the Management application database.

Before you install the Linux ODBC driver, download the ODBC RedHat Package Manager (RPM) file based on the Linux version.

TABLE 7 ODBC RedHat Package Manager (RPM) file requirements

Linux version	RedHat Package Manager file
SUSE	Rpm –I unixODBC-2.2.12-197.17.i586.rpm
RedHat or Oracle Enterprise	Rpm –i unixODBC-2.2.11-1.i386.rpm

Installing the ODBC driver on Linux systems

To install the ODBC driver, complete the following steps.

1. Execute the following command in the terminal:

```
> su
>chmod 777 edb_psqlodbc.bin
> ./edb_psqlodbc.bin
```

2. On the **Setup psqlODBC** screen, click **Next**.
3. Install the file to the usual location for your system's application files (for example, /opt/PostgreSQL/psqlODBC) on the **Installation Directory** screen and click **Next**.

   NOTE
   If you select an invalid location, the ODBC driver is installed in a different location than where the ODBC executable drivers are located.

4. On the **Ready to Install** screen, click **Next**.
5. On the **Completing the psqlODBC Setup Wizard** screen, click **Finish** to complete the installation.

Adding the datasource on Linux systems

Before you edit the INI files, make sure the PostgreSQL database is up and running.

NOTE
For RedHat and Oracle Enterprise systems, the odbc.ini and odbcinst.ini files are located in /etc. For SUSE systems, the odbc.ini and odbcinst.ini files are located in /etc/unixODBC.

1. Open the odbc.ini file in an editor and enter the datasource information as follows:

```
[TestDB]
Description = PostgreSQL 9.2
Driver = /opt/PostgreSQL/psqlODBC/lib/psqlodbcw.so
Database = dcmdb
Servername = 172.26.1.54
UserName = dcmadmin
Password = passw0rd
Port = 5432
```

2. Save and close the odbc.ini file.
3. Open the odbcinst.ini file in a text editor and make sure that the driver path information is correct.

   After you install the PostgreSQL ODBC driver, the odbcinst.ini should automatically update the driver path. If the driver path is not updated, add the following:

```
[psqlODBC]
Description=PostgreSQL ODBC driver
Driver=/opt/PostgreSQL/psqlODBC/lib/psqlodbcw.so
```

4. Save and close the odbcinst.ini file.

Testing the connection on Linux systems

To test the connection, complete the following steps.

1. Download and install Open Office.
2. Select **File > New > Database**.

   The **Database Wizard** displays.
3. On the **Select database** screen, complete the following steps.
   a. Select the **Connect to an existing database** option.
   b. Select **ODBC** from the list.
   c. Click **Next**.
4. On the **Set up ODBC connection** screen, complete the following steps.
   a. Click **Browse**.

      The datasource saved in the odbc.ini file is populated in the **Datasource** dialog box.
   b. Select the datasource and click **OK** on the **Datasource** dialog box.
   c. Click **Next**.
5. On the **Set up user authentication** screen, complete the following steps.
   a. Enter the database user name in the **User name** field.
   b. Select the **Password required** check box.
   c. Click **Test Connection** to test the connection.

      The **Authentication Password** dialog box displays.
   d. Enter the database password in the **Password** field and click **OK**.
   e. Click **OK** on the **Connection Test** dialog box.

      If an error message (file not found while testing the connection) displays, copy the lib files from the <postgresSQL path>/lib/* directory to the /usr/lib/ directory.
   f. Click **Next**.
6. On the **Save and proceed** screen, click **Finish**.

Executing SQL queries from the CLI

To execute SQL queries from the command line interface (CLI) , complete the following steps.

1. Choose one of the following options:
   - On Windows systems, launch the dbsql.bat script in the *Install_Home*\bin\ directory.
   - On UNIX systems, launch the dbsql script in the *Install_Home*\bin\ directory.
2. Execute your query from the command window.
3. Close the command window.

Changing the database user password

To change the read/write or read only database password, complete the following steps in the *Install_Home*/bin directory.

1. Open a command window.
2. Type **dbpassword** *User_Name Password New_Password Confirm_Password* and press **Enter**.

   Where *User_Name* is your user name, *Password* is your current password, and *New_Password* and *Confirm_Password* are your new password. The read/write user name and password defaults are dcmadmin and passw0rd (zero), respectively. The read-only user name and password defaults are dcmuser and password (all lowercase), respectively.

   If the password changed successfully, the following message displays:
   Password changed successfully.

   If an error occurs and the password did not change, the following message displays:
   Error while updating password. Please try again.
   Press any key to continue.

   If the current password and new password are the same, the following message displays:
   Old and New passwords cannot be same. Use different password and try again.
   Press any key to continue.

   If the new password and confirm password do not match, the following message displays:
   New password and confirm password do not match. Please try again.
   Press any key to continue.
3. Launch the Server Management Console.
4. Click the **Services** tab.
5. Click **Stop** to stop all services.
6. Click **Close** to close the Server Management Console.
7. Launch the Server Management Console.
8. Click **Start** to start all services.

   NOTE
   If the server is configured to use an external FTP server, the Server Management Console does not attempt to start the built-in FTP service.
9. Click **Close** to close the Server Management Console.

Supported open source software and third-party software products

Table 8 lists the open source software and third-party software products used in this release.

TABLE 8 Open source software and third-party software products

Component	License
Maverick Java SSH API (1.4.25)	SSHTools software license
powerdesigner (15.1)	Any License
SafeNet Sentinel RMS SDK (8.2.2)	SafeNet Sentinel RMS SDK license
SafeNet Sentinel Caffe (1.6.1)	safenet license

TABLE 8 Open source software and third-party software products (Continued)

Component	License
Quartz Enterprise Job Scheduler (2.2.0)	Apache License 2.0
Xalan-Java (2.7.1)	Apache License 2.0
WebNMS SNMP API (4.0.6a)	webnms SNMP API
atmosphere-runtime (2.0.4)	Apache License 2.0
VMware SDK (2.5.0)	Any License
VCEM SDK (7.4)	Any License
YourKit Java Profiler (9.5.1)	YourKit License
yguard 2.4 (2.4)	yWorks license for yGuard
yfiles 2.9 (2.9)	yWorks license for yFiles
primefaces (4.0)	Apache License 2.0
yworks (1.3)	yWorks license for yExport
Oracle-java update (51)	Oracle JRE license
JBoss Application Server (7.2)	LGPL
sigar (1.6.4)	Apache License 2.0
Java Service Wrapper (3.5.23)	Tanuki Software, Ltd. Development Software License Agreement, Version 1.1
Jasper Report (5.1.2)	GNU LGPL
ehcache (2.8.1)	Apache License 2.0
shiro-core (1.2.2)	Apache License 2.0
pgbadger (4.1)	BSD License
J2SSH Maverick (1.4.50)	SSHTools Professional License
javamelody (1.49)	Apache License 2.0
birt (4.3.0)	Eclipse Public License 1.0
Hermes JMS Console (1.14)	Apache License 2.0
Google Guice (3.0)	Apache License 2.0
Axis (1.4)	Apache License 2.0
hornetq (2.3)	Apache License 2.0
IBM ESS (Embedded Security Service) (6.2)	IBM SOW 7 Amendment 4
IAIK PKCS#11 WRAPPER (1.3)	IAIK PKCS#11 Wrapper License
Java Tar (2.5)	Public Domain
JbcParser (3.7)	Math Parser License
Install Anywhere (2012)	FLEXERA Software End-User License Agreement
iReasoning SNMP API (5.0.36)	IReasoning Inc SNMP API End User License Agreement
iBatis for Java (2.3.4.726)	Apache License 2.0
iBatis DAO Framework (2.2.0.638)	Apache License 2.0
Bean Scripting Framework (2.4.0)	Apache License 2.0
AXL Radius Client (3.29)	AXL Software® Library License Agreement
Castor Binding Framework (0.9.9.1)	Apache License 2.0

TABLE 8 Open source software and third-party software products (Continued)

Component	License
Bouncy Castle Crypto Provider (1.45)	Bouncy Castle License (an adaptation of MIT X11 License)
Checkstyle (5.0)	GNU Lesser General Public License.
Conf-M (1.9.7)	Tail-f license
DNS Java (2.0.7)	Sun Public License 1.0
dom4j (1.6.1)	dom4j License
Eclipse IDE (3.4.1)	Eclipse Public License Version 1.0
Emma (2.0.5312)	Emma Common Public License v1.0
FindBugs (1.3.9)	LGPL
JIDE 3.5.3 (3.5.3)	Jide 3.5.3 license
Microsoft VC++ Redistributable Package (2010 sp1)	Microsoft Visual Studio License
OpenSSL for Linux (1.0.0a)	OpenSSL License
OpenSSL for Windows (1.0.0)	OpenSSL license and SSLeay license
birt (4.2.1)	Eclipse Public License 1.0
Portlet API (2.0)	Apache License 2.0
XML RPC (1.2-B1)	Apache License 1.1
xdoclet (1.2.3)	BSD 3-clause "New" or "Revised" License
VI Java API (5.1 [20121126])	BSD License
testng (5.9)	Apache License 2.0
SNMP4J. (2.0.2)	Apache License 2.0
RockSaw Raw Socket Library (1.0.0)	Apache License 2.0
itextpdf (2.1.7)	GNU Library General Public License v2.0
Apache Commons DBCP (1.2.2)	Apache License 2.0
Apache Commons Configuration (1.6)	Apache License 2.0
jcifs (1.3.12)	GNU Lesser General Public License v2.1
Apache Commons Collections (3.2.1)	Apache License 2.0
Apache Commons Codec (1.4)	Apache License 2.0
jcalendar (1.3.3)	GNU Lesser General Public License v2.1
Apache Ant (1.7.1)	Apache License 2.0
jcommon (1.0.16)	GNU Lesser General Public License v2.1
7-Zip LZMA SDK (4.65)	LZMA SDK
jfreechart (1.0.13)	GNU Lesser General Public License v2.1
tartool (1.4)	GNU General Public License v2.0
Quality First Library (0.99.0)	Mozilla Public License 1.1
TableLayout (2009-06-10)	The Clearthought Software License, Version 2.0
Abator (1.2.1-681)	Apache License 2.0
Apache Commons Lang (2.6)	Apache License 2.0
Apache Commons Logging (1.1.1)	Apache License 2.0

TABLE 8 Open source software and third-party software products (Continued)

Component	License
Apache Commons Validator (1.3.1)	Apache License 2.0
edtftpj (2.3.0)	GNU Lesser General Public License v2.1
Apache Commons FileUpload (1.2.1)	Apache License 2.0
freehep-freehep-vectorgraphics	GNU Lesser General Public License v2.1
Apache Commons JXPath (1.3)	Apache License 2.0
infinispan (4)	GNU Lesser General Public License v2.1
Apache Commons Digester (2.0)	Apache License 2.0
glazedlists (1.8.0)	GNU Lesser General Public License v2.1
WS J WBEM Server 3 (3.9.5)	License - Jserver 3.x
Apache Commons Discovery (0.4)	Apache License 2.0
faenil-google-opensans-fonts (0.1)	Apache License 2.0
Apache Log4j (1.2.16)	Apache License 2.0
Apache SSHD (0.7.0)	Apache License 2.0
Apache FTP Server (1.0.3)	Apache License 2.0
Apache HttpComponents (4.2.1)	Apache HttpComponents
Oracle-java update (60)	Oracle JRE license
rrd4j (2.0.7)	Apache License 2.0
Postgresql-ODBC (09.02.0100)	Library General Public Licence
Google Guava (14.0)	Apache License 2.0
PostgreSQL-JDBC (9.2-1004)	BSD License
Apache Commons HttpCore (4.2.1)	Apache License 2.0
PostgreSQL (9.2.8)	PostgreSQL License
Apache Extras Companion for Apache log4j (1.1)	Apache License 2.0
jackson (2.0.5)	GNU Lesser General Public License v2.1
jaxen (1.1.1)	BSD 3-clause "New" or "Revised" License
alphanum-comparator (1.0)	GNU General Public License v3.0 or later
slf4j (1.7.2)	QOS.ch
javahelp (2.0.05)	GNU General Public License v2.0 with Classpath Exception
sblim (1.3.9.3)	Eclipse Public License 1.0
JDOM component (1.1.1)	JDOM License
JBoss Web (2.1.9)	LGPL v3
JGoodies Forms (1.2.1)	BSD
JGoodies Binding (2.0.6)	BSD
JBoss Drools (5.5)	Apache License 2.0
jmockit (1.2)	MIT License
Jsonrpc4j (0.24)	MIT License
JSON-RPC-Client (5.0)	Apache License 2.0

TABLE 8 Open source software and third-party software products (Continued)

Component	License
JGoodies Looks (2.2.2)	BSD
JGoodies Validation (2.0.1)	BSD
jgraph (5.13.0.1)	BSD License
jmesa (2.4.5)	Apache License 2.0
pmd (4.2.5)	BSD License
Report Ng (1.1.1)	Apache License 2.0
L2FProd.com Common Components (7.3)	Apache License 2.0
Mime Type Detection Utility (2.1.2)	Apache License 2.0
MyBatis Persistence Framework (3.0.2 GA)	Apache License 2.0
opensaml (2.3.0)	Apache License 2.0

SAN feature-to-firmware requirements

Use the following table to determine whether the Management application SAN features are only available with a specific version of the Fabric OS firmware as well as if there are specific licensing requirements.

TABLE 9 SAN feature-to-firmware requirements

Feature	Fabric OS
Access Gateway (AG)	AG connected to Fabric OS devices requires firmware 7.0 or later.
Call Home (Trial and Licensed version Only)	Requires Fabric OS 7.0 or later for supportSave. Requires Fabric Watch license for SNMP traps.
Configuration Management	Requires Fabric OS 7.0 or later
D-port	Requires Fabric OS 7.0 or later
Discarded Frames	Requires Fabric OS 7.0 or later for 16 Gbps-capable E_Ports, 10 Gbps-capable D_Ports or E_Ports. Requires Fabric OS 7.1 or later for 16 Gbps-capable F_Ports, ICL ports, and AG N_Ports.
Discovery	Requires Fabric OS 7.0 or later for the seed switch.
Encryption (Trial and Licensed version Only)	Requires Fabric OS 6.1.1_enc or 6.2 or later.
Enhanced Group Management (Trial and Licensed version Only)	Requires Enhanced Group Management license.
Fault Management	Requires Fabric OS 7.0 or later for SNMP traps
Fabric Binding (Trial and Licensed version Only)	Requires Fabric OS 7.0 or later.
Fabric Vision (Monitoring and Alerting Policy Suite (MAPS) and Flow Vision)	Flow Vision requires Fabric OS 7.2 or later and requires Fabric Vision license on the device. MAPS requires Fabric OS 7.1 or later and requires Fabric Watch and Advanced Performance Monitoring (APM) licenses on the device.
FCIP Management	Requires FCIP license. Requires Fabric OS 7.0 or later to enable the FICON Emulation tab on the FCIP Tunnel Advanced Settings dialog box.
FCoE Management	Requires FCoE license on the device. Requires Fabric OS version v6.1.2_CEE or later.

TABLE 9 SAN feature-to-firmware requirements (Continued)

Feature	Fabric OS
FICON (Trial and Licensed version Only)	Requires Fabric OS 7.0 or later. Requires FICON CUP license to allow CUP management features.
Firmware Management	Requires Fabric OS 7.0 or later. Requires Enhanced Group Management license to perform group actions.
High Integrity Fabric	Requires Fabric OS 7.0 or later.
Meta SAN	Requires Fabric OS 7.0 or later for FC router and router domain ID configuration. Requires Integrated Routing license.
Performance	Requires Advanced Performance Monitoring (APM) license for End-to-end Monitoring and Top Talkers. Requires Enhanced Group Management license for Historical graphs and tables. Requires Fabric Watch license for Performance thresholds.
Port Commissioning	Requires Fabric OS 7.1 or later
Port Fencing (Trial and Licensed version Only)	Requires Fabric OS 7.0 or later.
Security Management	Requires Fabric OS 7.0 and later.
Technical Support Data Collection	Requires Fabric OS 7.0 or later.
Troubleshooting and Diagnostics	Requires Fabric OS 7.0 or later.
Virtual Fabrics (Trial and Licensed version Only)	Requires at least one Virtual Fabrics-enabled physical chassis running Fabric OS 7.0 or later.
Zoning	Requires Adaptive Networking license for Quality of Service zones.

Patches

Installing a patch

The patch installer enables you to update the Management application between releases. Each patch installer includes the previous patches within a specific release. For example, patch F (11.X.Xf) includes the upgrades in the patch installers for A (11.X.Xa) through E (11.X.Xe).

To install a patch, complete the following steps.

1. Stop all services by completing the following steps.
   a. Launch the Server Console.
   b. Click the **Services** tab.
   c. Click **Stop** to stop all services.

      NOTE
      If you perform patch upgrade while services are running, an error message displays.
2. Go to the `/bin` directory.

   Install_Home/bin (Windows)

   /opt/*Application_Name*/bin (UNIX)
3. Execute the patch file for your operating system:

   `patch.bat` (Windows)

   `patch.sh` (UNIX)

   The **Upgrade** dialog box displays.
4. Browse to the patch file.

   The patch zip file uses the following naming convention:
   <*Application*>_<*Major_Version*><*Minor_Version*><*Revision_Number*><*Patch_Version*>_<*Company_Name*>.zip (for example na_1130a_<*Company_Name*>.zip).
5. Click **Upgrade**.

   If the patch process is interrupted (for example, loss of power), you must restart the patch process.

   The patch installer performs the following functions:

   - Extracts patch files to the *Install_Home* folder.
   - Creates a back up (zip) of the original files to be updated and copies the zip file to the *Install_Home*\patch-backup directory (for example, *Install_Home*\patch-backup\na_11-3-0a.zip).

     The first time you apply a patch, the back up patch zip file uses the following naming convention:
     <*Application*>_<*Major_Version*>-<*Minor_Version*>-<*Revision_Number*><*Patch_Version*>.zip (for example, *Install_Home*\patch-backup\na_11-3-0a.zip).

Each additional time you apply a patch, the back up patch zip file uses the following naming convention: <*Application*>_<*Major_Version*>-<*Minor_Version*>-<*Revision_Number*><*Patch_Version*>-<*Previous_Patch_Version*>.zip (for example, *Install_Home*\patch-backup\na_11-3-0-patch-a.zip).

- Generates a patch log.
- Updates the conf file (*Install_Home*\conf\patch.conf) to include the patch version applied and patch created date.
- Updates the patch version in the **About** dialog box (Select **Help > About** in the main window).

6. Start all services by completing the following steps.
   a. Launch the Server Console.
   b. Click the **Services** tab.
   c. Click **Start** to start all services.

Uninstalling a patch

Note that only one set of back up files are retained which enables you revert back to the previous version. You can only revert back one version. For example:

- If you upgrade from patch A to patch B, you can revert back to patch A.
- If you upgrade from patch A to patch B to patch C then to patch F, you can only revert back to patch C.

To uninstall a patch, complete the following steps.

1. Stop all services by completing the following steps.
   a. Launch the Server Console.
   b. Click the **Services** tab.
   c. Click **Stop** to stop all services.
2. Go to the *Install_Home*/patch-backup directory.
3. Extract the patch zip file (for example, na_1120a_<*Company_Name*>.zip).
4. Open the restore.xml file from the extracted files.

   The artifacts (jar files, war files, and so on) you need to replace display as separate file tags in the restore.xml file. The location of each artifact in the extracted folder is detailed in the src value under each file tag.
5. Go to the location of the first artifact (as shown in the src value under the file tag).
6. Copy the artifact from the extracted folder to the source folder in the *Install_Home*/patch-backup directory.
7. Repeat step 5 and 6 for all artifacts listed in the restore.xml folder.
8. Go to the *Install_Home*/conf directory.
9. Open the version.properties file in a text editor.
10. Change the patch version (patch.version) value to the reverted patch (for example, if you are reverting from patch F to patch C then `patch.version = c`).

    If the previous version is the initial version (no patches), change the patch version value to none (for example, `patch.version = None`).

11. Go to the *Install_Home*/patch-backup/conf directory.
12. Copy the patch.conf file in this directory to the *Install_Home*/conf directory.

    If the previous version is the initial version (no patches), delete the patch.conf file in the *Install_Home*/conf directory.
13. Start all services by completing the following steps.
    a. Launch the Server Console.
    b. Click the **Services** tab.
    c. Click **Start** to start all services.

Discovery

SAN discovery overview

Discovery is the process by which the Management application contacts the devices in your SAN. When you configure discovery, the application discovers devices connected to the SAN. The application illustrates each device and its connections on the Connectivity Map (topology).

When you discover a fabric, the Management application checks to confirm that the seed switch is running a supported Fabric OS version in the fabric, and if it is not, the Management application prompts you to select a new seed switch.

NOTE
Discovery of a Secure Fabric OS fabric in strict mode is not supported.

For a Fabric OS fabric, the seed switch must be the primary Fabric Configuration Server (FCS). If you use a non-primary FCS to discover the fabric, the Management application displays an error and will not allow the discovery to proceed. If the Management application has already discovered the fabric, but afterward you create the FCS policy and the seed switch is not a primary FCS, an event is generated during the next poll.

The Management application cannot discover a fabric that is in the process of actively configuring to form a fabric. Wait until the fabric is formed and stable, then re-attempt the fabric discovery.

After fabric discovery successfully completes, all clients are updated to display the newly discovered fabric.

During fabric discovery, you can define an IPv4 address or IPv6 address for the device; however, the Management application uses the preferred IP format to connect with the device. To configure the preferred IP format, refer to "Configuring the preferred IP format" on page 122.

Enabling or disabling auto enclosure

This batch file enables you to create auto enclosures from end devices present in the Management application database.

Once a fabric is discovered an enclosure is formed for the Host having FDMI with symbolic name enabled. When FDMI name is same for the adapters (HBA and CNA) which are displayed through fabric discovery, auto enclosure will be displayed for the fabric or fabrics.

To enable or disable auto enclosure, complete the following steps.

- On Windows systems, complete the following steps.

  Open a command prompt and navigate to the *Install_Home*\utilities directory.

  Enable auto enclosure by typing `Updateautoenclosure.bat dbusername dbpassword enable` and press **Enter**.

Disable auto enclosure by typing `Updateautoenclosure.bat dbusername dbpassword disable` and press **Enter**.

- On UNIX systems, complete the following steps.

  Open an SSH/Telnet session and navigate to the *Install_Home*\utilities directory.

  Enable auto enclosure by typing `Updateautoenclosure dbusername dbpassword enable` and press **Enter**.

  Disable auto enclosure by typing `Updateautoenclosure dbusername dbpassword disable` and press **Enter**.

NOTE
Professional Plus edition can discover, but not manage the Backbone chassis. Use the device's Element Manager, which can be launched from the Connectivity Map, to manage the device. This device cannot be used as a Seed switch.

FCS policy and seed switches

The Management application requires that the seed switch is the primary Fabric Configuration Server (FCS) switch at the time of discovery.

Setting time on the fabric will set the time on the primary FCS switch, which will then distribute the changes to other switches.

When FCS Policy is defined, **ConfigDownload** is allowed only from the primary FCS switch, but Management application does not check at the time of download that the switch is the primary FCS Switch.

NOTE
Switches running in Access Gateway mode cannot be used as the seed switch.

NOTE
The Backbone Chassis cannot be used as seed switch to discover and manage edge fabrics. You must discover a seed switch from each edge fabric to discover and manage the edge fabric.

NOTE
The Backbone Chassis can only discover and manage the backbone fabric.

Backbone Chassis discovery requirements

Table 10 details which Backbone Chassis models can be discovered by each version of the Management application and whether or not the model can be discovered as a seed switch or only as a member switch.

TABLE 10 Backbone Chassis discovery

Device	Professional	Professional Plus	Enterprise
8-slot Backbone Chassis as seed switch	No	No	Yes
8-slot Backbone Chassis as member switch	Yes for discovery; however, it cannot be managed.	Yes for discovery; however, it cannot be managed.	Yes
4-slot Backbone Chassis as seed switch	Yes	Yes	Yes
4-slot Backbone Chassis as member switch	Yes	Yes	Yes
16 Gbps 8-slot Backbone Chassis as seed switch	No	No	Yes
16 Gbps 8-slot Backbone Chassis as member switch	Yes for discovery; however, it cannot be managed.	Yes for discovery; however, it cannot be managed.	Yes

TABLE 10 Backbone Chassis discovery

Device	Professional	Professional Plus	Enterprise
16 Gbps 4-slot Backbone Chassis as seed switch	Yes	Yes	Yes
16 Gbps 4-slot Backbone Chassis as member switch	Yes	Yes	Yes
32 Gbps, 8-slot Backbone Chassis as seed switch	No	No	Yes
32 Gbps, 8-slot Backbone Chassis as member switch	Yes for discovery; however, it cannot be managed.	Yes for discovery; however, it cannot be managed.	Yes
32 Gbps, 4-slot Backbone Chassis as seed switch	Yes	Yes	Yes
32 Gbps, 4-slot Backbone Chassis as member switch	Yes	Yes	Yes

Discovering fabrics

NOTE
Fabric OS devices must be running Fabric OS 7.0 or later.

NOTE
Only one copy of the application should be used to monitor and manage the same devices in a subnet.

NOTE
When accessing additional data from the **SAN Inventory** or **SAN Status** widgets, it takes a few moments to populate newly discovered products in the **SAN Products - *Status*** dialog box (where ***Status*** is the section of the widget you selected).

To discover specific IP addresses or subnets, complete the following steps.

1. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays.

FIGURE 8 Discover Fabrics dialog box

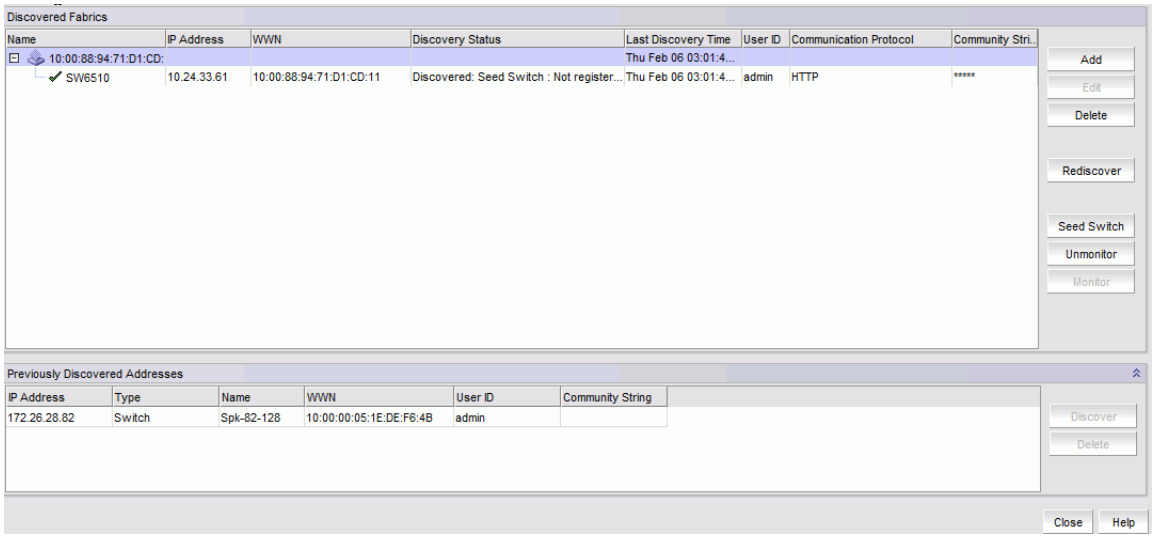

2. Click **Add** to specify the IP addresses of the devices you want to discover.

   The **Add Fabric Discovery** dialog box displays.

   An error message "Discovery Failed. Fabric is busy, try again after sometime." displays when the switch is busy. It is not recommended to continue with the other operations as the Management Application will not receive any updates from the fabric unless it is discovered. Refresh the Management Application and try again to discover the Fabric. The switch property dialog will display the discovered state as Busy when the switch is in busy state and already discovered. A Master log event will be displayed if you try to rediscover the switch in busy state.

FIGURE 9 Add Fabric Discovery dialog box (IP Address tab)

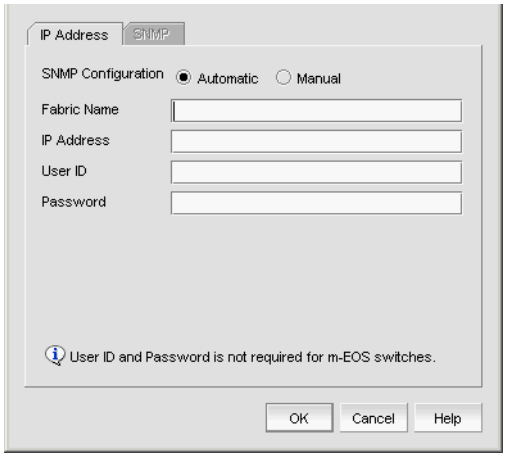

3. Enter a name for the fabric in the **Fabric Name** field.

4. Enter an IP address (IPv4 or IPv6) for a device in the **IP Address** field.

   To configure the preferred IP format for the Management application server to connect with Fabric OS devices, refer to "Configuring the preferred IP format" on page 122. If the product has both an IPv4 and IPv6 address, the Management server uses the preferred address. If a product does not have the preferred address type, the Management server uses the other IP type.

   For seed switch requirements, refer to "Seed switch requirements" on page 50.

   NOTE
   The Backbone Chassis cannot be used as seed switch to discover and manage edge fabrics. You must discover a seed switch from each edge fabric to discover and manage the edge fabric.

   NOTE
   The Backbone Chassis can only discover and manage the backbone fabric.

   NOTE
   Professional and Professional Plus editions cannot manage the Backbone Chassis.

   NOTE
   For Admin Domain (AD) devices, you must enable the AD configuration on the switch before discovery; otherwise, end devices associated with the user-configure AD display as missing in the topology. In addition, the Fabric OS switch must have Physical AD visibility.

   For Virtual Fabric discovery device requirements, refer to "Virtual Fabrics requirements" on page 604.

   To discover a Virtual Fabric device, you must have the following permissions:

   - Switch user account with Chassis Admin role permission on the physical chassis.
   - Switch and SNMPv3 user account with access rights to all logical switches (all Fabric IDs (1 - 128).

     For information about configuring permissions on a Fabric OS device, refer to the ***Fabric OS Administrator's Guide***.:

5. (Fabric OS devices only) Enter the user ID and password for the switch in the **User ID** and **Password** fields.
6. Choose one of the following options:

   - Select the **Automatic** option to use the default SNMPv3 profile.

     The default SNMPv3 profile uses the following attributes:

Attribute	Value
Timeout	5 seconds
Retries	3
User name	snmpadmin1
Context name	None
Auth Protocol	None
Priv Protocol	None

   - Select the **Manual** option to configure SNMP and complete the following steps.

a. Click the **SNMP** tab.

FIGURE 10 Add Fabric Discovery dialog box (SNMP - v1 tab)

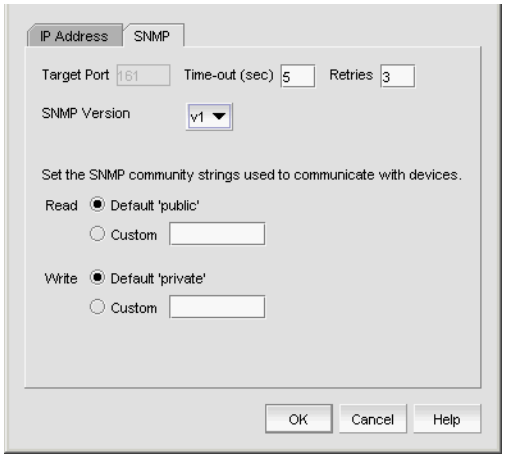

b. Enter the duration (in seconds) after which the application times out in the **Time-out (sec)** field.

c. Enter the number of times to retry the process in the **Retries** field.

d. Select the SNMP version from the **SNMP Version** list.

- If you selected v1, continue with step e.
- If you select v3, the SNMP tab displays the v3 required parameters. Go to step i.

  To discover a Fabric OS device (not virtual fabric-capable), you must provide the existing SNMPv3 username present in the switch.

  To discover a Virtual Fabric device, you must configure SNMPv3 and your SNMP v3 user account must be defined as a Fabric OS switch user.

  When you discovers Virtual Fabric-enabled switch with the SNMPv3 username “admin”, which is the same as the Fabric OS switch user, the Management application automatically creates an SNMP username “admin” in the switch by replacing the sixth username.

e. Specify the **Read** option by selecting **Default ‘public’** or **Custom**.

f. If you selected **Custom**, enter the community string in the **Custom** and **Confirm Custom** fields.

g. Specify the **Write** option by selecting **Default ‘private’** or **Custom**.

h. If you selected **Custom**, enter the community string in the **Custom** and **Confirm Custom** fields.

Go to step 7.

i. If you are configuring a 256-port director, select the **Configure for** *256-Port_Director_Name* check box.

- If you selected **Configure for** *256-Port_Director_Name*, go to step m.
- If you did not select **Configure for** *256-Port_Director_Name*, continue with step j.

j. Enter a user name in the **User Name** field.

k. Enter a context name In the **Context Name** field.

l. Select the authorization protocol in the **Auth Protocol** field.

m. Enter the authorization password in the **Auth Password** field.

- If you selected **Configure for** *256-Port_Director_Name*, go to step 7.
- If you did not select **Configure for** *256-Port_Director_Name*, continue with step n.

n. Select the privacy protocol in the **Priv Protocol** field.

o. Enter the privacy password in the **Priv Password** field.

7. Click **OK** on the **Add Fabric Discovery** dialog box.

   If the seed switch is partitioned, the **Undiscovered Seed Switches** dialog box displays.

   a. Select the **Select** check box for each undiscovered seed switch to discover their fabrics.

   b. Click **OK** on the **Undiscovered Seed Switches** dialog box.

8. Repeat step 2 through step 7 for each fabric you want to discover.

9. Click **Close** on the **Discover Fabrics** dialog box.

Editing the password for multiple devices

You can only edit password for Fabric OS devices in the same fabric.

To edit the password for multiple devices within the same fabric, complete the following steps.

1. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select multiple devices within the same fabric from the **Discovered Fabrics** table.

3. Click **Edit**.

   The *Fabric_Name* **Edit Switches** dialog box displays.

   FIGURE 11 Edit Switches dialog box

4. Enter the user ID for the switch in the **User ID** field.

5. Enter the password for the switch in the **Password** field.

6. Click **OK**. on the *Fabric_Name* **Edit Switches** dialog box.

   The **Credential Update Status** dialog box displays. This dialog box displays the status of the change on the selected devices. If you selected a logical switch, the updated credentials will be applied to the other logical switches in the same chassis.

   - **IP Address** — The IP address of the device.
   - **WWN** — The world wide name of the device.
   - **Name** — The name of the device.
   - **FID** — The fabric ID of the logical switch.

- **Fabric Name** — The name of the fabric where device is located.
- **Status** — The status of the update (such as Success, Failed, or Not Applicable).
- **Reason** — The reason for the status for Failed or Not Applicable.
  - Failed — Not Reachable
  - Not Applicable — Credentials not applied

7. Click **Close**. on the **Credential Update Status** dialog box.

Configuring SNMP credentials

1. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select an IP address from the **Discovered Fabrics** table.
3. Click **Edit**.

   The **Add Fabric Discovery** dialog box displays.

4. To revert to the default SNMPv3 settings, click the **Automatic** option. Go to step 19.
5. To manually configure SNMP, select the **Manual** option. Go to step 6.
6. Click the **SNMP** tab.

   FIGURE 12 Add Fabric Discovery dialog box (SNMP tab)

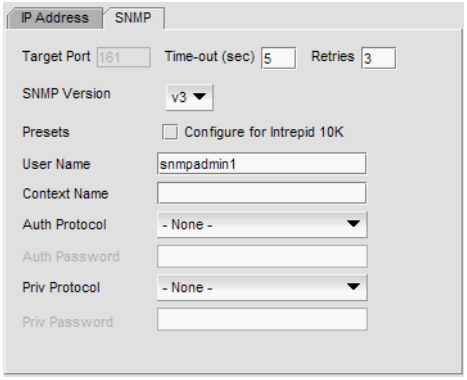

7. Select the SNMP version from the **SNMP Version** list.
   - If you selected v1, continue with step 8.
   - If you select v3, the **SNMP** tab displays the v3 required parameters. Go to step 12.

     To discover a Virtual Fabric device, you must configure SNMPv3 and your SNMP v3 user account must be defined as a Fabric OS switch user.

8. Specify the **Read** option by selecting **Default 'public'** or **Custom**.
9. If you selected **Custom**, enter the community string in the **Custom** and **Confirm Custom** fields.
10. Specify the **Write** option by selecting **Default 'private'** or **Custom**.
11. If you selected **Custom**, enter the community string in the **Custom** and **Confirm Custom** fields.

12. If you are configuring a 256-Port director, select the **Configure for** *256-Port_Director_Name* check box.
    - If you selected **Configure for** *256-Port_Director_Name*, go to step 16.
    - If you did not select **Configure for** *256-Port_Director_Name*, continue with step 13.
13. Enter a user name in the **User Name** field.
14. Enter a context name In the **Context Name** field.
15. Select the authorization protocol in the **Auth Protocol** field.
16. Enter the authorization password in the **Auth Password** field.
    - If you selected **Configure for** *256-Port_Director_Name*, go to step 19.
    - If you did not select **Configure for** *256-Port_Director_Name*, continue with step 17.
17. Select the privacy protocol in the **Priv Protocol** field.
18. Enter the privacy password in the **Priv Password** field.
19. Click **OK** on the **Add Fabric Discovery** dialog box.

    If the seed switch is not partitioned, continue with step 20.

    If the seed switch is partitioned, the **Undiscovered Seed Switches** dialog box displays.

    a. Select the **Select** check box for each undiscovered seed switch to discover their fabrics.
    b. Click **OK** on the **Undiscovered Seed Switches** dialog box.
20. Click **Close** on the **Discover Fabrics** dialog box.

Reverting to a default SNMP community string

To revert to the default SNMP parameters, complete the following steps.

1. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays.
2. Select an IP address from the **Discovered Fabrics** table.
3. Click **Edit**.

   The **Add Fabric Discovery** dialog box displays.
4. Select the **Automatic** option.
5. Click **OK** on the **Add Fabric Discovery** dialog box.
6. Click **Close** on the **Discover Fabrics** dialog box.

Rediscovering a fabric

To refresh discovery of a fabric, complete the following steps.

1. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays.
2. Select a fabric in the **Discovered Fabrics** table.

3. Click **Rediscover**.

   The application triggers all fabric and switch level collectors. The status of the refresh displays in the Master Log as an application event for the fabric as well as each switch in the fabric. For example, "Fabric information collection was successful for the fabric - *Fabric_Name*".

4. Click **Close** on the **Discover Fabrics** dialog box.

Removing a fabric from active discovery

If you decide you no longer want the Management application to discover and monitor a specific fabric, you can delete it from active discovery. Deleting a fabric also deletes the fabric data on the server (both system collected and user-defined data) except for user-assigned names for the device port, device node, and device enclosure information.

To delete a fabric from active discovery, complete the following steps.

1. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select the fabric you want to delete from active discovery in the **Discovered Fabrics** table.
3. Click **Delete**.
4. Click **OK** on the confirmation message.

   The deleted fabric displays in the **Previously Discovered Addresses** table.

5. Click **Close** on the **Discover Fabrics** dialog box.

Rediscovering a previously discovered fabric

To return a fabric to active discovery, complete the following steps.

1. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select the fabric you want to return to active discovery in the **Previously Discovered Addresses** table.
3. Click **Discover**.
4. Click **OK** on the confirmation message.

   The rediscovered fabric displays in the **Discovered Fabrics** table.

5. Click **Close** on the **Discover Fabrics** dialog box.

Deleting a fabric

To delete a fabric permanently from discovery, complete the following steps.

1. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select one or more switches that you want to delete permanently from discovery in the **Previously Discovered Addresses** table.

3. Click **Delete**.
4. Click **OK** on the confirmation message.
5. Click **Close** on the **Discover Fabrics** dialog box.

Viewing the fabric discovery state

The Management application enables you to view device status through the **Discover Setup** dialog box.

To view the discovery status of a device, complete the following steps.

1. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Right-click a fabric and select **Expand All** to show all devices in the fabric.

   The **Name** field displays the discovery status icons in front of the device name. The following table illustrates and describes the icons that indicate the current status of the discovered devices.

   TABLE 11 Discovery Status Icons

Icon	Description
	Displays when the fabric or host is managed and the management status is okay.
	Displays when the switch is managed and the switch management status is not okay.
	Displays when the fabric, switch, or host is not managed or not monitored.

   The **Discovery Status** field details the actual status message text, which varies depending on the situation. The following are samples of actual status messages:

   - Discovered: Seed Switch: Not registered for SNMP Traps
   - Discovered: Seed Switch: Not Manageable: Not registered for SNMP Traps
   - Discovered: Current seed switch is not recommended. Change Seed Switch. : Seed Switch: Not registered for SNMP Traps
   - New Discovery Pending

Troubleshooting fabric discovery

If you encounter discovery problems, complete the following checklist to ensure that discovery was set up correctly.

1. Verify IP connectivity by issuing a ping command to the switch.
   a. Open the command prompt.
   b. From the Server, type `ping` *Switch_IP_Address*.
2. Enter the IP address of the device in a browser to verify the http reachability.

   For example, *http://10.1.1.11*.

Managed count exceeded troubleshooting

The following section states possible issues and the recommended solution when you exceed your managed count limits.

Problem	Resolution
If you exceed your managed count limit, the Management application displays a "licensed exceeded" message on the topology.	Perform one or more of the following actions to • "Changing your network size" • "Remove a device from active discovery" • "Deleting a fabric" **Changing your network size** If you are at the maximum network size for your license, contact your preferred network provider. To change the size of your network, complete the following steps. 1 Select **Server > Options**. The **Options** dialog box displays. 2 Select **Memory Allocation** in the **Category** list to change the network size. 3 Select the size of the SAN (small, medium, or large) you need. 4 Click **OK** on the confirmation message. 5 Click **Apply** or **OK** to save your work. **NOTE:** Changes to this option take effect after an application restart. **NOTE:** You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**). 6 Click **OK** on the "changes take effect after application restart" message. **Remove a device from active discovery** To remove a fabric from active discovery, complete the following steps. 1 Select **Discover > Fabrics**. The managed count exceeded message displays. Managed counts that have been exceeded display with a light red background. Managed counts that are within the grace count limit display with a pale yellow background. 2 Click **OK** on the message. The **Discover Fabrics** dialog box displays. 3 Select the fabric you want to delete from active discovery in the **Discovered Fabrics** table. 4 Click **Delete**. 5 Click **OK** on the confirmation message. The deleted fabric displays in the **Previously Discovered Addresses** table. 6 Click **Close** on the **Discover Fabrics** dialog box. **Deleting a fabric** Before you can delete a fabric permanently from discovery, you must remove it from active discovery. Refer to "Remove a device from active discovery". To delete a fabric permanently from discovery, complete the following steps. 1 Select **Discover > Fabrics**. The managed count exceeded message displays. Managed counts that have been exceeded display with a light red background. Managed counts that are within the grace count limit display with a pale yellow background. 2 Click **OK** on the message. The **Discover Fabrics** dialog box displays. 3 Select one or more switches that you want to delete permanently from discovery in the **Previously Discovered Addresses** table. 4 Click **Delete**. 5 Click **OK** on the confirmation message. 6 Click **Close** on the **Discover Fabrics** dialog box.

Virtual Fabric discovery troubleshooting

The following section state possible issues and the recommended solutions for Virtual Fabric discovery errors.

Problem	Resolution
At the time of discovery, the seed switch is Virtual Fabric-enabled; however, the user does not have Chassis Admin role for the seed switch. At the time of discovery, the user does not have the Chassis Admin role for all other switches in the fabric. After discovery, a device is upgraded to Fabric OS 7.0 or later and is Virtual Fabric-enabled; however, the user does not have Chassis Admin role.	Make sure the user account has Chassis Admin role on the Fabric OS device.
At the time of discovery, the seed switch is Virtual Fabric-enabled; however, the user does not have access to all possible logical switches (access to all possible Fabric IDs 1 - 128). At the time of discovery, the user does not have access to all possible logical switches for all other devices in the fabric. After discovery, a device is upgraded to Fabric OS 7.0 or later and is Virtual Fabric-enabled; however, the user does not have access to all possible logical switches.	Make sure the user account has access rights to all logical switches (access to all possible Fabric IDs 1 - 128) on the Fabric OS device.
At the time of discovery, SNMP v3 is not configured. At the time of discovery, SNMP v3 is not configured for all other switches in the fabric. After discovery, a device is upgraded to Fabric OS 7.0 or later and is Virtual Fabric-enabled; however, SNMP v3 is not configured	Configure the SNMP v3 information for the Virtual Fabric-enabled device.
At the time of discovery or fabric refresh, the SNMP v3 user account does not have the Chassis Admin role.	Make sure the SNMP v3 user account has the Chassis Admin role on the Fabric OS device.
At the time of discovery or refresh, the SNMP v3 user account does not have access to all possible logical switches (access to all possible Fabric IDs 1 - 128). This access is required to obtain performance statistics from all logical switches.	Make sure the SNMP v3 user account has access rights to all logical switches (access to all possible Fabric IDs 1 - 128) on the Fabric OS device.
At the time of discovery or fabric refresh, the SNMP v3 user account does not have a matching Fabric OS switch user account. This is required to obtain performance statistics from all logical switches.	Make sure the SNMP v3 user account is also defined as a Fabric OS switch user.
At the time of fabric refresh, the physical chassis is reachable; however, a previously discovered logical switch is not reachable.	The logical switch has been deleted or the Fabric ID was changed. To find a logical switch, right-click the physical chassis within the **Chassis Group** in the **Product List** and select **Logical Switches**. All logical switches on the selected physical chassis display in a list.

SAN Fabric monitoring

NOTE
Monitoring is not supported on Hosts. The upper limit to the number of HBA and CNA ports that can be monitored at the same time is 32. The same upper limit applies if switch ports and HBA ports are combined. You can select switch ports and adapter ports from a maximum of ten devices.

Fabric monitoring enables discovery of and data collection for the specified fabric and all associated devices. The Management application enables you to view fabric monitoring status through the **Discover Fabrics** dialog box. The following table illustrates and describes the icons that indicate the current status of the discovered switches.

TABLE 12 Monitor Icons

Icon	Description
	Displays when the switch is managed and the switch management status is okay.
	Displays when the switch is managed and the switch management status is not okay.
	Displays when the fabric or switch is not managed or not monitored.

For Professional and Professional Plus, the default monitoring interval is 120 seconds (minimum interval is 120 seconds).

Table 6 details the default and minimum monitoring intervals used to query the monitored switches:

TABLE 13 Monitor Intervals

SAN Size	Default	Minimum
Small	120 seconds (2 minutes)	60 seconds (1 minute)
Medium	900 seconds (15 minutes)	120 seconds (2 minutes)
Large	1800 seconds (30 minutes)	180 seconds (3 minutes)

To change the monitoring interval, refer to "Configuring asset polling" on page 119.

Stop monitoring of discovered fabrics

NOTE
Monitoring is not supported on Hosts.

When you stop monitoring a fabric, the Management application performs the following actions:

- Stops all data collection for the fabric and all associated devices.
- Unregisters as SNMP trap recipient from the fabric and all associated devices.
- Unregisters as SYSLOG recipient from the fabric and all associated devices.
- Does not perform any scheduled or on demand operations (other than monitor) on the fabric and all associated devices.

- Removes the fabric and all associated devices from product list, topology, and all feature dialog boxes.
- Displays the fabric and all associated devices in the Discovery Fabrics dialog box with the unmonitored icon and prefixes “Unmonitored” to the discovery status

To stop monitoring a fabric and all associated devices, complete the following steps.

1. Select **Discovery** > **Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select the fabric you want to stop monitoring from the **Discovered Fabrics** table.
3. Click **Unmonitor**.
4. Click **Close** on the **Discover Fabrics** dialog box.

Stop monitoring of discovered switches

NOTE
You cannot stop monitoring the seed switch.

When you stop monitoring a switch, the Management application performs the following actions:

- Stops all data collection for the switch.
- Unregisters as SNMP trap recipient from the switch. For Virtual Fabric switches, only unregister as SNMP trap recipient when all Virtual Fabric switches of that chassis are unmonitored.
- Unregisters as SYSLOG recipient from the switch. For Virtual Fabric switches, only unregister as SYSLOG recipient when all Virtual Fabric switches of that chassis are unmonitored.
- Does not perform any scheduled or on demand operations (other than monitor) on the switch.
- Removes the switch from product list, topology, and all feature dialog boxes.
- Displays the switch in the Discovery Fabrics dialog box with the unmonitored icon and prefixes “Unmonitored” to the discovery status.

The following details the behavior that occurs when you unmonitor a switch:

- If you unmonitor a switch, the switch does not display in the topology, but end devices connected to the switch continue to display in the product list and topology (with no connections).
- If you segment an unmonitored switch, you cannot discover it separately until you accept changes in the original fabric.
- If you unmonitor a switch in Access Gateway mode, that switch is unmonitored from all fabrics in which it is participating.
- If you unmonitor a Virtual Fabric switch (logical switch in a chassis), only that partition is unmonitored, but end devices connected to the Virtual Fabric switch continue to display in the product list and topology (with no connections). Any other partitions of the associated chassis continue to be monitored.
- If fabric tracking is enabled and you unmonitor a switch, fabric tracking continues to track the unmonitored switch.
- If fabric tracking is enabled and the unmonitored switch segments out of the fabric, the switch is marked as “missing” in the **Accept Changes** dialog box. If an ISL connected to this switch is disconnected, the ISL is also marked as “missing” in the in the **Accept Changes** dialog box. If a device connected to this switch is disconnected, the device is also marked as “missing” in the product list and topology.
- If fabric tracking is enabled for two managed fabrics and you move an unmonitored switch from one fabric to the other, the unmonitored switches is marked as “missing” in the original fabric and marked as “untrusted” in the new fabric in the **Accept Changes** dialog box.

To stop monitoring a switch, complete the following steps.

1. Select **Discovery** > **Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select one or more switches in the same fabric that you want to stop monitoring from the **Discovered Fabrics** table.

   NOTE
   You cannot select switches in different fabrics.

3. Click **Unmonitor**.

   The **Unmonitor Status** dialog box displays with the following details:

   - **IP Address** — The IP address of the switch.
   - **WWN** — The WWN of the switch.
   - **Name** — The name of the switch.
   - **FID** — The FID of the switch.
   - **Fabric Name** — The name of the associated fabric.
   - **Status** — Whether the unmonitor was successful or failed.
   - **Reason** — The reason for the failure. Blank for success.

4. Click **Close** on the **Unmonitor Status** dialog box.
5. Click **Close** on the **Discover Fabrics** dialog box.

Resume monitoring of discovered fabrics

NOTE
Monitoring is not supported on Hosts.

To monitor a fabric and all associated devices, complete the following steps.

1. Select **Discovery** > **Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select the fabric you want to monitor from the **Discovered Fabrics** table.
3. Click **Monitor**.

   The **Monitor Status** dialog box displays with the status.

   NOTE
   If there is a unmonitored switch in the fabric, it stays unmonitored.

   The monitor function fails if the fabric has user-defined Admin Domains created or if the fabric is merged with another fabric already in the monitored state.

4. Click **Close** on the **Monitor Status** dialog box.
5. Click **Close** on the **Discover Fabrics** dialog box.

Resume monitoring of discovered switches

NOTE
Monitoring is not supported on Hosts.

NOTE
You can only monitor a switch that is reachable and has valid credentials.

To monitor a switch, complete the following steps.

1. Select **Discovery** > **Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select one or more switches that you want to monitor from the **Discovered Fabrics** table.
3. Click **Monitor**.

   The **Monitor Status** dialog box displays with the status.

4. Click **Close** on the **Monitor Status** dialog box.
5. Click **Close** on the **Discover Fabrics** dialog box.

SAN Seed switch

The seed switch must be running a supported Fabric OS version and must be HTTP-reachable.

Sometimes, the seed switch is auto-selected, such as when a fabric segments or when two fabrics merge. Other times, you are prompted (an event is triggered) to change the seed switch, such as in the following cases:

- If, during fabric discovery, the Management application detects that the seed switch is not running a supported version, you are prompted to change the seed switch.
- When one or more switches join the fabric or if the switch firmware is changed on any of the switches in the fabric, the Management application checks to make sure that the seed switch is still running a supported version. If it is not, then you are prompted to either upgrade the firmware on the seed switch or to change the seed switch to a switch running a supported firmware.

If a fabric of switches running only Fabric OS 7.0 or later is created due to segmentation, the Management application continues to monitor that fabric, but if any switch with a later Fabric OS version joins the fabric, an event is triggered informing you that the seed switch is not running the latest firmware and you should change to the seed switch running the highest firmware.

ATTENTION
If a seed switch is segmented or merged, historical data such as offline zone DB, profile and reports, and Firmware Download Profile can be lost. Segmentation of a seed switch does not result in formation of a new fabric. If a merge occurs, the historical data is lost only from the second fabric.

You can change the seed switch as long as the following conditions are met:

- The new seed switch is HTTP-reachable from the Management application.
- The new seed switch is a primary FCS.
- The new seed switch is running the latest Fabric OS version in the fabric.

This operation preserves historical and configuration data, such as performance monitoring and user-customized data for the selected fabric.

If, during the seed switch change, the fabric is deleted, but the rediscovery operation fails (for example, if the new seed switch becomes unreachable using HTTP), then you must rediscover the fabric again. If you rediscover the fabric using a switch that was present in the fabric before the change seed switch operation was performed, then all of the historical and configuration data is restored to the rediscovered fabric. If you rediscover the fabric using a switch that was added to the fabric after the fabric was deleted, then the historical and configuration data is lost.

If multiple users try to change the seed switch of the same fabric simultaneously, only the first change seed switch request is executed; subsequent requests that are initiated before the first request completes will fail.

If another user changes the seed switch of a fabric you are monitoring, and if you have provided login credentials for only that seed switch in the fabric, then you lose connection to the seed switch.

Seed switch requirements

The seed switch must be running Fabric OS 7.0 or later. For a complete list of all supported Fabric OS hardware, refer to "Supported hardware and software" on page lxi.

Seed switch failover

The Management application collects fabric-wide data (such as, fabric membership, connectivity, name server information, zoning, and so on) using the seed switch. Therefore when a seed switch becomes unreachable or there is no valid seed switch, the fabric becomes unmanageable.

When the seed switch cannot be reached for three consecutive fabric refresh cycles, the Management application looks for another valid seed switch in the fabric, verifies that it can be reached, and has valid credentials. If the seed switch meets this criteria, the Management application automatically fails over to the recommended seed switch.

Note that it is possible that auto-failover may occur to a seed switch not running the latest firmware version. In this instance, any functionality which has a direct dependency on the firmware version of the seed switch is affected and restricted by the failover seed switch capabilities.

Changing the seed switch

When you change the seed switch for a fabric, the Management application performs the following checks in the order they are listed:

- Identifies all switches and removes those running unsupported firmware version.
- Identifies which of the remaining switches are running the latest firmware versions.
- Filters out those switches that are not reachable.
- Identifies which switches are Virtual Fabric-enabled switches (Fabric OS only).

  If there are Virtual Fabric-enabled switches, the Management application only uses these switches as recommended seed switches. If there are no Virtual Fabric-enabled switches, continue with the next check.
- Identifies which switches are Virtual Fabric-capable devices (Fabric OS only).

  If there are Virtual Fabric-capable switches, the Management application only uses these switches as recommended seed switches. If there are no Virtual Fabric-capable switches, the Management application uses the list from the second check.

To change the seed switch, complete the following steps.

1. Select **Discovery** > **Fabrics**.

   The **Discover Fabrics** dialog box displays.

2. Select the fabric for which you want to change the seed switch from the **Discovered Fabrics** table.

   If a device joins or merges with a fabric and fabric tracking is active, you must accept changes to the fabric before the new devices display in the **Seed Switch** dialog box. For more information about fabric tracking, refer to "Fabric tracking" on page 130.

3. Click **Seed Switch**.

   If the fabric contains other switches that are running the latest version and are also HTTP-reachable from the Management application, the **Seed Switch** dialog box appears. Otherwise, a message displays that you cannot change the seed switch.

4. Select a switch to be the new seed switch from the **Seed Switch** dialog box.

   You can select only one switch. Only switches that are running the latest Fabric OS version in the fabric are displayed. The current seed switch is not displayed in this list.

5. Click **OK** on the **Seed Switch** dialog box.

   If you are not already logged in to the seed switch, the **Fabric Login** dialog box displays.

   If you are successfully authenticated, the fabric is deleted from the Management application without purging historical data, and the same fabric is rediscovered with the new seed switch.

6. Click **Close** on the **Discover Fabrics** dialog box.

Host discovery

The Management application enables you to discover individual hosts, import a group of hosts from a comma-separated values (CSV) file, or import all hosts from discovered fabrics or VM Managers.

NOTE
Host discovery requires HCM Agent 2.0 or later.

CIM and WMI host discovery requirements

ESXi host adapter discovery requires a vendor-specific HBA CIM provider to be installed on the ESXi host. The Management application supports CIMOM-based discovery for third-party adapters irrespective of the operating system and firmware version. Perform the following steps to configure HTTPS certificate validation for ESXi host discovery.

1. Download the host certificate to the Management application trust store from the following path.

   `/etc/vmware/ssl`

2. Refer to "Certificates" on page 105 to import the host certificate to the Management application server from the Management application. If the certificate is not imported properly, an error message displays.

3. Restart the server after importing the certification properly. The server must be restarted each time a host certificate is configured.

4. Refer to "Discovering Hosts by network address or host name" on page 52 to configure discovery authentication in the **Add Host Adapters** dialog box.

   NOTE
   CIM HTTPS Certificate validation occurs whether the **Enable Certificate Validation** check box is selected or not in the **Server** dialog box.

For Windows, the third-party adapter discovery is based on Windows Management Instrumentation (WMI). Perform the following steps to configure HTTPS certificate validation.

1. Import the host certificate when the **Enable Certificate Validation** check box is selected. Discovery will occur successfully even without importing the certificate when the **Enable Certificate Validation** check box is not selected.
2. Restart the Server Management Console (SMC) whether the **Enable Certificate Validation** check box is selected or not.
3. Refer to "Discovering Hosts by network address or host name" on page 52 to configure discovery authentication in the **Add Host Adapters** dialog box.

Discovering Hosts by network address or host name

To discover a Host by network address or host name, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

   FIGURE 13 Discover Host Adapters dialog box

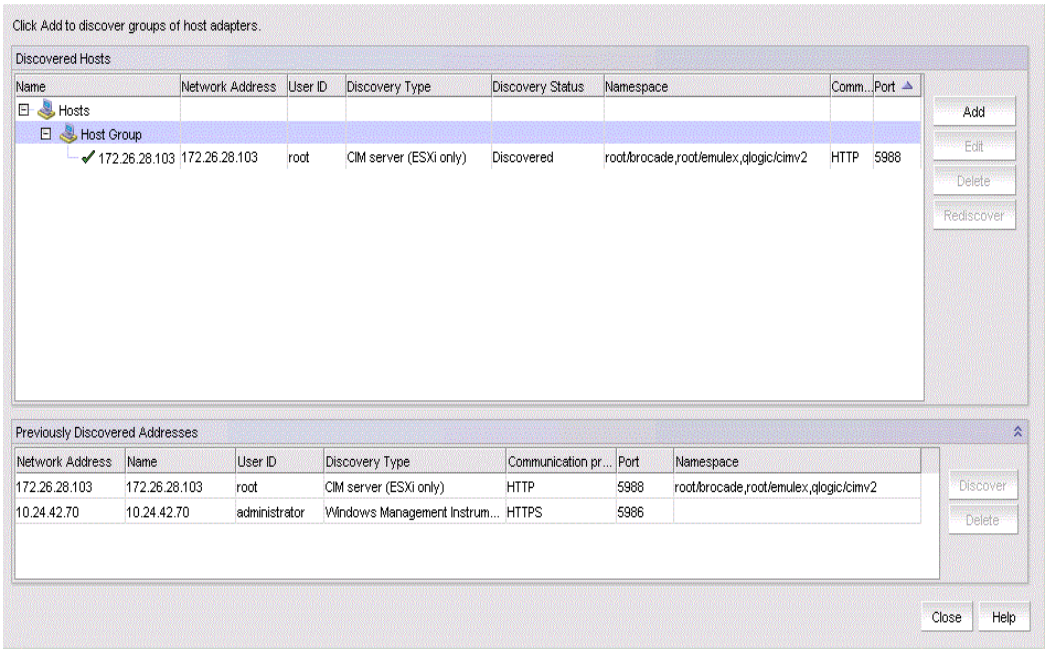

2. Click **Add**.

   The **Add Host Adapters** dialog box displays.

FIGURE 14 Add Host Adapters dialog box

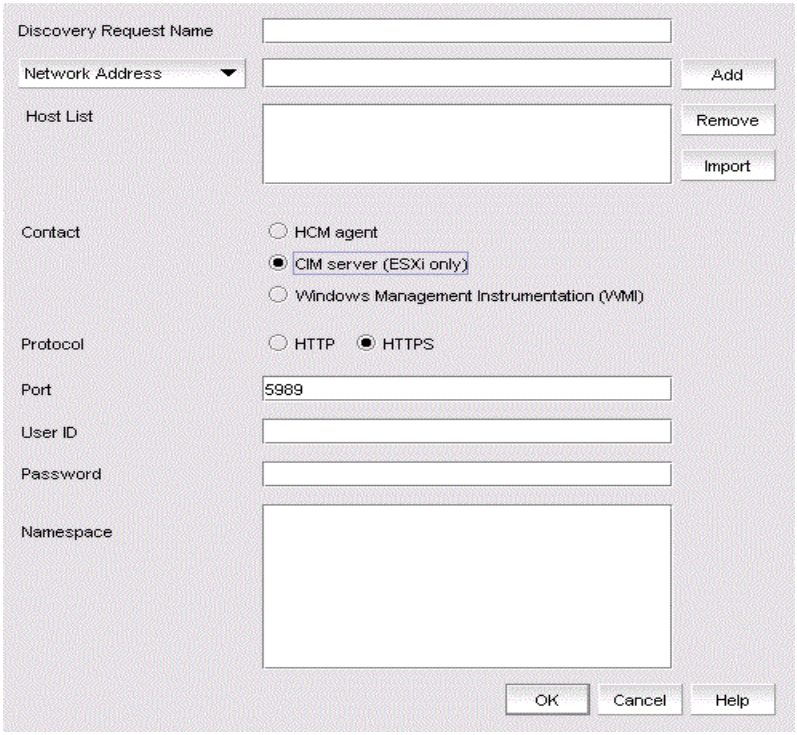

3. (Optional) Enter a discovery request name (such as Manual 06/12/2009) in the **Discovery Request Name** field.
4. Select **Network Address** from the list.
5. Enter the IP address (IPv4 or IPv6 formats) or host name in the **Network Address** field.
6. Click **Add**.

   The IP address or host name of the Host displays in the **Host List**.
7. Configure Host credentials by choosing one of the following options:
   - To configure HCM agent credentials, select the **HCM agent** option. Go to step 9.
   - To configure CIM server credentials, select the **CIM server (ESXi only)** option. Continue with step 8.
   - To configure WMI server credentials, select the **Windows Management Instrumentation (WMI)** option. Continue with step 8.

   If you do not need to configure Host credentials, skip to step 14.
8. Configure discovery authentication by choosing one of the following options:
   - To configure discovery with authentication, select the **HTTPS** option in **Protocol**.
   - To configure discovery without authentication, select the **HTTP** option in **Protocol**.
9. Enter the port number in the **Port** field.

   The HCM agent default is 34568. The CIM server HTTPS default is 5989. The CIM server HTTP default is 5988. The WMI HTTPS default is 5986. The WMI HTTP default is 5985.

10. Enter your user name in the **User ID** field.

    The HCM agent default is admin. For CIM/WMI server, enter the user ID used during the CIM/WMI configuration, otherwise leave this field blank.

11. Enter your password in the **Password** field.

    The HCM agent default is password. For CIM/WMI server, enter the password used during the CIM/WMI configuration, otherwise leave this field blank.

12. Enter a namespace in the **Namespace** field. Enter namespaces separated by commas if multiple adapters are connected to a host.

    The default namespace is root/brocade, root/emulex, and qlogic/cimv2.

13. Repeat step 5 through step 12 for each Host you want to discover.

14. Click **OK** on the **Add Host Adapters** dialog box.

    If an error occurs, a message displays. Click **OK** to close the error message and fix the problem.

    A host group displays in the **Discovered Hosts** table of the **Discover Host Adapters** dialog box. The discovery status is updated dynamically.

15. Click **Close** on the **Discover Host Adapters** dialog box.

Importing Hosts from a CSV file

To discover Hosts by importing a CSV file, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

2. Click **Add**.

   The **Add Host Adapters** dialog box displays.

FIGURE 15 Add Host Adapters dialog box

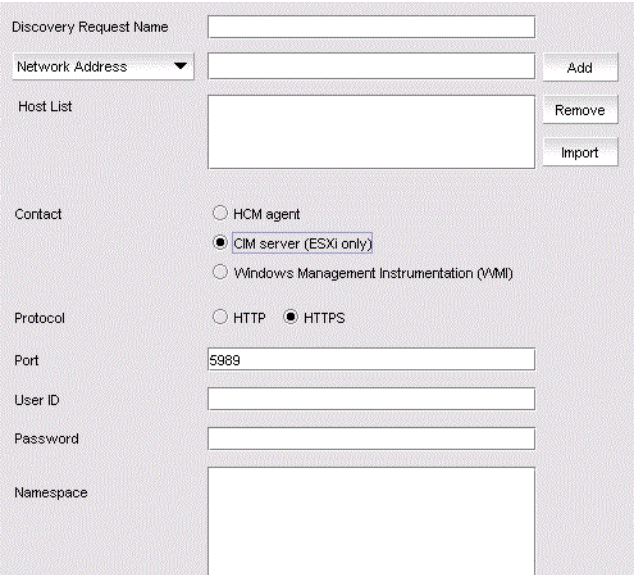

3. Enter a discovery request name (such as MyFabric) in the **Discovery Request Name** field.
4. Click **Import**.

   The **Open** dialog box displays.
5. Browse to the CSV file location.

   The CSV file must meet the following requirements:

   - Comma-separated IP addresses or host names
   - No commas within the values
   - No escaping supported

     For example, XX.XX.XXX.XXX, XX.XX.X.XXX, computername.company.com
6. Click **Open**.

   The CSV file is imported to the **Add Host Adapters** dialog box. During import, duplicate values are automatically dropped. When import is complete, the imported values display in the **Host List**. If the file cannot be imported, an error displays.
7. Verify the imported values in the **Host List**.
8. Configure Host credentials by choosing one of the following options:

   - To configure HCM agent credentials, select the **HCM agent** option. Go to step 10.
   - To configure CIM server credentials, select the **CIM server (ESXi only)** option. Continue with step 9.
   - To configure WMI server credentials, select the **Windows Management Instrumentation (WMI)** option. Continue with step 9.

   If you do not need to configure Host credentials, skip to step 14.
9. Configure discovery authentication by choosing one of the following options:

   - To configure discovery with authentication, select the **HTTPS** option in **Protocol**.
   - To configure discovery without authentication, select the **HTTP** option in **Protocol**.

10. Enter the port number in the **Port** field.

    The HCM agent default is 34568. The CIM server HTTPS default is 5989. The CIM server HTTP default is 5988. The WMI HTTPS default is 5986. The WMI HTTP default is 5985.

11. Enter your user name in the **User ID** field.

    The HCM agent default is admin. For CIM/WMI server, enter the user ID used during the CIM/WMI configuration, otherwise leave this field blank.

12. Enter your password in the **Password** field.

    The HCM agent default is password. For CIM/WMI server, enter the password used during the CIM/WMI configuration, otherwise leave this field blank.

13. Enter a namespace in the **Namespace** field. Enter namespaces separated by commas if multiple adapters are connected to a host.

    The default namespace is root/brocade, root/emulex, and qlogic/cimv2.

14. Click **OK** on the **Add Host Adapters** dialog box.

    If an error occurs, a message displays. Click **OK** to close the error message and fix the problem.

    A host group displays in the **Discovered Hosts** table of the **Discover Host Adapters** dialog box. The discovery status is updated dynamically.

15. Click **Close** on the **Discover Host Adapters** dialog box.

Importing Hosts from a fabric

To discover a Host from a discovered fabric, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

2. Click **Add**.

   The **Add Host Adapters** dialog box displays.

FIGURE 16 Add Host Adapters dialog box

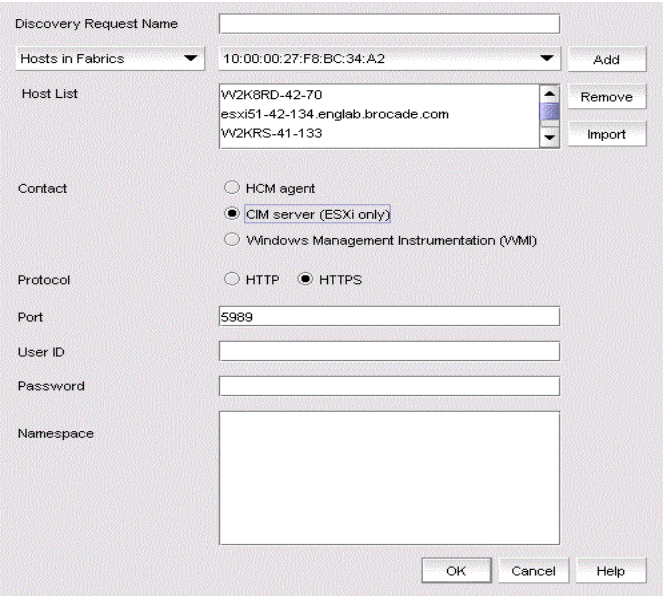

3. Enter a discovery request name (such as MyFabric) in the **Discovery Request Name** field.
4. Select **Hosts in Fabrics** from the list.
5. Select **All fabrics** or an individual fabric from the host list.
6. Click **Add**.

   All hosts that are part of a managed fabric and have a registered host name display in the host list. If no host with a registered host name exists, an error message displays. Click **OK** to close the error message.

7. Configure Host credentials by choosing one of the following options:

   - To configure HCM agent credentials, select the **HCM agent** option. Go to step 9.
   - To configure CIM server credentials, select the **CIM server (ESXi only)** option. Continue with step 8.
   - To configure WMI server credentials, select the **Windows Management Instrumentation (WMI)** option. Continue with step 8.

   If you do not need to configure Host credentials, skip to step 13.

8. Configure discovery authentication by choosing one of the following options:

   - To configure discovery with authentication, select the **HTTPS** option in **Protocol**
   - To configure discovery without authentication, select the **HTTP** option in **Protocol**.

9. Enter the port number in the **Port** field.

   The HCM agent default is 34568. The CIM server HTTPS default is 5989. The CIM server HTTP default is 5988. The WMI HTTPS default is 5986. The WMI HTTP default is 5985.

10. Enter your user name in the **User ID** field.

    The HCM agent default is admin. For CIM/WMI server, enter the user ID used during the CIM/WMI configuration, otherwise leave this field blank.

11. Enter your password in the **Password** field.

    The HCM agent default is password. For CIM/WMI server, enter the password used during the CIM/WMI configuration, otherwise leave this field blank.

12. Enter a namespace in the **Namespace** field. Enter namespaces separated by commas if multiple adapters are connected to a host.

    The default namespace is root/brocade, root/emulex, and qlogic/cimv2.

13. Click **OK** on the **Add Host Adapters** dialog box.

    If an error occurs, a message displays. Click **OK** to close the error message and fix the problem.

    A host group displays in the **Discovered Hosts** table of the **Discover Host Adapters** dialog box. The discovery status is updated dynamically.

14. Click **Close** on the **Discover Host Adapters** dialog box.

Importing Hosts from a VM Manager

To discover Hosts from a discovered VM Manager, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

2. Click **Add**.

   The **Add Host Adapters** dialog box displays.

   FIGURE 17 Add Host Adapters dialog box

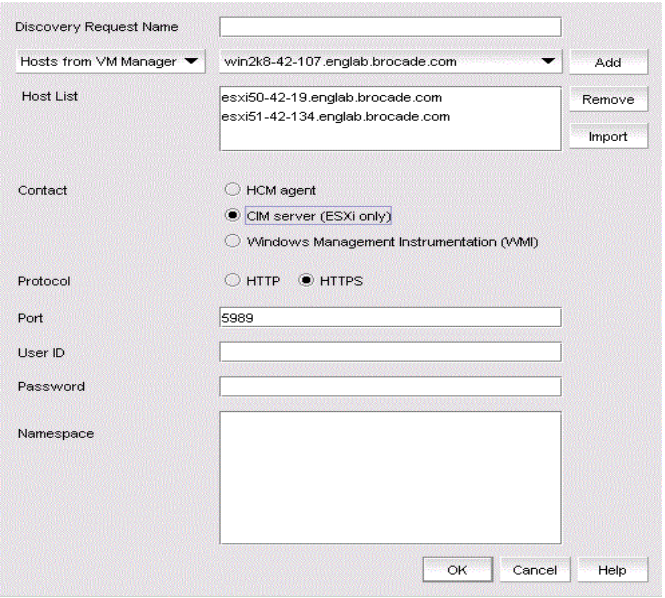

3. Enter a discovery request name (such as MyVMManager) in the **Discovery Request Name** field.
4. Select **Hosts from VM Manager** from the list.
5. Select **All VM** or an individual VM from the list.

6. Click **Add**.

   All hosts that are part of a discovered VM Manager and have a registered host name display in the list. If no host with a registered host name exists, an error message displays. Click **OK** to close the error message.

7. Configure Host credentials by choosing one of the following options:

   - To configure HCM agent credentials, select the **HCM agent** option. Go to step 9.
   - To configure CIM server credentials, select the **CIM server (ESXi only)** option. Continue with step 8.
   - To configure WMI server credentials, select the **Windows Management Instrumentation (WMI)** option. Continue with step 8.

   If you do not need to configure Host credentials, skip to step 13.

8. Configure discovery authentication by choosing one of the following options:

   - To configure discovery with authentication, select the **HTTPS** option in **Protocol**.
   - To configure discovery without authentication, select the **HTTP** option in **Protocol**.

9. Enter the port number in the **Port** field.

   The HCM agent default is 34568. The CIM server HTTPS default is 5989. The CIM server HTTP default is 5988. The WMI HTTPS default is 5986. The WMI HTTP default is 5985.

10. Enter your user name in the **User ID** field.

    The HCM agent default is admin. For CIM/WMI server, enter the user ID used during the CIM/WMI configuration, otherwise leave this field blank.

11. Enter your password in the **Password** field.

    The HCM agent default is password. For CIM/WMI server, enter the password used during the CIM/WMI configuration, otherwise leave this field blank.

12. Enter a namespace in the **Namespace** field. Enter namespaces separated by commas if multiple adapters are connected to a host.

    The default namespace is root/brocade, root/emulex, and qlogic/cimv2.

13. Click **OK** on the **Add Host Adapters** dialog box.

    If an error occurs, a message displays. Click **OK** to close the error message and fix the problem.

    A host group displays in the **Discovered Hosts** table of the **Discover Host Adapters** dialog box. The discovery status is updated dynamically.

14. Click **Close** on the **Discover Host Adapters** dialog box.

Editing host adapter credentials

To edit host credentials, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

2. Select the Host in the **Discovered Hosts** table and click **Edit**.

The **Edit Host Adapters** dialog box displays.

FIGURE 18 Edit Host Adapters dialog box

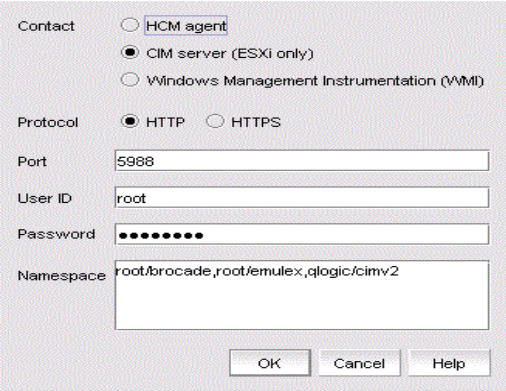

3. Configure Host credentials by choosing one of the following options:
   - To configure HCM agent credentials, select the **HCM agent** option. Go to step 5.
   - To configure CIM server credentials, select the **CIM server (ESXi only)** option. Continue with step 4.
   - To configure WMI server credentials, select the **Windows Management Instrumentation (WMI)** option. Continue with step 4.

   If you do not need to configure Host credentials, skip to step 9.
4. Configure discovery authentication by choosing one of the following options:
   - To configure discovery with authentication, select the **HTTPS** option in **Protocol**.
   - To configure discovery without authentication, select the **HTTP** option in **Protocol**.
5. Enter the port number in the **Port** field.

   The HCM agent default is 34568. The CIM server HTTPS default is 5989. The CIM server HTTP default is 5988. The WMI HTTPS default is 5986. The WMI HTTP default is 5985.
6. Enter your user name in the **User ID** field.

   The HCM agent default is admin. For CIM/WMI server, enter the user ID used during the CIM/WMI configuration, otherwise leave this field blank.
7. Enter your password in the **Password** field.

   The HCM agent default is password. For CIM/WMI server, enter the password used during the CIM/WMI configuration, otherwise leave this field blank.
8. Enter a namespace in the **Namespace** field. Enter namespaces separated by commas if multiple adapters are connected to a host.

   The default namespace is root/brocade, root/emulex, and qlogic/cimv2.
9. Click **OK** on the **Edit Host Adapters** dialog box.

   If an error occurs, a message displays. Click **OK** to close the error message and fix the problem.
10. Click **Close** on the **Discover Host Adapters** dialog box.

Removing a host from active discovery

If you decide you no longer want the Management application to discover and monitor a specific host, you can delete it from active discovery. Deleting a host also deletes the host data on the server (both system-collected and user-defined data) except for user-assigned names for the device port, device node, and device enclosure information.

To delete a host from active discovery, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

2. Select the host you want to delete from active discovery in the **Discovered Hosts** table.
3. Click **Delete**.
4. Click **OK** on the confirmation message.

   The deleted host displays in the **Previously Discovered Addresses** table.

5. Click **Close** on the **Discover Host Adapters** dialog box.

Rediscovering a host to active discovery

To rediscover a host to active discovery, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

2. Select the host you want to rediscover in the **Discovered Hosts** table.
3. Click **Rediscover**.
4. Click **OK** on the confirmation message.

   The rediscovered host displays in the **Discovered Hosts** table. The discovery status is updated dynamically.

5. Click **Close** on the **Discover Host Adapters** dialog box.

Rediscovering a previously discovered host

To return a host to active discovery, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

2. Select the host you want to return to active discovery in the **Previously Discovered Addresses** table. The table displays the Network Address, Name, User ID, Discovery Type, Communication Protocol, and Port details of the selected host.
3. Click **Discover**.
4. Click **OK** on the confirmation message.

   The rediscovered host displays in the **Discovered Hosts** table.

5. Click **Close** on the **Discover Host Adapters** dialog box.

Deleting a host from discovery

To delete a host permanently from discovery, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

2. Select the host you want to delete permanently from discovery in the **Previously Discovered Addresses** table.
3. Click **Delete**.
4. Click **OK** on the confirmation message.
5. Click **Close** on the **Discover Host Adapters** dialog box.

Viewing the host discovery state

The Management application enables you to view device discovery status through the **Discover Host Adapters** dialog box.

To view the discovery status of a device, complete the following steps.

1. Select **Discover > Host Adapters**.

   The **Discover Host Adapters** dialog box displays.

2. Right-click the **Hosts** node and select **Expand All** to show all devices.

   The **Name** field displays the discovery status icons in front of the device name. The following table illustrates and describes the icons that indicate the current status of the discovered devices.

Icon	Description
	Displays when the fabric or host is managed and the management status is okay.
	Displays when the fabric or host is not managed.
	Displays when any operation is pending.
	Displays when the namespace operation is partially failed.

   The **Discovery Status** field details the actual status message text, which varies depending on the situation. The following are samples of actual status messages:

   - CIM Server authentication failed
   - CIM Server connection failed
   - CIM Server unknown failure
   - HCM Agent connection failed
   - HCM Agent authentication failed
   - HCM Agent unknown failure

- WMI authentication failed
- WMI connection failed
- WMI Unknown Error
- Discovery ignored. One or more adapters in the host are already a part of Host group {}
- Discovery ignored. One or more adapters in the host are already a part of auto/manual enclosure {}. Please delete the enclosure and try again.
- Discovered
- Discovering....
- Rediscovering...
- Deleting....
- Partially Failed. Reason: Invalid Namespace <name>
- Failed.Reason: Invalid Namespace <name>
- Internal Error
- No Adapters found

Troubleshooting host discovery

If you encounter discovery problems, complete the following checklist to ensure that discovery was set up correctly. For more complete information about troubleshooting adapters, refer to the *Adapters Troubleshooting Guide*.

1. Verify IP connectivity by issuing a **ping** command to the host.
   a. Open the command prompt.
   b. From the server, type `ping` ***Host_IP_Address*** **.**
2. If the host is responding to ping, but discovery still fails, verify that HCM agent is up or not by browsing to the following URL:

   https://***Host_IP_Address***:34568/JSONRPCServiceApp/JSON-RPC

   If HCM agent is running and reachable, you should receive a prompt of credentials and then show an Error 500 (No Reason) result page.
3. Verify that firewall port 34568 is open.

   There are firewall issues with the HCM Agent on Windows 2008 and VMware systems. When installing the driver package on these systems, open TCP/IP port 34568 to allow agent communication with the Management application.

   - For VMware, use the following commands to open port 34568:
     - `esxcfg-firewall –o 34568,tcp,in,https`
     - `esxcfg-firewall –o 34568,udp,out,https`
   - For Windows, use Windows Firewall and Advanced Service (WFAS) to open port 34568.

VM Manager discovery

The Management application enables you to discover VM managers. VM Manager discovery requires vCenter Server 4.0 or later.

NOTE
vCenter discovery time is dynamically determined based on the number of hosts being managed by the vCenter. For every 50 hosts managed, the vCenter collection period increases 30 minutes.
For 0-50 hosts managed, the collection duration is 30 minutes; for 50-100 hosts managed, the collection duration is one hour, and so on.

VM Manager discovery requirements

- Discovery of a vCenter server (refer to "Discovering a VM manager" on page 64, step 4 and step 5), requires a vCenter user with read-only or read-write privilege on the vCenter server node and all objects in the inventory below the vCenter server.
- Enabling the vSphere client plug-in registration (refer to "Discovering a VM manager" on page 64, step 6), requires a vCenter user with the vCenter server Admin privileges.

Discovering a VM manager

Before you discover a VM Manager, make sure you meet the discovery requirements (refer to "VM Manager discovery requirements" on page 64).

To discover a VM manager, complete the following steps.

1. Select **Discover > VM Managers**.

   The **Discover VM Managers** dialog box displays.

   FIGURE 19 Discover VM Managers dialog box

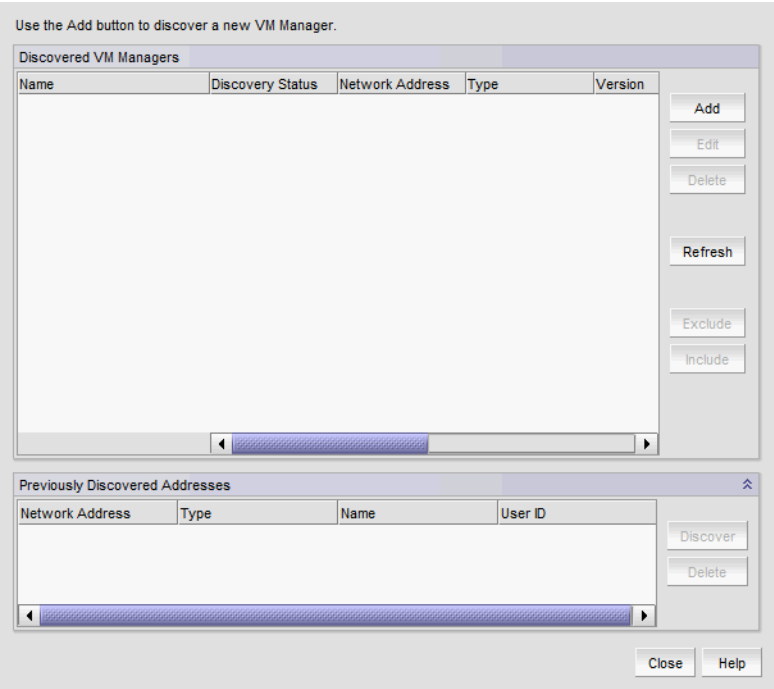

2. Click **Add**.

   The **Add VM Manager** dialog box displays.

FIGURE 20 Add VM Manager dialog box

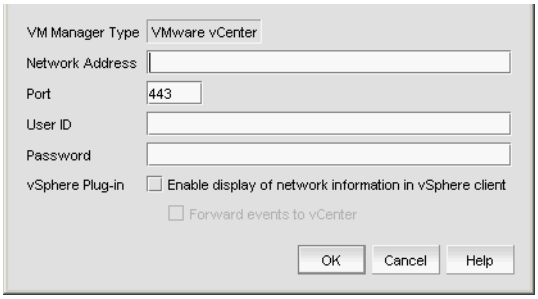

3. Enter the IP address or host name in the **Network Address** field.
4. Enter the VM manager port number in the **Port** field.
5. Enter the VM manager username in the **User ID** field.
6. Enter the VM manager password **Password** field.
7. Select the **Enable display of network information in vSphere client** check box to enable vSphere client plug-in registration.

   Clear to disable vSphere client plug-in registration.
8. Select the **Forward event to vCenter** check box to enable event forwarding from the Management application to vCenter.

   Clear to disable event forwarding.
9. Click **OK** on the **Add VM Manager** dialog box.

   If an error occurs, a message displays. Click **OK** to close the error message and fix the problem.

   A VM manager displays in **Discovered VM Managers** table with pending status. To update the status from pending you must close and reopen the **Discover VM Managers** dialog box.
10. Refresh the **Discover VM Managers** list by clicking **Refresh**.
11. Click **Close** on the **Discover VM Managers** dialog box.

Editing a VM manager

To edit VM manager discovery, complete the following steps.

1. Select **Discover > VM Managers**.

   The **Discover VM Managers** dialog box displays.
2. Select the Host in the **Discovered VM Managers** list and click **Edit**.

   The **Edit VM Manager** dialog box displays.

FIGURE 21 Edit VM Manager dialog box

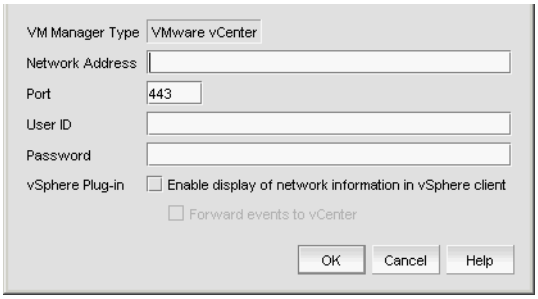

3. Change the VM manager port number in the **Port** field.
4. Enter the VM manager username in the **User ID** field.
5. Enter the VM manager user password **Password** field.
6. Select the **Enable display of network information in vSphere client** check box to enable vSphere client plug-in registration.

   Clear to disable vSphere client plug-in registration.
7. Select the **Forward event to vCenter** check box to enable event forwarding from the Management application to vCenter.

   Clear to disable event forwarding.
8. Click **OK** on the **Edit VM Manager** dialog box.

   If an error occurs, a message displays. Click **OK** to close the error message and fix the problem.
9. Refresh the **Discover VM Managers** list by clicking **Refresh**.
10. Click **Close** on the **Discover VM Managers** dialog box.

Excluding a host from VM manager discovery

To exclude host from VM manager discovery complete the following steps.

1. Select **Discover > VM Managers**.

   The **Discover VM Managers** dialog box displays.
2. Select the Host you want to exclude in the **Discovered VM Managers** list and click **Exclude**.
3. Click **Close** on the **Discover VM Managers** dialog box.

Including a host in VM manager discovery

To include host in VM manager discovery complete the following steps.

1. Select **Discover > VM Managers**.

   The **Discover VM Managers** dialog box displays.
2. Select a Host you want to include in the **Discovered VM Managers** list and click **Include**.
3. Click **Close** on the **Discover VM Managers** dialog box.

Removing a VM manager from active discovery

If you decide you no longer want the Management application to discover and monitor a specific VM manager, you can delete it from active discovery. Deleting a VM manager also deletes the data on the server (both system collected and user-defined data) except for user-assigned names for the device port, device node, and device enclosure information.

To delete a VM manager from active discovery, complete the following steps.

1. Select **Discover > VM Managers**.

   The **Discover VM Managers** dialog box displays.

2. Select the VM manager you want to delete from active discovery in the **Discovered VM Managers** table.
3. Click **Delete**.
4. Click **OK** on the confirmation message.

   The deleted VM manager displays in the **Previously Discovered Addresses** table.

5. Refresh the **Discover VM Managers** list by clicking **Refresh**.
6. Click **Close** on the **Discover VM Managers** dialog box.

Rediscovering a previously discovered VM manager

To return a VM manager to active discovery, complete the following steps.

1. Select **Discover > VM Managers**.

   The **Discover VM Managers** dialog box displays.

2. Select the VM manager you want to return to active discovery in the **Previously Discovered Addresses** table.
3. Click **Discover**.
4. Click **OK** on the confirmation message.

   The rediscovered VM manager displays in the **Discovered VM Managers** table.

5. Refresh the **Discover VM Managers** list by clicking **Refresh**.
6. Click **Close** on the **Discover VM Managers** dialog box.

Deleting a VM manager from discovery

To delete a host permanently from discovery, complete the following steps.

1. Select **Discover > VM Managers**.

   The **Discover VM Managers** dialog box displays.

2. Select the VM manager you want to delete permanently from discovery in the **Previously Discovered Addresses** table.
3. Click **Delete**.
4. Click **OK** on the confirmation message.
5. Refresh the **Discover VM Managers** list by clicking **Refresh**.

6. Click **Close** on the **Discover VM Managers** dialog box.

Viewing the VM manager discovery state

The Management application enables you to view device discovery status through the **Discover VM Managers** dialog box.

To view the discovery status of a device, complete the following steps.

1. Select **Discover > VM Managers**.

   The **Discover VM Managers** dialog box displays.

2. Right-click the Hosts node select **Expand All** to show all devices.

   The **Discovery Status** field details the actual status message text, which varies depending on the situation.

   The following are samples of actual VMM status messages:

   - Active
   - Failed – Not reachable
   - Failed – Authentication failure

   The following are samples of actual ESX host status messages:

   - Active
   - Discovery pending,
   - Excluded,
   - Conflict – Existing Host <hostname>

3. Refresh the **Discover VM Managers** list by clicking **Refresh**.
4. Click **Close** on the **Discover VM Managers** dialog box.

Troubleshooting VM manager discovery

If you encounter discovery problems, complete the following checklist to ensure that discovery was set up correctly.

Verify IP connectivity by issuing a ping command to the switch.

1. Open the command prompt.
2. From the Server, type `ping` *Device_IP_Address* .

Application Configuration

Configurable preferences

You can use the **Options** dialog box to configure the following preferences in the Management application:

- Event Storage — Use to configure the maximum number of historical events saved to the repository as well as the retention period for the events. For more information, refer to "Event storage settings" on page 87.
- Flyovers — Use to customize the properties display in product and connection flyovers. For more information, refer to "Flyover settings" on page 88.
- Inventory Upload — Use to send SAN inventory information with firmware levels and licenses on devices to a particular E-mail address. For more information, refer to "Inventory Upload settings" on page 80.
- Look and Feel — Use to customize the Management application interface to mimic your system settings as well as define the size of the font. For more information, refer to "Look and feel customization" on page 16.
- OUI Mapping — Use to import OUI file from a specified location. For more information, refer to "Inventory Upload settings" on page 80.
- OUI Mapping — Use to import OUI file from a specified location. For more information, refer to "OUI mapping settings" on page 82.
- Performance Graph Style — Use to configure the color scheme, to display data points, and to configure the error units for all performance graphics in the management application. For more information, refer to "Performance Data" on page 959.
- SAN Display — Use to configure the display for FICON and to reset the display to the default settings. For more information, refer to "SAN display settings" on page 84.
- SAN End Node Display — Use to display (or turn off display of) end nodes on the Connectivity map for newly discovered fabrics. Disabling end node display limits the Connectivity map to switch members only. For more information, refer to "SAN End node display" on page 85.

- SAN Ethernet Loss Events — Use to enable events for a loss of ethernet connection to SAN switches. For more information, refer to "SAN Ethernet loss events" on page 86.
- SAN Names — Use to set whether unique names are required. For more information, refer to "Name settings" on page 91.
- Miscellaneous Security — Use to configure server security configurations and the login banner. For more information, refer to "Miscellaneous security settings" on page 99.
- Server Backup — Use to configure backup settings. Backup is a service process that periodically copies and stores application files to an output directory. The output directory is relative to the server and must use a network share format to support backup to the network. If you use a network path as the output directory, you must add network credentials. For more information, refer to "Server Data backup" on page 71 and "Server Data restore" on page 76.
- Syslog Registration — Use to automatically register the server as the syslog recipient on products. For more information, refer to "Syslog Registration settings" on page 102.
- Trap Registration — Use to automatically register the server as the trap recipient on products. If SAN products have Informs enabled, the registration is for the Informs. For more information, refer to "SNMP Trap Registration settings" on page 103.
- Trap Forwarding Credentials — Use to configure SNMP credentials for the traps forwarded by the server. For more information, refer to "SNMP Trap forwarding credential settings" on page 103.
- Certificates — Use to manage keystore and truststore certificates as well as enable or disable certificate validation. For more information, refer to "Certificates" on page 105.
- Client Export Port — Use to assign a communications port between the client and server. For more information, refer to "Client export port settings" on page 112.
- Client/Server IP — Use to configure IP address of the Management application server. For more information, refer to "Client/Server IP" on page 112.
- Memory Allocation — Use to configure memory allocation for the client and server. You can also use this option to configure asset polling. For more information, refer to "Memory allocation settings" on page 117.
- Product Communication — Use to configure HTTP or HTTP over SSL for connecting to the server. For more information, refer to "Product communication settings" on page 120.
- FTP/SCP/SFTP servers — Use to configure internal or external FTP, SCP, or SFTP server settings. For more information, refer to "FTP/SCP/SFTP server settings" on page 122.
- Server Port — Use to configure server port settings. For more information, refer to "Server port settings" on page 127.
- Support Mode — Use to configure support settings to enable enhanced diagnostics. For more information, refer to "Support mode settings" on page 128.

Server Data backup

The Management application helps you to protect your data by backing it up automatically. Backup is a service process that periodically copies and stores application files to an output directory. The output directory is relative to the server and must use a network share format to support backup to the network. The data can then be restored, as necessary.

> NOTE
> The backup data process performance is optimized when you backup large database. By using parallel processing capabilities provided by PostgreSQL the backup process is optimized and reduced by 50 percentage of the time consumed. For example, a backup of 4GB of data approximately takes 25 to 30 minutes.

> NOTE
> Sometimes when you backup large data, it is possible that, in a disaster recovery situation, configuration changes made after the last backup interval will be missing from the backup.

The Management application allows you to view the backup status at a glance, initiate immediate backup, enable or disable automatic backup, reconfigure the backup directory, interval, and start time, and retrieve backup events.

What is backed up?

The data is backed up to the following directories:

- Backup\databases — contains database and log files.
- Backup\data — contains Element Manager data files (including Dump files, Data collection progress files, Director/Switch firmware files FAF files, Switch technical supportSave, and Switch backup files) and miscellaneous files.
- Backup\conf – contains the Management application configuration files.
- Backup\cimom – contains the SMIA configuration files.

Management server backup

There are two options for backing up data to the Management server:

- Configuring backup to a hard drive (internal or external)
- Configuring backup to a network drive

The Management server is backed up to D:\Backup (Windows systems) by default. Make sure you have an internal or external hard disk configured to ensure that backup can occur. Critical information from the Management application is automatically backed up to the internal or external hard disk when the data directory contents change or when you restart the Management application.

> NOTE
> For networks with large amounts of data to backup, the Management application's performance is degraded during the daily scheduled backup. To avoid performance degradation, configure backup to an external hard drive or use Backup Now on demand.

Backup directory structure overview

The Management server backs up data to two alternate folders. For example, if the backup directory location is D:\Backup, the backup service alternates between two backup directories, D:\Backup\Backup and D:\Backup\BackupAlt. The current backup is always D:\Backup and contains a complete backup of the system. The older backup is always D:\BackupAlt.

If a backup cycle fails, the cause is usually a full hard drive. When the backup cycle fails, there may only be one directory, D:\Backup. There may also be a D:\BackupTemp directory. Ignore this directory because it may be incomplete.

Configuring backup

To configure backup, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Server Backup** in the **Category** list.

   The **Server Backup** pane displays (Figure 22) with the currently defined directory displays in the **Backup Output Directory** field.

   FIGURE 22 Options dialog box (Server Backup pane)

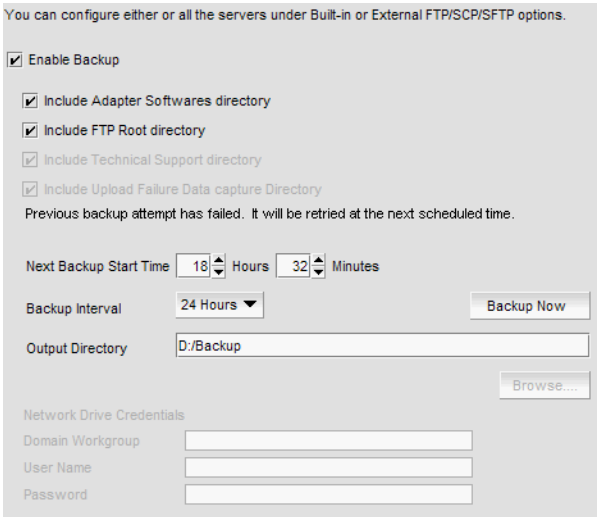

3. Select the **Enable Backup** check box, if necessary.
4. Select what information you want to include in the backup by choosing one or more of the following options:
   - Select the **Include Adapter Boot Image directory** check box.
   - Select the **Include FTP Root directory** check box.

     If you select the FTP Root directory, the FTP Root sub-directories, Technical Support and Trace Dump, are selected automatically and you cannot clear the sub-directory selections. If you do not select the FTP Root directory, the sub-directories can be selected individually.
   - Select the **Include Technical Support directory** check box, if necessary. Only available if the **Include FTP Root directory** check box is clear.
   - Select the **Include Upload Failure Data Capture directory** check box, if necessary. Only available if the **Include FTP Root directory** check box is clear.

5. Enter the time (using a 24-hour clock) you want the backup process to begin in the **Next Backup Start Time Hours** and **Minutes** fields.
6. Select an interval from the **Backup Interval** drop-down list to set how often backup occurs.
7. Backup data to a hard drive by browsing to the hard drive and directory to which you want to backup your data.

   NOTE
   This requires a hard drive. The drive should not be the same physical drive on which your Operating System or the Management application is installed.

8. Backup data to a network drive by completing the following steps.

   To backup to a network drive, your workstation can be either in the same domain or in the same workgroup. However, you must have rights to copy files for the network drive.

   NOTE
   The Management application should not directly access local or network resources through mapped drive letters. When the Management application must access a remote resource (or any process that is running in a different security context), you should use the Universal Naming Convention (UNC) name to access the resource. For more information about services and redirected drives, refer to http://support.microsoft.com/kb/180362/en-us.

   NOTE
   Configuring backup to a network drive is not supported on UNIX systems.

   NOTE
   It is recommended that this configuration be completed on the Local client (the client application running on the Server) so that the backup path and location can be confirmed.

   a. Browse to the network share and directory to which you want to backup your data.

      NOTE
      You must specify the directory in a network share format (for example, \\network-name\share-name\directory). Do not use the drive letter format (C:\directory).

   b. (Windows only) Enter the name of the Windows domain or workgroup in which you are defined in the **Domain Workgroup** field.

      NOTE
      You must be authorized to write to the network device.

   c. (Windows only) Enter your Windows login name in the **User Name** field.
   d. (Windows only) Enter your Windows password in the **Password** field.

9. Click **Apply** or **OK**.

   The application verifies that the backup device exists and that the server can write to it.

   For backup to a network drive, if the device does not exist or you are not authorized to write to the network drive, an error message displays that states you have entered an invalid device path or invalid network credentials.

   Click **OK** to go back to the **Options** dialog box and fix the error.

   Backup occurs, if needed, at the interval you specified.

Enabling backup

Backup is enabled by default. However, if it has been disabled, complete the following steps to enable the function.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Server Backup** in the **Category** list.
3. Select the **Enable Backup** check box.
4. Click **Apply** or **OK**.

Disabling backup

Backup is enabled by default. If you want to stop the backup process, you need to disable backup. To disable the backup function, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Server Backup** in the **Category** list.
3. Clear the **Enable Backup** check box.
4. Click **Apply** or **OK**.

Viewing the backup status

The Management application enables you to view the backup status at a glance by providing a backup status icon on the Status Bar. The following table illustrates and describes the icons that indicate the current status of the backup function.

TABLE 14 Backup status

Icon	Description
	Backup in Progress — displays the following tooltip: "Backup started at hh:mm:ss, in progress... *XX* directories are backed up."
	Countdown to Next Scheduled Backup — displays the following tooltip: "Next backup scheduled at hh:mm:ss."
	Backup Disabled — displays the following tooltip: "Backup is disabled."
	Backup Failed — displays the following tooltip: "Backup failed at hh:mm:ss mm/dd/yyyy."

Changing the backup interval

When the backup feature is enabled, your SAN is protected by automatic backups. The backups occur every 24 hours by default. However, you can change the interval at which backup occurs.

> **NOTE**
> Do NOT modify the backup.properties file.

To change the backup interval, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Server Backup** in the **Category** list.
3. Select an interval from the **Backup Interval** drop-down list to set how often backup occurs.
4. Click **Apply** or **OK**.

   The minimum value is 6 hours and the maximum value is 24 hours.

Starting immediate backup

> **NOTE**
> You must have backup privileges to use the Backup Now function. For more information about privileges, refer to "User Privileges" on page 1333.

To start the backup process immediately, complete one of the following procedures:

Using the Backup Icon, right-click the **Backup** icon and select **Backup Now**.

The backup process begins immediately.

OR

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Server Backup** in the **Category** list.
3. Click **Backup Now**.

   Click **Yes** on the confirmation message. The backup process begins immediately.

4. Click **Apply** or **OK**.

Reviewing backup events

The Master Log, which displays in the lower left area of the main window, lists the events that occur on the Fabric.

If you do not see the Master Log, select **View > Show Panels > All Panels**.

The following backup events appear in the Master Log:

- Backup started
- Backup error

- Backup Enabled
- Backup Disabled
- Backup Now
- Backup destination change
- Backup interval change
- Backup start time change
- Domain workgroup change
- User name change
- User password change
- Number of files backed up on completion
- Network share access problem when backup starts or during backup (not when the backup configuration is changed)

Server Data restore

NOTE
You cannot restore data from a previous version of the Management application.

NOTE
You cannot restore data from a higher or lower configuration (Trial or Licensed version) of the Management application.

NOTE
You cannot restore data from a different package of the Management application.

The Management application helps you to protect your data by backing it up automatically. The data can then be restored, as necessary.

The data in the following directories is automatically backed up to disk. The data includes the following items:

- Backup\databases — contains database and log files.
- Backup\data — contains Element Manager data files (including Dump files, Data collection progress files, Director/Switch firmware files FAF files, Switch technical supportSave, and Switch backup files) and miscellaneous files. .
- Backup\conf – contains the Management application configuration files.
- Backup\cimom – contains the SMIA configuration files.

In a disaster recovery situation, it is possible that configuration changes made less than 45 minutes before Server loss (depending on the backup interval you set) could be missing from the backup.

Restoring data

NOTE
The restore data files must use the exact directory structure as the backup directory structure (refer to "Backup directory structure overview" on page 72).

1. (Windows) Open the **Server Management Console** from the **Start** menu on the Management application server.

   OR

   (UNIX) Open *Install_Home/bin* from the Management application server and type `./smc.sh` at the command line.

2. Click the **Services** tab.

   The tab lists the Management application services.

3. Click **Stop Services** to stop all of the services.
4. Click the **Restore** tab.
5. Browse to the backup location.

   Browse to the location specified in the **Output Directory** field on the **Options** dialog box - Backup pane.

6. Click **Restore**.

   Upon completion, a message displays the status of the restore operation. Click OK to close the message and the Server Management Console. For the restored data to take effect, re-launch the Configuration Wizard using the instructions in "Launching the Configuration Wizard" on page 6.

Restoring data to a new server

NOTE
The restore data files must use the exact directory structure as the backup directory structure (refer to "Backup directory structure overview" on page 72).

If your Management application server fails and you must recover information to a new server, restore the data (Refer to "Restoring data" on page 76 for complete instructions).

SAN data collection

The Management application uses collectors to gather data from switches, persist the switches in the database, and to publish the collected data to the client. Each collector polls data for one feature area using HTTP or HTTPS (web pages or CAL calls) to communicate with the switch. For a given switch, only one collector runs at a time. When you first discover a switch, all collectors associated with that switch type run to gather data on the switch. After that, the Management application schedules each collector to run independently. Since this is a fairly repetitive task, the Management application has a collection framework which schedules these collectors to run periodically (using lazy polling).

When a data collector fails, the Management application automatically retries the collection after the short tick interval. However, there are two exceptions to this retry rule:

- If collection failure is due to an ACL rule blocking access to the switch, the Management application retries collection after the lazy polling interval (not the short tick interval).
- If collection failure is due to incorrect switch credentials in the Management application, the Management application retries collection three times after which all communication with the switch is stopped to prevent lock out of the Fabric OS user due to too many failed login attempts.

In addition to the automatic collection retry, you can configure adaptive asset collection to trigger specific collectors to run when a particular event occurs. For example, when the Management application receives an SNMP trap that a port has been disabled, the Management application triggers the TopologyCollector (which collects ISL information) and the SwitchAssetCollector to make sure that the client reflects the changes due to the port going down. Adaptive asset collection occurs within the short tick interval.

The Management application uses the short tick interval to ping the switch for a periodic reachability check. If the reachability check succeeds, then the Management application runs pending collectors triggered by an event. When no SNMP traps or syslog events occur, the Management application uses the lazy polling interval to schedule collection of configuration and status changes. The lazy polling interval process schedules any pending collectors for the next short tick. Therefore, the interval between collections when there are no SNMP traps or syslog events is the lazy polling interval plus the short tick interval. To increase polling efficiency, you can configure both the short tick interval (**Check for state change every** option) and the lazy polling interval (**If no state change, poll switch every** option) on the **Options** dialog box. For step-by-step instructions, refer to "Configuring asset polling" on page 119.

There are two types of collectors, fabric-level collectors and switch-level collectors. Fabric-level collectors gather fabric-level information. The Management application collects fabric-level data from the seed switch, for example, the NameServerCollector gathers data about all end devices present in the fabric.

The Management application uses the following Fabric-level collectors:

- DeviceFDMICollector – Collects FDMI-related information for end devices in the fabric.
- NameServerInfoCollector – Collects data about end devices in the fabric.
- ActiveZoneInfoCollector – Collects the active zone configuration in the fabric.
- ZoneInfoCollector – Collects the defined zone configuration in the fabric.
- TopologyCollector – Collects data about the ISLs in the fabric.
- TrunkInfoCollector – Collects data about trunks in the fabric.
- WtJarsCollector – Downloads the jar files needed to launch WebTools from the Management application.

Switch-level collectors gather individual switch-level information (such as, port details and so on). The Management application also uses specialty collectors which run only for switches that have a particular feature. For example the EncryptionBaseCollector only runs for encryption switches. The Management application uses the following Switch-level collectors:

- BottleneckConfigCollector – Collects data about bottleneck configuration on the switch.
- BottleneckStatusCollector – Collects data about the bottleneck status (whether or not a port is bottlenecked) for each port on the switch.
- EncryptionBaseCollector – Collects all encryption related data.
- GroupConfigChangeCollector – Collects encryption data related to HA Cluster, Target Containers, Crypto Host, and Crypto LUN.
- GroupConfigCollector – Collects group member and group collection data.
- FabricCollector – Collects the fabric members (switches) and persists the members in the application. This is the main collector that organizes fabric discovery.
- CeeSwitchCalCollector – Collects the association of a device port to a 10 G physical port on the DCB switch.
- DCBCollector – Collects data specific to the DCB switch.
- FportTrunkCollector – Collects data about F-port trunks present on the switch.
- GigePortCollector – Collects GigE-port data on the switch.
- LicenseCollector – Collects data about licenses on the switch.
- LiteSwitchAssetCollector – Collects the FMS mode setting on the switch.
- SwitchAssetCollector – Collects data about the switch including, inventory details, port level data, any blades that may be present (on directors), AG-port mapping, and auto trace dump settings. This is the major collector for switch data.
- FCIPCollector – Collects FCIP-related data on the older FCIP switches.
- XFCIPCollector – Collects extended FCIP-related data on the newer FCIP switches.

- MapsPolicyCollector – Collects data about MAPS policies configured on the switch.
- MetaSANCollector – Collects data about the IFLs (Inter Fabric Links) on the switch.
- FlowCollector – Collects data about the flow definitions on the switch. Also collects the subflows for each flow definition. This collector requires the Fabric Insight license on the switch.
- VPWwnInfoCollector – Collects data about the VPWWN (Virtual Port World Wide Name) on the switch.

The Management application collects performance monitoring data via SNMP. Performance monitoring data is collected asynchronously and is not affected by the collector scheduling. Performance data (mostly port statistics) is collected every 5 minutes.

Product communication protocols

Table 15 details the protocols that the Management application uses for communication between products and the Management application server.

TABLE 15 Product communication protocols

Protocol	Description	Management application use	Communicates with device type
ICMP	The Internet Control Message Protocol (ICMP) is part of the Internet Protocol Suite, as defined in RFC 792. ICMP messages are typically used for diagnostic or control purposes or generated in response to errors in IP operations.	Used during initial discovery.	Fabric OS
SNMP	Simple Network Management Protocol (SNMP) is an internet-standard protocol for managing devices on IP networks. SNMP is a component of the Internet Protocol Suite as defined by the Internet Engineering Task Force (IETF). It consists of a set of standards for network management, including an application layer protocol, a database schema, and a set of data objects.	Used during discovery, performance polling, and operation status polling. Note that performance polling (including data collection for dashboards) completely relies on SNMP. For Historical data collection, the minimum time interval is 1 minute and 5 minutes for Fabric OS devices. For real time graphs, the minimum time interval is 10 seconds.	Fabric OS
HTTP/HTTPS	The Hypertext Transfer Protocol (HTTP) is an application protocol for distributed, collaborative, hypermedia information systems. HTTP is the foundation of data communication for the World Wide Web. Hypertext Transfer Protocol Secure (HTTPS) is a communications protocol for secure communication over a computer network, with especially wide deployment on the Internet. Technically, it is not a protocol in and of itself; rather, it is the result of simply layering the Hypertext Transfer Protocol (HTTP) on top of the SSL/TLS protocol, thus adding the security capabilities of SSL/TLS to standard HTTP communications.	Used to collect all information required to manage Fabric OS devices. You can configure the HTTP/HTTPS protocol from the **Options** dialog box (refer to "Product communication settings" on page 120).	Fabric OS
FTP	File Transfer Protocol (FTP) is a standard network protocol used to transfer files from one host to another host over a TCP-based network	Used for firmware download. For Fabric OS devices, used to collect technical support information. For more information, refer to "FTP/SCP/SFTP server settings" on page 122.	Fabric OS Network OS

TABLE 15 Product communication protocols

Protocol	Description	Management application use	Communicates with device type
SCP	Secure copy (SCP) is a means of securely transferring computer files between a local host and a remote host or between two remote hosts. It is based on the Secure Shell (SSH) protocol.	Used for firmware download. For Fabric OS devices, used to collect technical support information. For more information, refer to "FTP/SCP/SFTP server settings" on page 122.	Fabric OS IronWare Network OS
SFTP	Secure File Transfer Protocol (SFTP) or SSH File Transfer Protocol is a network protocol that provides file access, file transfer, and file managementfunctionalities over any reliable data stream.	Used for firmware download. For Fabric OS devices, used to collect technical support information. For more information, refer to "FTP/SCP/SFTP server settings" on page 122.	Fabric OS Network OS

Inventory Upload settings

You can configure inventory upload option to send SAN inventory information along with the firmware levels and licenses on the switches to specific e-mail addresses.

The inventory details are sent immediately to the specified e-mail address if the e-mail server is configured ("Configuring e-mail notification" on page 156) and the data is available. If the data is not available or the e-mail server is not configured then the upload will be triggered the next day and continue to send the details as scheduled (once in a week).

NOTE
You have to configure the E-mail server before configuring the Inventory upload settings.

Enabling Inventory Upload settings

To send data of the device, complete the following steps.

1. Select **Server** > **Options**.

   The **Options** dialog box displays.

2. Select **Inventory Upload** in the **Category** list.

   The **Inventory Upload** pane displays.

FIGURE 23 Options dialog box -Inventory Upload pane

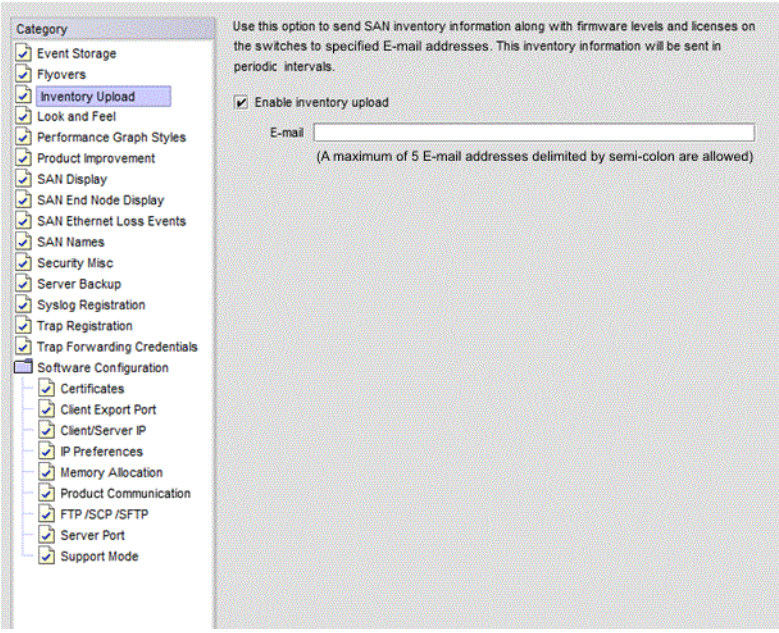

3. Select **Enable inventory upload** check box to configure the settings.
4. Enter a valid E-mail address in the **E-mail** field to which you want to send the inventory details.

   Enter more than one e-mail address, separating each with a semicolon. You can enter up to five e-mail addresses.
5. Click **OK**.

   The inventory details (zip file) is sent to the specified e-mail addresses. The zip file uses the following naming convention: IBM_InventoryDetails_2015-01-yyyy-mm-dd.zip.

FIGURE 24 Inventory Details example

IBM Network Advisor_InventoryDetails_2017-04-19 06-00-18.xml

```
<?xml version="1.0" encoding="UTF-8"?><Inventory_Upload>
<customer_details>
 <property name="Company Name" value=""/>
 <property name="Customer Name" value=""/>
</customer_details>
<chassis identifier="10:00:00:27:F8:BC:37:47">
 <property name="Firmware Version" value="v8.1.0_bld44"/>
 <property name="Contact Name" value="Field Support."/>
 <property name="IP Address" value="10.24.44.12"/>
 <property name="Product Type" value="109"/>
 <property name="Call Home" value="Disabled"/>
 <property name="Fabric Watch" value=""/>
 <property name="Ethernet IP Mask" value="255.255.248.0"/>
 <property name="Factory Serial number" value="BRW2539J09T"/>
 <property name="Product Manufacturer" value="BRD"/>
 <property name="Type" value="109"/>
 <property name="Sub Type" value="0"/>
 <property name="Chassis WWN" value="10:00:00:27:F8:BC:37:47"/>
 <property name="RNID Sequence" value="0BRW2539J09T"/>
 <property name="Product Type Number" value="BROCAD"/>
 <property name="Model number" value="Brocade 6510"/>
 <property name="Ethernet IP" value="10.24.44.12"/>
 <property name="FCIP" value="0.0.0.0"/>
 <property name="Licenses installed" value=""/>
 <property name="Customer Name" value="yraj@brocade.com"/>
 <property name="Chassis Service Tag" value="NOT_AVAILABLE"/>
```

OUI mapping settings

You can import the OUI (Organizationally Unique Identifier) information and assign the product type as Initiator, Target, or Default. The latest OUI can be imported from the following site. http://standards.ieee.org/develop/regauth/oui/oui.txt

> NOTE
> Discovery overrides the product type based on the OUI mapping. On discovery, the preference is given to the switch returned value of the switch only when the OUI mapping is not available.

You can change the product type of OUI on the **Product Type Mapping** dialog box, and the end devices which have the same OUI will be overridden. When you change the product type from **Default** to **Initiator** or **Target**, it will be reflected in the application immediately. However, changing the product type from **Initiator** or **Target** to **Default** will be reflected only when there is a change in the name server data.

Importing the OUI file

To import the OUI file, complete the following steps.

1. Select **Server** > **Options**.

   The **Options** dialog box displays.

2. Select **OUI Mapping** in the **Category** list.

   The **OUI Mapping** pane displays.

   FIGURE 25 Options dialog box - OUI Mapping pane

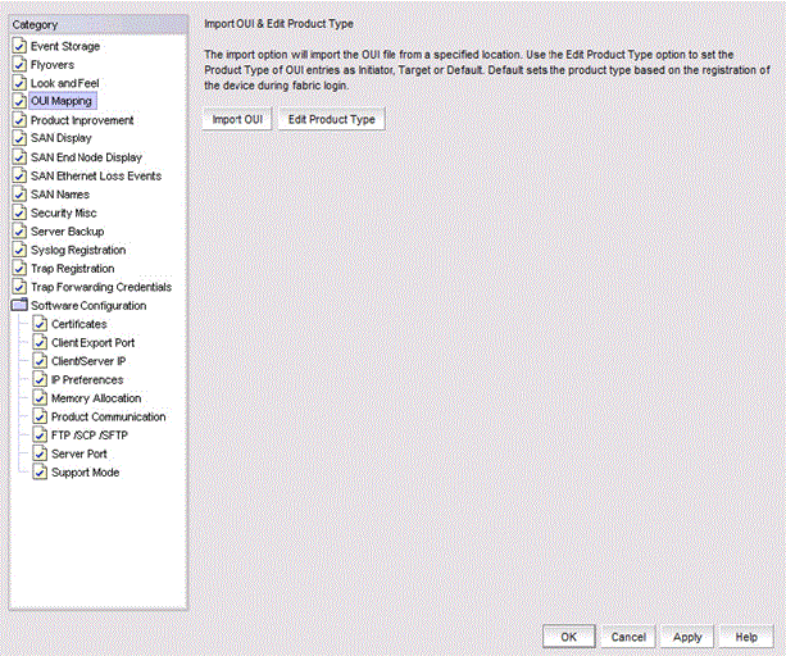

3. Click **Import OUI**.

4. Select the folder where you want the OUI file to be uploaded.

   NOTE
   If you select the correct OUI file and new OUI files are found, the **Product Type Mapping** dialog box is displayed. The product type for each OUI will be **Default** and can be changed to **Target** or **Initiator.**

   The following example formats are supported in the OUI.text file:

   - 00-00-88 (hex) Brocade Communications Systems, Inc.
   - 00:00:88 (hex) Brocade Communications Systems, Inc.
   - 000088 (hex) Brocade Communications Systems, Inc.

5. Click **OK**.

Editing the product type

You can edit the product type of the OUI by setting it as **Initiator**, **Target,** or **Default.**

To edit the product type, complete the following steps.

1. Select **Server** > **Options**.

   The **Options** dialog box displays.

2. Select **OUI Mapping** in the **Category** list.

   The **OUI Mapping** pane displays.

3. Click **Edit Product Type**.

   The **Product Type Mapping** dialog box displays.

   FIGURE 26 Product Type Mapping dialog box

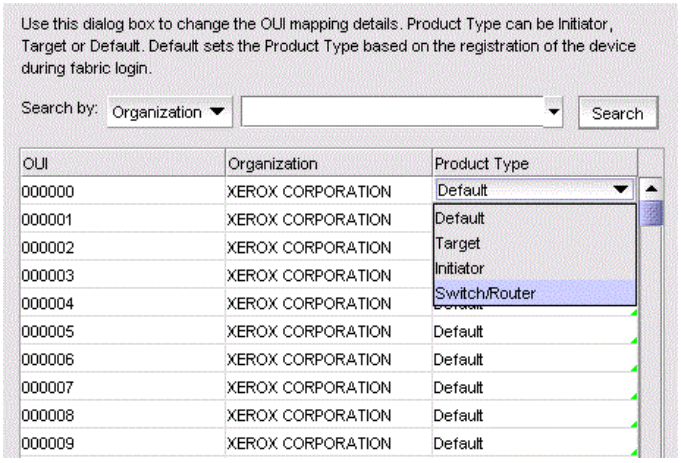

   NOTE
   You can search for an OUI by using a search string in the **Search** list or with the **Organization** list.

4. Select the **product type** for a particular OUI file and change to **Target**, **Initiator, Switch/Router**, or **Default**.
5. Click **OK**.

SAN display settings

You can configure the display for FICON and reset the display to the default settings.

Setting your FICON display

FICON display setup rearranges the columns of any table that contains end device descriptions to move the following columns to be the first columns: Attached Port#, FC Address, Serial #, Tag, Product Type, Model, Vendor, Port Type, and WWN.

> **NOTE**
> You cannot set the FICON display for Professional and Professional Plus software.

To set the FICON display, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **SAN Display** in the **Category** list.

   The **SAN Display** pane displays (Figure 27).

   FIGURE 27 Options dialog box (SAN Display pane)

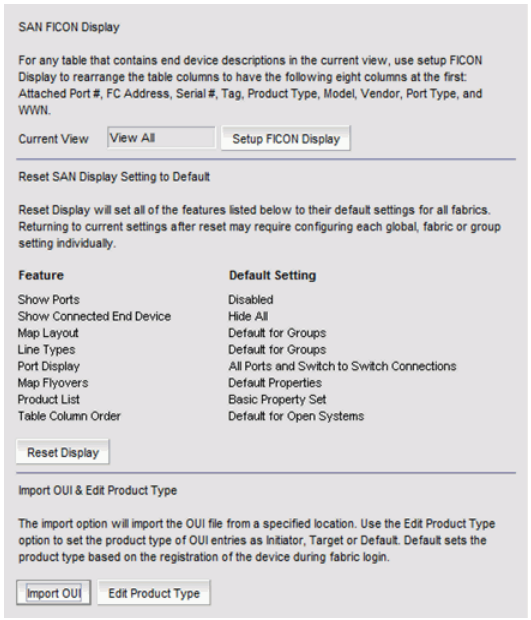

3. Click **Set Up FICON Display**.

   Any table that contains end device descriptions move the following nine columns to the beginning of the table: Attached Port #, FC Address, Serial #, Tag, Device Type, Model, Vendor, Port Type, and WWN.

4. Click **Apply** or **OK** to save your work.

Resetting your display

You can reset your system to display the default display settings for all fabrics. Note that returning to current settings after a reset may require configuring each global fabric or group setting individually. The following table (Table 16) details the settings that change with reset and the associated default state.

TABLE 16 Default display settings

Settings	Default State
Show Ports	Disabled
Show Connected End Device	Hide All
Map Layout	Default for Groups
Line Types	Default for Groups
Port Display	All Ports and Switch to Switch Connections
Map Flyovers	Default Properties — includes the following properties: • Product Display — Name, Device Type, WWN, IP Address, and Domain ID. • Connection Display — Name (port), Address, Node WWN, Port WWN, and Port #.
Product List	Basic Poperty Set
Table Column Order	Default for Open System

To reset the Management application to the default display and view settings, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **SAN Display** in the **Category** list.
3. Click **Reset Display**.
4. Click **Yes** on the reset confirmation message.

   The display and view settings are immediately reset to the default display settings (as detailed in the Default display settings table (Table 16).

5. Click **Apply** or **OK** to save your work.

SAN End node display

The connectivity map can be configured to display or not display end nodes. This option enables you to set the end node display for all newly discovered fabrics. Note that disabling end node display limits the connectivity map to emphasize switch members only.

Displaying end nodes

To display end nodes when discovering a new fabric, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays (Figure 28).

FIGURE 28 Options dialog box (SAN End Node Display pane)

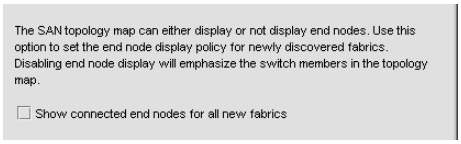

2. Select **SAN End Node Display** in the **Category** list.
3. Select the **Show connected end nodes when new fabric is discovered** check box to display end nodes on your system.

   NOTE
   Before changes can take effect, the topology must be rediscovered.

4. Click **Apply** or **OK** to save your work.

SAN Ethernet loss events

An Ethernet event occurs when the Ethernet link between the Management Server and the managed SAN device is lost. You can configure the application to enable events when the Ethernet connection is lost.

Enabling SAN Ethernet loss events

The **Options** dialog box enables you to configure the Management application to generate an Ethernet event after a device is offline for a specific period of time.

To enable Ethernet loss events, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

   FIGURE 29 Options dialog box (SAN Ethernet Loss Event pane)

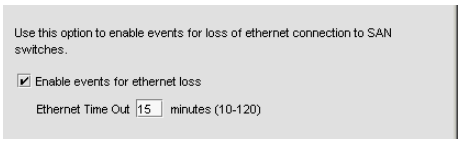

2. Select **SAN Ethernet Loss Events** in the **Category** list.
3. Select the **Enable events for ethernet loss** check box.
4. Enter the Ethernet time out value (10 to 120 minutes).
5. Click **Apply** or **OK** to save your work.

Disabling SAN Ethernet loss events

To disable Ethernet loss events, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **SAN Ethernet Loss Events** in the **Category** list.

3. Clear the **Enable events for ethernet loss** check box.
4. Click **Apply** or **OK** to save your work.

Event storage settings

You can configure the maximum number of historical events save to the repository, how long the events will be retained, as well as whether to store historical events to a file before purging them from the repository.

Configuring event storage

To configure event storage, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Event Storage** in the **Category** list (Figure 30).

   FIGURE 30 Options dialog box (Event Storage pane)

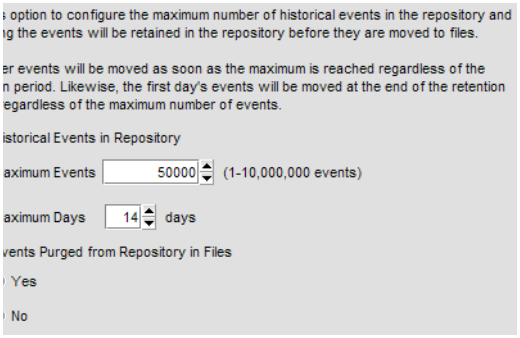

3. Enter the maximum number of events you want to be retained in the repository in the **Maximum Events** field.

   Depending on your installation, the maximum number of events stored are as follows:

   - Professional — 1 through 100,000
   - Professional Plus — 1 through 1,000,000
   - Enterprise — 1 through 10,000,000

   Default is 50,000. Older events are purged at midnight on the date the maximum event limit is reached regardless of the retention days.

4. Enter then number of days (1 through 365) you want to store events in the **Maximum Days** field.

   The events are purged at midnight on the last day of the retention period regardless of the number of maximum events.

5. Choose one of the following options:

   - Select the **Yes** option to store all historical events from the repository to a file while purging occurs.
   - Select the **No** option to purge historical events from the repository without storing them as a file.

6. Click **OK**.

Storing historical events purged from repository

To store historical events purged from the repository, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Event Storage** in the **Category** list.
3. Select the **Yes** option.
4. Click **OK**.

   Purged events from the master log table are stored in the ***Install_Home***\data\archive\events directory using the format event_***MMDDYYY***.zip (for example, event_04052011.zip. These files are retained for a maximum of 30 days. The zip file contains multiple archive text files that use the format event_***MMDDYYY_N***.txt (for example, event_04052011_1.txt).

Flyover settings

You can configure your system to display information for products and connections in a pop-up window on the Connectivity Map.

Configuring flyovers

To display product and connection information in a pop-up window, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Flyovers** in the **Category** list.
3. Select the **Enable flyover display** check box to enable flyover display on your system.
4. Select the **Include labels** check box to include labels on flyover displays.
5. Add product properties you want to display on flyover by selecting the **Product** tab (Figure 31) and completing the following steps.

FIGURE 31 Options dialog box (Flyovers pane, Product tab)

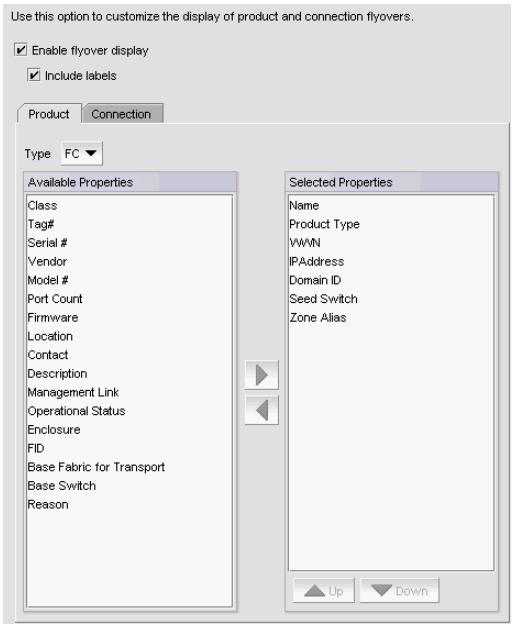

a. Select the protocol type from the **Type** list, if necessary.

b. Select each property you want to display in the product flyover from the **Available Properties** table.

Depending on which protocol you select, some of the following properties may not be available:

FC (default)

- Name
- Device Type
- WWN
- IP Address
- Domain ID
- Class
- Tag#
- Serial #
- Vendor
- Model #
- Port Count
- Seed Switch
- Firmware
- Location
- Contact
- Description
- Management Link
- Operational Status
- Enclosure
- Reason
- FID
- Base Fabric for Transport
- Base Switch
- Zone Alias

c. Click the right arrow to move the selected properties to the **Selected Properties** table.

d. Use the **Move Up** and **Move Down** buttons to reorder the properties in the **Selected Properties** table, if necessary.

The properties displayed in the **Selected Properties** table appear in the flyover display.

6. Remove product properties you do not want to display on flyover by selecting the property in the **Selected Properties** table and clicking the left arrow.

7. Add connection properties you want to display on flyover by selecting the **Connection** tab (Figure 32) and completing the following steps.

FIGURE 32 Options dialog box (Flyovers pane, Connection tab)

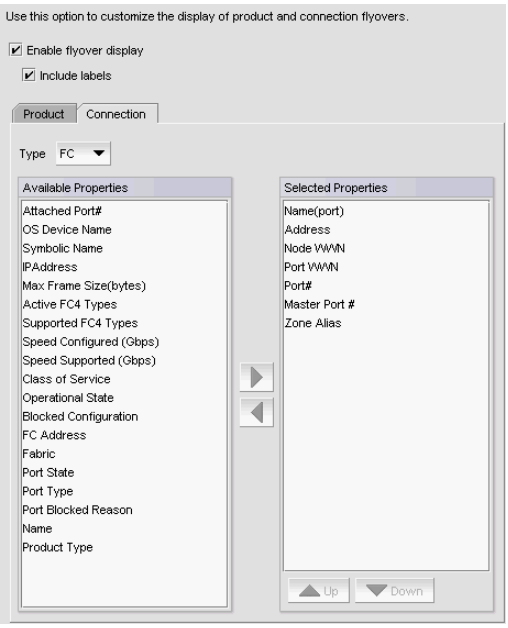

a. Select the protocol type from the **Type** list, if necessary.

   Depending on which protocol you select, some properties may not be available for all protocols.

b. Select each property you want to display in the connection flyover from the **Available Properties** table.

   Depending on which protocol you select, some of the following properties may not be available for all protocols:

   FC (default)

- Active FC4 Types
- Address
- Attached Port#
- Blocked Configuration
- Class of Service
- Device Type
- Fabric
- FC Address
- IP Address
- Master Port #
- Max Frame Size (bytes)
- Name
- Name (port)
- Node WWN
- Operational State
- OS Device Name
- Port #
- Port Blocked Reason
- Port State
- Port Type
- Port WWN
- Speed Configured (Gbps)
- Speed Supported (Gbps)
- Symbolic Name
- Supported FC4 Types
- Zone Alias

   FCoE

- Name
- Node WWN
- MAC
- Port#
- Port Type
- FCoE Index #

c. Click the right arrow to move the selected properties to the **Selected Properties** table.

d. Use the **Move Up** and **Move Down** buttons to reorder the properties in the **Selected Properties** table.

   The properties displayed in the **Selected Properties** table appear in the flyover display.

8. Remove connection properties you do not want to display on flyover by selecting the property in the **Selected Properties** table and clicking the left arrow.
9. Click **Apply** or **OK** to save your work.

Turning flyovers on or off

Flyovers display when you place the cursor on a product. They provide a quick way to view a product's properties.

To turn flyovers on or off, select **Enable Flyover Display** from the **View** menu.

Viewing flyovers

On the Topology Map, rest the pointer over a product icon, port, or connection.

The pop-up window containing the product, port, or connection information displays.

For the product icon, the pop-up window displays the display name and IP address of the device.

For the connection, the pop-up window displays the IP address and port number for each device at either end of the connection. If one of the connections is a cloud, the port number does not display.

Name settings

You can use Names as a method of providing familiar simple names to products and ports in your SAN. Using your Management application you can:

- Set names to be unique or non-unique.
- Fix duplicate names.
- Associate a name with a product, port WWN, or Fabric Assigned WWN currently being discovered.
- Add a WWN and an associated name for a product or port that is not yet being discovered.
- Remove or disassociate a name from a WWN.

Setting names to be unique

You can edit duplicate names so that each device has a unique name. Note that the **Duplicated Names** dialog box only displays when you set names to be unique and there are duplicate names in the system.

To edit duplicate names, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **SAN Names** in the **Category** list.

   The **SAN Names** pane displays (Figure 33).

FIGURE 33 Options dialog box (SAN Names pane)

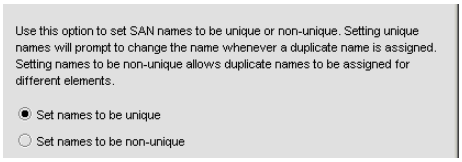

3. Select **Set names to be unique** to require that names be unique on your system.
4. Click **OK** on the **Options** dialog box.
5. Click **OK** on the "duplicate names may exist" message.
   To fix duplicated names, refer to "Fixing duplicate names" on page 92.

Setting names to be non-unique

You can choose to allow duplicate names in your fabric.

To set names to be non-unique, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.
2. Select **SAN Names** in the **Category** list.
3. Select **Set names to be non-unique** to allow duplicate names on your system.
4. Click **OK** on the **Options** dialog box.

Fixing duplicate names

To fix duplicated names, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.
2. Click **Fix Duplicates**.

   The **Duplicated Names** dialog box displays (Figure 34).

FIGURE 34 Duplicated Names dialog box

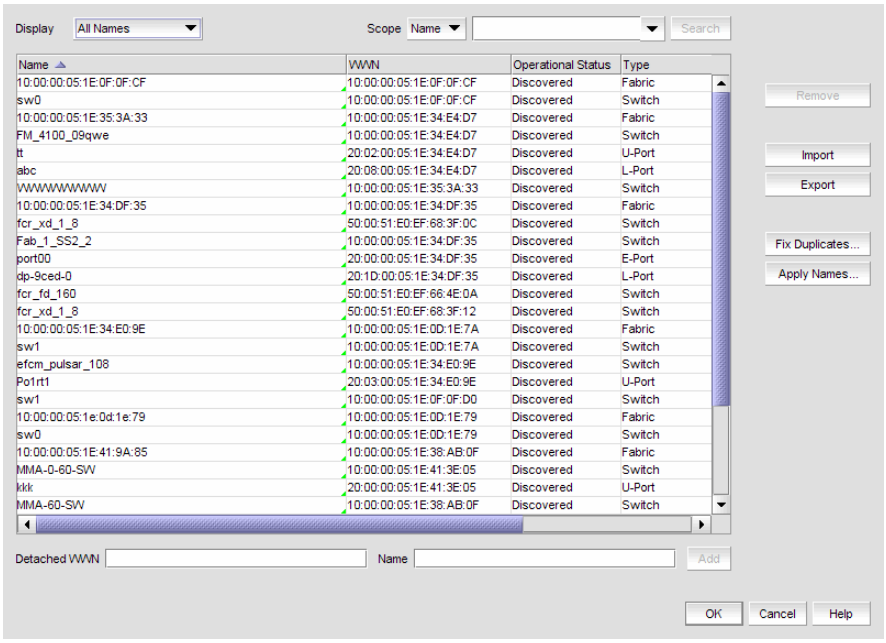

The **Duplicated Names** dialog box contains the following information:

- **Description** — A description of the device.
- **Duplicate Names** table — Every instance of duplicate names.
  - **Fabric** — The fabric name.
  - **FC Address** — The Fibre Channel address.
  - **Names** — The current name of the device.
    If you selected the **Append Incremental numbers for all repetitive names** option, the names display with the incremental numbering.
    If you selected the **I will fix them myself** option, this field becomes editable.
  - **Operational Status** — The operational status of the device. There are four possible values:
    - Up — Operation is normal.
    - Down — The port is down or the route to the remote destination is disabled.
    - Disabled — The connection has been manually disabled.
    - Backup Active — The backup TCP port is active due to a failover.
  - **Port #** — The port number.
  - **Type** — The type of device.

3. Select one of the following options.
   - If you select **Append Incremental numbers for all repetitive names**, the names are edited automatically using incremental numbering.
   - If you select **I will fix them myself**, edit the name in the **Name** field.
4. Click **OK** on the **Duplicated Names** dialog box.
5. Click **OK** to close the **Configure Names** dialog box.
6. Click **OK** on the confirmation message.

Viewing names

To view names associated with devices, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays (Figure 35).

   FIGURE 35 Configure Names dialog box

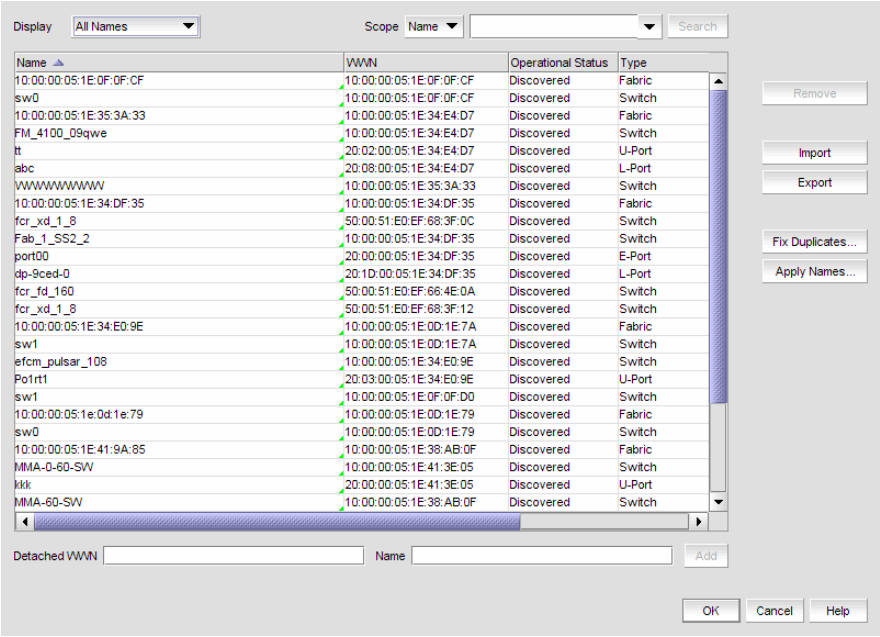

2. Select **All Names** from the **Display** list.

   Only devices with a name display. The table displays the following information.

   - **Scope** list — Select a search value (Name or WWN) from the list.
   - **Search text box —** Enter the name or WWN of the device for which you are searching.
   - **Search** button — Click to search on the value in the Search field. For more information, refer to "Searching for a device by name" on page 98.
   - **Display** table — This table displays the following information:
     - **Description–**A description of the device.
     - **Name–**The name of the device. Enter a name for the device.
     - **Operational Status–**The operational status of the device (discovered, operational, and unknown).
     - **Type–**The type of device (port, node, Fabric Assigned WWN, and unknown).
     - **WWN–**The world wide node (WWN) of the device. Enter a WWN for the device.

     Click a column head to sort the list. Click a column head again to reverse the sort order.
   - **Remove** button — Click to remove a device from the Display table. For more information, refer to "Removing a name from a device" on page 96.
   - **Import** button — Click to import name data. For more information, refer to "Importing Names" on page 97.
   - **Export** button — Click to export the name data. Depending on your operating system, the default export location are as follows:
     - Desktop\My documents (Windows)
     - \root (Linux)

For more information, refer to "Exporting names" on page 97.

- **Fix Duplicates** button — Click to launch the Fix Duplicates dialog box. For more information, refer to "Fixing duplicate names" on page 92.
- **Apply Names** button — Click to apply unassigned (detached) names to newly discovered devices. For more information, refer to "Applying a name to a detached WWN" on page 96.
- **Detached WWN text box —** Enter the WWN of the device you want to add.
- **Name text box —** Enter a name for the device you want to add.
- **Add** button — Click to add a device by detached WWN and Name to the table. For more information, refer to "Adding a name to a new device" on page 95.

3. Click **OK** to close the **Configure Names** dialog box.

Adding a name to an existing device

To add a name to an existing device, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.

2. Select how you want to display devices from the **Display** list.

   You can display devices by **All Names**, **All WWNs**, **Fabric Assigned WWNs**, **Only Fabrics**, **Only Products**, **Only Ports**, or **Switch and N Ports**.

   All discovered devices display.

3. Select the device to which you want to assign a name in the **Display** table.
4. Double-click in the **Name** column for the selected device or port and enter a name for the device or port.

   If you set names to be unique on the **Options** dialog box and the name you entered already exists, the entry is not accepted. To search for the device already using the name, refer to "Searching for a device by name" on page 98 or "Searching for a device by WWN" on page 99 in the **Configure Names** dialog box or "Searching for a device" on page 314 in the connectivity map.

   NOTE
   If you segment a fabric, the Fabric's name follows the assigned principal switch.

5. Click **OK** on the confirmation message.
6. Click **OK** to close the **Configure Names** dialog box.

Adding a name to a new device

To add a new device and name it, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.

2. Enter the WWN of the device in the **Detached WWN** field.
3. Enter a name for the device in the **Name** field.
4. Click **Add**.

The new device displays in the table.

If you set names to be unique on the **Options** dialog box and the name you entered already exists, a message indicating the name already in use displays. Click **OK** to close the message and change the name.

5. Click **OK** to close the **Configure Names** dialog box.
6. Click **OK** on the confirmation message.

Applying a name to a detached WWN

To apply a name to a detached wwn, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.

2. Click **Apply Names**.

   If there are any detached WWNs in a discovered state, the **Apply Names** dialog box displays.

3. Select or clear the check box for the associated switch or switch port.

   Select a check box to apply the detached name as the switch or switch port name and remove the duplicated WWN entry (detached) in the **Configure Names** dialog box.

   Clear a check box to remove the duplicated WWN entry (detached) in the **Configure Names** dialog box.

4. Click **OK** on the **Apply Names** dialog box.
5. Click **OK** on the **Configure Names** dialog box.

Removing a name from a device

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.

2. In the **Display** table, select the name you want to remove.
3. Click **Remove**.

   An application message displays asking if you are sure you want clear the selected name.

4. Click **Yes**.
5. Click **OK** to close the **Configure Names** dialog box.
6. Click **OK** on the confirmation message.

Editing names

To edit the name associated with a device, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.

2. Select **All Names** from the **Display** list.

   Only devices with a name display. The table displays the Name, WWN, Operational Status, Type, and a Description of the device.

3. Click the name you want to edit in the **Name** column.
4. Edit the name and press **Enter**.
5. Click **OK** to close the **Configure Names** dialog box.
6. Click **OK** on the confirmation message.

Exporting names

To export the names associated with devices, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.

2. Click **Export.**

   The **Export Files** dialog displays.

3. Browse to the location where you want to save the export file.

   Depending on your operating system, the default export location are as follows:

   - Desktop\My documents (Windows)
   - \root (Linux)

4. Enter a name for the file and click **Save.**
5. Click **OK** to close the **Configure Names** dialog box.

Importing Names

If the name length exceeds the limitations detailed in the following table, you must edit the name (in the CSV file) before import. Names that exceed these limits will not be imported. If you migrated from a previous version, the .properties file is located in the *Install_Home*\migration\data folder.

TABLE 17 Name length limitations

Device	Character limit
Fabric OS switch 6.2 or later	30 (24 character limit when in FICON mode)
Fabric OS switch 6.1.X or earlier	15
Fabric OS switch port 7.0 or later	128 (24 character limit when in FICON mode)
Fabric OS switch port 6.4.X or earlier	32 (24 character limit when in FICON mode)

TABLE 17 Name length limitations

Device	Character limit
HBA	256
HBA port	256
Others names	128

To import names, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.

2. Click **Import.**

   The **Import Files** dialog displays.

3. Browse to the import (.csv) file location.
4. Select the file and click **Import.**
5. Click **OK** to close the **Configure Names** dialog box.
6. Click **OK** on the confirmation message.

Searching for a device by name

You can search for objects (switch, fabric, product, ports, or N Ports) by name. To search for a name in the Connectivity Map, refer to "Searching for a device" on page 314.

To search by name, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.

2. Select **All Names** from the **Display** list.
3. Select **Name** from the **Scope** list.
4. Enter the name you want to search for in the **Search** field.

   You can search on partial names.

   NOTE
   To search for a device, the device must be discovered and display in the topology.

5. Click **Search**.

   All devices with the specified name (or partial name) are highlighted in the **Display** table. You may need to scroll to see all highlighted names.

   If the search finds no devices, a 'no item found' message displays.

6. Click **OK** to close the **Configure Names** dialog box.

Searching for a device by WWN

You can search for objects (switch, fabric, product, ports, or N Ports) by WWN (world wide name). To search for a WWN in the Connectivity Map, refer to "Searching for a device" on page 314.

To search by WWN, complete the following steps.

1. Select **Configure > Names**.

   The **Configure Names** dialog box displays.

2. Select **All Names** from the **Display** list.
3. Select **WWN** from the **Scope** list.
4. Enter the WWN you want to search for in the **Search** field.

   You can search on partial WWNs.

   NOTE
   To search for a device, the device must be discovered and display in the topology.

5. Click **Search**.

   All devices with the specified WWN (or partial WWN) are highlighted in the **Display** table. You may need to scroll to see all highlighted WWNs.

   If the search finds no devices, a 'no item found' message displays.

6. Click **OK** to close the **Configure Names** dialog box.

Miscellaneous security settings

You can configure the Server Name, login banner, modify whether or not to allow clients to save passwords, and modify whether or not to enforce the MD5 checksum during import. When the login banner is enabled, each time a client connects to the server, the login banner displays with a legal notice provided by you. The client's users must acknowledge the login banner to proceed, otherwise they are logged out.

NOTE
M-EOS device support is no longer available in the Management application; therefore, the **CHAP Secret** and **Retype Secret** fields are no longer required.

Configuring the server name

To configure the server name, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Security Misc** in the **Category** list.

   The **Security Misc** pane displays (Figure 36).

FIGURE 36 Options dialog box (Security Misc pane)

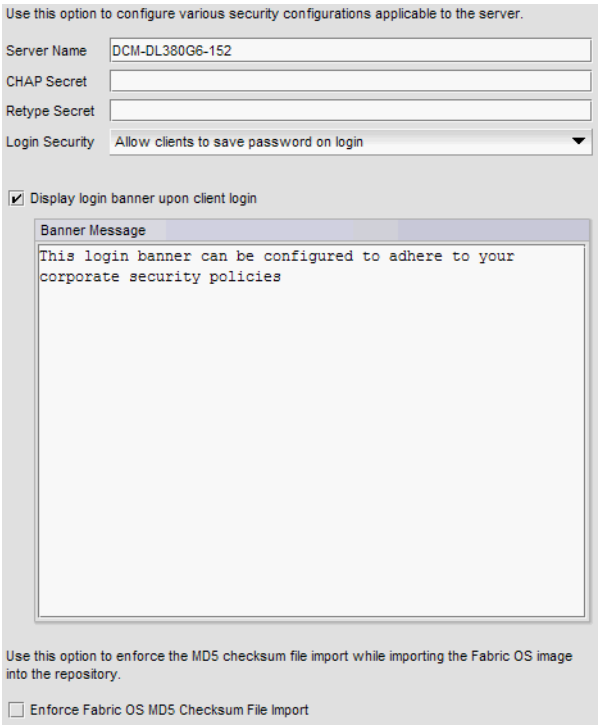

3. Enter the server name in the **Server Name** field.

   The **Server Name** field cannot be empty.

4. Click **OK** on the confirmation message.
5. Click **Apply** or **OK** to save your work.

Enforcing MD5 file during import

NOTE
The MD5 checksum file is required when you load Fabric OS firmware into the Management application version 12.0 or later.

You can configure the Management application to enforce the MD5 checksum file import during the import of the Fabric OS image into the firmware repository.

The MD5 checksum file can be obtained from the Fabric OS product download site in the same location as the firmware file. The MD5 checksum file cannot be downloaded directly from the site; however, you can open the file, copy and paste the contents into a new file, and save the file with the .md5 extension in the same directory as the firmware file.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Security Misc** in the **Category** list.
3. Select the **Enforce Fabric OS MD5 Checksum File Import** check box.
4. Click **Apply** or **OK** to save your work.

Configuring login security

To configure login security, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Security Misc** in the **Category** list.
3. Choose one of the following options:
   - To allow users to save their password in the **Login Security** list, select **Allow clients to save password on login**.
   - To not allow users to save their password in the **Login Security** list, select **Do NOT allow clients to save password on login**.
4. Click **Apply** or **OK** to save your work.

Configuring the login banner display

To configure the login banner display, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Security Misc** in the **Category** list.
3. Select the **Display login banner upon client login** check box.
4. Enter the message you want to display every time a user logs into this server in the **Banner Message** field.

   This field contains a maximum of 2048 characters.

5. Click **Apply** or **OK** to save your work.

Disabling the login banner

To disable the login banner display, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Security Misc** in the **Category** list.
3. Clear the **Display login banner upon client login** check box.

   NOTE
   Users logging into the client will not see the banner when logging in to this Server.

4. Click **Yes** on the confirmation message.
5. Click **Apply** or **OK** to save your work.

Syslog Registration settings

You can automatically register the server as the syslog recipient on products.

Registering a server as a Syslog recipient automatically

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Syslog Registration** in the Category pane.

   The **Syslog Registration** pane displays (Figure 37).

   FIGURE 37 Options dialog box (Syslog Registration pane)

   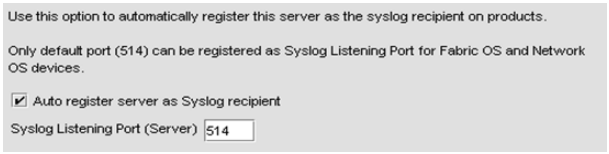

3. Select the **Auto register server as Syslog recipient** check box, if necessary.

   This check box is selected by default.

4. Click **Apply** or **OK** to save your work.

Configuring the Syslog listing port number

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Syslog Registration** in the Category pane.

   The **Syslog Registration** pane displays (Figure 37).

3. Enter the Syslog listening port number of the Server in the **Syslog Listening Port (Server)** field, if necessary.

   The default Syslog listening port number is 514 and is automatically populated.

   For Fabric OS and devices, only the default port (514) can be registered as the Syslog Listening Port.

4. Click **Apply** or **OK** to save your work.

SNMP Trap Registration settings

You can automatically register the server as the trap recipient on products. If SAN products have Informs enabled, the registration is for the Informs.

Registering a server as a SNMP trap recipient automatically

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Trap Registration** in the Category pane.

   The **Trap Registration** pane displays (Figure 38).

   FIGURE 38 Options dialog box (Trap Registration pane)

   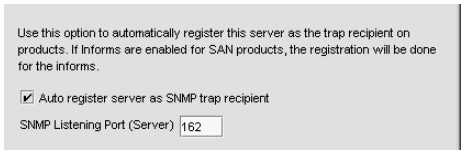

3. Select the **Auto register server as SNMP trap or informs recipient** check box, if necessary.

   This check box is selected by default.

4. Click **Apply** or **OK** to save your work.

Configuring the SNMP trap listing port number

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Trap Registration** in the Category pane.
3. Enter the SNMP listening port number of the Server in the **SNMP Listening Port (Server)** field, if necessary.

   The default SNMP listening port number is 162 and is automatically populated.

4. Click **Apply** or **OK** to save your work.

SNMP Trap forwarding credential settings

You can configure SNMP credentials for the traps forwarded by the server.

Configuring SNMP v1 and v2c credentials

To configure a SNMP v1 or v2c credentials, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Trap Forwarding Credentials** in the Category pane.

   The **Trap Forwarding Credentials** pane displays (Figure 39).

   FIGURE 39 Options dialog box (Trap Forwarding Credentials pane)

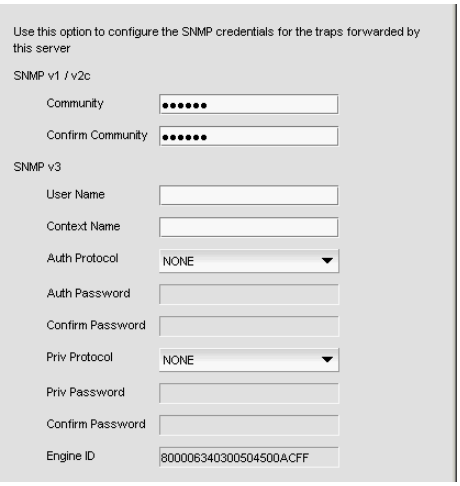

3. Enter the unique community string (case sensitive, 1 to 64 characters). in the **Community** and **Confirm Community** fields.

   Displays as asterisks. Allows all printable ASCII characters.

4. Click **Apply** or **OK** to save your work.

Configuring SNMP v3 credentials

To configure a SNMP v3 credentials, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Trap Forwarding Credentials** in the Category pane.

   The **Trap Forwarding Credentials** pane displays (Figure 39).

3. Enter the SNMP v3 name (case sensitive, 1 to 16 characters) to identify the credentials in the **User Name** field.

   Allows all printable ASCII characters.

4. Select on of the following authentication protocols from the **Auth Protocol** list.

   - HMAC_MD5 (continue with step 5)
   - HMAC_SHA (continue with step 5)
   - NONE (go to step 6)

5. Enter the SNMP v3 authentication password (case sensitive, 1 to 16 characters) in the **Auth Password** and **Confirm Password** fields.

   Displays as asterisks. Allows all printable ASCII characters.

6. Select one of the following privacy protocol types from the **Priv Protocol** list.
   - CBC-DES (continue with step 7)
   - CFB_AES-128 (continue with step 7)
   - CFB_AES_256 (continue with step 7)
   - NONE (go to step 8)
7. Enter the privacy password (case sensitive, 8 to 16 characters) in the **Priv Password** and **Confirm Password** fields.

   Displays as asterisks. Allows all printable ASCII characters.
8. Click **Apply** or **OK** to save your work.

Software Configuration

The Management application allows you to configure the following software settings:

- Certificates — Support settings to allow enhanced diagnostics.
- Client export port settings — A port for communication between the client and server.
- Client/Server IP — IP configuration settings.
- Memory allocation settings — Memory allocation for the client and server.
- Product communication settings — Connections between the server and SAN switches or IP products.
- FTP/SCP/SFTP server settings — Internal or external FTP or SCP server settings.
- Server port settings — Server port settings.
- Support mode settings — Support settings to allow enhanced diagnostics.

Certificates

Certificate management allows you to enable certificate validation between the Management application server and products when HTTPS is enabled and between server and client when SSL is enabled on server. For more information about product communication, refer to "Product communication settings" on page 120.

Certificate management also allows you to manage the Management application server truststore as well as the Management application client truststore. On the Management application server, the truststore is maintained as two separate files: truststore and keystore. A truststore contains certificates from other third-parties with which the Management application server communicates. The truststore file is used when making decisions on what to trust. The server truststore (truststore.jks) is stored in the *Install_Home*/conf/security/ directory. A keystore file stores the Management server's identity and its private key. The server keystore file (keystore.jks) is stored in the *Install_Home*/conf/security directory.

When SSL is enabled on the server, the server presents the keystore certificate to authenticate itself with the client. The Management application client truststore contains certificates from the Management application servers with which the client communicates. The Management application client truststore does not have a private key.

Viewing certificates

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays (Figure 40).

FIGURE 40 Options dialog box (Certificates pane)

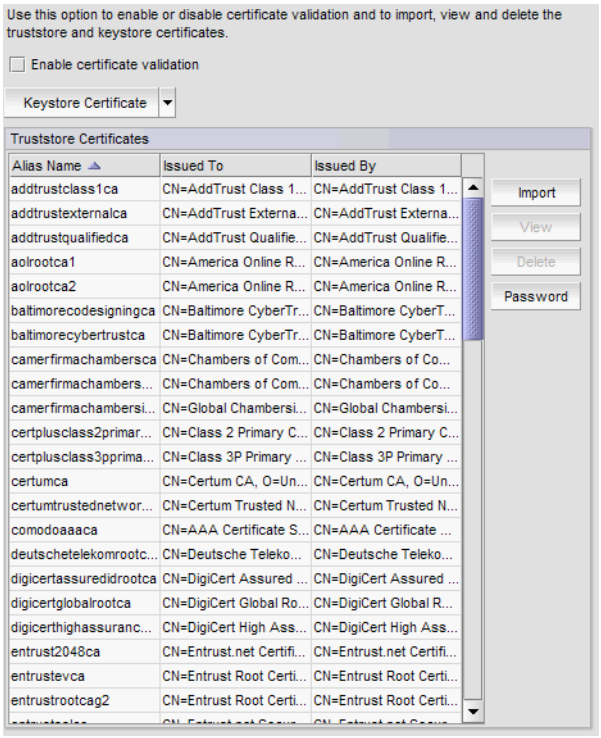

The **Certificates** pane contains the following fields and components:

- **Enable certificate validation** check box — Select to enable certificate validation. Clear to disable certificate validation
- **Keystore Certificates** drop-down list — Select one of the following options:
  - **View** — Click to view the keystore certificate details. For more information, refer to "Viewing a truststore certificate" on page 107.
  - **Export** — Click to export a keystore certificate. For more information, refer to "Importing a truststore certificate" on page 108.
  - **Replace —** Click to replace the keystore certificate. For more information, refer to "Deleting a truststore certificate" on page 108.
  - **Change Password —** Click to change the password for the keystore. For more information, refer to "Changing the keystore password" on page 111.

- **Truststore Certificates** table — Contains the following fields and components:
  - **Alias Name** — Unique alias of the certificate.
  - **Issued To** — To whom the certificate was issued.
  - **Issued By** — Author of the certificate.
  - **Import** button — Click to import a certificate. For more information, refer to "Importing a truststore certificate" on page 108.
  - **View** button — Click to view the certificate details. For more information, refer to "Viewing a truststore certificate" on page 107.
  - **Delete** button — Click to delete the certificate. For more information, refer to "Deleting a truststore certificate" on page 108.
  - **Password** button — Click to change the password for the trusstore. For more information, refer to "Changing the password for the truststore repository" on page 109.

3. Click **Apply** or **OK** to save your work.

Viewing a truststore certificate

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays.

3. Select a truststore in the **Truststore Certificates** table.

4. Click **View**.

   The **Details - Certificate** *Name* dialog box displays (Figure 41).

FIGURE 41 Details - Certificate *Name* dialog box

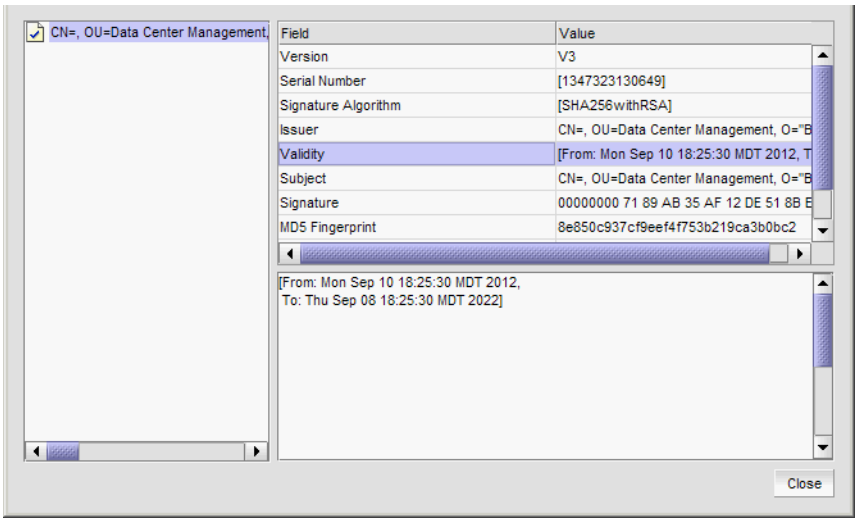

The **Details - Certificate** *Name* dialog box contains the following fields:

- Left-side text box — Name of the Issuer.

- Right-side table — Displays the following certificate details:
  - Version — Version of the certificate.
  - Serial Number — Serial number of the certificate.
  - Signature Algorithm — Signature algorithm used to sign the certificate. The signature algorithm is derived from the algorithm of the underlying private key. For example, if the underlying private key is of type "RSA", the default signature algorithm is "SHA256withRSA".
  - Issuer — Entity that signed the certificate.
  - Validity — Dates that the certificate is valid.
  - Subject — Name of the entity whose public key the certificate identifies.
  - Signature — Digital signature of the certificate.
  - MD5 Fingerprint — MD5 fingerprint used to authenticate the public key.
  - SHA1 Fingerprint — SHA1 fingerprint used to authenticate the public key
- Right-side text box — Displays the value for the field selected in the table above.

5. Click **Close**.
6. Click **OK** on the **Options** dialog box.

Importing a truststore certificate

NOTE
The Managment application supports importing the PKCS#12 certificate format which uses the ".pfx" file extention.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays.

3. Click **Import**.
4. Browse to the location of the new certificate.
5. Enter a unique alias for the certificate in the Alias Name field.
6. Click **OK**.
7. Click **Apply** or **OK** to save your work.

Deleting a truststore certificate

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays.

3. Select the truststore you want to delete in the **Truststore Certificates** table.
4. Click **Delete**.
5. Click **Yes** on the confirmation message.

   The truststore is deleted from the **Truststore Certificates** table.

6. Click **Apply** or **OK** to save your work.

   The truststore is deleted from the server truststore.

Changing the password for the truststore repository

To change the keystore password, refer to "Changing the keystore password" on page 111.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays.

3. Select a truststore in the **Truststore Certificates** table.
4. Click **Password**.

   The **Truststore Password** dialog box displays.

5. Enter the current password in the **Old Password** field.
6. Enter the new password in the **New Password** and **Confirm New Password** fields.

   The password can be from 6 through 256 characters long, case sensitive, and allows all printable ASCII characters. The password displays as asterisks.

7. Click **OK**.

   The password is cached locally in the client.

8. Click **Apply** or **OK** to save your work.

   The password is saved to the server.

Viewing a keystore certificate

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays.

3. Select **View** from the **Keystore Certificate** list.

   The **Details - Certificate** *Name* dialog box displays with the following fields:

   - Left-side text box — Name of the Issuer.

- Right-side table — Displays the following certificate details:
  - Version — Version of the certificate.
  - Serial Number — Serial number of the certificate.
  - Signature Algorithm — Signature algorithm used to sign the certificate. The signature algorithm is derived from the algorithm of the underlying private key. For example, if the underlying private key is of type "RSA", the default signature algorithm is "SHA256withRSA".
  - Issuer — Entity that signed the certificate.
  - Validity — Dates that the certificate is valid.
  - Subject — Name of the entity whose public key the certificate identifies.
  - Signature — Digital signature of the certificate.
  - MD5 Fingerprint — MD5 fingerprint used to authenticate the public key.
  - SHA1 Fingerprint — SHA1 fingerprint used to authenticate the public key
  - Public Key — Public key used for the certificate.
- Right-side text box — Displays the value for the field selected in the table above.

4. Click **Close**.
5. Click **OK** on the **Options** dialog box.

Exporting a keystore certificate

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays.

3. Select **Export** from the **Keystore Certificate** list.

   The **Export Keystore Certificate** - *Name* dialog box displays.

4. Browse to the location to which you want to export the certificate.
5. Click **OK**.
6. Click **Apply** or **OK** to save your work.

Replacing a keystore certificate

NOTE
Changes to this option take effect after an application restart.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays.

3. Select **Replace** from the **Keystore Certificate** list.

   The **Replace Keystore Certificate** dialog box displays.

4. To replace the current certificate with a new self-signed certificate, select the **A new self signed certificate** option.

5. To replace the current certificate with a certificate file, select the **Certificate File** option and complete the following steps.
   a. Browse to the location of the new certificate.
   b. Enter the password for the new certificate in the **Password** field.

   The new certificate is cached locally in the client.
6. Click **Apply** or **OK** to save your work.

   The new certificate is saved to the server.
7. Click **OK** on the "changes take effect after application restart" message.

Changing the keystore password

NOTE
Changes to this option take effect after an application restart.

1. Select **Server > Options**.

   The **Options** dialog box displays.
2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays.
3. Select **Change Password** from the **Keystore Certificate** list.

   The **Keystore Password** dialog box displays.
4. Enter the current password in the **Old Password** field.
5. Enter the new password in the **New Password** and **Confirm New Password** fields.
6. Click **OK**.
7. Click **Apply** or **OK** to save your work.

Enabling and disabling certificate validation

The Management application server only validates the certifying authority and the date in the certificate.

Certificate validation requires HTTPS connections between the server and the switches. To configure product communication to HTTPS, refer to "Product communication settings" on page 120.

1. Select **Server > Options**.

   The **Options** dialog box displays.
2. Select **Certificates** to in the **Category** list.

   The **Certificates** pane displays.
3. Select the **Enable certificate validation** check box.

   Clear the check box to disable certificate validation.
4. Click **Apply** or **OK** to save your work.

Client export port settings

You can configure a port for communication between the client and server.

Configuring the client export port

To configure client export port settings, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Client Export Port** to assign a communications port between the client and server in the **Category** list.

   The **Client Export Port** pane displays (Figure 42).

   FIGURE 42 Options dialog box (Client Export Port pane)

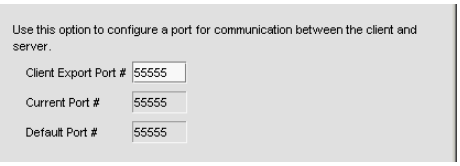

3. Enter the client export port number to set a fixed port number for the client in the **Client Export Port** field.

   The current port number displays in the **Current Port #** field.

   The default port number (55555) displays in the **Default Port #** field.

4. Click **Apply** or **OK** to save your work.

   NOTE
   Changes to this option take effect after a client restart.

5. Click **OK** on the "changes take effect after client restart" message.

Client/Server IP

You can configure connections between the client or switches and the Management application server.

Configuring the server IP address

If your Operating System is IPv4-enabled or IPv6-enabled (running in dual mode), the server binds using an IPv4 address. IPv6 only mode does not support server to client communication (the IPv6 address cannot be bound to the server).

NOTE
If the Management server or client has multiple Network Interface Cards and if any of these interfaces are not plugged in, you must disable them; otherwise, the following features do not work properly:

Server impact

- Configuration wizard (does not display all IP addresses)
- Trap and Syslog auto registration
- Report content (Ipconfiguration element does not display all server IP addresses)

- Network OS configuration backup through FTP
- Trace dump through FTP

Client impact

- Options dialog box (does not display all IP addresses)
- Firmware import and download dialog box
- Firmware import for Fabric OS and Network OS products
- FTP button in Technical Support Repository dialog box
- Technical supportSave of Fabric OS, Network OS, and Host products through FTP

To configure the IP address used by the server for client-server communications, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Client/Server IP** in the **Category** list to set the IP address.

   The **Client/Server IP** pane displays (Figure 43).

   FIGURE 43 Options dialog box (Client/Server IP option)

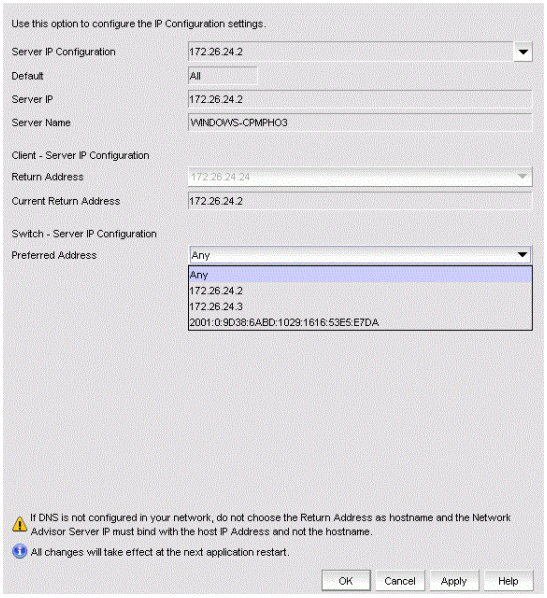

3. Choose one of the following options in the **Server IP Configuration** list.

   - Select **All**. Go to step 4.
   - Select a specific IP address. Continue with step 5.
   - Select **localhost**. Continue with step 5.

   When **Server IP Configuration** is set to **All**, you can select any available IP address as the **Return Address**. If you select a specific IP address, the **Return Address** list shows the same IP address and you cannot change it.

   If DNS is not configured for your network, do not select the 'hostname' option from either the **Return Address** or **Preferred Address** list. Selecting the 'hostname' option prevents clients and devices from communicating with the Server.

4. Select the return IP address in the **Client - Server IP Configuration Return Address** list.
5. Select one of the following options from the **Switch - Server IP Configuration Preferred Address** list:
   - **Any** - Any one of the reachable Server IP address is used for the FFDC, FTP, SCP, SFTP, SNMP trap, Syslog registration, Technical Support Save, and Firmware Download functions. It is not mandatory that the same IP address will be used for all the functions.
   - **Specific IP address** - The IP address selected is used for the FFDC, FTP, SCP, SFTP, SNMP trap, Syslog registration, Technical Support Save, and Firmware Download functions. When a specific IP address is selected, the Management application will not check reachability for any product or device.

     NOTE
     When the Management application is installed, the IP address selected by the user under **Switch - Server IP Configuration Preferred Address** list in the **Server IP Configuration** screen is taken as the default IP address for the FFDC, FTP, SCP, SFTP, SNMP trap, Syslog registration, Technical Support Save, and Firmware Download function.

   If DNS is not configured in your network, do not choose the Return Address as hostname and the the Management application IP must bind with the host IP address and not the host name.
6. Click **Apply** or **OK** to save your work.

   NOTE
   Changes to this option take effect after an application restart.

   NOTE
   You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).
7. Click **OK** on the "All changes will take effect at the next application restart" message.

Configuring an explicit server IP address

If you selected a specific IP address from the **Server IP Configuration** screen during installation and the selected IP address changes, you will not be able to connect to the server. To connect to the new IP address, you must manually update the IP address information.

To change the IP address, complete the following steps.

1. Choose one of the following options:
   - On Windows systems, select **Start > Programs >** *Management_Application* **12.X.X >** *Management_Application* **Configuration**.
   - On UNIX systems, execute `sh` *Install_Home*`/bin/configwizard` on the terminal.
2. Click **Next** on the **Welcome** screen.
3. Click **Yes** on the confirmation message.
4. Click **Next** on the **FTP Server** screen.
5. Complete the following steps on the **Server IP Configuration** screen (Figure 44).

FIGURE 44 Server IP Configuration screen

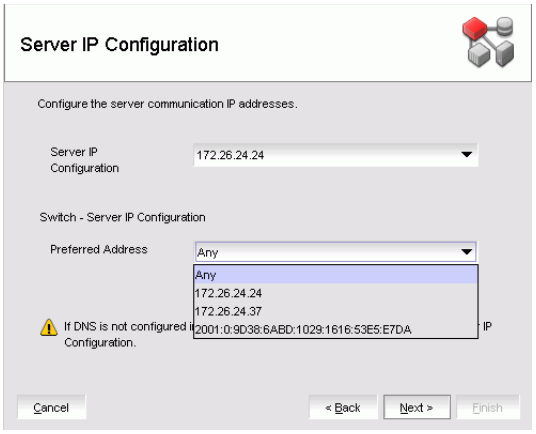

a. Select an address from the **Server IP Configuration** list.

NOTE
The host name does not display in the list if it contains invalid characters. Valid characters include alphanumeric and dash (-) characters. The IP address is selected by default.

If DNS is not configured for your network, do not select the "hostname" option from the **Server IP Configuration** list. Selecting the "hostname" option prevents clients and devices from communicating with the server.

b. Select an IP address from the **Switch - Server IP Configuration Preferred Address** list. The preferred IP address is used for switch and server communication.

or

Select **Any** from the **Switch - Server IP Configuration Preferred Address** list to enable switch and server communication with one of the reachable IP address present in the server. By default, **Any** option is selected.

c. Click **Next**.

6. Click **Next** on the **Server Configuration** screen.
7. Click **Next** on the **SMI Agent Configuration** screen.
8. Verify the IP address on the **Server Configuration Summary** screen and click **Next**.
9. Click **Finish** on the **Start Server** screen.
10. Click **Yes** on the restart server confirmation message.
11. Choose one of the following options:
    - If you configured authentication to CAC, enter your PIN in the CAC PIN field.
    - If you configured authentication to the local database, an external server (RADIUS, LDAP, or TACACS+), or a switch, enter your user name and password.

      The defaults are **Administrator** and **password**, respectively.

      NOTE
      Do not enter *Domain\User_Name* in the **User ID** field for LDAP server authentication.
12. Click **Login**.

13. Click **OK** on the **Login Banner**.

    NOTE
    When you launch the Management application or navigate to a new view, the **SAN** tab displays with a gray screen over the Product List and Topology Map while data is loading.

Configuring the application to use dual network cards

Issues with Client-to-Server connectivity can be due to different reasons. Some examples are:

- The computer running the Server has more than one network interface card (NIC) installed.
- The computer running the Server is behind a firewall that performs network address translation.

To make sure that Clients can connect to the Server, you may need to edit the IP configuration setting in the **Options** dialog to manually specify the IP address that the Server should use to communicate to its Clients.

NOTE
If your Operating System is IPv4-enabled or IPv6-enabled (dual mode), the server binds using IPv4 address by default.

NOTE
IPv6 only mode does not support server to client communication (the IPv6 address cannot be bound to the server).

To configure the IP address to override the default RMI server host IP address, complete the following steps.

NOTE
This configuration option replaces the -Djava.rmi.server.hostname value used in previous releases.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Client/Server IP** in the **Category** list to set the IP address.
3. Choose one of the following options in the **Server IP Configuration** list.
   - Select **All**. Go to step 4.
   - Select a specific IP address. Continue with step 5.
   - Select **localhost**. Continue with step 5.
4. Select the return IP address in the **Client - Server IP Configuration Return Address** list.

   When **Server IP Configuration** is set to **All**, you can select any available IP address as the **Return Address**. If you select a specific IP address, the **Return Address** field shows the same IP address and you cannot change it.

5. Click **Apply** or **OK** to save your work.

   NOTE
   Changes take effect after you restart the Management Server.

   NOTE
   You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).

6. Click **OK** on the "changes take effect after "application restart" message.

Memory allocation settings

You can configure memory allocation for the client and server to improve performance. You can trigger switch polling when a state changes or you can poll at intervals when no state change occurs.

NOTE
SAN size is a consideration in the selection of polling periods.

Configuring memory allocation settings

To configure memory allocation settings, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Memory Allocation** in the **Category** list to set the memory allocation for the server and client.

   The **Memory Allocation** pane displays (Figure 89).

3. (Enterprise only) In the **SAN Network Size** list, complete the following steps.

   For other editions, the SAN Network size is small. You cannot change the SAN Network size.

   a. Select the size of the SAN (small, medium, or large) you want to configure.

      Product and port recommended counts change to the new default values when you change the SAN Network size. Recommended counts are as follows:

      - Small SAN — 40 products, 2,000 ports
      - Medium SAN — 100 products, 5,000 ports
      - Large SAN — 400 products, 15,000 ports

        NOTE
        For full performance management and dashboard functionality, the **Large** option of the SAN Enterprise edition only supports 5000 switch ports on a 32-bit system.

      Memory and asset polling values change to the new default values when you change the SAN Network size. You may increase these values. For default values, refer to step 4 and step 5.

   b. Click **OK** on the confirmation message.

4. Enter the memory allocation (MB) for the client in the **Client Memory Allocation** field.

   If you enter an invalid value, an error message displays with the minimum value allowed. Click **OK** and edit the value again.

   The current configured number of megabytes for client memory allocation displays in the **Current Value** field. The default minimum number of megabytes for client memory allocation displays in the **Default Minimum** field.

   The default minimum Client Heap Size is 950 MB. To manage more than 9000 ports, increase the memory allocation for the client to 2048 MB.

   NOTE
   There is no restriction on the Client Heap Size value. The correct Client Heap Size value should be given according to the RAM present in the server when it is launched.

   NOTE
   For a 32-bit server, configuring a value higher than 1024 MB impacts the client launch.

5. Enter the memory allocation (MB) for the server in the **Server Memory Allocation** field.

   If you enter an invalid value, an error message displays with the minimum value allowed. Click **OK** and edit the value again.

   The current configured number of megabytes for server memory allocation displays in the **Current Value** field. The default minimum number of megabytes for server memory allocation displays in the **Default Minimum** field. The IP address of the server displays in the **Server IP** field. The server name displays in the **Server Name** field.

   For 64-bit servers, the minimum Server Heap Size is 2048 MB and the maximum is 3072 MB for all network sizes. The default Server Heap Size is 3072 MB. To manage more than 9000 ports, increase the memory allocation for the server to 6144 MB.

   NOTE
   There is no restriction on the maximum value for server heap size in a 64-Bit server. The correct server heap size value must be given according to the RAM present in the server.

   For 64-bit servers (Professional Plus and Enterprise), the default values for the CIMOM Heap Size are as follows:

   - Small: 1536 MB
   - Medium: 2048 MB
   - Large: 3072 MB

   For 32-bit servers (Professional only), the default Server Heap Size is 768 MB (minimum and maximum).

6. Click **Apply** or **OK** to save your work.

   NOTE
   Changes to this option take effect after an application restart.

   NOTE
   You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).

7. Click **OK** on the “changes take effect after application restart” message.

Configuring asset polling

Asset polling allows you set the length of time between state change polling. To maximize the efficiency of the polling feature (balancing the amount of possible information with any possible performance impact), base your settings on the size of the SAN.

To configure asset polling, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Memory Allocation** in the **Category** list to set the memory allocation for the server and client.

   The **Memory Allocation** pane displays (Figure 89).

3. Enter how often you want to check for state changes in the **Check for state change every** field.

   This is the short tick interval which is used for adaptive asset collection (for more information, refer to "SAN data collection" on page 77). Valid values are from 1 through 600 seconds. You cannot enter a value lower than the default minimum value.

   Default minimum values are as follows:

   - Small/0–2000 ports (Professional): 60 seconds
   - Medium/2000–5000 ports: 120 seconds
   - Large/5000 or more ports: 180 seconds

4. Enter how often you want to check for state changes in the **If no state change, Poll switch every** field.

   This is the lazy polling interval which is used to schedule data collection when not triggered by an event (for more information, refer to "SAN data collection" on page 77). Valid values are from 1 through 3,600 seconds. Default values are as follows:

   - Small/0–2000 ports (Professional): 120 seconds
   - Medium/2000–5000 ports: 900 seconds
   - Large/5000 or more ports: 1800 seconds

5. Click **Apply** or **OK** to save your work.

   NOTE
   Changes to this option take effect after an application restart.

   NOTE
   You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).

6. Click **OK** on the "changes take effect after application restart" message.

Viewing the network size status

The Management application enables you to view the network size status at a glance by providing a status icon on the status bar. Double-click the icon to launch the **Memory Allocation** pane of the **Options** dialog box.

NOTE
If you exceed the recommended count, the network size status icon refreshes when the license is refreshed (every three hours) or after a client restart.

NOTE
The recommended count is the supported scalability limit based on the network size. If the maximum license count is less than the recommended count, the license count displays as the recommended count.

Table 18 illustrates and describes the icons that indicate the current network size status.

TABLE 18 Network size status

Icon	Description
	This icon displays when the network size is within the recommended count.
	This icon displays when the network size exceeds the recommended count. This icon displays when any of the following counts are exceeded: • SAN Product Count • SAN Port Count • Fabric Count

Product communication settings

You can configure HTTP or HTTPS connections between the products and the Management application server.

Configuring SAN communication

To configure connections between the SAN devices and the Management application server, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Product Communication** from the **Software Configurations** list in the **Category** pane.

   The **Product Communication** pane displays (Figure 45).

FIGURE 45 Options dialog box (Product Communication pane)

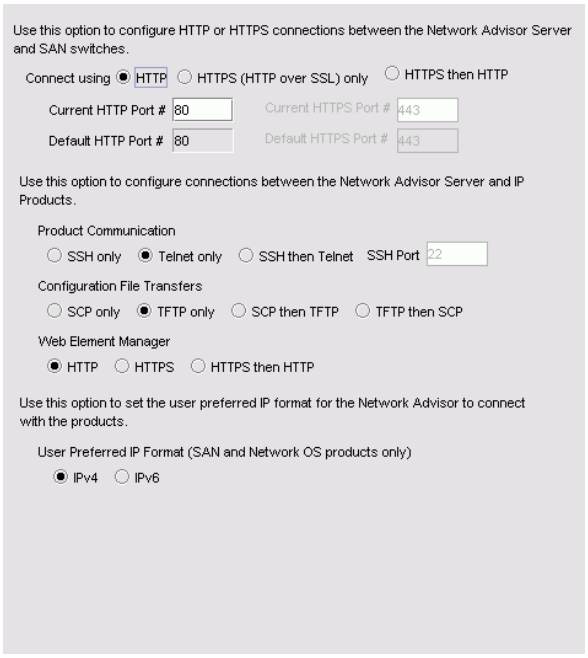

3. To connect using HTTP, complete the following steps.
   a. Select the **Connect using HTTP** option.
   b. Enter the connection port number in the **Port #** field. Continue with step 6.

      The default HTTP port number is 80.

      NOTE
      To manage FIPS-enabled Fabric OS fabrics, you must configure Product Communication using the **Connect using HTTPS (HTTP over SSL) only** or **HTTPS then HTTP** option.

4. To connect using HTTPS (HTTP over SSL), complete the following steps.
   a. Select the **Connect using HTTPS (HTTP over SSL) only** option.
   b. Enter the connection port number in the **Port #** field. Continue with step 6.

      The default HTTPS port number is 443.

5. To connect using HTTPS then HTTP, complete the following steps.
   a. Select the **Connect using HTTPS then HTTP** option.
   b. Enter the connection port number in the **Current Port #** field. Continue with step 6.

      The default HTTPS port number is 443 and the default HTTP port number is 80.

6. (Fabric OS and Network OS products only) Select **IPv4** or **IPv6** to set the preferred IP format.
7. Click **Apply** or **OK** to save your work.

Configuring the preferred IP format

To configure the preferred IP format for the Management application server to connect with Fabric OS and Network OS devices, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Product Communication** from the **Software Configurations** list in the **Category** pane.

   The **Product Communication** pane displays (Figure 45).

3. (Fabric OS and Network OS products only) Select **IPv4** (default) or **IPv6** to set the preferred IP format.

4. Click **Apply** or **OK** to save your work.

FTP/SCP/SFTP server settings

> NOTE
> For FIPS-enabled Fabric OS switches, you must configure the FTP/SCP/SFTP server communication to an external SCP server to download firmware and allow technical support.

File Transfer Protocol (FTP) is a network protocol used to transfer data from one computer to another over a TCP computer network. During installation, a built-in FTP server and its services are installed. Other FTP servers on your system are recognized by the application as external FTP servers.

For Windows systems, the built-in FTP server is the default configuration and installation starts the FTP service if port 21 is not used by any other FTP server. For UNIX systems, built-in FTP is the default for UNIX systems during installation; the external FTP server is the default only if port 21 is busy.

Note that when uninstalling the application the built-in FTP server is removed with all other services even if the FTP service is used by firmware upgrade or supportSave features.

> NOTE
> FTP is supported on all Fabric OS devices.

Secure Copy (SCP) is a means of securely transferring computer files between a local and a remote host or between two remote hosts, using the Secure Shell (SSH) protocol. You must configure SCP on your machine to support Technical Support and firmware download.

> NOTE
> SCP is supported on Fabric OS devices running 7.0 and later.

SSH File Transfer Protocol (SFTP) is a network protocol used to transfer data from one computer to another over a secure channel. You must configure SCP on your machine to support Technical Support and firmware management.

> NOTE
> SFTP is supported on Fabric OS devices running 7.0 and later.

The built-in SCP/SFTP servers use the port 22 by default.

To view the port status for the FTP and SCP/SFTP servers, refer to "Viewing port status" on page 12.

Accessing the FTP server folder

Choose from one of the following options to access the FTP server folder:

- To access the internal FTP folder, select **Monitor > Techsupport > View Repository**.
- To access the external FTP folder, type the following in a browser window: ftp://*Username*@*External_FTP_Server_IP_Address* (for example, ftp://admin@10.1.1.1) and press **Enter**. Type your password in the pop-up window and press **Enter**. The external FTP folder displays.

Configuring an internal FTP server

To configure the internal FTP server settings, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **FTP/SCP/SFTP** in the **Category** list.

   The **FTP/SCP/SFTP** pane displays (Figure 46).

   FIGURE 46 Options dialog box (FTP/SCP/SFTP pane)

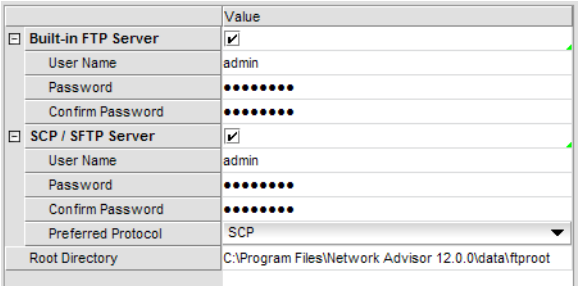

3. Select the **Use built-in FTP/SCP/SFTP Server** option to use the default built-in FTP server.

   All active fields are mandatory. The default user name is admin. The full path to the built-in FTP directory displays in the **Root Directory** field.

4. Select the **Built-in FTP Server** check box.
5. Change your password by entering a new password in the **Password** and **Confirm Password** fields.

   The default password is passw0rd (where 0 is a zero).

6. Click **Test** to test the FTP server.

   An "FTP Server running successfully" or an error message displays.

   If you receive an error message, make sure your credentials are correct, the server is running, the remote directory path exists, and you have the correct access permission; then try again.

7. Click **Apply** or **OK** to save your work.

Configuring an internal SCP or SFTP server

NOTE
SCP is supported on Fabric OS devices running 7.0 and later.

NOTE
SFTP is supported on Fabric OS devices running 7.0 and later.

To configure the internal SCP or SFTP server settings, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **FTP/SCP/SFTP** in the **Category** list.

   The **FTP/SCP/SFTP** pane displays (Figure 46).

   FIGURE 47 Options dialog box (FTP/SCP/SFTP pane)

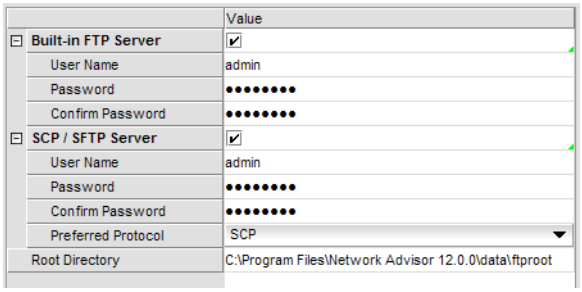

3. Select the **Use built-in FTP/SCP/SFTP Server** option to use the default built-in SCP or SFTP server.

   All active fields are mandatory. The default user name is admin. The full path to the built-in SCP or SFTP directory displays in the **Root Directory** field.

4. Select the **SCP/SFTP Server** check box.
5. Change your password by entering a new password in the **Password** and **Confirm Password** fields.

   The default password is passw0rd (where 0 is a zero).

6. Select the protocol (**SCP** or **SFTP**) from the **Preferred Protocol** list.
7. Click **Test** to test the server.

   An "SCP/SFTP Server running successfully" or an error message displays.

   If you receive an error message, make sure your credentials are correct, the SCP/SFTP server is stopped, the remote directory path exists, and you have the correct access permission; then try again.

8. Click **Apply** or **OK** to save your work.

Configuring an external FTP, SCP, or SFTP server

NOTE
For FIPS-enabled Fabric OS switches, you must configure the FTP/SCP/SFTP server communication to an external SCP or SFTP server to download firmware and allow technical support.

NOTE
SCP is supported on Fabric OS devices running 7.0 and later.

NOTE
SFTP is supported on Fabric OS devices running 7.0 and later.

To configure external FTP, SCP, or SFTP server settings, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **FTP/SCP/SFTP** in the **Category** list.

   FIGURE 48 The **FTP/SCP/SFTP** pane displays (Figure 48).Options dialog box (FTP/SCP/SFTP pane)

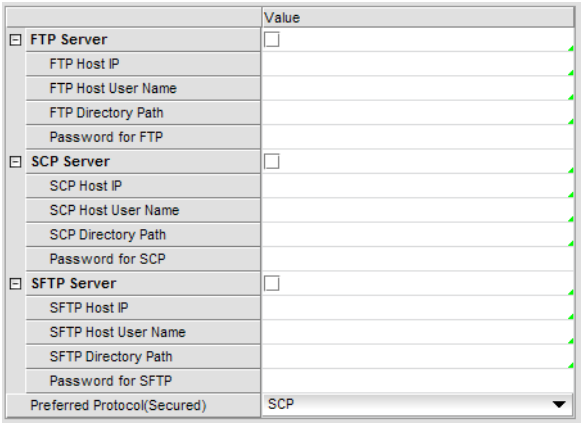

3. Select the **Use External FTP Server and/or SCP Server** option.
4. To configure an external FTP server, complete the following steps.
   a. Select the **FTP Server** check box to configure the external FTP server.

      All fields are mandatory.

   b. Enter the IP address for the remote host in the **Remote Host IP** field.
   c. Enter a user name in the **Remote Host User Name** field.
   d. Enter the path to the remote host in the **Remote Directory Path** field.

      Use a slash (/) or period (.) to denote the root directory.

   e. Enter the password in the **Password Required for FTP** field.
5. To configure an external SCP server, complete the following steps.
   a. Select the **SCP Server** check box to configure the external SCP server.

      All fields are mandatory.

b. Enter the IP address for the remote host in the **SCP Host IP** field.
c. Enter a user name in the **SCP Host User Name** field.
d. Enter the path to the remote host in the **SCP Directory Path** field.

   Use a slash (/) or period (.) to denote the root directory.
e. Enter the password in the **Password Required for SCP** field.
f. Select **SCP** from the **Preferred Protocol (Secured)** list.

6. To configure an external SFTP server, complete the following steps.
   a. Select the **SFTP Server** check box to configure the external SCP server.

      All fields are mandatory.
   b. Enter the IP address for the remote host in the **SFTP Host IP** field.
   c. Enter a user name in the **SFTP Host User Name** field.
   d. Enter the path to the remote host in the **SFTP Directory Path** field.

      Use a slash (/) or period (.) to denote the root directory.
   e. Enter the password in the **Password Required for SFTP** field.
   f. Select **SFTP** from the **Preferred Protocol (Secured)** list.
7. Click **Test** to test the server.

   A "Server running successfully" or an error message displays.

   If you receive an error message, make sure your credentials are correct, the server is running, the remote directory path exists, and you have the correct access (read and write) permissions; then try again.
8. Click **OK** on the message.
9. Click **Apply** or **OK** to save your work.

Testing the FTP, SCP, and SFTP server

To test the FTP, SCP, or SFTP server, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.
2. Select **FTP/SCP/SFTP** in the **Category** list.
3. Choose one or more of the following options:
   - If you are using the internal FTP server, select the **Use built-in FTP/SCP/SFTP Server** option.

     For step-by-step instructions about configuring the built-in server, refer to "Configuring an internal FTP server" on page 123.
   - If you are using the external FTP server, select the **Use external FTP/SCP/SFTP Server** option.

     For step-by-step instructions about configuring the built-in server, refer to "Configuring an external FTP, SCP, or SFTP server" on page 125.

4. Click **Test**.

   An "FTP, SCP, or SFTP Server running successfully" or an error message displays.

   If you receive an error message, make sure your credentials are correct, the server is running, the remote directory path exists, and you have the correct access permission; then try again.

5. Click **OK** on the message.
6. Click **OK** to close the **Options** dialog.

Server port settings

You can configure the server port settings so that you can assign a web server port number and set the server port to be SSL-enabled.

Configuring the server port

To configure server settings, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Server Port** in the **Category** list.

   The **Server Port** pane displays (Figure 49).

   FIGURE 49 Options dialog box (Server Port pane)

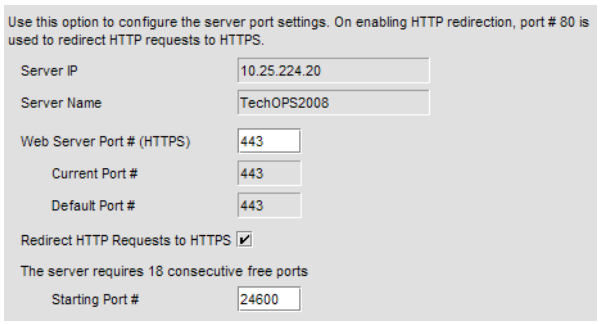

3. Enter a port number in the **Web Server Port # (HTTPS)** field.

   The default is 443.

4. Enable HTTP redirection to HTTPS by selecting the **Redirect HTTP Requests to HTTPS** check box.

   When you enable HTTP redirection, the server uses port 80 to redirect HTTP requests to HTTPS. Make sure that port 80 is available before you enable HTTP redirection.

5. Enter a port number in the **Starting Port #** field. The default is 24600.

   For Professional, the server requires 15 consecutive free ports beginning with the starting port number.

   For Trial and Licensed versions, the server requires 18 consecutive free ports beginning with the starting port number.

6. Click **Apply** or **OK** to save your work.

   NOTE
   Changes to this option take effect after application restart.

7. Click **OK** on the “changes take effect after application restart” message.

Support mode settings

You can configure support settings to allow enhanced diagnostics.

Configuring support mode settings

To configure support mode settings, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays (Figure 50).

2. Select **Support Mode** in the **Category** list.

   NOTE
   Only use this option when directed to by customer support.

   The **Support Mode** pane displays (Figure 49).

   FIGURE 50 Options dialog box (Support Mode pane)

   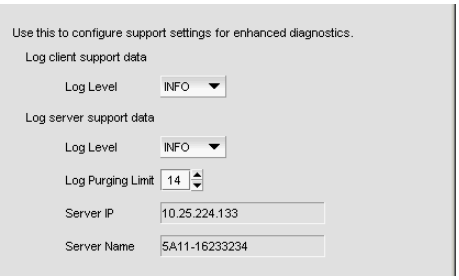

3. Select the **Log client support data - Log Level** list, and select the type of log data you want to configure.

   Log level options include: **All**, **Fatal**, **Error**, **Warn**, **Info**, **Debug**, **Trace**, and **Off**. Default is **Info**.

4. Select the **Log server support data - Log Level** list, and select the type of log data you want to configure.

   Log level options include: **All**, **Fatal**, **Error**, **Warn**, **Info**, **Debug**, **Trace**, and **Off**. Default is **Info**.

5. Click **Apply** or **OK** to save your work.

   NOTE
   Changes to the server log levels reset to the default (INFO) after a server restart.

   NOTE
   Changes to the **Log client support data** log level is persisted on all clients launched from the same machine for the same server.

   client. log file properties

   - Client logs are collected separately for each server. After successful login, a log file is created and prefixed with the network address provided in the **Login** dialog box.

     For example, 172.26.1.1.client.log or localhost.client.log

Each log file is limited to 5 MB. When a file reaches the maximum size, and there are less than 5 log files for the Client, a new file is created.

- For local clients, log files (*network_address*.client.log.1 through *network_address*.client.log.5) are created in the *User_Home*/*Product_Name*/*localhost* directory.
- For web start clients, log files (*network_address*.client.log.1 through *network_address*.client.log.5) are created in the *User_Home*/*Product_Name*/*Server_IP_Address* directory.

server. log file properties

- There is only one server.log file each day with no log size limit.
- The server.log file rolls over at 12:00 midnight everyday.
- When the log file rolls over, it is compressed and renamed using the following file name format:

  server.yyyy-mm-dd.log.zip
  for example, server.2010-04-14.log.zip, server.2010-04-15.log.zip, and so on
- For servers, log files are created in the *Install_Home*/logs/server directory.

Configuring the server log file purge limit

To configure server log file purging, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Support Mode** in the **Category** list.

   NOTE
   Only use this option when directed to by customer support.

3. Select the maximum number of days to retain the server log file in the **Log Purging Limit** field.

   Valid values are 1 through 90. Default is 14.

   The log files are purged at 1:00 AM on the day after the retention period ends.

4. Click **Apply** or **OK** to save your work.

FIPS Support

To manage FIPS-enabled Fabric OS fabrics and switches, make sure you complete the following configuration requirements:

- Configure Product Communication to HTTPS (refer to "Configuring SAN communication" on page 120) to allow communication between the server and the Fabric OS switches.
- Configure an external SCP server (refer to "Configuring an external FTP, SCP, or SFTP server" on page 125) to allow firmware download, product technical support, and supportSave.

Fabric tracking

When you discover a new fabric and initial discovery is complete, fabric tracking is automatically enabled. Subsequently, if a switch or end-device is added to or removed from the fabric, a plus (+) or minus (-) icon displays (see table below) next to the product icon. Connections are also tracked. A new connection displays a solid gray line with an added icon and missing connections display a yellow dashed line with a removed icon.

	Device Added
	Device Removed

When you enable fabric tracking and a switch is missing from the fabric, a warning level call home event (Switch ***Switch_WWN*** is missing from the Fabric ***Fabric_Name***) is generated in the Master Log and a call home alert is sent to the corresponding call center for this event.

To avoid call home events for missing switches, create a call home event filter and clear the **Switch is missing from the Fabric** check box in the Available Call Home Event Types table. Once you create the call home event filter, assign it to the appropriate call center. To create a call home event filter, refer to "Defining an event filter" on page 341.

Enabling fabric tracking

1. Enable fabric tracking by choosing one of the following options:
   - Select a fabric on the Product List or Connectivity Map and select **Monitor > Track Fabric Changes**.
   - Right-click a fabric on the Product List or Connectivity Map and select **Track Fabric Changes**.

   The accept changes summary message displays. This message includes the following information:
   - **Do not show me this again** check box — Select if you do not want to see this dialog box again when you enable or disable fabric tracking or accept changes for a switch or fabric.
   - **Switches** — This table shows a brief summary of the switches including status (whether the device port will be added () or removed ( ) from the fabric), name, fabric name, IP address, WWN, and domain ID. This table includes unmonitored switches which becomes segmented from the fabric.
   - **Device Ports** — This table shows a brief summary of the device ports including status (whether the device port will be added () or removed () from the fabric), reason (why the device is missing), product type, port, fabric name, port WWN, node WWN, and attached port number.
   - **Connections** — This table shows a brief summary of the switch connections including the status (whether the device port will be added () or removed () from the fabric), reason (why the connection is missing), and connection type as well as the fabric name, WWN, domain ID, IP address, and port number of the connected switches.

     The reason for the missing device or connection requires devices running Fabric OS 7.2 or later.
2. Click **Yes** to accept changes.

Disabling fabric tracking

1. Disable fabric tracking by choosing one of the following options:
   - Select the fabric on which you want to disable fabric tracking on the Product List or Connectivity Map and select **Monitor > Track Fabric Changes**.
   - Right-click the fabric on which you want to disable fabric tracking on the Product List or Connectivity Map and select **Track Fabric Changes**.

   The accept changes summary message displays. This message includes the following information:

- **Do not show me this again** check box — Select if you do not want to see this dialog box again when you enable or disable fabric tracking or accept changes for a switch or fabric.
- **Switches** — This table shows a brief summary of the switches including status (whether the device port will be added ( ) or removed ( ) from the fabric), name, fabric name, IP address, WWN, and domain ID. This table includes unmonitored switches which becomes segmented from the fabric.
- **Device Ports** — This table shows a brief summary of the device ports including status (whether the device port will be added ( ) or removed ( ) from the fabric), reason (why the device is missing), product type, port, fabric name, port WWN, node WWN, and attached port number.
- **Connections** — This table shows a brief summary of the switch connections including the status (whether the device port will be added ( ) or removed ( ) from the fabric), reason (why the connection is missing), and connection type as well as the fabric name, WWN, domain ID, IP address, and port number of the connected switches.

2. Click **Yes**.

Accepting changes for a fabric

1. Accept the changes to a fabric by choosing one of the following options:
   - Select a fabric on the Product List or Connectivity Map and select **Monitor > Accept Changes**.
   - Right-click a fabric on the Product List or Connectivity Map and select **Accept Changes**.

   The accept changes summary message displays. This message includes the following information:

FIGURE 51 Accept changes summary message

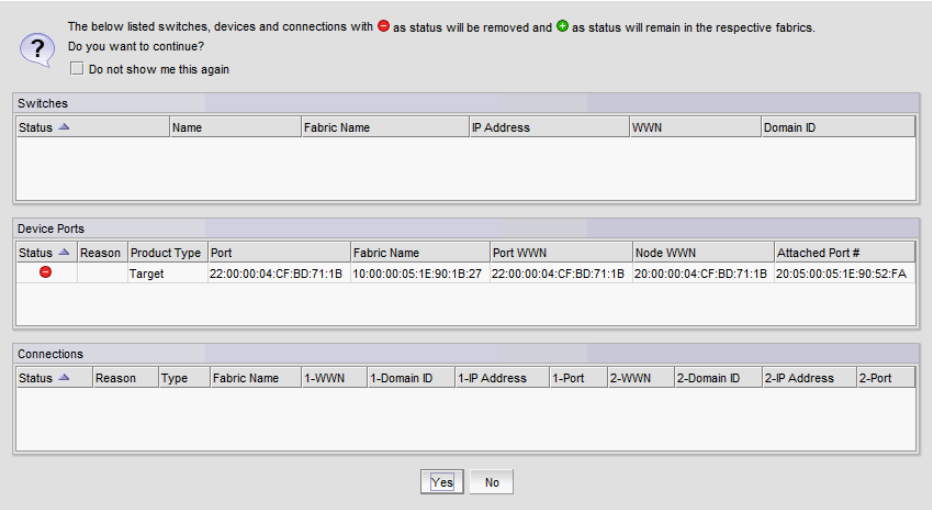

- **Do not show me this again** check box — Select if you do not want to see this dialog box again when you enable or disable fabric tracking or accept changes for a switch or fabric.
- **Switches** — This table shows a brief summary of the switches including status (whether the device port will be added ( ) or removed ( ) from the fabric), name, fabric name, IP address, WWN, and domain ID. This table includes unmonitored switches which becomes segmented from the fabric.
- **Device Ports** — This table shows a brief summary of the device ports including status (whether the device port will be added ( ) or removed ( ) from the fabric), reason (why the device is missing), product type, port, fabric name, port WWN, node WWN, and attached port number.

- **Connections** — This table shows a brief summary of the switch connections including the status (whether the device port will be added ( ) or removed ( ) from the fabric), reason (why the connection is missing), and connection type as well as the fabric name, WWN, domain ID, IP address, and port number of the connected switches.

2. Click **Yes** to accept changes.

Accepting changes for all fabrics

1. Accept the changes to all fabrics by choosing one of the following options:
   - Click in the white space on the Connectivity Map and select **Monitor > Accept All Changes**.
   - Right-click in the white space on the Connectivity Map and select **Accept All Changes**.

   The accept changes summary message displays. This message includes the following information:

FIGURE 52 Accept all changes summary message

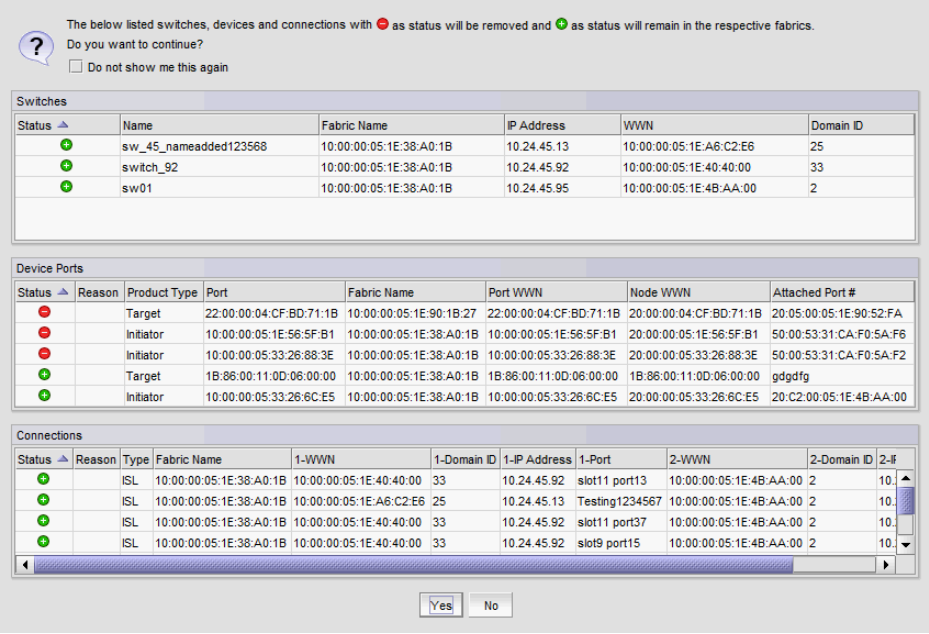

- **Do not show me this again** check box — Select if you do not want to see this dialog box again when you enable or disable fabric tracking or accept changes for a switch or fabric.
- **Switches** — This table shows a brief summary of the switches including status (whether the device port will be added ( ) or removed ( ) from the fabric), name, fabric name, IP address, WWN, and domain ID. This table includes unmonitored switches which becomes segmented from the fabric.
- **Device Ports** — This table shows a brief summary of the device ports including status (whether the device port will be added ( ) or removed ( ) from the fabric), reason (why the device is missing), product type, port, fabric name, port WWN, node WWN, and attached port number.
- **Connections** — This table shows a brief summary of the switch connections including the status (whether the device port will be added ( ) or removed ( ) from the fabric), reason (why the connection is missing), and connection type as well as the fabric name, WWN, domain ID, IP address, and port number of the connected switches.

2. Click **Yes** to accept changes.

Accepting changes for a switch, access gateway, or phantom domain

1. Accept the changes to a switch, access gateway, or phantom domain by choosing one of the following options:

   - Select the switch, access gateway, or phantom domain on the Product List or Connectivity Map and select **Monitor > Accept Changes**.
   - Right-click the switch, access gateway, or phantom domain on the Product List or Connectivity Map and select **Accept Change**.

   The accept changes summary message displays. This message includes the following information:

   - **Do not show me this again** check box — Select if you do not want to see this dialog box again when you enable or disable fabric tracking or accept changes for a switch or fabric.
   - **Switches** — This table shows a brief summary of the switches including status (whether the device port will be added () or removed () from the fabric), name, fabric name, IP address, WWN, and domain ID. This table includes unmonitored switches which becomes segmented from the fabric.
   - **Device Ports** — This table shows a brief summary of the device ports including status (whether the device port will be added () or removed () from the fabric), reason (why the device is missing), product type, port, fabric name, port WWN, node WWN, and attached port number.
   - **Connections** — This table shows a brief summary of the switch connections including the status (whether the device port will be added () or removed () from the fabric), reason (why the connection is missing), and connection type as well as the fabric name, WWN, domain ID, IP address, and port number of the connected switches.

2. Click **Yes** to accept changes.

User Account Management

Users overview

The Management application allows you to manage accounts of users who manage devices on the network. When a user logs in to the Management application, the user name and password can be authenticated and authorized by the local server or by a supported external server.

User accounts are assigned privileges, which you define within roles. Each privilege provides access to a specific feature of the Management application. This enables you to maintain privileges common to a group of administrators within a role, instead of in individual accounts.

You can group devices, access points, and their groups in areas of responsibilities (AORs), then assign one or more AORs to a user's privilege. When you assign a user an AOR, that user will be able to manage only the devices in that AOR. Devices in a user's AOR are the only devices that user sees in device trees and on the **Dashboard** tab. You can place selected devices, device groups, port groups, access points, access point groups, and access point port groups in an AOR.

Users who create a device group are the only users who can manage the devices in that group. Other users may view the groups, but do not have the ability to add, delete, or modify the groups.

Configuration requirements

To administer accounts on the Management application server, you must have an administrative login on the platform on which the Management application is running. Use the "Administrator" login to create other logins with administrative permissions.

Viewing configured users

To view configured users, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click the **Users** tab, if necessary.

FIGURE 53 Users dialog box - Users tab

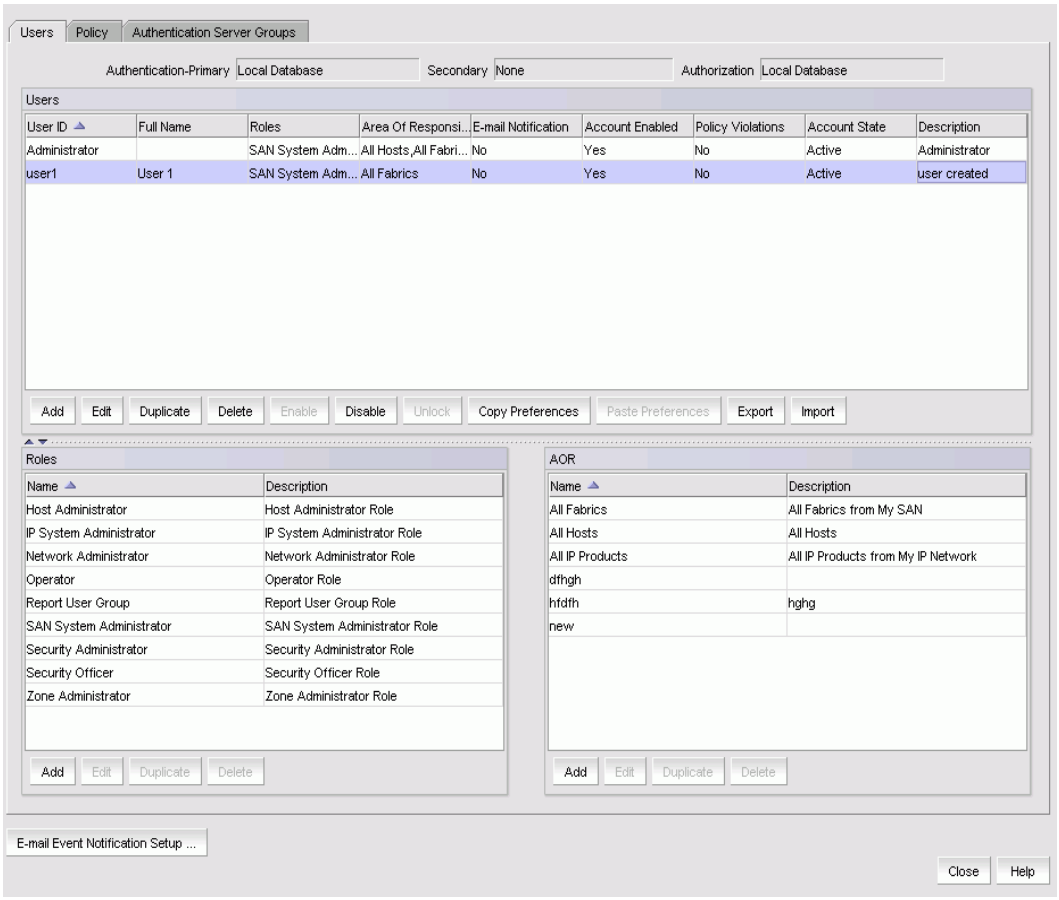

The **Users** dialog box contains the following fields and components:

- Authentication-Primary — The primary authentication server type configured through the Server Management Console.
- Secondary — The secondary authentication server type configured through Server Management Console.
- Authorization — The authorization source configured through the Server Management Console.
- **Users** table — The configured users.
  - **User ID** — The unique name used to identity a user.
  - **Full Name** — The user's full name.
  - **Roles** — List of roles the user belongs to separated by commas.
  - **Area Of Responsibility** — List of Area Of Responsibility (AORs) the user belongs to separated by commas.
  - **E-mail Notification** — Whether e-mail notification is enabled for the user.
  - **Account Enabled** — Whether the user account status is enabled.
  - **Policy Violations** — Whether there is a current policy violation for the user.
  - **Account State** — The current account state for the user. Options include:
    - Active
    - Locked by User manager
    - Password Expired
    - Password format policy violated
    - Password history policy violated

  - Locked out threshold reached
- **Description** — The description of the user.
- **Add** button — Click to launch the **Add Users** dialog box and configure a new user (refer to "Creating a new user account" on page 138).
- **Edit** button — Click to launch the **Edit Users** dialog box for the selected user (refer to "Editing a user account" on page 139).
- **Duplicate** button — Click to launch the **Duplicate Users** dialog box for the selected user (refer to "Copying a user account" on page 140).
- **Delete** button — Click to delete the selected users (refer to "Deleting a user account" on page 143).
- **Enable** button — Select to enable the selected users (refer to "Enabling a user account" on page 143). Disabled if the selected user is already enabled.
- **Disable** button — Select to disable the selected users (refer to "Disabling a user account" on page 142). Disabled if the selected user is already disabled.
- **Unlock** button — Select to unlock the selected user account (refer to "Unlocking a user account" on page 143).
- **Copy User Preferences** button — Select to copy user preferences from the selected user account (refer to "Copying and pasting user preferences" on page 140).
- **Paste User Preferences** button — Select to paste user preferences from the selected user account (refer to "Copying and pasting user preferences" on page 140).
- **Export** button — Select to export the selected user profile (refer to "Exporting a user account" on page 141).
- **Import** button — Select to import the selected user profile (refer to "Importing a user account" on page 141).

- **Roles** table — Lists the default system roles and any user-defined roles.
  - **Name** — The unique name of the role.

  Default system roles for SAN only environments include:

    - SAN System Administrator
    - Network Administrator
    - Security Administrator
    - Zone Administrator
    - Operator
    - Security Officer
    - Host Administrator
  - **Description** — A description of the role.
  - **Add** button — Click to add a new role (refer to "Creating a new role" on page 144).
  - **Edit** button — Click to edit the selected role (refer to "Editing a role" on page 145).
  - **Duplicate** button — Click to copy the selected role (refer to "Copying a role" on page 145).
  - **Delete** button — Click to delete the selected role (refer to "Deleting a role" on page 146).
- **AOR** table — Lists the default system AOR and any user-defined AORs.
  - **Name** — The unique name of the AOR. Default system AORs include:
    - **All Fabrics** — All discovered SAN devices.
    - **All Hosts** — All discovered Hosts devices.
    - **All IP Products** — All discovered IP devices.
  - **Description** — A description of the AOR.
  - **Add** button — Click to launch the **Add AOR** dialog box.
  - **Edit** button — Click to launch the **Edit AOR** dialog box for the selected AOR. You cannot edit system AORs.
  - **Duplicate** button — Click to launch the **Duplicate AOR** dialog box for the selected AOR. You cannot duplicate system AORs.
  - **Delete** button — Click to delete the selected AOR. You cannot delete system AORs.

- **E-mail Event Notification Setup** button — Click to configure e-mail event notification (refer to "Configuring e-mail notification" on page 156).

3. Click **Close** to close the **Users** dialog box.

User accounts

NOTE
You must have User Management Read and Write privileges to add new accounts, set passwords for accounts, and apply roles to the accounts. For a list of privileges, refer to "User Privileges" on page 1333.

Management application user accounts contain the identification of the Management application user, as well as privileges, roles, and AORs assigned to the user. Privileges provide access to the features in Management application. A role is a group of selected privileges. A role can be assigned to one or more Management application users who need access to the same menu options.

An AOR contains selected fabrics and devices that a Management application user is allowed to manage.

Creating a new user account

To create a new user account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click **Add** under the **Users** table.

   The **Add User** dialog box displays.

   FIGURE 54 Add User dialog box

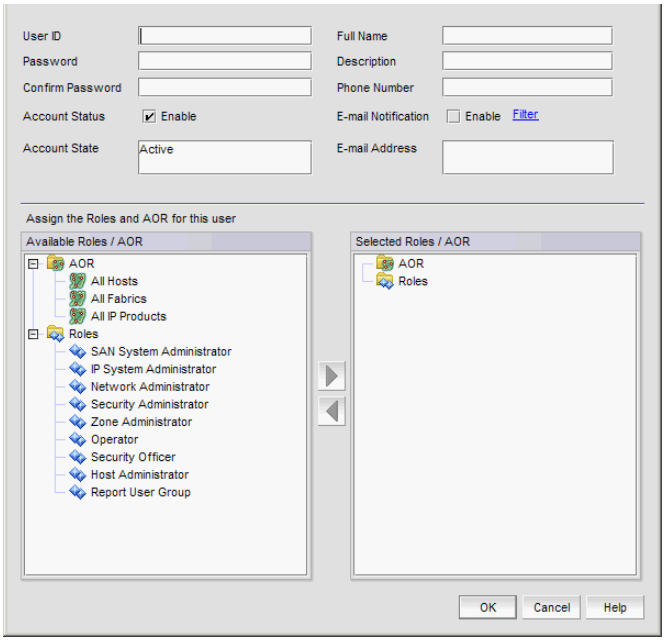

3. Enter a unique name to identify the user in the **User ID** field.

4. Enter a password for the user in the **Password** and **Confirm Password** fields.

   Passwords displays as dots (.). For password policy details, refer to "Viewing your password policy" on page 155.

5. Select the **Account Status - Enable** check box to enable the account of the user.

   Account Status is enabled by default.

6. (Optional) Enter the full name of the user in the **Full Name** field.
7. (Optional) Enter a description for the user in the **Description** field.
8. (Optional) Enter the phone number of the user in the **Phone Number** field.
9. Select the **E-mail Notification - Enable** check box to enable e-mail notification for the user.

   E-mail Notification is disabled by default.

10. Click **Filter** to set up basic event filters for the user.

    For step-by-step instructions about setting up basic event filters, refer to "Setting up basic event filtering" on page 1133.

11. Enter the e-mail address of the user in the **E-mail Address** field.

    Enter more than one e-mail address, separating each with a semi-colon. To send a text message or page via e-mail, use the following format *number@carrier.com*, where *number* is your phone number and *carrier.com* is the SMS server. For example, 3035551212@txt.att.net (text message) or 3035551212@page.att.net (page).

    NOTE
    Check with your carrier for the exact e-mail address.

12. Assign roles and AORs by selecting the role or AOR in the **Available Roles / AOR** table and click the right arrow button to move the role or AOR to the **Selected Roles / AOR** table.

    Select multiple roles or AORs by holding down the CTRL key and clicking more than one role or AOR.

13. Remove roles and AORs by selecting the role or AOR in the **Selected Roles / AOR** table and click the left arrow button to move the role or AOR to the **Available Roles / AOR** table.

    Select multiple roles or AORs by holding down the CTRL key and clicking more than one role or AOR.

14. Click **OK** to save the new user and close the **Add User** dialog box.

    The new user account displays in the **Users** table of the **Users** dialog box. You must assign at least one role to a user account. Users without an assigned role cannot log in to the client.

15. Click **Close** to close the **Users** dialog box.

Editing a user account

To make changes to an existing user account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the user account you want to edit and click **Edit** under the **Users** table.

The **Edit User** dialog box displays.

3. Complete step 3 through step 13 in "Creating a new user account" on page 138.
4. Click **OK** to save the user account and close the **Edit User** dialog box.

   If you make changes to the user's role or AOR while the user is logged in, a confirmation message displays. When you click **OK** on the confirmation message, the user is logged out and must log back in to see the changes.
5. Click **Close** to close the **Users** dialog box.

Copying a user account

You can create a user account by copying an existing one. When you copy an account, you copy the selected roles and AORs of that account. You can then enter a new user name, ID, e-mail address, and telephone number.

To create a new user account from an existing account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.
2. Select the user account you want to copy and click **Duplicate** under the **Users** table.

   The **Duplicate User** dialog box displays.
3. Complete step 3 through step 13 in "Creating a new user account" on page 138.
4. Click **OK** to save the new user and close the **Duplicate User** dialog box.

   The new user account displays in the **Users** table of the **Users** dialog box.
5. Click **Close** to close the **Users** dialog box.

Copying and pasting user preferences

You can copy user preference settings, such as window and dialog box sizes, table column and sort order, as well as other customizations, and all the user-defined views (including fabrics and hosts) from the selected user account to one or more other user accounts.

If the fabric and hosts from the original user account are not included in the other user's AOR, then the copied fabrics and hosts do not display in the other user's views. To include fabrics and hosts from the original user account, you must add them to the other user's account (refer to "Exporting a user account" on page 141).

If a user-created view with the same name already exists in the other user's views, user-defined views with the same name are ignored. For example, user_acct1 (copy) has the following user-defined views: Fabric1, Fabric2, and Host1 and user_acct2 (paste) has the following user-defined views: Fabric1, Fabric_CO, and Hosts. When you paste the user_acct1 user preferences to user_acct2, user_acct2 now has the following user-defined views: Fabric1, Fabric2, Fabric_CO, Host1, and Hosts.

NOTE
You cannot copy user preferences to user accounts that are currently logged in to the Management application.

NOTE
You cannot copy user preferences to the original user account.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the user account you want to copy user preferences from and click **Copy User Preferences** under the **Users** table.
3. Select the user account you want to copy user preferences to and click **Paste User Preferences** under the **Users** table.

   If you need to make any other changes to this user account, refer to "Editing a user account" on page 139.
4. Click **Yes** on the confirmation message.
5. Click **Close** to close the **Users** dialog box.

Exporting a user account

To export a user account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.
2. Select the user account you want to export and click **Export** under the **Users** table.

   The **Export User** dialog box displays.
3. Browse to the location you want to save the file.
4. Enter a name for the exported user profile data in the **File Name** field, if needed.

   Export uses the following naming convention: <*Flavor*>-UserProfile-<*Time stamp*>.zip.
5. Click **Save**.The file is saved to the location you selected.

   If the export is successful, the following message displays:

   User profile data exported successfully to <*Flavor*>-UserProfile-<*Time stamp*>.zip

Importing a user account

To import a user account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.
2. Click **Import** under the **Users** table.

   The **Import User** dialog box displays.
3. Select the file you want to import.

   The file will be imported to the **Users** table.

   Import uses the following naming convention: <*Flavor*>-UserProfile-<*Time stamp*>.zip.
4. Click **OK**.

   If the import is successful, the following message displays:

   User profile data imported successfully. Restart the Server for the changes to take effect.

Assigning roles and areas of responsibility to a user account

To assign roles and AORs to an existing user account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the user account you want to edit and click **Edit** under the **User**s table.

   The **Edit User** dialog box displays.

3. Assign roles and AORs by selecting the role or AOR in the **Available Roles / AOR** table and click the right arrow button to move the role or AOR to the **Selected Roles / AOR** table.

   Select multiple roles or AORs by holding down the CTRL key and clicking more than one role or AOR.

4. Click **OK** to save the user account and close the **Edit User** dialog box.

   If you make changes to the user's role or AOR while the user is logged in, a confirmation message displays. When you click **OK** on the confirmation message, the user is logged out and must log back in to see the changes.

5. Click **Close** to close the **Users** dialog box.

Removing roles and areas of responsibility from a user account

To remove roles and AORs from an existing user account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the user account you want to edit and click **Edit** under the **User**s table.

   The **Edit User** dialog box displays.

3. Remove roles and AORs by selecting the role or AOR in the **Selected Roles / AOR** table and click the left arrow button to move the role or AOR to the **Available Roles / AOR** table.

   Select multiple roles or AORs by holding down the CTRL key and clicking more than one role or AOR.

4. Click **OK** to save the user account and close the **Edit User** dialog box.

   If you make changes to the user's role or AOR while the user is logged in, a confirmation message displays. When you click **OK** on the confirmation message, the user is logged out and must log back in to see the changes.

5. Click **Close** to close the **Users** dialog box.

Disabling a user account

To make the user account inactive, but keep it in the database, you can disable the user account.

NOTE
You cannot disable the default "Administrator" account.

To disable a user account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the enabled user account you want to disable in the **Users** table and click **Disable**.
3. Click **Yes** on the confirmation message.

   If currently accessing the server, the user will be logged out once the user account is disabled. The user cannot log back in until you re-enable the user account.

4. Click **Close** to close the **Users** dialog box.

Enabling a user account

To re-activate a user account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the disabled user account you want to enable in the **Users** table and click **Enable**.
3. Click **Yes** on the confirmation message.
4. Click **Close** to close the **Users** dialog box.

Deleting a user account

NOTE
You cannot delete the default "Administrator" user account.

To permanently delete a user account from the server, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the user you want to delete in the **Users** table and click **Delete**.
3. Click **Yes** on the confirmation message.

   If currently accessing the server, the user will be logged out once the user account is deleted.

4. Click **Close** to close the **Users** dialog box.

Unlocking a user account

NOTE
You must have User Management Read and Write privileges to unlock a user account.

You can unlock a user account when a user is locked out of the system because of too many invalid login attempts.

To unlock a user account, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the locked user account you want to unlock in the **Users** table and click **Unlock**.
3. Click **Yes** on the confirmation message.
4. Click **Close** to close the **Users** dialog box.

Roles

NOTE
You must have User Management Read and Write privileges to view, add, modify, or delete roles.

A role is a group of Management application tasks or privileges that can be assigned to several users who have similar functions.

When you create a role, it immediately becomes available in the **Users** dialog box.

Creating a new role

To create a new role, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click **Add** under the **Roles** table.

   The **Add Role** dialog box displays.

   FIGURE 55 Add Role dialog box

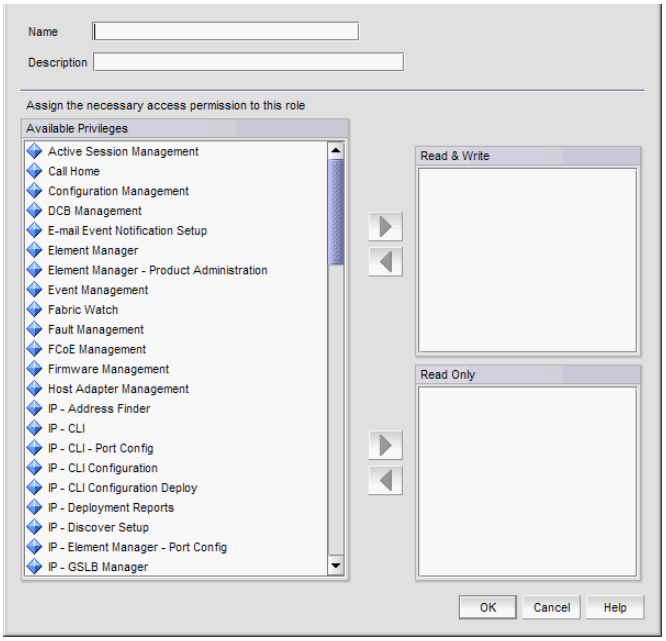

3. Enter a name of the role in the **Name** field.

4. (Optional) Enter a short description for the role in the **Description** field.
5. Add or remove privileges as needed.

   For step-by-step instructions, refer to "Adding privileges to a role" on page 146 or "Removing privileges from a role" on page 146.
6. Click **OK** to save the new role and close the **Add Role** dialog box.

   The new role displays in the **Roles** list of the **Users** dialog box. To add users to this role, follow the instructions in "Exporting a user account" on page 141.
7. Click **Close** to close the **Users** dialog box

Editing a role

To make changes to an existing role, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.
2. Select the role you want to edit in the **Roles** table and click **Edit**.

   The **Edit Role** dialog box displays.
3. Complete step 3 through step 5 in "Creating a new role" on page 144.
4. Click **OK** to save the role and close the **Edit Role** dialog box.

   If you make changes to the user's role or AOR while the user is logged in, a confirmation message displays. When you click **OK** on the confirmation message, the user is logged out and must log back in to see the changes.
5. Click **Close** to close the **Users** dialog box.

Copying a role

You can create a new role by copying an existing one. When you copy a role, you copy the selected privileges in that role.

To copy an existing role, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.
2. Select the role you want to copy in the **Roles** table and click **Duplicate**.

   The **Duplicate Role** dialog box displays.
3. Complete step 3 through step 5 in "Creating a new role" on page 144.
4. Click **OK** to save the role and close the **Duplicate Role** dialog box.

   The new role displays in the **Roles** list of the **Users** dialog box. To add users to this role, follow the instructions in "Exporting a user account" on page 141.
5. Click **Close** to close the **Users** dialog box.

Deleting a role

To delete a role, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the role you want to delete in the **Roles** table and click **Delete**.
3. Click **Yes** on the confirmation message.
4. Click **Close** to close the **Users** dialog box.

Adding privileges to a role

Each option under the Management application main menu corresponds to a privilege. By adding a privilege to a role and assigning that role to a user, you give the user access to a feature of the Management application. When a user logs in to the Management application, the user sees only the options that correspond to the privileges listed in the **Add Roles**, **Edit Roles**, or **Duplicate Roles** dialog box.

To add privileges to a role, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click **Add**, **Edit**, or **Duplicate** under the **Roles** table.

   The **Add Roles**, **Edit Roles**, or **Duplicate Roles** dialog box displays.

3. Add read and write access by selecting the features to which you want to allow read and write access in the **Available Privileges** list and click the right arrow button to move the features to the **Read & Write Privileges** list.

   Select multiple features by holding down the CTRL key and clicking more than one privilege. The features are moved to the **Read & Write Privileges** list.

4. Add read-only access by selecting the features to which you want to allow read-only access in the **Available Privileges** list and click the right arrow button to move the features to the **Read Only Privileges** list.

   Select multiple features by holding down the CTRL key and clicking more than one privilege. The features are moved to the **Read Only Privileges** list.

5. Click **OK** to save your work.
6. Click **Close** to close the **Users** dialog box.

Removing privileges from a role

You remove privileges from the **Edit** or **Duplicate Users** dialog boxes.

To remove privileges from role, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the role you want to edit in the **Roles** table and click **Edit** or **Duplicate** under the **Roles** table.

   The **Edit Roles** or **Duplicate Roles** dialog box displays.

3. Remove read and write access by selecting the features to which you want to remove read and write access in the **Read & Write Privileges** list and click the left arrow button to move the features to the **Available Privileges** list.

   Select multiple features by holding down the CTRL key and clicking more than one privilege. The features are moved to the **Available Privileges** list.

4. Remove read-only access by selecting the features to which you want to remove read-only access in the **Read Only Privileges** list and click the right arrow button to move the features to the **Available Privileges** list.

   Select multiple features by holding down the CTRL key and clicking more than one privilege. The features are moved to the **Available Privileges** list.

5. Click **OK** to save your work.
6. Click **Close** to close the **Users** dialog box.

Areas of responsibility

NOTE
You must have User Management Read and Write privileges to view, add, modify, or delete operational areas of responsibility.

An area of responsibility (AOR) allows you to place Fabrics and Hosts into management groups that can be assigned to an Management application user. Users can manage only the Fabrics and Hosts in the AOR assigned to them, because only devices their AOR display in the Product List and Topology Map.

For example, devices 10.10.10.1, 10.10.10.2, and 10.10.14.3 may be placed in AOR Group 1. This AOR group can then be assigned to UserA. When using the Management application, UserA will be able to create configurations, generate reports, and perform backups only to entries in AOR Group 1 (which consists of devices 10.10.10.1, 10.10.10.2, and 10.10.14.3).

Creating an AOR

When creating an AOR, you assign devices or groups to that AOR. After you save the AOR, it can be assigned to one or more user account. Users of those accounts can then view the devices or groups in their assigned AOR. Users can deploy configurations and payloads only to devices in assigned AORs.

When you create an AOR, it immediately becomes available in the **Users** dialog box.

To create an AOR, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click **Add** under the **AOR** table.

   The **Add AOR** dialog box displays.

FIGURE 56 Users dialog box - Users tab

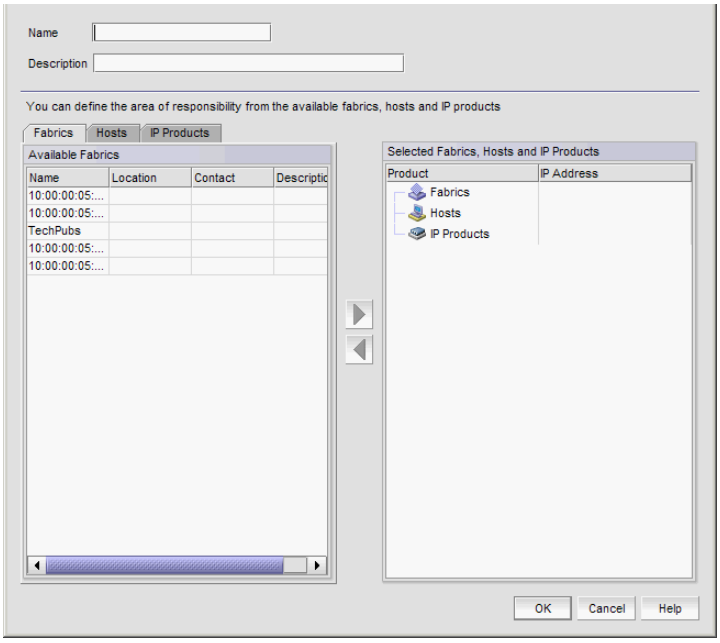

3. Enter a name of the AOR in the **Name** field.
4. (Optional) Enter a short description for the AOR in the **Description** field.
5. Assign or remove products as needed.

   For step-by-step instructions, refer to "Assigning products to an AOR" on page 149 or "Removing products from an AOR" on page 150.
6. Click **OK** to save the new AOR and close the **Add AOR** dialog box.

   The new AOR displays in the **AOR** list of the **Users** dialog box.
7. Click **Close** to close the **Users** dialog box.

Editing an AOR

NOTE
You cannot edit system AORs.

To make changes to an existing AOR, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.
2. Select the AOR you want to edit in the **AOR** table and click **Edit**.

   The **Edit AOR** dialog box displays.
3. Complete step 3 through step 5 in "Creating an AOR" on page 147.
4. Click **OK** to save the AOR and close the **Edit AOR** dialog box.

If you make changes to the user's role or AOR while the user is logged in, a confirmation message displays. When you click **Yes** on the confirmation message, the user is logged out and must log back in to see the changes.

5. Click **Close** to close the **Users** dialog box.

Copying an AOR

NOTE
You cannot duplicate system AORs.

To create a new AOR by copying an existing one, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the AOR you want to copy in the **AOR** table and click **Duplicate**.

   The **Duplicate AOR** dialog box displays.

3. Complete step 3 through step 5 in "Creating an AOR" on page 147.
4. Click **OK** to save the new AOR and close the **Duplicate AOR** dialog box.

   The new AOR displays in the **AOR** table of the **Users** dialog box. To add this AOR to a user, follow the instructions in "Exporting a user account" on page 141.

5. Click **Close** to close the **Users** dialog box

Deleting an AOR

NOTE
You cannot delete system AORs.

To delete an AOR, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select the AOR you want to delete in the **AOR** table and click **Delete**.
3. Click **Yes** on the confirmation message.
4. Click **Close** to close the **Users** dialog box.

Assigning products to an AOR

You can assign fabrics and hosts to an AOR from the **Add**, **Edit**, or **Duplicate AOR** dialog box.

To assign fabrics and hosts to an AOR, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click **Add**, **Edit**, or **Duplicate** under the **AOR** table.

   The **Add AOR**, **Edit AOR**, or **Duplicate AOR** dialog box displays.

3. Click the **Fabrics** tab.
4. Select the fabrics you want to assign to the AOR in the **Available Fabrics** table and click the right arrow button to move the products to the **Selected Products** table.

   Select multiple fabrics by holding down the CTRL key and clicking more than one fabric.

5. Click the **Hosts** tab.
6. Select the hosts you want to assign to the AOR in the **Available Hosts** table and click the right arrow button to move the products to the **Selected Products** table.

   Select multiple hosts by holding down the CTRL key and clicking more than one host.

7. Click **OK** to save your work
8. Click **Close** to close the **Users** dialog box.

Removing products from an AOR

You can remove fabrics and hosts from and AOR from the **Edit AOR** or **Duplicate AOR** dialog box.

To remove fabrics and hosts from the AOR, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click **Edit** or **Duplicate** under the **AOR** table.

   The **Edit AOR** or **Duplicate AOR** dialog box displays.

3. In the **Selected Products** table, select the products or groups you want to remove and click the left arrow button.

   Select multiple products or groups by holding down the CTRL key and clicking more than one item.

4. Click **OK** to save your work.
5. Click **Close** to close the **Users** dialog box.

Password policies

> NOTE
> You must have User Management Read and Write privileges to configure password policy.

Passwords are an important aspect of computer security. They are the front line of protection for user accounts. The purpose of the password policy is to establish a standard for the creation of strong passwords, the protection of those passwords, and the frequency of change.

Configuring a password policy

To configure password policies for all user accounts, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click the **Policy** tab.

3. Configure the password expiration by completing the following steps.

   a. Enter the maximum number of days that can elapse before a password must be changed by the user in the **Password Age** field.

      Valid values are 0 through 999. The default is 0, which means the policy is disabled.

   b. Enter the number of days to warn the user prior to password expiration in the **Warning Period** field.

      Only enabled when the **Password Age** value is greater than zero. Valid values are 0 through 998. The default is 0. The **Warning Period** value must be less than the **Password Age** value.

4. Enter the number of unique passwords you must use before you can reuse a password in the **History Count** field.

   Valid values are 1 through 24. The default is 1. When you update the **History Count** value, the current password history is not cleared.

5. Configure the password format by completing the following steps.

   a. Select the **Empty Password - Allow** check box to allow user accounts to be created or edited with empty passwords or to allow passwords with any format.

      Empty Password is enabled by default.

   b. Enter the minimum password length in the **Minimum Length** field.

      Only enabled when the **Empty Password - Allow** check box is clear. Valid values are 4 through 127. The default is 8.

   c. Enter the minimum number of uppercase characters required in the **Upper Case Characters** field.

      Only enabled when the **Empty Password - Allow** check box is clear. Valid values are 0 through 127. The default is 0.

   d. Enter the minimum number of lowercase characters required in the **Lower Case Characters** field.

      Only enabled when the **Empty Password - Allow** check box is clear. Valid values are 0 through 127. The default is 0.

   e. Enter the minimum number of digits required in the **Number of Digits** field.

      Only enabled when the **Empty Password - Allow** check box is clear. Valid values are 0 through 127. The default is 0.

   f. Enter the minimum number of punctuation characters required in the **Punctuation Required** field.

      Only enabled when the **Empty Password - Allow** check box is clear. Valid values are 0 through 127. The default is 0.

   g. Enter the maximum number that the same character can repeat without a different intervening character in the **Maximum Repeat** field.

      Only enabled when the **Empty Password - Allow** check box is clear. Valid values are 0 through 127. The default is 2.

   h. Enter the maximum number of sequence characters from the ASCII collating series or keyboard sequences in the **Maximum Sequence** field.

      For example, 'ab' is a sequence of 2 and '456' is a sequence of 3.

      Only enabled when the **Empty Password - Allow** check box is clear. Valid values are 0 through 127. The default is 1.

6. Configure the password lockout support by completing the following steps.

a. Enter the number of failed login attempts allowed before the user account is locked out in the **Lockout Threshold** field.

   Valid values are 0 through 999. The default is 0 (disabled).

b. Enter the time frame after which the account automatically unlocks and resumes normal operation in the **Lockout Duration** field.

   Only enabled when the **Lockout Threshold** is greater than zero. If you specify zero, the user account is locked out indefinitely until an administrator manually unlocks it. Valid values are 0 through 99999. The default is 30.

7. Configure the password login policy by completing the following steps.

   a. Select **Concurrent Login** or **Single Login** from the **Login Mode** list.

      Single Login allows only one user to login at a time. If you selected **Single Login**, continue with step b.

      Concurrent Login allows multiple users to login at the same time. If you selected **Concurrent Login**, go to step 8.

   b. Select **Reject New Sessions** or **Logout Existing Sessions** from the **Action** list.

8. Click **View Policy Violators** to view the user accounts affected by any policy violations caused by your changes to the **Policy** tab before you save your work.

   If none of the user accounts violate the updated password policy, an empty **View Policy Violators** dialog box displays.

9. Click **Apply**.
10. Click **Yes** on the confirmation message.
11. Click **Close** to close the **Users** dialog box.

Viewing password policy violators

To view password policy violators, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click the **Policy** tab.
3. Click **View Policy Violators**.

   The **View Policy Violators** dialog box displays.

4. Review the password policy violator details.

   The **View Policy Violators** dialog box includes the following details:

   - **User ID —** Displays the identifier of the user who violated the password policy.
   - **Full Name —** Displays the full name of the user who violated the password policy.
   - **Reason —** Displays the reason the user violated the password policy.

5. Click **Close** on the **View Policy Violators** dialog box.
6. Click **Close** on the **Users** dialog box.

User profiles

User profiles contain the standard identification information of the user account, such as name, password, phone number, and e-mail address. The Management application enables you to make the following changes to your user profile:

- Change your name
- Change your password
- Change your user account description
- Change your phone number
- Change your e-mail address
- View your account state
- View your password policy
- Reset Management application messages
- Enable e-mail notification
- Configure e-mail notification

Viewing your user profile

To view your user profile, complete the following steps. To edit your user profile, refer to "Editing your user profile" on page 154.

1. Select **Server > User Profile**.

   The **User Profile** dialog box displays the following information:

   - **User ID —** Displays your user identifier.
   - **Full Name —** Displays the name if entered while adding a user; otherwise, this field is blank.
   - **Password —** Displays your password as dots (.). If the password policy is configured for an empty password, this field is blank. To change your password, refer to "Changing your password" on page 154.
   - **Confirm Password —** Displays your password as dots (.). If the password policy is configured for an empty password, this field is blank.
   - **Description —** Displays your description if entered while adding a user; otherwise, this field is blank.
   - **Phone Number —** Displays your phone number if entered while adding a user; otherwise, this field is blank.
   - **Account State —** Displays the current state of the account. Valid states include:
     - Active
     - Locked out by user manager
     - Locked out threshold reached
     - Password expired
     - Password format policy violated
     - Password history policy violated
   - **E-mail Notification Enable** check box — Select to enable e-mail notification.
   - **Filter —** Click to configure e-mail notification (refer to "Configuring e-mail notification" on page 156).
   - **E-mail Address —** Displays your e-mail, text message, or page addresses if entered while adding a user; otherwise, this field is blank.
   - **Password Age —** Displays the age of the password in days. Default is zero.
   - **Password Policy View** button — Click to display the current password policy (refer to "Viewing your password policy" on page 155).

- **Optional Messages Reset** button — Click to reset all optional messages to the default behavior. For more information, refer to "Resetting optional messages" on page 155.

2. Click **OK** on the **User Profile** dialog box.

Editing your user profile

To edit your user profile, complete the following steps.

1. Select **Server > User Profile**.

   The **User Profile** dialog box displays.

2. Change your name in the **Full Name** field.
3. Change your password in the **Password** and **Confirm Password** fields.

   Passwords display as dots (.).

4. Change your user profile description in the **Description** field.
5. Change your phone number in the **Phone Number** field.
6. Select the **E-mail Notification Enable** check box to enable e-mail notification.

   Clear the **E-mail Notification Enable** check box to disable e-mail notification.

7. Click **Filter** to set up basic event filters.

   For step-by-step instructions about setting up basic event filters, refer to "Setting up basic event filtering" on page 1133.

8. Change your e-mail, text message, or page address in the **E-mail Address** field.

   Enter more than one e-mail address, separating each with a semi-colon. To send a text message or page via e-mail, use the following format *number@carrier.com*, where *number* is your phone number and *carrier.com* is the SMS server. For example, 3035551212@txt.att.net (text message) or 3035551212@page.att.net (page).

   NOTE
   Check with your carrier for the exact e-mail address.

9. Click **OK** on the **User Profile** dialog box to save your changes.

Changing your password

To change your password from your user profile, complete the following steps.

1. Select **Server > User Profile**.

   The **User Profile** dialog box displays.

2. Change your password in the **Password** and **Confirm Password** fields.

   Passwords display as dots (.).

3. Click **OK** on the **User Profile** dialog box to save your changes.

If your password expires or your current password violates the password policy, you will be prompted to change your password from the **Change Password** dialog box. To view your password policy, click **Password Policy - View**.

To change your password from the **Change Password** dialog box, complete the following steps.

1. Enter your current password in the **Existing Password** field.
2. Enter your new password in the **New Password** and **Confirm Password** fields.

   Passwords display as dots (.).
3. Click **OK** to save your new password.

Viewing your password policy

To view your password policy, complete the following steps.

1. Select **Server > User Profile**.

   The **User Profile** dialog box displays.
2. Click **Password Policy - View** to display your password policy.

   The **View Password Policy** dialog box displays.

   - **Password History Count** — The number of unique passwords you must use before you can reuse a password.
   - **Empty Password** — Whether or not to allow empty passwords.
   - **Minimum Length** — The minimum length allowed for the password.
   - **Upper Case Characters** — The minimum number of uppercase characters required in the password.
   - **Lower Case Characters** — The minimum number of lowercase characters required in the password.
   - **Number of Digits** — The minimum number of digits required in the password.
   - **Punctuation Required** — The minimum number of punctuation characters required in the password.
   - **Maximum Repeat** — The maximum number that the same character can repeat without a different intervening character in the password.
   - **Maximum Sequence** — The maximum number of sequence characters from the ASCII collating series or keyboard sequences in the password.
3. Click **OK** on the **Password Policy** dialog box.
4. Click **OK** on the **User Profile** dialog box.

Resetting optional messages

To reset all Management application optional messages to their default behaviors, complete the following steps.

1. Select **Server > User Profile**.

   The **User Profile** dialog box displays.
2. Click **Optional Messages Reset**.

   The **Password Policy** dialog box displays.
3. Click **Yes** on the confirmation message.

   A successful reset message displays.
4. Click **OK** on the **User Profile** dialog box.

Configuring e-mail notification

To configure and enable e-mail notification, complete the following steps.

1. Select **Server > User Profile**.

   The **User Profile** dialog box displays.

2. Select the **E-mail Notification - Enable** check box to enable e-mail notification.

3. Click **Filter** to set up basic event filter.

   For step-by-step instructions about setting up basic event filters, refer to "Setting up basic event filtering" on page 1133.

4. Enter your e-mail, text message, or page address in the **E-mail Address** field.

   Enter more than one e-mail address, separating each with a semi-colon. To send a text message or page via e-mail, use the following format *number@carrier.com*, where *number* is your phone number and *carrier.com* is the SMS server. For example, 3035551212@txt.att.net (text message) or 3035551212@page.att.net (page).

   NOTE
   Check with your carrier for the exact e-mail address.

5. Click **OK** on the **User Profile** dialog box.

Fabric Insight Portal

Fabric Insight Portal overview

The Fabric Insight Portal provides high-level overview of a network, as well as quick access to dashboard monitors and reports. The Fabric Insight Portal's main window contains a number of areas. The following are the various areas, and descriptions listed below.

- Fabric Insight Portal login page
- Dashboard
- Collections
- Inventory
- Events
- Preferences

Management application license

License keys are encoded form of supported configuration or features. License keys verify ownership of the Management application software as well as determine the maximum port count allowed or any additional features that you receive as part of the license.

The Fabric Insight Portal is available for Professional, Professional Plus and Enterprise editions; however, the Professional edition does not support the Performance widgets in the dashboard that do not contain data. The Fabric Insight Portal support is not available for SMI Agent only installation.

If you exceed the maximum port count for your version, software functionality is impacted and you must reduce the port count through discovery or contact your vendor to purchase an additional license for your version.

If your License expires or changes, an error message displays.

Prerequisites

The following sections explain the prerequisites for the Fabric Insight Portal.

Browser requirements

The Fabric Insight Portal Management application is supported in the following browsers:

- Internet Explorer 11 and later (Windows only, except Windows 8 and Windows 2012)
- Edge 13 (Windows 10 only)
- Firefox 41 and later (Linux only)
- Chrome 46 and later (Windows, MAC OS)

Opening the Fabric Insight Portal

You must log into a Management application server to monitor the network. To launch the Fabric Insight Portal, complete the following steps.

1. Open a web browser and enter the IP address or host name of the Management application server in the **Address** bar.

   If the web server port number does not use the default (443 if is SSL Enabled; otherwise, the default is 80), you must enter the web server port number in addition to the IP address. For example, IP_Address:Port_Number.

   The Fabric Insight Portal login page displays. You can launch the Desktop client from login page of the Fabric Insight Portal by clicking **Desktop Client**. You can download the client bundle (64-bit OS only), JRE, or MIB files by clicking JRE and MIB files.

   FIGURE 57 Fabric Insight Portal log in page

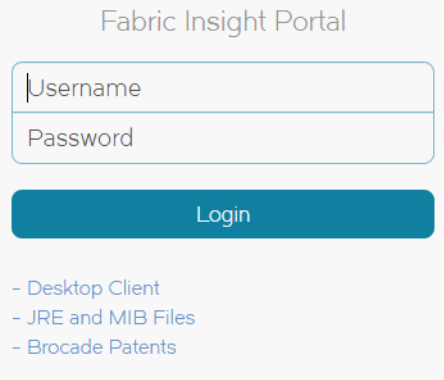

2. Enter your user name and password.

   Do not enter Domain\User_Name in the User ID field for LDAP server authentication.

3. Press **Enter** or click **Login**.

   If the Administrator disconnects the Fabric Insight Portal using the Active Sessions dialog box (Server > Active Sessions), the Fabric Insight Portal redirects to the login page after three minutes or as soon as you make a selection.

Launching the Desktop client from Fabric Insight Portal login page

You can launch the Desktop client from the Fabric Insight Portal login page.

NOTE
It is recommended to use JRE version 1.8.0_121 to launch the remote client from Fabric Insight Portal login page.

To launch the Desktop client, complete the following steps.

1. Open a web browser and enter the IP address or host name of the Management application server in the **Address** bar.

   If the web server port number does not use the default (443 if is SSL Enabled; otherwise, the default is 80), you must enter the web server port number in addition to the IP address. For example, IP_Address:Port_Number.

   The Fabric Insight Portal login page displays.

2. Click **Desktop Client**.

   The Fabric Insight Portal page displays.

3. Open the download.

   The Log In dialog box displays. Log into another server by entering the IP address to the other server in the Network Address field.

   The server must be the exact same version, edition, starting port number, and network size as the client.

   Remove a server from the Network Address list by selected the IP address and clicking Delete.

4. Choose one of the following options:

   - If you configured authentication to CAC, enter your PIN in the CAC PIN field.
   - If you configured authentication to the local database, an external server (RADIUS, LDAP, or TACACS+), or a switch, complete the following steps.

     a. Enter your user name and password.

     The defaults are Administrator and password, respectively.

     Do not enter Domain\User_Name in the User ID field for LDAP server authentication.

     b. Select or clear the **Save password** check box to choose whether you want the application to remember your password the next time you log in.

   To change your password, refer to your user manual or online help.

5. Click **Login**.
6. Click **OK** on the **Login Banner** dialog box.

   The Management application displays

   When you launch the Management application or navigate to a new view, the SAN tab displays with a gray screen over the Product List and Topology Map while data is loading.

Global Filter

Global filter is available on any page in the Fabric Insight Portal. You can search for the following using search criteria such as the name, WWN, IP address, Zone alias, model, port speed, and firmware version. The search criteria is based on a free text keyword for an exact match.

- A network resource (fabric, device, port, interface, host/target or blade)
- Widgets
- Events
- Filters
- Collections
- Flows

You can perform global filter to view the following details:

- Flow list
- Inventory list
- Filter list
- Collection list
- Collection management list

To perform a global filter, complete the following steps:

1. Click the desired icon (for example Inventory icon) for which you wish to view data.

   The Inventory page displays with all available SAN products.

2. Enter your search criteria in the **Global Filter** search box.

   FIGURE 58 Global filter

3. Press **Enter**.

   The search result displays.

   FIGURE 59 Global filter search

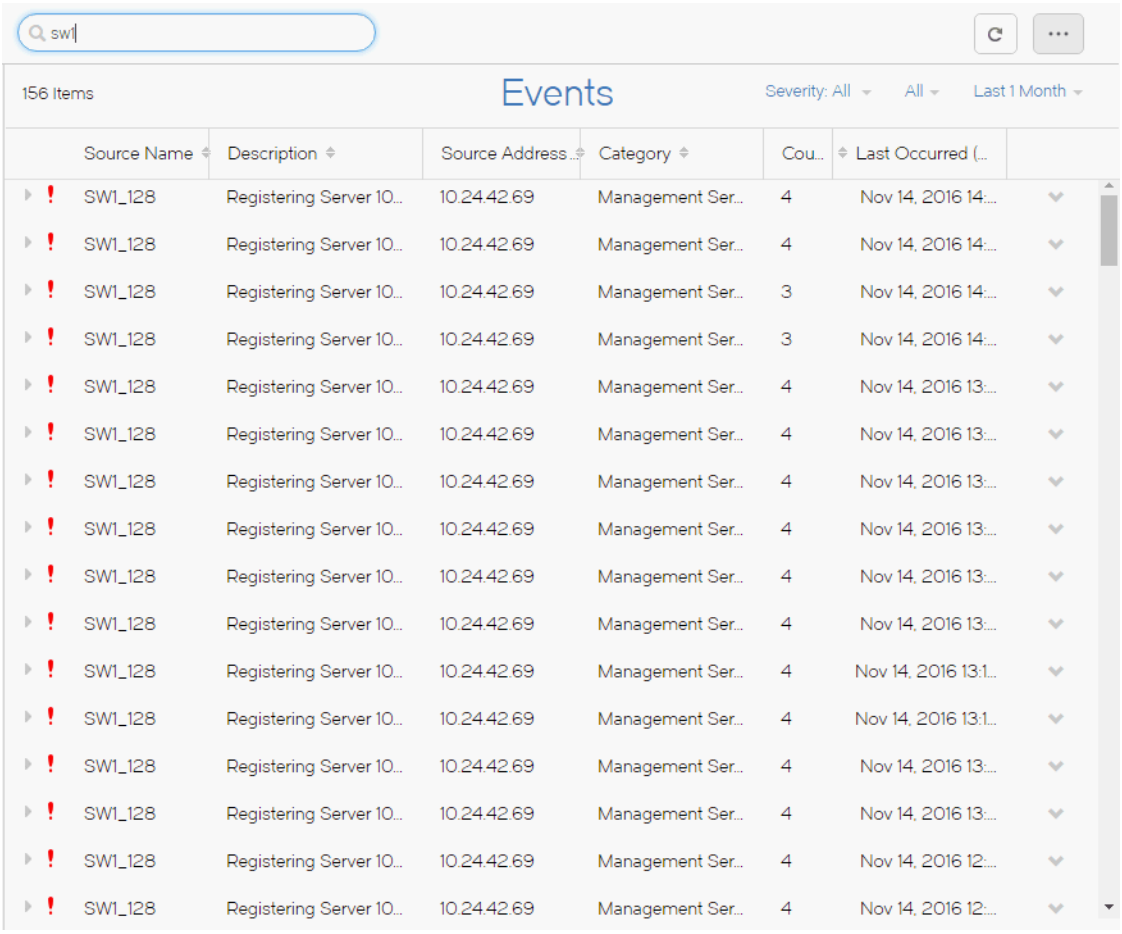

Logging off the Fabric Insight Portal

Click Logout, to log off the Fabric Insight Portal.

The Fabric Insight Portal login page displays.

Dashboard

NOTE
You must have the Dashboard Management privilege with read permission to view dashboards.

NOTE
You can only create and add customized widgets from the Desktop client.

NOTE
Only devices in your area of responsibility (AOR) display in the dashboard.

The Dashboard displays the performance monitors. You can also display additional performance monitors, as needed. The Fabric Insight Portal has the following default dashboards: SAN Port Health.

The dashboard provides a high-level overview of the network and the current states of managed devices. This allows you to easily check the status of the devices on the network. The dashboard also provides several features to help you quickly access reports, device configurations, and system event logs.

The dashboard refreshes the flow widgets every five minutes and refreshes the status and performance widgets every 10 minutes regardless of the size of your network. Note that data may become momentarily out of sync between the dashboard and other areas of the application. For example, if you remove a product from the network while another user navigates from the dashboard to a more detailed view of the product, the product may not appear in the detailed view.

Dashboard toolbar

The dashboard toolbar (as shown in the following figure) is located above the performance monitors and provides a information about the selected dashboard as well as buttons to perform various functions.

FIGURE 60 Dashboard toolbar

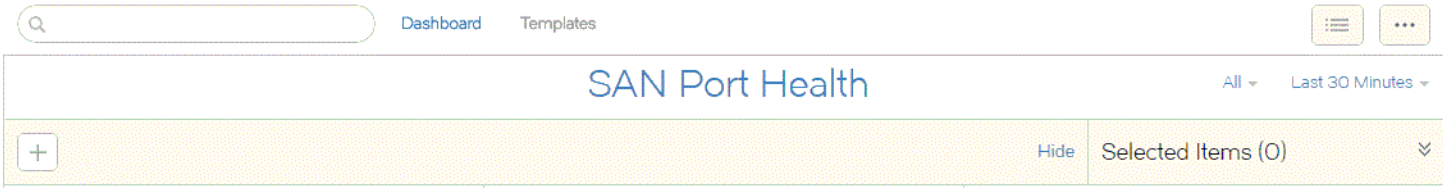

The dashboard toolbar contains the following fields and components:

1. Global filter—Use to search for a network resource (fabric, switch, port, interface, host/target or blade), widgets, events, filters, collection, or flows on the Dashboard page. For more information, refer to "Global Filter" on page 159.
2. Dashboard tab—Click to display the dashboard.
3. Templates tab—Click to view the templates.
4. Select Template icon—Click to display the Select Templates dialog box.
5. Create Template icon (not shown)—Click to create a new template. For more information, refer to "Creating a user-defined dashboard" on page 165.
6. More icon—Click to edit the current dashboard (template), export the dashboard, access the user guide, and view the application information.

7. Dashboard title—The title of the selected dashboard.
8. Show Filters link (not shown)—Click to show filters and Selected Items.
9. Network scope list—Use to select the network scope for which you want to display data in the dashboard. For more details refer to "Setting the network scope" on page 163.
10. Time scope—Click to select network scope and time scope for which you want to display data in the dashboard. For more details refer to "Setting the time interval" on page 164.
11. Add filter icon—Click to add a filter to or create a filter for the dashboard. For more details refer to "Adding a filter" on page 811.
12. Hide link—Click to hide filters and Selected Items.
13. Selected Items count and arrow icon—Click to investigate performance for an object. For more information, refer to "Investigating flow performance" on page 822 and "Investigating port performance" on page 172.

Accessing a dashboard

To access a specific dashboard, complete the following steps.

1. Click the **Select Template** icon.

   The **Select Template** dialog box displays. Sort the table by clicking the **Name** column head to sort the list. Click the **Name** column head again to reverse the sort orders.

2. Select the dashboard you want to view from the list.

   Search for a dashboard by entering search criteria in the search field and pressing enter.

3. Click **OK**.

   The dashboard you selected displays.

Dashboard customization

You can customize the dashboard display by setting the network scope and time scope.

Setting the network scope

You can configure the Dashboards page to display all objects in your area of responsibility (AOR) or a subset of objects (fabrics, devices, or groups) using the network scope selection. Default network scopes are visible to all users. User-defined scopes are visible only to the user who created it. Whenever scope is changed, all widgets corresponding to the network scope in the dashboard automatically refresh.

1. Click the **Scope** arrow.

   FIGURE 61 Select Network Scope dialog box

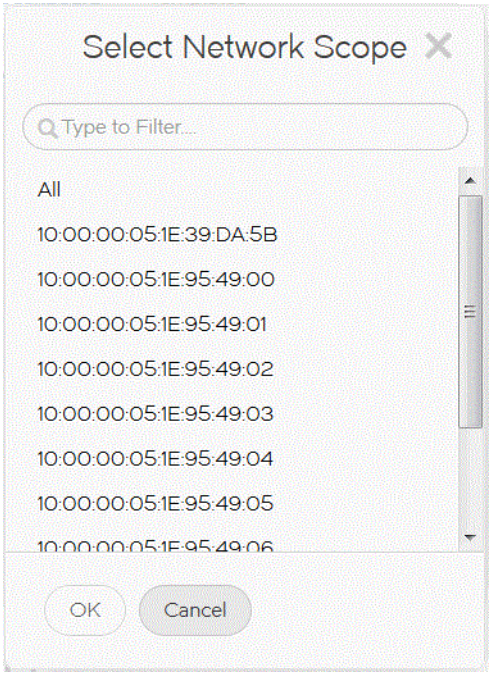

2. Select a network from the **Network Scope** list.

   The default network scope is All. It includes all managed and monitored fabrics or groups in your AOR. If the selected fabric or group is deleted from discovery, the widget refreshes and returns to the default network scope (All).

3. Click **OK**.

   Search for a product or fabric by entering search criteria in the Search field. The Network Scope list automatically filters out any products that do not match the search criteria.

   If you select a fabric scope, dashboard widgets displays data for all products and ports in the fabric.

   If you select a product scope, dashboard widgets displays data for the selected products and the ports that belong to the selected products.

   If you select a port scope, dashboard widgets displays data for the specified ports and the products to which the ports belong. If any of the selected ports are initiator or target ports, dashboard widgets displays data for the attached switch port.

Setting the time interval

Setting the global time interval in the Dashboards page configures the data display time range for all the applicable widgets. Time interval in the Scope list allows you to select a specific time range for which you want to display data in the Dashboards page.

1. Click the **Scope** arrow.

FIGURE 62 Select Date Range dialog box

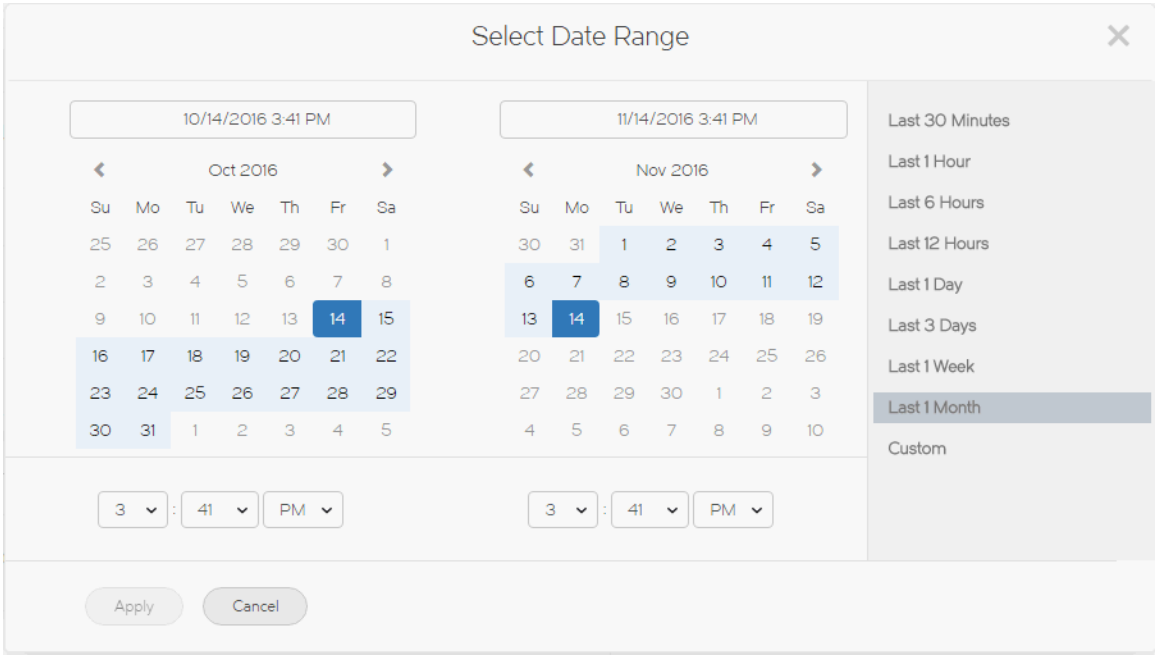

2. To display data for the current date and time based on the selected time scope, select one of the following options:
   - 30 Minutes—Displays data for 30 minutes.
   - 1 Hour—Displays data for 1 hour.
   - 6 Hours—Displays data for 6 hours.
   - 12 Hours—Displays data for 12 hours.
   - 1 Day—Displays data for 24 hours.
   - 3 Days—Displays data for 3 days.
   - 1 Week—Displays data for 1 week.
   - 1 Month—Displays data for 30 days.

   Go to step 9.
3. To display data for data for a custom date and time based on the selected time scope, select **Custom** from the Time Scope list. Continue with step 4.
4. Select a start date from the left Calendar pop-up window.
5. Select a start time from the left Choose Time pop-up window
6. Select a end date from the right Calendar pop-up window.

7. Select a end time from the right Choose Time pop-up window
8. Select a duration from the Time Scope list.

   The displayed data changes to the new time range for all the applicable widgets.
9. Click **Apply**.

User-defined dashboard templates

You can create and edit user-defined dashboard templates and customize them with the performance monitors you need to monitor your network.

Creating a user-defined dashboard

You can create a dashboard template and customize it with the performance monitors you need to monitor your network.

1. Click the Templates tab.

   The list of preconfigured and user-defined templates display.
2. Choose **Select Widgets** from the + (Create Template) button.

   FIGURE 63 Create Template options

   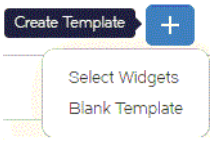

   The Select Widgets dialog box displays.
3. Select **Status** from the list, if necessary and select the check box for each status widget you want to add in this dashboard template.
4. Select **Performance** from the list and select the check box (up to 30) for each performance widget you want to add in this dashboard.

   Filter the list of widgets by entering filter criteria in the **Filter** field and pressing Enter.
5. Click **OK**.

   The new dashboard displays with the title “Create New Template”.
6. Select **Save** from the **Save** list.

   The **Template Details** dialog box displays.
7. Enter a name (up to 50 characters), tags, and description (up to 256 characters) for the template.
8. Select the Share check box to share the template with other users.
9. Click **Save**.
10. Display the template in the dashboard by selecting **View in Dashboard** from the ··· (More) icon.

Creating a unique layout user-defined dashboard

You can create a dashboard template and customize it with the performance monitors you need to monitor your network.

1. Click the **Templates** tab.

   The list of preconfigured and user-defined templates display.

2. Choose **Blank Template** from the + (Create Template) button.

   The **Create New Template** page displays.

   FIGURE 64 Create New Template page

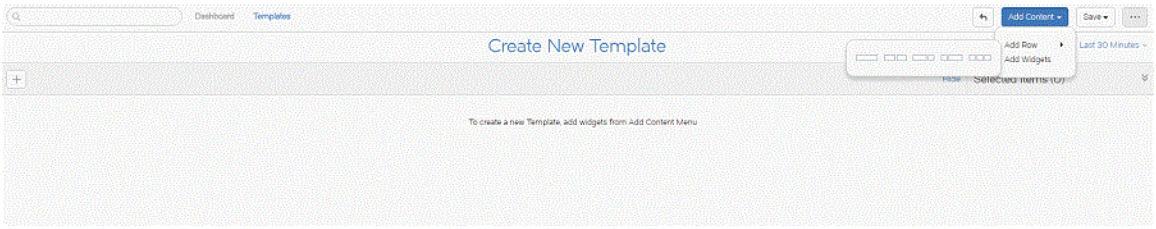

3. Select **Add Row >** row_format from the **Add Content** list, where row_format includes the following options:

   - one large widget
   - two equal size widgets
   - one large and one small widget
   - one small and one large widget
   - three equal size widgets

   The new row displays in the **Create New Template** page.

4. Click the Select Widget button (+) in the blank widget to define the widget.

   The **Select Widget** dialog box displays.

5. Choose one of the following options:

   - Add a status widget by selecting Status from the list, if necessary, and selecting the status widget you want to add
   - Add a performance widget by selecting Performance from the list and selecting the performance widget you want to add.

     Filter the list of widgets by entering filter criteria in the **Filter** field and pressing **Enter**.

6. Click **OK**.

   The selected widget displays. To change a widget, click the Select Widget icon in the widget and repeat step 5 and step 6.

7. Repeat step 3 through step 6 for each widget you want to include in the template.

8. Select **Save** from the **Save** list.

   The **Template Details** dialog box displays.

9. Enter a name (up to 50 characters), tags, and description (up to 256 characters) for the template.

10. Select the **Share** check box to share the template with other users.

11. Click **Save**.
12. Display the template in the dashboard by selecting **View in Dashboard** from the ··· (More) icon.

Editing a user-defined dashboard

You can edit a user-defined dashboard name, description, and customize it with the performance monitors you need to monitor your network.

1. Click the **Templates** tab.

   The list of preconfigured and user-defined templates display.

2. Select the user-defined dashboard you want to edit in list of templates.
3. Edit the name, tags, description, and share status by selecting **Edit** Info from the ··· (More) icon.

   The **Template Details** dialog box displays.

4. Edit the name, tags, and description for the dashboard.
5. Select the Share check box to share the template with other users.
6. Add new widgets to the template by selecting **Add Row >** row_format from the **Add Content** list, where row_format includes the following options:
   - one large widget
   - two equal size widgets
   - one large and one small widget
   - one small and one large widget
   - three equal size widgets

   The new row displays in the dashboard.

7. Define the blank widget by clicking the **Select Widget** button (+) in the blank widget.

   The **Select Widget** dialog box displays.

8. Select the new widget by choosing one of the following options:
   - Add a status widget by selecting **Status** from the list, if necessary, if necessary and selecting the status widget you want to add
   - Add a performance widget by selecting Performance from the list and selecting the performance widget you want to add.

   Filter the list of widgets by entering filter criteria in the **Filter** field and pressing **Enter**.

9. Click **Save**.

   The selected widget displays.

10. Change an existing widget by clicking the **Select Widge**t icon in the widget you want to change.
11. Select the new widget by repeating step 8 and step 9.
12. Select **Save As** from the **Save** list.
13. Display the template in the dashboard by selecting **View in Dashboard** from the ··· (More) icon.

    The template displays in the dashboard with live data.

Deleting a user-defined dashboard

You can only delete a user-defined dashboard template.

1. Click the **Templates** tab.

   The list of preconfigured and user-defined templates display.

2. Select the dashboard you want to delete
3. Select **Delete** from the ··· (More) icon.
4. Click **OK** on the confirmation message.

   The user-defined dashboard template is deleted **Templates** tab.

Exporting the dashboard or widget

You can export the current dashboard display (all widgets and monitors) or a selected widget or monitor in PDF, PNG, or JPEG format.

To export the dashboard, select **Export >** ***output_type*** from the ··· (More) icon.

To export a widget, select **Export >** ***output_type*** from the ☰ (Other actions) icon.

Where output types is PDF, PNG, or JPEG:

The file is saved to the default download location (such as, C:\Users***user_name***\Downloads).

Default dashboard templates

Fabric Insight Portal provides preconfigured dashboard templates which provide high-level overview of the network, the current states of managed devices, and performance of devices, ports, and traffic on the network.

SAN Ports Health

The SAN Ports Health dashboard provides the following preconfigured status and performance widgets for the ISL, Host, and Target ports:

- Top Port Utilization widget
- Top Port Utilization Percentage widget
- Bottom Port Utilization Percentage widget
- Bottom Port Utilization widget
- Top Port CRC Errors widget
- Top Port Sync Losses widget
- Top Port Link Failures widget
- Top Port C3 Discards RX TO widget
- Top Port Link Resets widget
- Top Port PCS Block Errors widget

Performance

Fabric Insight Portal provides the following additional preconfigured performance widgets. These performance widgets are available when you create a user-defined dashboard (refer to "Dashboard customization" on page 163):

- Top Port C3 Discards widget
- Top Port Traffic widget
- Top Port Invalid Transmissions widget

SAN Port Health widgets

Fabric Insight Portal provides nine preconfigured SAN Port Health widgets.

Top Port Utilization widget

The Top Port Utilization widget displays the top port utilization in a bar chart. There are three port widgets: ISL, Initiator, and Target.

FIGURE 65 Top Port Utilization widget

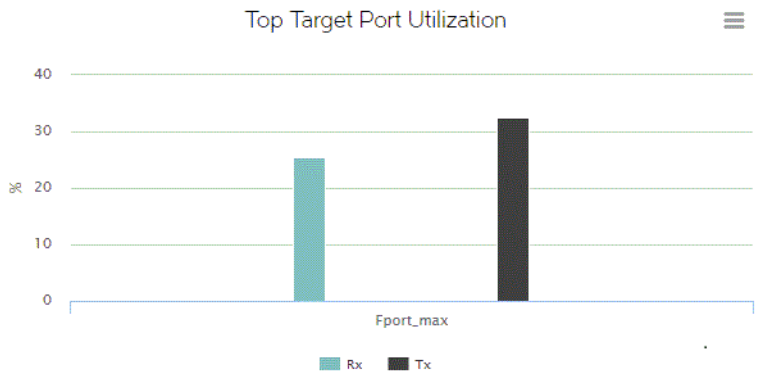

The Top Port Utilization widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- RX—The top receive port utilization for each affected port.
- TX—The top transmit port utilization for each affected port.

  Pause on a bar to view utilization.

  You can perform the following functions for both RX and TX bars.

  - View additional details by clicking the bar and selecting Show Details (refer to "Viewing additional details for the Top Port Utilization widget" on page 170).
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Viewing additional details for the Top Port Utilization widget

To view additional details for the widget, complete the following steps.

1. Click a bar in the chart and select Show Details.

   FIGURE 66 Top Port Utilization detailed view

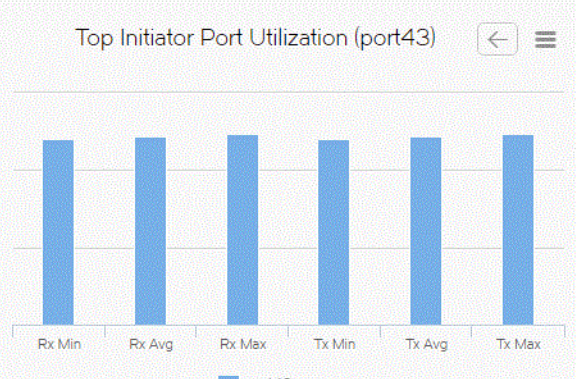

   A more detailed widget displays which includes the following data:

   - Widget Name (Port Name)—The name of the widget and the port affected by this widget.
   - Back icon ( )—Click to go back to the Top Port Utilization widget.
   - Other actions icon—Select to export a graphic of the widget.
   - RX Min—The minimum top receive port utilization.
   - RX Avg—The average top receive port utilization.
   - RX Max—The maximum top receive port utilization.
   - TX Min—The minimum top transmit port utilization.
   - TX Avg—The average top transmit port utilization.
   - TX Max—The maximum top transmit port utilization.

   Pause on a bar to view utilization.

2. Click the back ( ) icon to go back to the Top Port Utilization widget.

Viewing port properties

1. Click a bar in the chart and select Properties.

FIGURE 67 Properties dialog box

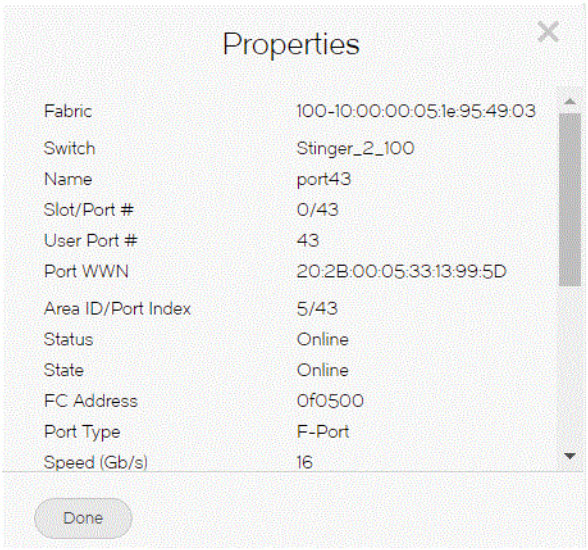

A Properties dialog box displays which includes the following data:

- Fabric—The IP address of the fabric.
- Switch—The name of the switch.
- Name—The port name.
- Slot/Port #—The slot and port number.
- User Port #—The number of the user port.
- Port WWN—The port's world wide name.
- Area ID /Port Index—The area identifier and port index number.
- Status—The status of the port.
- State—The state (such as Online) of the port.
- FC Address—The FC address of the port.
- Port Type—The type of port, for example, F-port.
- Speed (Gb/s)—The port speed, in Gigabits per second.
- Protocol—The network protocol, for example, Fibre Channel.
- Long Distance Settings—Whether the connection is considered to be normal or longer distance.
- Forward Error Correction—Whether FEC is enabled or disabled.
- Encryption—Whether encryption is enabled or disabled.
- Compression—Whether compression is enabled or disabled.
- NPIV Enabled—Whether the port is NPIV enabled or not.
- Connected to—The name of the connected switch.

- Attached Port #—The port number of the attached product.
- Product Type—The product type.
- Calculated Status—The calculated operational status.
- Zone Alias—The zone alias of the port.

2. Click **Done** to go back to the widget.

Investigating port performance

1. Chose one of the following options:
   - Click a data point on the flow chart and select Investigate.
   - Click the Selected Items down arrow icons, select one or more flows, and click Investigate.

     The Investigate Performance page displays. The top right displays the port performance in a flow/bar chart.

     Pause on a data point to view data about the selected measure.

     The x-axis displays the date and time.

     The bottom right displays the port performance in a table format.
   - Name—The port name.
   - Type—The port type.
   - WWN—The port world wide name.
   - vTap—The vTap.
   - Status—The status of the port (such as No_Module).
   - State—The state (such as Online) of the port.
2. Click the close (X) button to return to the dashboard.

Top Port Utilization Percentage widget

The Top Port Utilization widget displays the bottom port utilization percentages in a bar chart.

FIGURE 68 Top Port Utilization Percentage widget

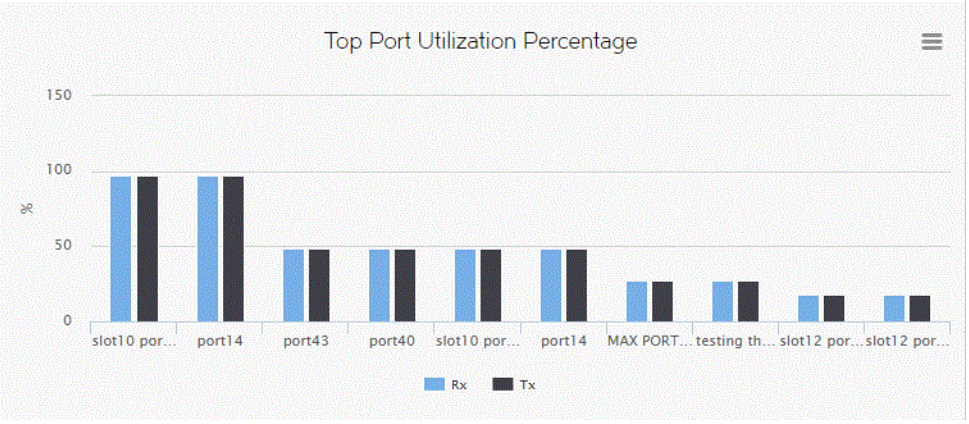

The Top Port Utilization Percentage widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- RX—The bottom receive port utilization percentages for each port.
- TX—The bottom transmit port utilization percentages for each port.

  Pause on a bar to view utilization.

  You can perform the following functions for both RX and TX bars.

  - View additional details by clicking the bar and selecting Show Details (refer to "Viewing additional details for the Top Port Utilization Percentage widget" on page 173).
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Viewing additional details for the Top Port Utilization Percentage widget

To view additional details for the widget, complete the following steps.

1. Click a bar in the chart and select Show Details.

   FIGURE 69 Top Port Utilization Percentage details widget

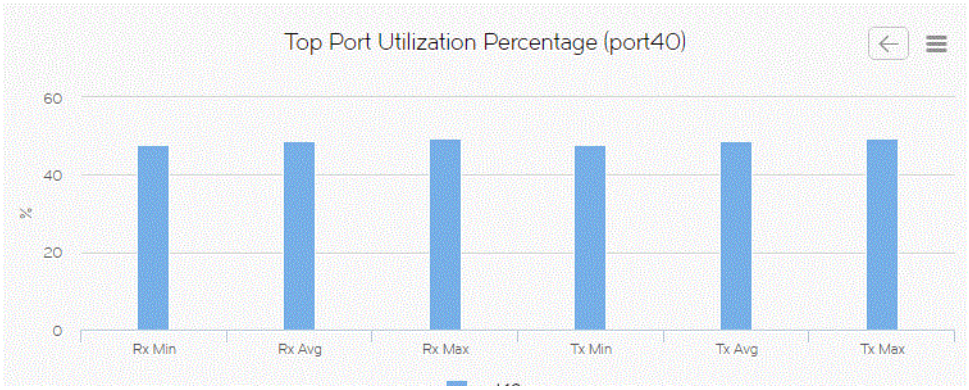

   A more detailed widget displays which includes the following data:

   - Widget Name (Port Name)—The name of the widget and the port affected by this widget.
   - Back icon—Click to go back to the Top Port Utilization Percentage widget.
   - Other actions icon—Select to export a graphic of the widget.
   - RX Min—The minimum top receive port utilization percentages.
   - RX Avg—The average top receive port utilization percentages.
   - RX Max—The maximum top receive port utilization percentages.
   - TX Min—The minimum top transmit port utilization percentages.
   - TX Avg—The average top transmit port utilization percentages.
   - TX Max—The maximum top transmit port utilization percentages.

     Pause on a bar to view utilization.

2. Click the back ( ) button to go back to the Top Port Utilization Percentage widget.

Bottom Port Utilization Percentage widget

The Bottom Port Utilization Percentage widget displays the bottom port utilization percentages in a bar chart. There are four port widgets: All, ISL, Initiator, and Target.

FIGURE 70 Bottom Port Utilization Percentage widget

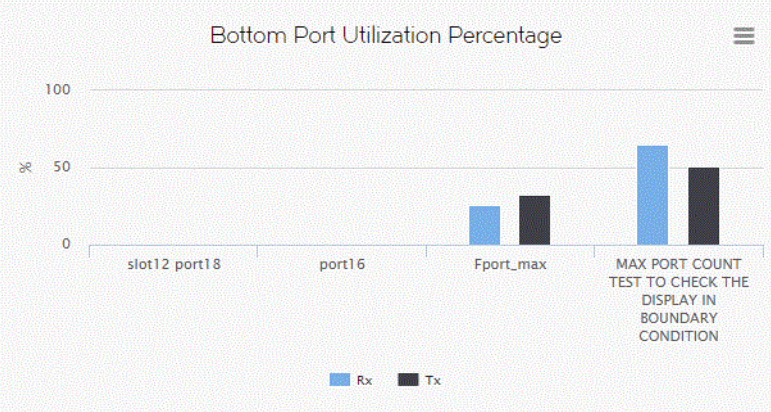

The Bottom Port Utilization Percentage widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- RX—The bottom receive port utilization percentages for the port.
- TX—The bottom transmit port utilization percentages port.

  Pause on a bar to view utilization.

  You can perform the following functions for both RX and TX bars.

  - View additional details by clicking the bar and selecting Show Details (refer to "Viewing additional details for the Bottom Port Utilization Percentage widget" on page 175).
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Viewing additional details for the Bottom Port Utilization Percentage widget

To view additional details for the widget, complete the following steps.

1. Click a bar in the chart and select Show Details.

FIGURE 71 Bottom Port Utilization Percentage details widget

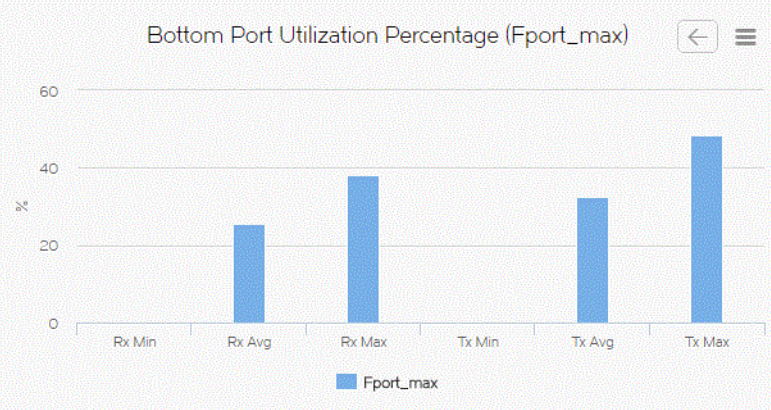

A more detailed widget displays which includes the following data:

- Widget Name (Port Name)—The name of the widget and the port affected by this widget.
- Back icon—Click to go back to the Bottom Port Utilization Percentage widget.
- Other actions icon—Select to export a graphic of the widget.
- RX Min—The minimum bottom receive port utilization percentages.
- RX Avg—The average bottom receive port utilization percentages.
- RX Max—The maximum bottom receive port utilization percentages.
- TX Min—The minimum bottom transmit port utilization percentages.
- TX Avg—The average bottom transmit port utilization percentages.
- TX Max—The maximum bottom transmit port utilization percentages.

  Pause on a bar to view utilization.

2. Click the back ( ) button to go back to the Bottom Port Utilization Percentage widget.

Bottom Port Utilization widget

The Bottom Port Utilization widget displays the top port utilization in a bar chart. There are three port widgets: ISL, Initiator, and Target.

FIGURE 72 Bottom Port Utilization widget

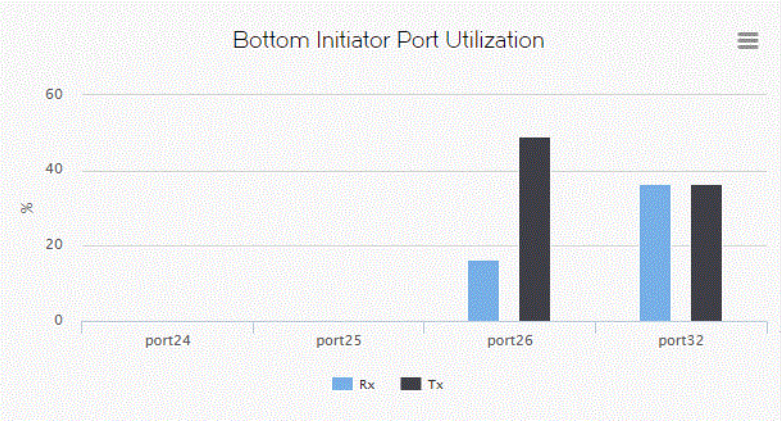

The Bottom Port Utilization widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- RX—The bottom receive port utilization for each affected port.
- TX—The bottom transmit port utilization for each affected port.

  Pause on a bar to view utilization.

  You can perform the following functions for both RX and TX bars.

  - View additional details by clicking the bar and selecting Show Details (refer to "Viewing additional details for the Bottom Port Utilization widget" on page 177).
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Viewing additional details for the Bottom Port Utilization widget

To view additional details for the widget, complete the following steps.

1. Click a bar in the chart and select Show Details.

   FIGURE 73 Bottom Port Utilization widget

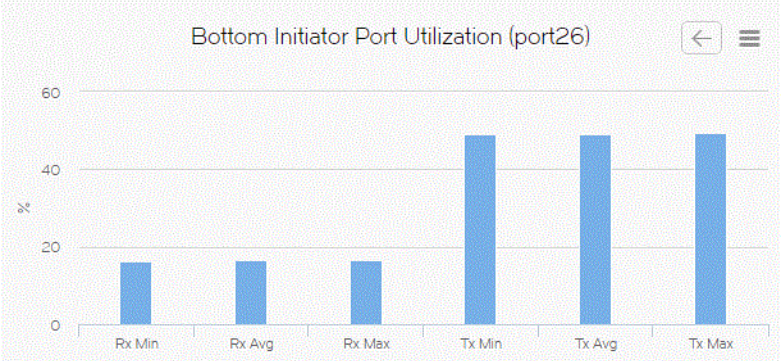

   A more detailed widget displays which includes the following data:

   - Widget Name (Port Name)—The name of the widget and the port affected by this widget.
   - Back icon—Click to go back to the Top Port Utilization widget.
   - Other actions icon—Select to export a graphic of the widget.
   - RX Min—The minimum bottom receive port utilization.
   - RX Avg—The average bottom receive port utilization.
   - RX Max—The maximum bottom receive port utilization.
   - TX Min—The minimum bottom transmit port utilization.
   - TX Avg—The average bottom transmit port utilization.
   - TX Max—The maximum bottom transmit port utilization.

     Pause on a bar to view utilization.

2. Click the back ( ) icon to go back to the Bottom Port Utilization widget.

Top Port CRC Errors widget

The Top Port CRC Errors widget displays the top ports with frames that contain cyclic redundancy check (CRC) errors in a table. There are three top CRC Errors port widgets: ISL, Initiator, and Target.

The Top Port CRC Errors widget includes the following data:

- Widget title—The name of the widget.
- Port—The port affected by this widget. Click to launch the Port page (refer to "Port Summary View" on page 194). When you launch the Port page, the detailed view closes.
- CRC Errors/sec—The number (error rate) of cyclic redundancy check (CRC) errors per second for the duration specified in the widget.
- CRC Errors—The number (error count) of cyclic redundancy check (CRC) errors for the duration specified in the widget.

Viewing additional details for the Top Port CRC Errors widget

1. Click the Show Details icon.

   A more detailed widget displays which includes the following data:

   - Scope—The scope configured for the dashboard.
   - Port—The port affected by this widget. Click to launch the Port page (refer to "Port Summary View" on page 194). When you launch the Port page, the detailed view closes.
   - Connected_Port (where Connected_Port is Connected Port, Initiator, or Target)—Displays the address of the port:
   - CRC Errors/sec—The number (error rate) of cyclic redundancy check (CRC) errors per second for the duration specified in the widget.
   - CRC Errors—The number (error count) of cyclic redundancy check (CRC) errors for the duration specified in the widget.

2. Click the close (X) button.

Top Port Sync Losses widget

The Top Port Sync Losses widget displays the top ports with synchronization failures in a bar chart. There are three top Sync Losses port widgets: ISL, Initiator, and Target.

FIGURE 74 Top Port Sync Losses widget

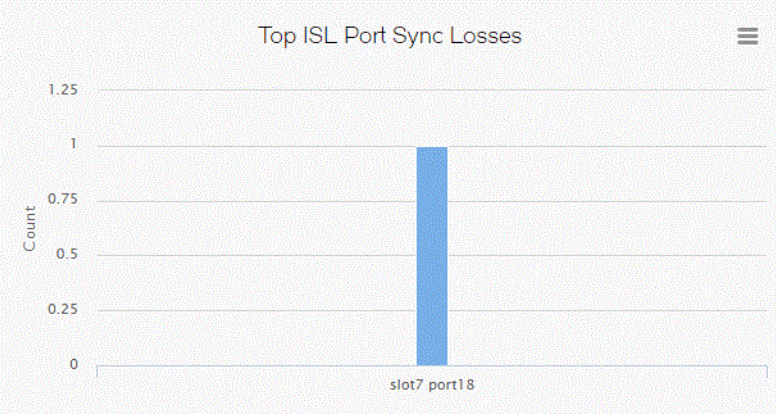

The Top Port Sync Losses widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- Count—The x-axis displays the count.
- Slot / Port—The y-axis displays the slot number and port number. Pause on the bar to display the slot and port number and the number of synchronization failures for the port. You can also perform the following functions:
  - View additional details by clicking the bar and selecting Show Details (refer to "Viewing additional details for the Top Port Sync Losses widget" on page 179).
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Viewing additional details for the Top Port Sync Losses widget

To view additional details for the widget, complete the following steps.

1. Click a bar in the chart and select Show Details.

   A more detailed widget displays which includes the following data:

   - Widget Name (Port Name)—The name of the widget and the port affected by this widget.
   - Back icon—Click to go back to the Top Port Sync Losses widget.
   - Other actions icon—Select to export a graphic of the widget.
   - CRC Errors—The number of CRC errors for the port.
   - Sync Losses—The number of synchronization failures for the port.
   - Invalid Transmissions—The number of invalid transmissions for the port.

2. Click the back (←) button to go back to the Top Port Utilization Percentage widget.

Top Port Link Failures widget

The Top Port Link Failures widget displays the top ports with link failures in a bar chart. There are three top Link Failures port widgets: ISL, Initiator, and Target.

FIGURE 75 Top Port Link Failures widget

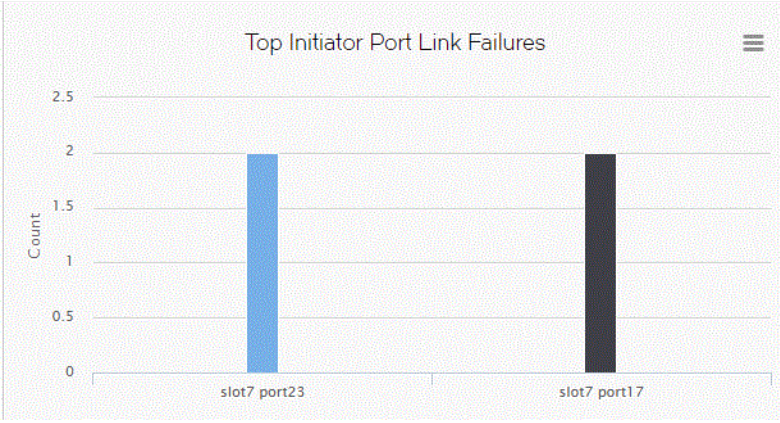

The Top Port Link Failures widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- RX—The number (error count) of receive link failure errors. The slot and port number display on the y-axis. Pause on the bar to display the slot and port number as well as the error count.
- TX—The number (error count) of transmit link failure errors. The slot and port number display on the y-axis. Pause on the bar to display the slot and port number as well as the error count.

  You can perform the following functions for both RX and TX bars.

  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Top Port C3 Discards RX TO widget

The Top Port C3 Discards RX TO widget displays the top ports with receive Class 3 frames received at this port and discarded at the transmission port due to timeout in a bar chart. There are three C3 Discards RX TO port widgets: ISL, Initiator, and Target.

FIGURE 76 Top Port C3 Discards RX TO widget

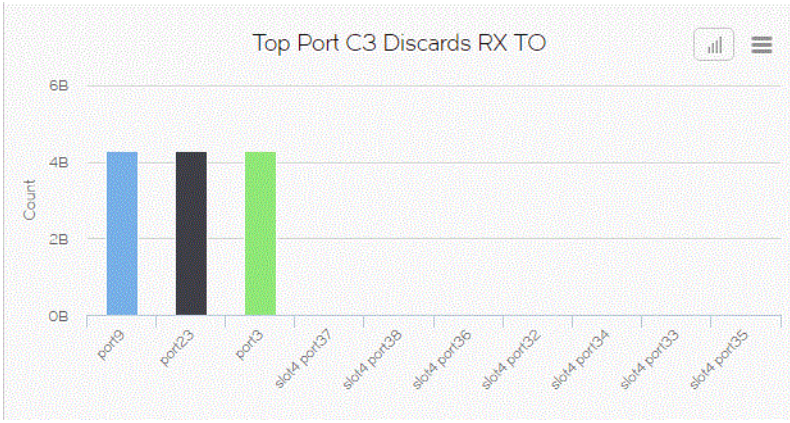

The Top Port C3 Discards RX TO widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- Count—The x-axis displays the error count.
- Slot / Port—The y-axis display the slot number and port number. Pause on the bar to display the slot and port number and the number (error count) of Class 3 frames received at this port and discarded at the transmission port due to timeout for the duration specified in the widget. You can also perform the following functions:
  - View additional details by clicking the bar and selecting Show Details (refer to "Viewing additional details for the Top Port C3 Discards RX TO widget" on page 180).
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Viewing additional details for the Top Port C3 Discards RX TO widget

To view additional details for the widget, complete the following steps.

1. Click a bar in the chart and select the Show Details icon.

   A more detailed widget displays which includes the following data:

   - Widget Name (Port Name)—The name of the widget and the port affected by this widget.
   - Back icon—Click to go back to the Top Port C3 Discards RX TO widget.
   - Other actions icon—Select to export a graphic of the widget.
   - Count—The x-axis displays the error count.
   - C3 Discards RX TO—The number (error count) of Class 3 frames received at this port and discarded at the transmission port due to timeout errors for the duration specified in the widget.
   - C3 Discards TX TO—The number (error count) of Class 3 frames transmitted from this port and discarded at the transmission port due to timeout errors for the duration specified in the widget.

2. Click the back ( ) button to go back to the Top Port C3 Discards RX TO widget.

Top Port Link Resets widget

The Top Port Link Resets widget displays the top ports with link resets in a bar chart. There are three Link Resets port widgets: ISL, Initiator, and Target.

FIGURE 77 Top Port Link Resets widget

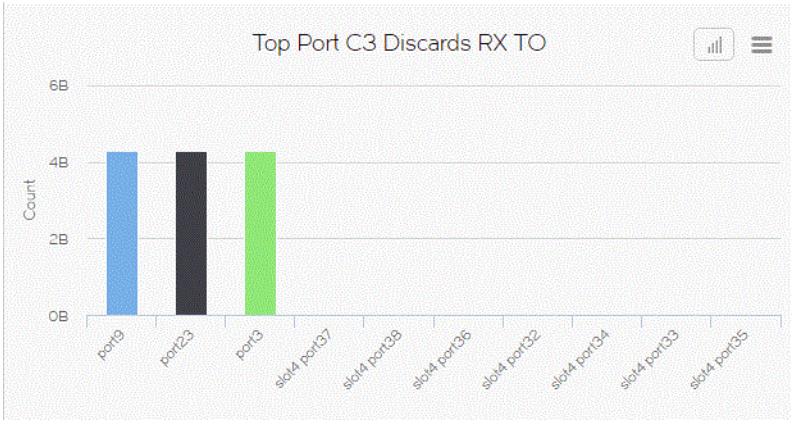

The Top Port Link Resets widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- Count—The x-axis displays the count.
- Slot / Port—The y-axis display the slot number and port number for receive and transmit link reset errors.
- RX—The number (error count) of receive link reset errors. The slot and port number display on the y-axis. Pause on the bar to display the slot and port number as well as the error count.
- TX—The number (error count) of transmit link reset errors. The slot and port number display on the y-axis. Pause on the bar to display the slot and port number as well as the error count.

  You can perform the following functions for both RX and TX bars.

  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Top Port PCS Block Errors widget

The Top Port PCS Block Errors widget displays the top ports with Physical Coding Sublayer (PCS) block errors outside of frames in a bar chart. There are three PCS Block Errors port widgets: ISL, Initiator, and Target.

NOTE
PCS block errors are only applicable on 10 and 16 Gbps ports.

FIGURE 78 Top Port PCS Block Errors widget

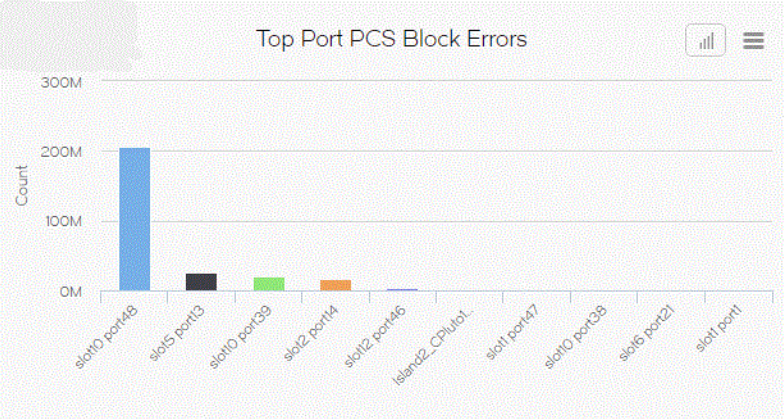

The Top Port PCS Block Errors widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- Count—The x-axis displays the count.
- Slot / Port—The y-axis display the slot number and port number for the number (error count) of PCS block errors outside of frames for the duration specified in the widget.

  You can perform the following functions.

  - View additional details by clicking the bar and selecting Show Details (refer to "Viewing additional details for the Top Port PCS Block Errors widget" on page 182).
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Viewing additional details for the Top Port PCS Block Errors widget

To view additional details for the widget, complete the following steps.

1. Click a bar in the chart and select the Show Details icon.

   A more detailed widget displays which includes the following data:

   - Back icon—Click to go back to the Top Port PCS Block Errors widget.
   - Other actions icon—Select to export a graphic of the widget.
   - Count—The x-axis displays the error count.
   - PCS Block Errors—The number (error count) of PCS block errors outside of frames for the duration specified in the widget.

2. Click the back (←) button to go back to the Top Port PCS Block Errors widget.

Performance widgets

Fabric Insight Portal provides additional preconfigured performance widgets. These performance widgets are available when you create a user-defined dashboard.

Top Port C3 Discards widget

The Top Port C3 Discards widget displays the top ports with Class 3 frames discarded in a bar chart.

The Top Port C3 Discards widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- Count—The x-axis displays the error count.
- Slot / Port—The y-axis display the slot number and port number. Pause on the bar to display the slot and port number and the number (error count) of Class 3 frames discarded for the duration specified in the widget. You can also perform the following functions:
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Top Port Traffic widget

The Top Port Traffic widget displays the top ports with receive and transmit traffic in a table.

FIGURE 79 Top Port Traffic widget

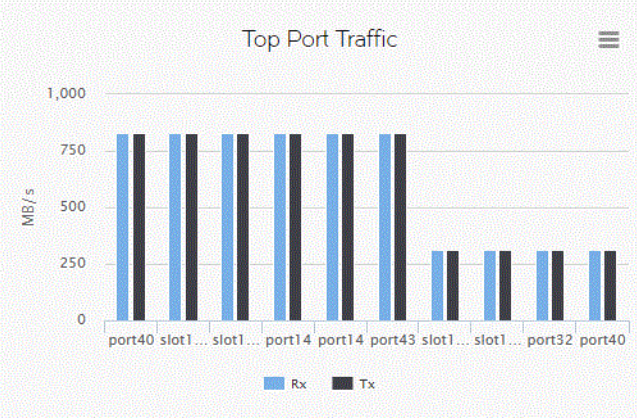

The Top Port Traffic widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- MB/s—The x-axis displays the MB/s value.
- Port/Slot—The y-axis displays the port number or slot and port number. Pause on the bar to display the slot and port number and the receive and transmission speed for the port. You can also perform the following functions:
  - View additional details by clicking the bar and selecting Show Details (refer to "Viewing additional details for the Top Port Traffic widget" on page 184).
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Viewing additional details for the Top Port Traffic widget

To view additional details for the Top Port Traffic widget, complete the following steps.

1. Click a bar in the chart and select Show Details.

   A more detailed widget displays which includes the following data:

   - Widget Name (Port Name)—The name of the widget and the port affected by this widget.
   - Back icon—Click to go back to the Top Port Traffic widget.
   - Other actions icon—Select to export a graphic of the widget.
   - RX Min—The minimum receive traffic rate in megabits per second. Pause on the bar to view the receive traffic rate.
   - RX Avg—The average receive traffic rate in megabits per second. Pause on the bar to view the receive traffic rate.
   - RX Max—The maximum receive traffic rate in megabits per second. Pause on the bar to view the receive traffic rate.
   - TX Min—The minimum transmit traffic rate in megabits per second. Pause on the bar to view the transmit traffic rate.
   - TX Avg—The average transmit traffic rate in megabits per second. Pause on the bar to view the transmit traffic rate.
   - TX Max—transmit traffic rate in megabits per second. Pause on the bar to view the receive traffic rate.

2. Click the back (←) button to go back to the Top Port Traffic widget.

Top Port Invalid Transmissions widget

The Top Port Invalid Transmissions widget displays the top ports with receive and transmit traffic in a table. The Top Port Invalid Transmissions widget includes the following data:

- Widget title—The name of the widget.
- Other actions icon—Select to export a graphic of the widget.
- Count—The x-axis displays the number of invalid transmissions.
- Port/Slot—The y-axis displays the port number or slot and port number. Pause on the bar to display the slot and port number and the number of invalid transmissions for the port. You can also perform the following functions:
  - View port properties by clicking the bar and selecting Properties (refer to "Viewing port properties" on page 171).
  - Investigate port performance by clicking the bar and selecting Investigate (refer to "Investigating port performance" on page 172).

Events

The Events page lists the events and alerts that have occurred on discovered devices in your network as well as on the Management server. The Events page displays a maximum of 5,000 events within the selected time scope.

Viewing events

Click the Events ( ) icon to display the Events page.

FIGURE 80 Events page

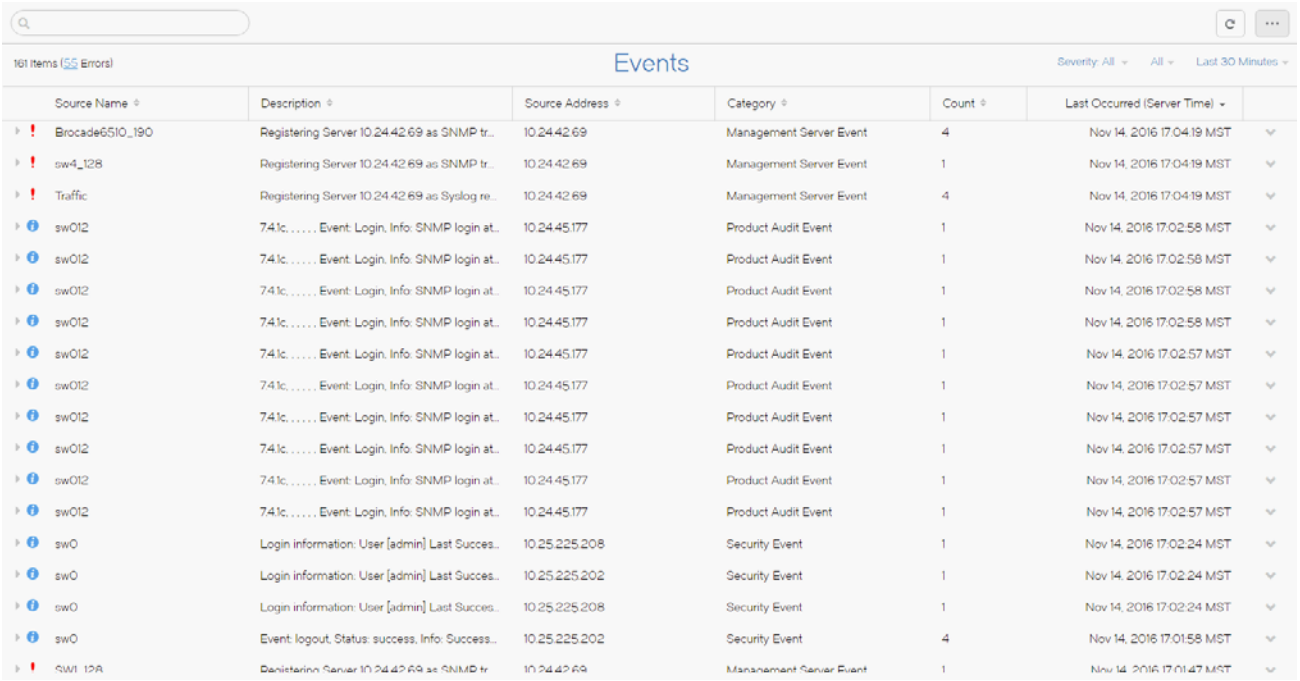

The Events page contains the following fields and components:

- Events filter—Use to filter the list of events. For more information, refer to "Global Filter" on page 159.
- Refresh icon—Click to refresh the Events page.
- Ellipsis icon—Click to access the user guide or information about the Management application release.
- Events count—The total number of events triggered and the total number of Alert and Error events triggered. Click the Alert or Error event count to display a filtered events list.
- Severity list—Select an option to filter events by severity group. Options include All (default), Alert, Error, Info, and Warning. The Alert group includes (Emergency, Alert, and Critical events. The Info group includes Notice, Info, and Debug events. For more information, refer to "Displaying events by severity" on page 188.
- Network scope list—The selected network scope. Click to select the network scope by which you want to filter events. Options include each discovered product and fabric or All (default). For more information, refer to "Setting the network scope" on page 188.
- Date range list—The selected date range. Click to define the date range by which you want to filter events. Options include 30 Minutes (default), 1 Hour, 6 Hours, 12 Hours, 1 Day, 3 Days, 1 Week, and 1 Month or you can customize a date range. For more information, refer to "Setting the time interval" on page 189.

- Expand Event Details icon—Click the expand icon to display event details. For more information, refer to "Viewing event details" on page 186.
- Severity icon—The severity group icon (Alert, Error, Info, and Warning) associated with the event. There are four severity groups including Alert, Error, Info, and Warning. The Alert group includes Emergency, Alert, and Critical events. The Info group includes Notice, Info, and Debug events.
- Source Name—The product on which the event occurred.
- Description—A description of the event.
- Source Address—The IP address (IPv4 or IPv6 format) of the product on which the event occurred.
- Category—The type of event that occurred (for example, client/server communication events).
- Count—The number of times that the event occurred.
- Last Occurred (Server Time) —The time and date on which the event last occurred on the server.
- Actions list—Click to view product or fabric data (refer to "Viewing the FC product summary" on page 191).

Event page functions

The Events page displays a maximum of 5,000 events within the selected time scope.

- Refresh the table by clicking the refresh icon (). Refresh returns the table to its original state. If you sorted the table, it returns to the default state. If you expanded an event, it closes the event. Refresh does not affect filter options such as severity, network scope, or date range.
- Sort the table by clicking a column head. Click a column head again to reverse the sort order.

Viewing event details

The Events page displays the most important fields for an event. You can expand an event to view additional event information. You can select only one event detail at a time. To view additional event details, click the expand icon (next to the severity icon).

FIGURE 81 Expanded event details

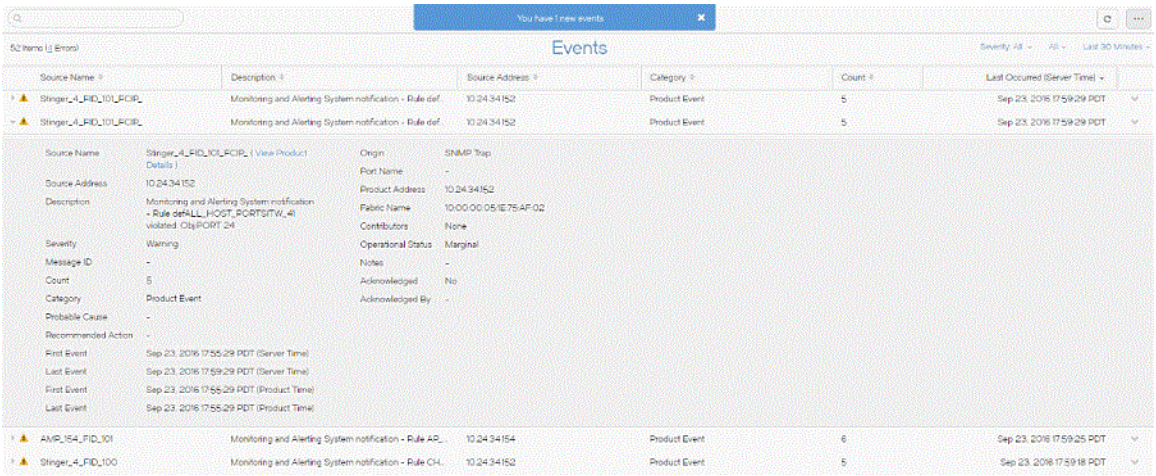

The following list details the fields included in event details. Data displays only if the field have values.

- Source Name—The product on which the event occurred. Click View Product Details to view additional information about the device. For more information, refer to "Viewing FC Fabric properties" on page 192.

- Source Address—The IP address (IPv4 or IPv6 format) of the product on which the event occurred.
- Description—A description of the event.
- Severity—The severity of the event. Options include Emergency, Alert, Critical, Error, Warning, Notice, Info, and Debug events.
- Message ID—The message ID of the event.
- Count—The number of times the event occurred.
- Category—The type of event that occurred (for example, client/server communication events).
- Probable Cause—The most likely reason the event occurred.
- Recommended Action—The recommended action to take to remedy the event
- First Event —The day, date, and time the first event occurred on the server.
- Last Event —The day, date, and time the last event occurred on the server.
- First Event —The day, date, and time the first event occurred on the product.
- Last Event —The day, date, and time the last event occurred on the product.
- Origin—The event source type.
- Port Name—The port name associated with the event.
- Product Address—The IP address of the product on which the event originated.
- Fabric Name—The VCS fabric name.
- Contributors—The contributor to this event.
- Operational Status—The product's operational status.
- Notes—Any notes entered for the event.
- Acknowledged—Whether the event is acknowledged.
- Acknowledged By—The name of the user who acknowledged the event.

To close event details, click the expand icon (next to the severity icon).

When you refresh the table, the event details automatically close.

Refreshing the Events page

The Events page does not update automatically. When new events occurs in your area of responsibility (AOR), a notification panel displays at the top of the Events page. The Notification panel displays the new event count as well as the new severity Alert and Error event counts with links to the filtered event list. The notification panel continues to update as new event notifications are received from the server.

To refresh the Events page, select one of the following options:

- Click the Refresh icon ().
- Click the notification panel to display the new events, with the default search criteria, in the Events page.

The notification panel closes on selection.

To close the notification panel without refreshing the Events page, click the close (x) icon on the panel.

The notification panel redisplays after a minute if any new events have occurred on any of the devices with in your AOR.

Displaying events by severity

You can configure the Events page to display events based on their severity group. There are four severity groups including Alert, Error, Info, and Warning. The Alert group includes Emergency, Alert, and Critical events. The Info group includes Notice, Info, and Debug events.

1. Click the Severity arrow icon.
2. Select the severity group from the list.

   Options include Alert, Error, Info, and Warning. The Alert group includes Emergency, Alert, and Critical events. The Info group includes Notice, Info, and Debug events.

   The Events page displays only events with the selected severity group.

Setting the network scope

You can configure the Events page to display all objects in your area of responsibility (AOR) or a subset of objects (fabrics, devices, or groups) using the network scope selection. Default network scopes are visible to all users. User-defined scopes are visible only to the user who created them.

1. Click the **Scope** arrow.

   FIGURE 82 Select Network Scope dialog box

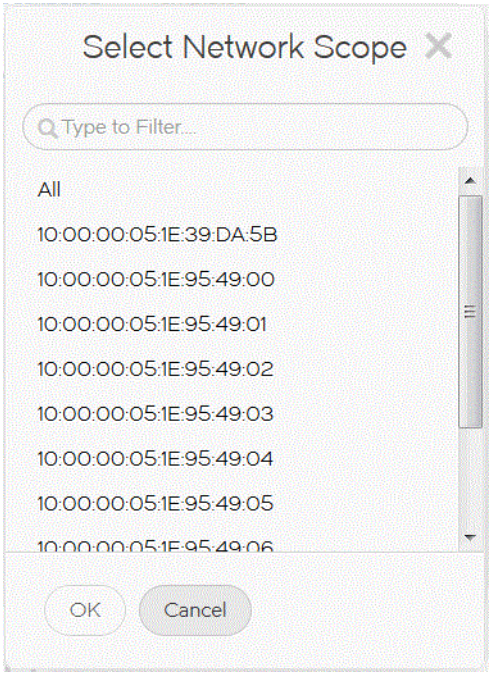

2. Select a network from the **Network Scope** list.

   The default network scope is All. It includes all managed and monitored fabrics or groups in your AOR. If the selected fabric or group is deleted from discovery, the widget refreshes and returns to the default network scope (All).

   Search for a product or fabric by entering search criteria in the Search field. The Network Scope list automatically filters out any products that do not match the search criteria.

3. Click OK.

   If you select a fabric scope, the Events page displays only events triggered in the fabric.

   If you select a product scope, the Events page displays only events triggered in the selected product.

   If you select a port scope, the Events page displays only events triggered in the specified ports.

Setting the time interval

Setting the global time interval in the Events page configures the data display time range. The time interval in the Scope list allows you to select a specific time range for which to display data in the Events page.

1. Click the **Scope** arrow.

   FIGURE 83 Select Date Range dialog box

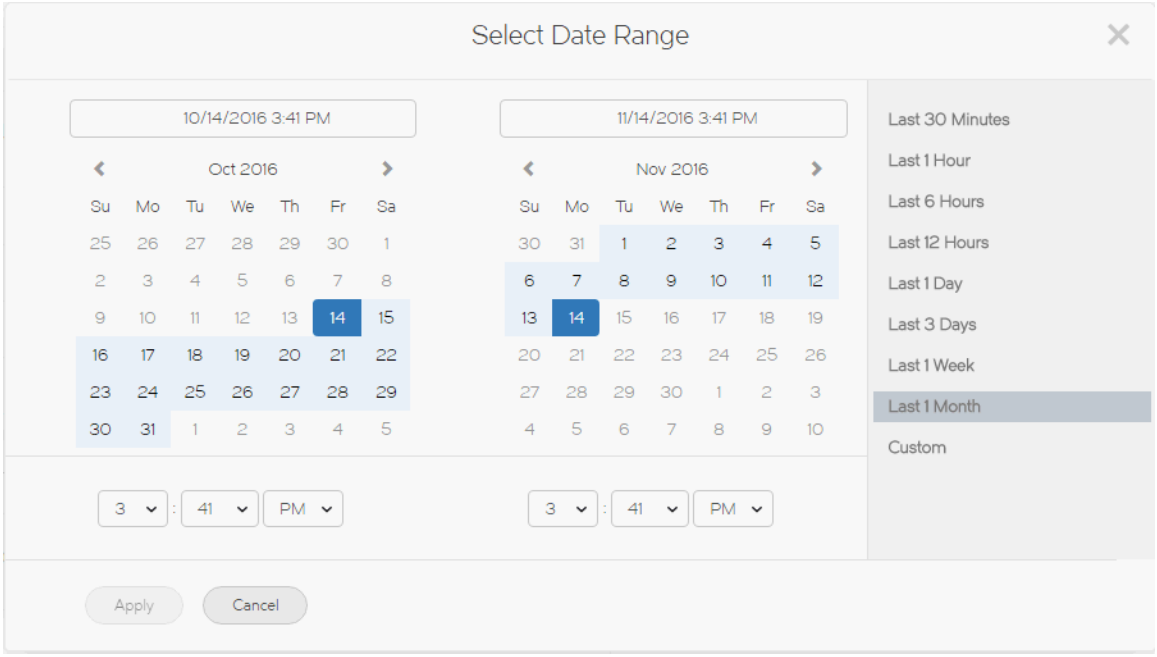

2. To display data for the current date and time based on the selected time scope, select one of the following options:

   - 30 Minutes—Displays data for 30 minutes.
   - 1 Hour—Displays data for 1 hour.
   - 6 Hours—Displays data for 6 hours.
   - 12 Hours—Displays data for 12 hours.
   - 1 Day—Displays data for 24 hours.
   - 3 Days—Displays data for 3 days.
   - 1 Week—Displays data for 1 week.
   - 1 Month—Displays data for 30 days.

   Go to step 9.

3. To display data for data for a custom date and time based on the selected time scope, select **Custom** from the Time Scope list. Continue with step 4.
4. Select a start date from the left Calendar pop-up window.
5. Select a start time from the left Choose Time pop-up window
6. Select a end date from the right Calendar pop-up window.
7. Select a end time from the right Choose Time pop-up window
8. Select a duration from the Time Scope list.

   The displayed data changes to the new time range for all the applicable widgets.
9. Click **Apply**.

Inventory

The Inventory page displays a detailed view of the discovered devices and fabrics. You can view the products based on the Product groups selected from the list. The following are the supported product groups:

- All Products (default)—Lists all the discovered fabrics and switches.
- FC Products—Lists only the FC products.

Click the Inventory icon ( ) to view the All Products inventory. The Inventory page supports infinite scrolling.

FIGURE 84 All Products inventory page

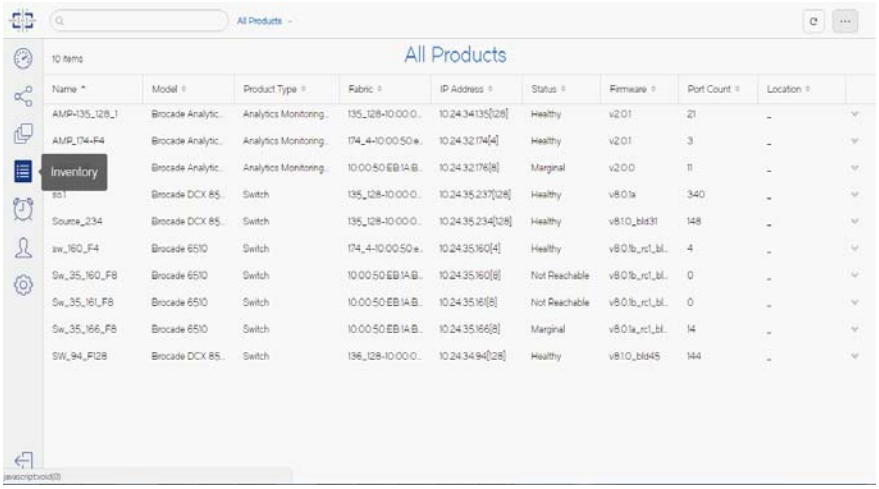

The Inventory page displays the devices in your inventory. This list can include all switches in your inventory or a subset of switches based on filtering criteria.

- Inventory count—The total number of switches available in the inventory.
- Inventory table—All the available switches in the inventory. The following details displays:
  - Name—The name of the switch.
  - Model—The model number of the switch.
  - Product Type—The type of product.

- Fabric—Specifies to which fabric the switch belongs.
- IP Address—The IP address (IPv4 or IPv6 format) of the switch.
- Status—The status of the product.
- Firmware—The firmware version of the product.
- Port Count—The available ports in the switch.
- Location—The location of the device.
- Action menu—

The following columns are supported only on the FC Products group:

- WWN—The WWN of the product.
- State—The state of the product.

Fabric summary view

The Fabric summary view displays the product list, fabric summary, and properties panes for a selected fabric.

The Fabric Insight Portal supports the following contextual menu only for Fabric OS products.

- View
- Fabric properties

Viewing the FC product summary

To view the details for a FC product, complete the following steps.

1. Click the Inventory ( ) icon.
2. Select FC Products from the drop-down list.

   The FC Products page displays all the discovered fabrics.

   FIGURE 85 FC Products Inventory page

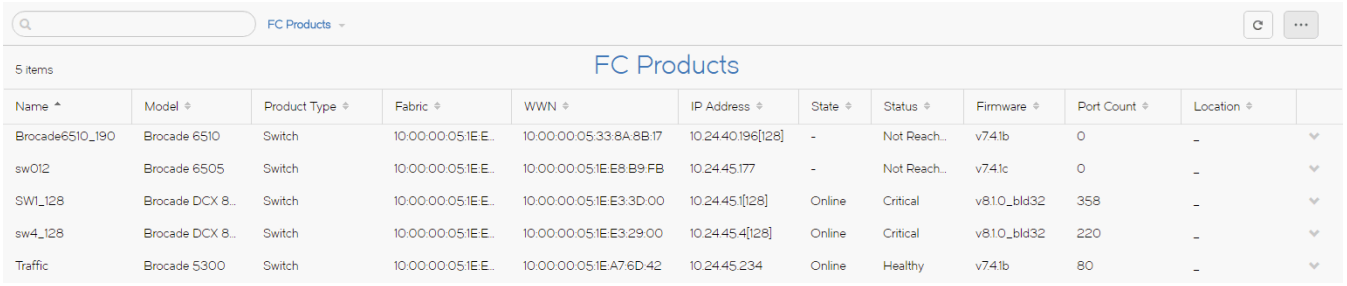

3. Select View from the Action menu.

   The detailed view page displays the following details:

   - Device properties
   - CPU Utilization
   - Memory Utilization
   - Up Time

- Port properties

FIGURE 86 FC Products detailed view page

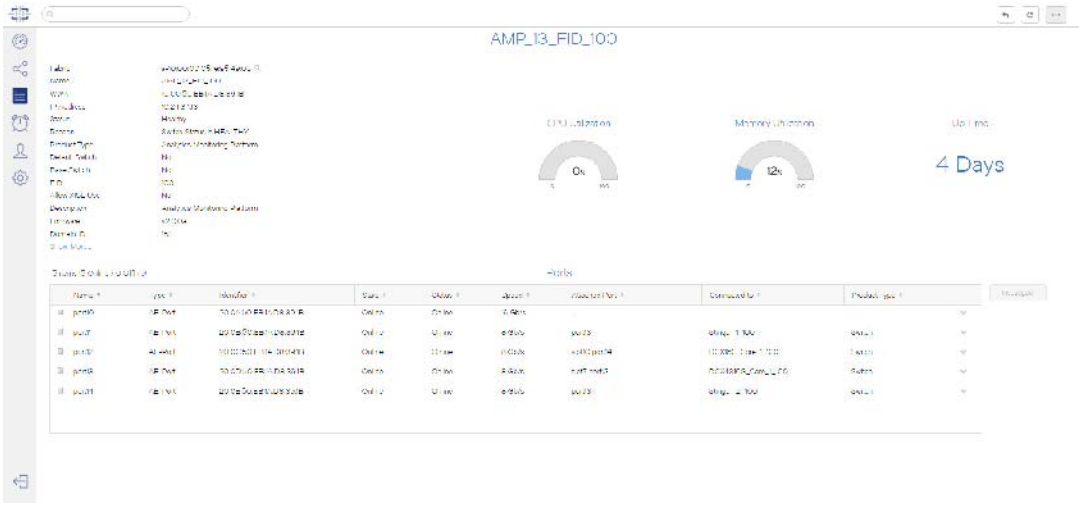

You can filter the Port properties view by selecting FC or VE/VEx or GigE to view the ports based on the port types only for the FC Products. The default port type is FC. You can further investigate the port or view port properties by selecting the action menu of the port.

4. Click Back ( ) to return to the main page.

Viewing FC Fabric properties

To view Fabric properties, complete the following steps.

1. Click the Inventory ( ) icon.

   The All Products Summary page is displayed.

2. Select FC Products from the drop-down list.

   The FC Products page displays all the discovered fabrics.

3. Select Fabric Properties from the Action menu for the fabric or device.

   The Fabric Properties dialog box displays.

FIGURE 87 SAN Fabric Properties

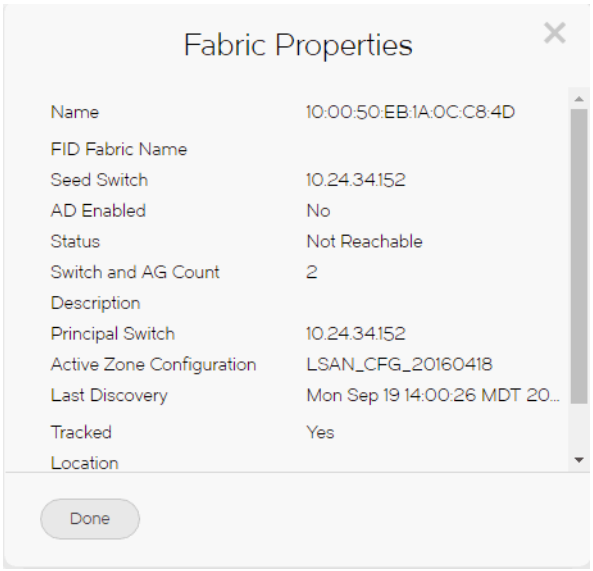

The fabric properties dialog box contains the following fields:

- Name—The name of the selected fabric.
- FID Fabric Name—The Fabric ID name of the selected fabric.
- Seed Switch—The IP address of the seed switch for the selected fabric.
- AD Enabled—The AD Enabled fabrics.
- Status—The worst status for the discovered products in the selected fabric.
- Switches and AG Count—The total number of switches and Access Gateways in the fabric.
- Description—The description of the product.
- Principal Switch—The IP address of the principal switch.
- Active Zone Configuration—Whether fabric tracking is on or off.
- Last Discovery—The physical location of the fabric.
- Tracked—The name of the person or group you should contact about the fabric.
- Location—The description of the fabric.
- Contact—The name of the contact person.

4. Click Done to close and return to the products page.

Port Summary View

The Port summary view displays the product summary and port panes for a selected devices and fabrics.

Viewing the port summary

To view data for a fabric, complete the following steps.

1. Click the Inventory ( ) icon.

   The All Products page displays all the available devices and fabrics.

2. Select a device or fabric name to view the port summary details.

   The port summary displays at the bottom of the device or fabric summary page.

   FIGURE 88 FC Port summary view

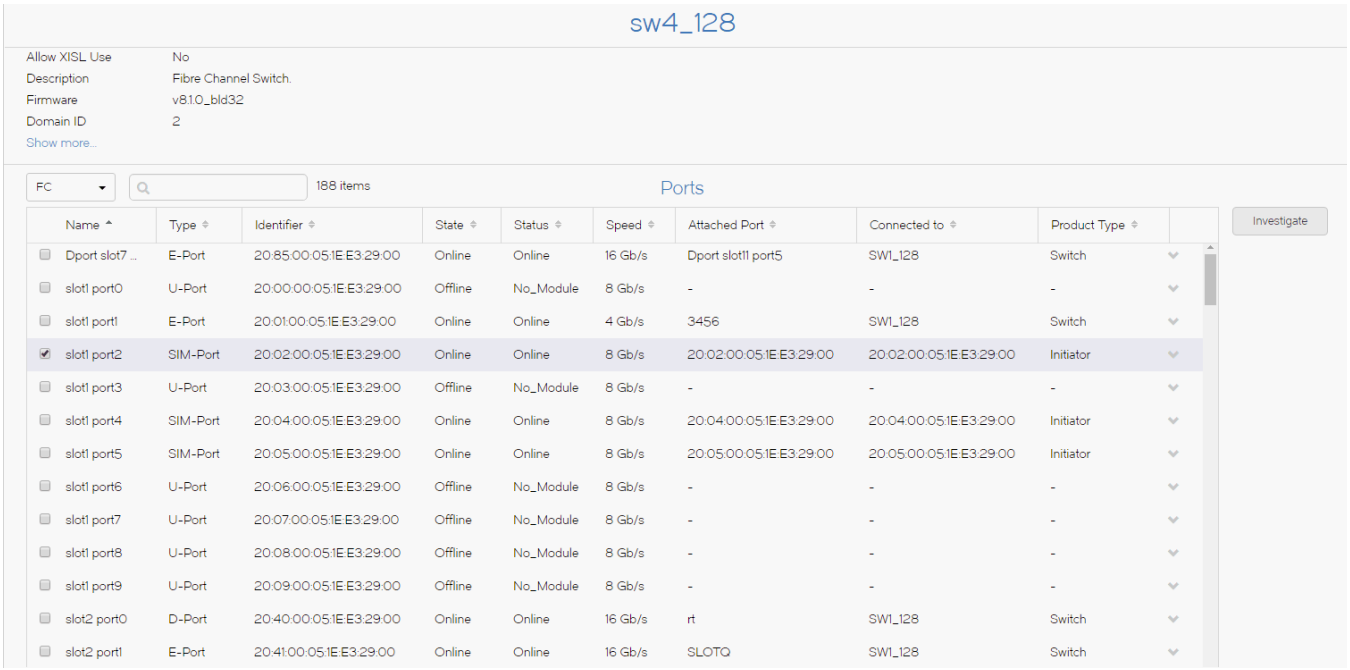

3. Click Back ( ) to return to the All Products page.

Viewing port properties

To view data for a port, complete the following steps.

1. Click the Inventory ( ) icon.

   The All Products page displays all the available devices and fabrics.

2. Select a device or fabric name to view the port summary details.

   The port summary displays at the bottom of the device or fabric summary page.

3. Select **Properties** from the Action menu for the port.

The port Properties dialog box displays.

FIGURE 89 FC Port properties dialog box

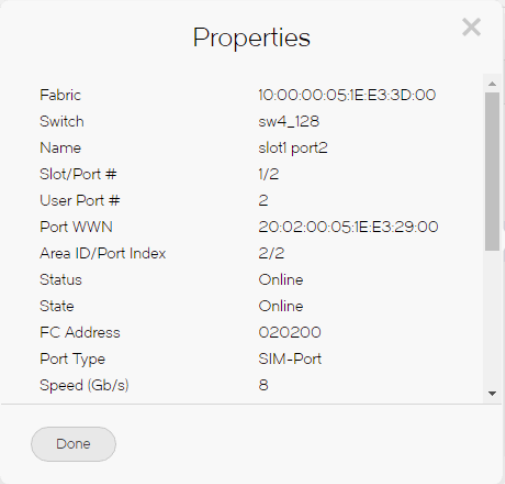

NOTE
Depending on the device type, some of the properties listed below may not be available for all products.

The SAN port properties dialog box contains the following fields:

- Fabric—The name of the fabric.
- Switch—The name of the switch.
- Name—The name of the port.
- Slot/Port #—The slot and port number of the selected fabric. The port number includes the type of port (FC, TE, GE, or XGE).
- User Port #—The user's port number of the switch.
- Port WWN—The world wide name of the port.
- Area ID/Port Index— The port index of the connected switch.
- Status—The operational status of the switch.
- State—The operational state of the port.
- FC Address—The FC address of the port.
- Port Type—The port type. For example, F-Port, L-Port, and so on.
- Speed (Gb/s)—The speed in gigabytes per second.
- Protocol—The network protocol, for example, Fibre Channel.
- Long Distance Setting—Whether the connection is considered to be normal or longer distance.
- Forward Error Correction—Whether FEC is enabled or disabled.
- Encryption—Whether encryption is enabled or disabled
- Compression—Whether compression is enabled or disabled.
- NPIV Enabled—Whether the port is NPIV enabled or not.
- Calculated Status—The calculated operational status.

FIGURE 90 IP Port properties dialog box

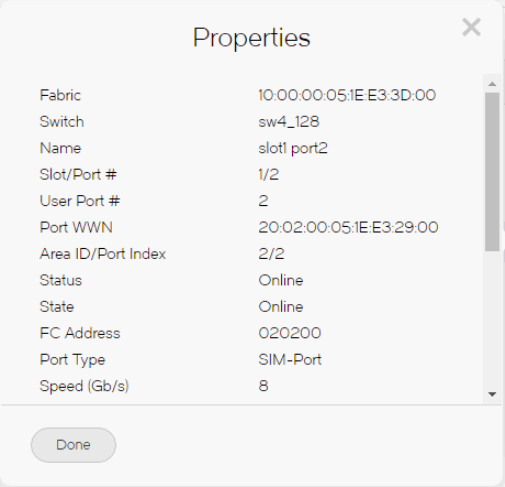

4. Click **Done** to close and return to the All Products page.

Investigating Historical performance

You can investigate the Historical performance for the selected device and port for a selected time range.

1. Click the Inventory ( ) icon.
2. Select a device to view the port view summary details.
3. Select a port under **Ports** table and click **Investigate**.

   The **Investigate Performance** page displays with following radio-button options.

   - RealTime
   - Last 30 Minutes drop-down list.
4. Select the measures from the **Measures** pane.
5. Select the ports from the **Show** drop-down list.
6. Select Last 30 Minutes option and click the drop-down list.

   The **Select Date Range** dialog box displays.

7. Set the date range for performance investigation and click **Apply**.

   The performance graph displays for the selected port and date range.

FIGURE 91 Historical performance graph

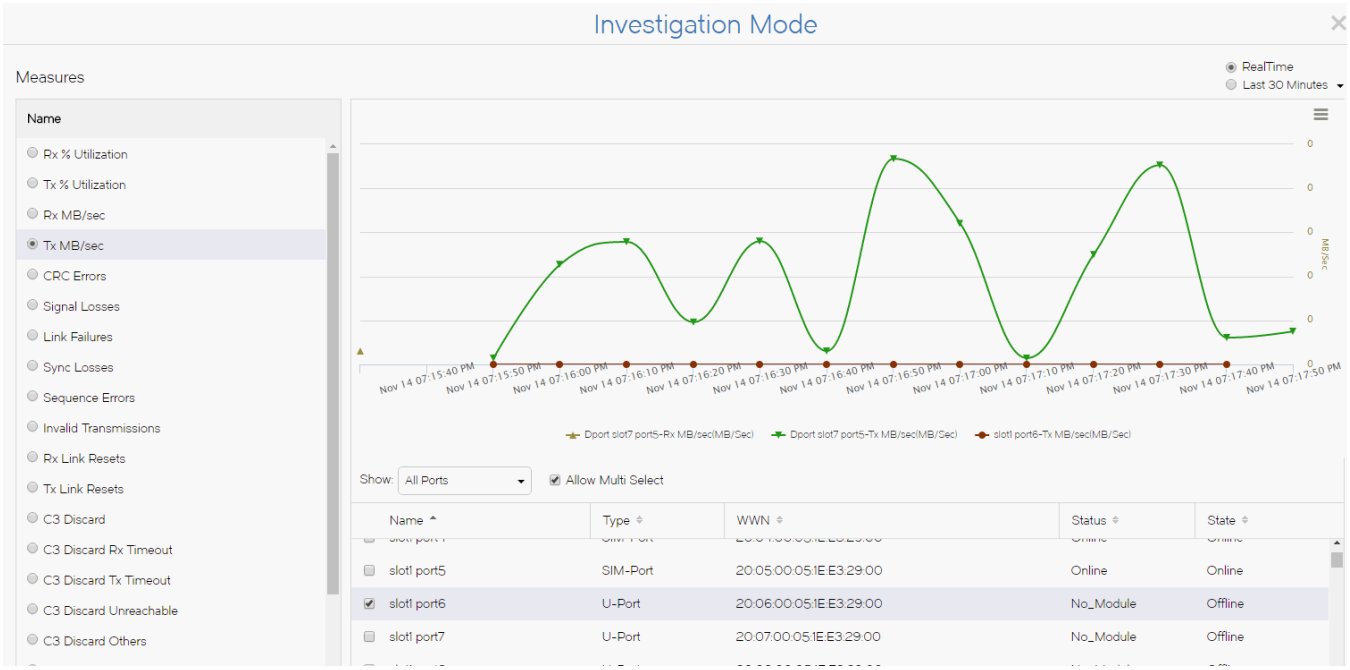

Investigating RealTime performance

You can investigate the RealTime performance for selected device and port.

1. Click the Inventory ( ) icon.
2. Select a device to view the device and port view summary details.
3. Select a port under **Ports** table and click **Investigate**.

   The **Investigate Performance** page displays with following radio-button options.

   - RealTime
   - Last 30 Minutes drop-down list
4. Select the **RealTime** option.
5. Select the measures from the **Measures** pane.

6. Select the ports type from the **Show** drop-down list.

   The realtime investigate graph displays for the selected port.

FIGURE 92 RealTime performance graph

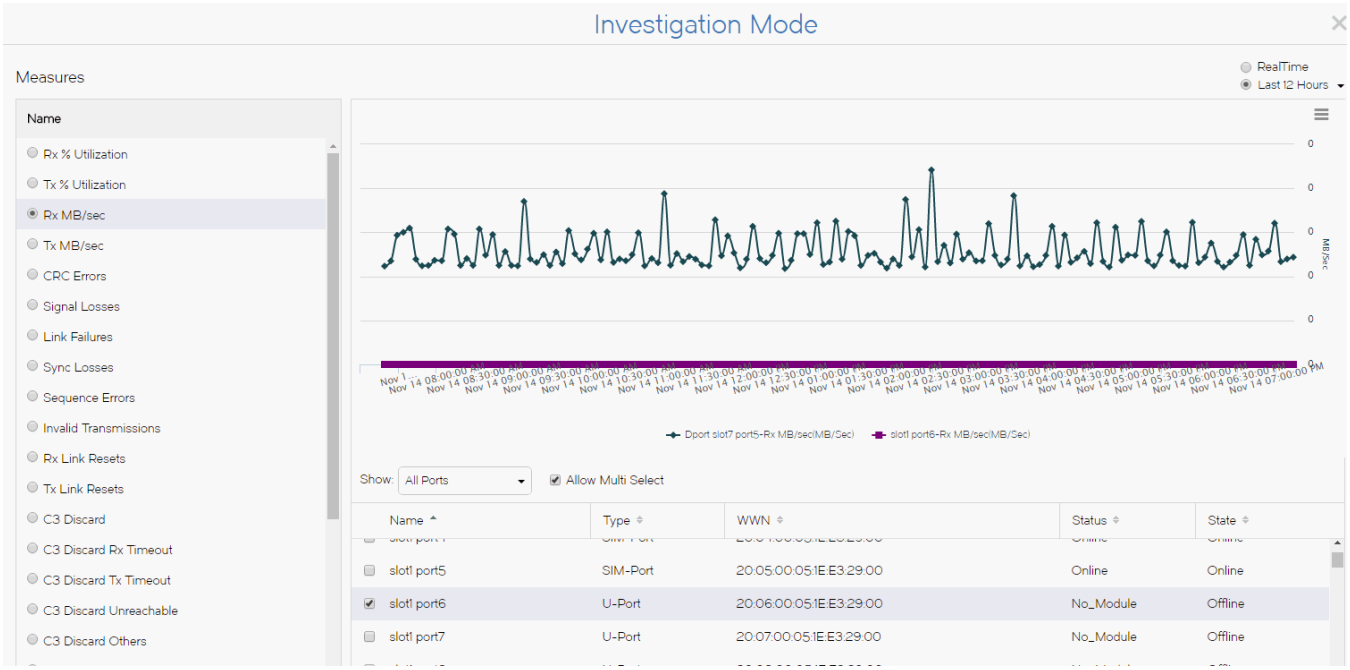

Dashboard Management

Dashboard overview

NOTE
Only devices in your area of responsibility (AOR) display in the dashboard.

The **Dashboard** tab displays the status widgets, performance monitors, and the Master Log. You can also display additional status widgets and performance monitors, as needed. The Management application has the following default dashboards: Product Status and Traffic, SAN Port Health, Storage Port Health, and WAN Vision.

The dashboard provides a high-level overview of the network and the current states of managed devices. This allows you to easily check the status of the devices on the network. The dashboard also provides several features to help you quickly access reports, device configurations, and system logs.

The dashboard updates regardless of the currently selected tab (**SAN**or **Dashboard**) or the SAN size. However, data may become momentarily out of sync between the dashboard and other areas of the application. For example, if you remove a product from the network while another user navigates from the dashboard to a more detailed view of the product, the product may not appear in the detailed view.

FIGURE 93 **Dashboard tab**

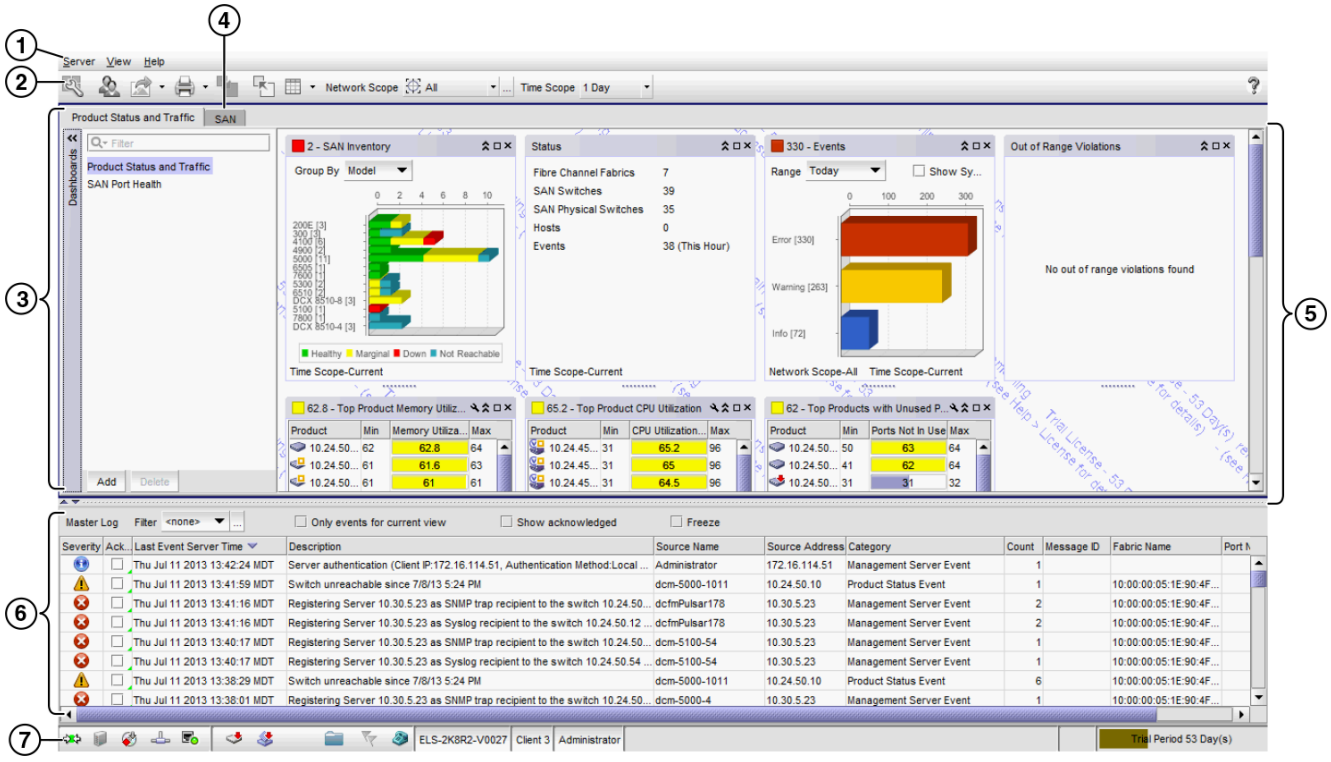

1. **Menu bar —** Lists commands you can perform on the dashboard. For a list of **Dashboard** tab menu commands, refer to "Dashboard main menus" on page 1295.

   The dashboard also provides a shortcut menu to reset the dashboard back to the defaults. Reset the dashboard back to the default settings by right-clicking in the white space and selected **Reset to Default**.

2. **Toolbar —** Provides buttons that enable quick access to dialog boxes and functions. For a list of **Dashboard** tab toolbar options, refer to "Dashboard toolbar" on page 201.

3. **Dashboard** tab and **expand navigation bar** — Provides a high-level overview of the network managed by Management application server. The expand navigation bar is located left of the status widgets or performance monitors and provides a list of dashboards to choose from as well as buttons to perform add and delete functions. For more information, refer to "Dashboards expand navigation bar" on page 202.

4. **SAN** tab — Displays the Master Log, Minimap, Connectivity Map (topology), and Product List. For more information, refer to the "SAN tab overview".

5. **Widgets —** Displays operational status, inventory status, event summary, and overall network or fabric status as well as performance monitors. For more information, refer to "Status widgets" on page 217 and "Performance monitors" on page 233.

6. **Master Log —** Displays all events that have occurred on the Management application. For more information, refer to "Master Log" on page 300.

7. **Status bar —** Displays the connection, port, product, fabric, special event, Call Home, and backup status, as well as Server and User data. For more information about the status bar, refer to "Status bar" on page 302.

Dashboard toolbar

The toolbar (Figure 94) is located beneath the menu bar and provides icons and buttons to perform various functions.

FIGURE 94 Toolbar

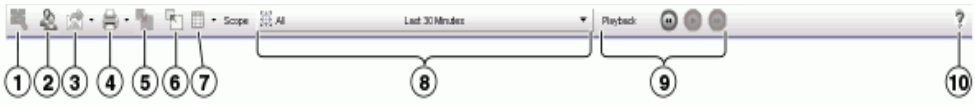

The toolbar contains the following icons and buttons:

1. **Customize Dashboard** — Displays the **Customize Dashboard** dialog box. Use to configure which status widgets and performance monitors display on the **Dashboard** tab. For more information, refer to **"Customizing the dashboard widgets and monitors"** on page 207.
2. **Users** — Displays the **Users** dialog box. Use to configure users, user groups, and permissions. For more information, refer to "User accounts" on page 138.
3. **Export** list — Saves the current dashboard display (all widgets) or a selected widget in a .png format. For more information, refer to "Exporting the dashboard display" on page 208.
4. **Print** list — Prints the dashboard display (all widgets) or a selected widget. For more information, refer to "Printing the dashboard display" on page 209.
5. **Attach** — Returns the dashboard to the main window. For more information, refer to "Attaching and detaching the Dashboard tab" on page 209.
6. **Detach** — Detaches the dashboard to a separate window. For more information, refer to "Attaching and detaching the Dashboard tab" on page 209.
7. **Dashboard display** list — Use to select how to display the status widgets and performance monitors in the dashboard. For more information, refer to "Filtering the dashboards list" on page 203.
8. **Scope** list — Use to select network scope and time scope for which you want to display data in the dashboard. For more information, refer to "Customizing the dashboard scope" on page 209.
9. **Dashboard Playback** — Use to play, pause, rewind, and forward the dashboard and widgets in playback mode. For more information, refer to "Dashboard playback" on page 214.
10. **Help** — Displays the online help.

Dashboard messages

The dashboard message bar (Figure 95) only displays when the **Scope** list (network scope and time scope) has changed in other clients. You can also view all dashboard messages and clear them.

FIGURE 95 Dashboard message bar

The toolbar contains the following components:

1. **Details** button — Use to view dashboard messages.
2. **Close** button — Use to close the dashboard message bar.

Dashboards expand navigation bar

The expand navigation bar (Figure 96) is located left of the status widgets or performance monitors and provides a list of dashboards to choose from as well as buttons to perform add and delete functions.

FIGURE 96 Expand navigation bar

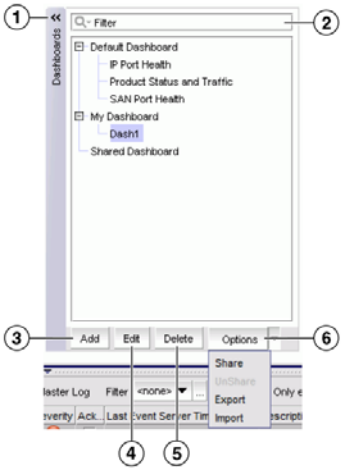

The toolbar contains the following fields and components:

1. **Dashboards** expand navigation bar — Use to select the dashboard you want to view from the list. For more information, refer to "Accessing a dashboard" on page 203.
2. **Filter** — Use to search for the dashboard you want to view.
3. **Add** button — Use to create a dashboard. For more information, refer to "Creating a user-defined dashboard" on page 204.
4. **Edit** button — Use to edit a user-defined dashboard name. For more information, refer to "Editing a user-defined dashboard" on page 204.
5. **Delete** button — Use to delete the selected user-defined dashboard. For more information, refer to "Deleting a user-defined dashboard" on page 205.
6. **Options** button — Use to share, unshare, export, and import a user-defined dashboard. For more information, refer to "Sharing a user-defined dashboard definition" on page 205, "Unsharing a user-defined dashboard definition" on page 205, "Exporting a user-defined dashboard definition" on page 206, and "Importing a user-defined dashboard definition" on page 206.

General dashboard functions

The Management application also provides the following general functions which are applicable to all widgets and monitors:

- **Preference persistence —** Any customization you make to the dashboards persists in the dashboard when you log off and log back in again.
- **Severity** — Most widgets display a severity icon (worst severity of the data shown) next to the widget title. The SAN Status and SANand Host Inventory widgets also indicate the number of products with that severity. The Events widget displays a severity icon with the highest severity event color. The Status widget does not display the severity icon.
- **Title bar buttons** — Status widgets have the following three (left to right) title bar buttons: expand/collapse, maximize/minimize, and close. Performance monitors are editable and have the following four (left to right) title bar buttons: edit, expand/collapse, maximize/minimize, and close.

- **Resizing** — All widgets can be resized by dragging the grab bars. Use the vertical grab bars between widget columns to adjust the width of widgets in the adjacent columns. Use the horizontal grab bars to adjust the height of adjacent widget rows.

  Reset the dashboard back to the default size by right-clicking in the white space and select **Reset to Default**.
- **Zoom in** — Only widgets with a bar graph enable you to zoom in using your mouse. To zoom in, click the upper left of the widget area on which you want to zoom in, drag the mouse to the lower right, and release the mouse button.
- **Zoom out** — Only widgets with a bar graph enable you to zoom out using your mouse. To zoom out, click the lower right widget area on which you want to zoom out, drag the mouse to the upper left, and release the mouse button.
- **Tooltips** — Only widgets with a pie chart or bar graph display tooltips when you pause on a section or bar.
  - For the pie chart widgets, the tooltip displays the name of the category, number of items in that category, and the percentage.
  - For the bar graph widgets, the tooltip displays the count represented by the selected bar.

Dashboard privileges

You must have dashboard privileges to perform various dashboard operations. For a new installation, all out-of-box roles have the dashboard **Read and Write** privilege. After migration, all out-of-box and custom roles have the dashboard **Read and Write** privilege. A user with the User Management privilege can create a new role with the dashboard Read Only, Read and Write, or No privilege, and assign it to new users. For more information, refer to "User privileges" on page 1333.

Accessing a dashboard

From the **Dashboards** expand navigation bar, double-click the dashboard you want to view. Options include:

- **Product Status and Traffic** — Displays preconfigured status widgets and performance monitors. You can display additional widgets and monitors in this dashboard.
- **SAN Port Health** — Displays preconfigured SAN performance monitors. You can display additional status widgets and performance monitors in this dashboard.
- **Storage Port Health** — Displays preconfigured status widgets and monitor ports that are connected to the storage devices.
- **My Dashboard** — Displays a user-defined dashboard. The dashboard name created is case- sensitive and unique to the user. The user-defined dashboard can be shared with other users.
- **Shared Dashboard** — Displays the user-defined dashboard shared by other users. A shared icon () is displayed for shared dashboards. For more information, refer to "Sharing a user-defined dashboard definition" on page 205, "Unsharing a user-defined dashboard definition" on page 205.

Filtering the dashboards list

You can filter the list of dashboards to display only the dashboard you need.

1. Click the **Dashboards** expand navigation bar.
2. Enter your filter criteria in the **Filter** text box.
3. To make the filter case-sensitive or insensitive, choose one of the following options from the filter icon list:
   - **Case sensitive** — Select to make the filter case-sensitive.
   - **Case insensitive** — Select to make the filter case-insensitive.
4. To allow wild cards or regular expressions, choose one of the following options from the filter icon list:

- **Use wildcards** — Select to use wildcards in the **Filter** text box.
- **Use regular expression** -- Select to use a unicode regular expression. Enter a Unicode regular expression in the **Filter** text box.

5. To determine how to match the filter text, choose one of the following options from the filter icon list:
   - **Match from start** — Select to match from the start of the dashboard name.
   - **Match exactly** — Select to match the dashboard name exactly.
   - **Match anywhere** — Select to match text anywhere in the dashboard name.
6. To determine how to handle leaf nodes as well as parent and children nodes, choose one of the following options from the filter icon list:
   - **Match leaf node only** — Select to only include leaf nodes in the filter.
   - **Hide nodes without children** — Select to exclude nodes without children from the filter.
   - **Keep the children if any of their ancestors match** — Select to include children in the filter when any of their ancestors match.
7. Press **Enter**.

   The filter results display in the **Dashboards** expand navigation bar. To stop the filter, click the stop filter (X) icon in the **Filter** text box.

Creating a user-defined dashboard

You can create a dashboard and customize it with the status widgets and performance monitors you need to monitor your network.

1. Click the **Dashboards** expand navigation bar.
2. Click **Add**.

   The **Add Custom Dashboard** dialog box displays.
3. Enter a name and description for the dashboard.Select **Copy active dashboard widgets** to include all widgets in the current dashboard in this dashboard.
4. Click **OK**.

   The new dashboard displays in the **Dashboards** expand navigation bar under **My Dashboard** and becomes the active dashboard.

   To cutomize the dashboard, refer to "Customizing the dashboard widgets and monitors" on page 207.

Editing a user-defined dashboard

You can edit a user-defined dashboard name.

1. Click the **Dashboards** expand navigation bar**.**
2. Select the dashboard you want to edit and click **Edit**.

   The **Edit Dashboard** dialog box displays.
3. Edit the name and description of the dashboard.

4. Click **OK.**

   The edited name displays in the **Dashboards** expand navigation bar under **My Dashboard**.

   To customize the dashboard, refer to "Customizing the dashboard widgets and monitors" on page 207.

Deleting a user-defined dashboard

You can delete a user-defined dashboard.

1. Click the **Dashboards** expand navigation bar.
2. Select the dashboard you want to delete and click **Delete**.
3. Click **Yes** on the confirmation message.

   The user-defined dashboard is deleted in the **Dashboards** expand navigation bar under **My Dashboard.**

Sharing a user-defined dashboard definition

You can share the user-defined dashboard with other users. The changes made in the shared dashboard will reflect to all the shared users. When the owner deletes a shared dashboard, it is unshared from all the shared users and removed from the Shared Dashboard list.

1. Click the **Dashboards** expand navigation bar.
2. Select a user-defined dashboard you want to share under **My Dashboard**, click **Options,** and then select **Share**.

   The following message displays: "The selected dashboard will be shared with all users".

3. Click **OK**.

   The selected user-defined dashboard is shared with all users and is displayed with a shared icon.

   NOTE
   You cannot share default dashboards and shared dashboards.

Unsharing a user-defined dashboard definition

You can unshare a user-defined dashboard from other users. When the owner deletes a shared dashboard, it is unshared from all the shared users and removed from the shared dashboard list. If the unshared dashboard is the active dashboard, the shared user is notified with a warning message and the Product Status and Traffic dashboard becomes the active dashboard.

1. Click the **Dashboards** expand navigation bar.
2. Select a user-defined dashboard you want to unshare under **My Dashboard**, click **Options,** and then select **Share**.

   The following message displays: "Are you sure of unsharing the dashboard <*Dashboard name*>".

3. Click **OK**.

The selected user-defined dashboard is unshared from all the users.

Exporting a user-defined dashboard definition

You can export a user-defined dashboard definition and save it in the .zip file format. The exported dashboard should include the dashboard details, widget definition, and user setting details.

1. Click the **Dashboards** expand navigation bar.
2. Select a user-defined dashboard you want to share under **My Dashboard**, click **Options,** and then select **Export.**

   The **Export** dialog box displays.
3. Browse the path to save the details and click **Select**. A successfully exported information message displays.
4. Click **OK**.

   The user-defined dashboard definition details are saved in a .zip file in a location that you specify.

   NOTE
   You cannot export an empty dashboard and published widgets.

Importing a user-defined dashboard definition

You can import a user-defined dashboard definition from the file system to the Management application.

1. Click the **Dashboards** expand navigation bar.
2. Select a user-defined dashboard you want to share under **My Dashboard**, click **Options,** and then select **Import.**

   The **Import** dialog box displays.
3. Browse the path and click **Select**. A successfully imported information message displays.
4. Click **OK**.

   The user-defined dashboard definition is imported. If the imported dashboard name already exists, the imported dashboard name is renamed as <*Dashboard name_1*>, <*Dashboard name_2*>, and so on in an incremental sequence. The imported dashboard becomes the active dashboard.

Setting the dashboard display

You can set the dashboard to minimize or expand all status widgets and performance monitors as well as return to the default settings.

Select one of the following options from the dashboard display list:

- **Collapse All** — Select to minimize all widgets and monitors on the dashboard.
- **Expand All** — Select to expand all widgets and monitors on the dashboard.
- **Reset to Default** — Select to reset the dashboard to the default display settings.

Customizing the dashboard widgets and monitors

NOTE
You cannot customize a default (system-defined) dashboard. When you view a default dashboard, the **Customize Dashboard** icon is not available. For a list of default dashboards, refer to "Default dashboards" on page 215.

1. From the dashboard, click the **Customize Dashboard** icon.

   The **Customize Dashboard** dialog box displays.

2. Click the **Status** tab.

   The preconfigured general status widgets display.

3. Select the **Display** check box in the **General Status Widgets** list for each status widget you want to add to the dashboard.

   Clear the check box to remove the associated status widget from the dashboard.

   The **General Status Widgets** list contains the following additional information:

   - **Title** — The name of the status widget. For more information, refer to "Status widgets" on page 217.
   - **Description** — A general description of the status widget.

4. Click the **Performance** tab (Figure 97).

   The preconfigured performance monitors display. You can create up to 100 performance monitors; however, you can only display up to 30 performance monitors. For more information about performance monitors, refer to "Performance monitors" on page 233.

FIGURE 97 Customize Dashboard dialog box, Performance tab

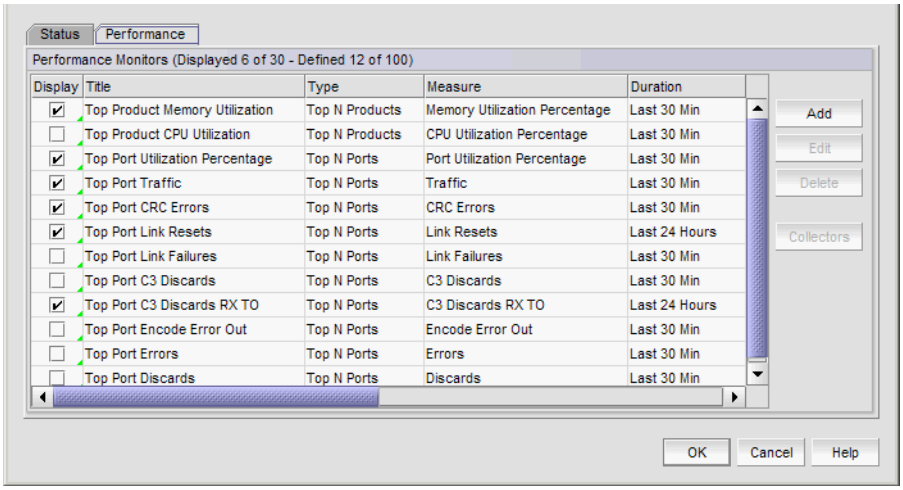

5. Select the **Display** check box in the **Performance Monitors** list for each performance monitor you want to add to the dashboard.

   Clear the check box to remove the associated performance monitor from the dashboard.

   The **Performance Monitors** list contains the following additional information:

   - **Title** — The name of the performance monitor. For more information, refer to "Performance monitors" on page 233
   - **Type** — The type of monitor.
   - **Measure** — The performance measures included in the monitor.
   - **Data Collectors —** The data collectors that provide data for the monitor.

6. Click **Add** to add a new performance monitor. For more information, refer to "Configuring a user-defined product performance monitor" on page 275.

7. Click **Edit** to edit an existing performance monitor. For more information, refer to "Configuring a user-defined product performance monitor" on page 275 or "Editing a preconfigured performance monitor" on page 267.

8. Select one or more user-defined monitors and click **Delete** to delete the user-defined performance monitors.

9. Click **OK** to close the **Customize Dashboard** dialog box.

Exporting the dashboard display

You can export the current dashboard display (all widgets and monitors) or a selected widget or monitor in .png format.

1. Select one of the following options from the **Export** list:

   - **Dashboard** — Exports the current dashboard.
   - ***Name*** — Exports the selected widget (where ***Name*** is the name of the widget or monitor on the dashboard).

   The **Export Dashboard to PNG File** or **Export *Name* to PNG File** dialog box displays.

2. Browse to the location you want to save the file.

3. Enter a name for the snapshot in the **File Name** field, if needed.

   Export uses the following naming convention: ***Name_yyyy_mm_dd_hh_mm_ss***.png.

4. Click **Save**.

   The file is saved to the location you selected.

Printing the dashboard display

You can print the current dashboard display (all widgets and monitors) or a selected widget or monitor.

1. Select one of the following options from the **Print** list:
   - **Dashboard** — Prints the current dashboard.
   - *Name* — Prints the selected widget (where *Name* is the name of the widget or monitor on the dashboard).

   The **Page Setup** dialog box displays.
2. Change the page setup options, as needed.
3. Click **OK**.

Attaching and detaching the Dashboard tab

You can detach the **Dashboard** tab from the main application to display in a separate window.

To detach the **Dashboard** tab, click the Detach icon. The **Dashboard - *Dashboard_Name* – *Application_Name*** window displays.

Reattach the **Dashboard** to the main application by clicking the Attach icon or by closing the **Dashboard - *Dashboard_Name* – *Application_Name*** window. The **Dashboard** tab displays in the main application window.

Customizing the dashboard scope

You can customize the dashboard display by setting the network scope and time scope in the **Scope** list (Figure 98).

FIGURE 98 Scope list

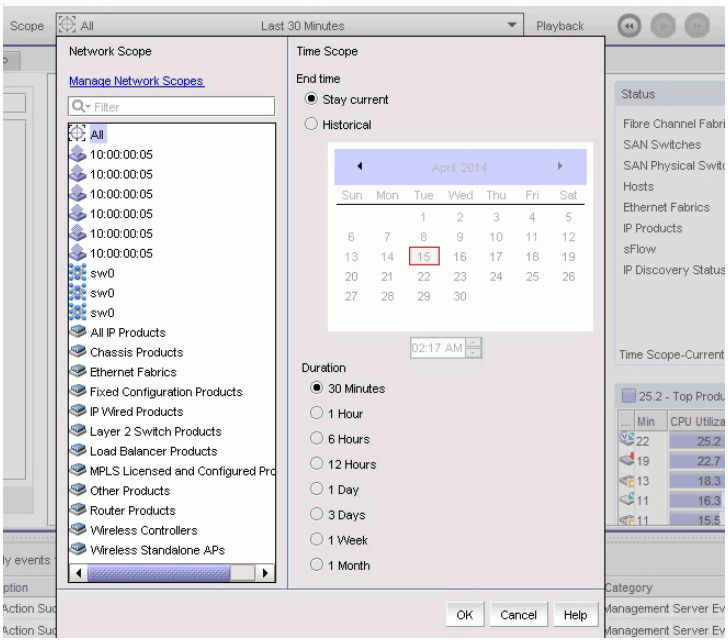

Setting the network scope

You can configure the dashboard to display all objects in your area of responsibility (AOR) or a subset of objects (fabrics, devices, or groups) using the network scope selection. Default network scopes are visible to all users. User-defined scopes are visible only to the user who created them. When a scope is changed, all widgets corresponding to the network scope in the dashboard get refreshed.

You can either select a network scope from the **Available Network Scopes** or create a user-defined network scope by clicking the **Manage Network Scopes** in the **Scope** list.

The available network scopes include the following list of options:

- All
- Any SAN fabric
- Any Ethernet fabric
- Any system-defined group
- Any user-defined group (IP product and port group)
- Any user-defined customized network

If you select a fabric scope, dashboard widgets display data for all products and ports in the fabric.

If you select a product scope, dashboard widgets display data for the selected products and the ports that belong to the selected products.

If you select a port scope, dashboard widgets display data for the specified ports and the products to which the ports belong. If any of the selected ports are initiator or target ports, dashboard widgets display data for the attached switch port.

In case of MAPS widgets, violation counts will be displayed based on the selected scope.

The default network scope is **All**. It includes all managed and monitored fabrics or groups in your AOR.

If the selected fabric or group is deleted from discovery, the widget refreshes and returns to the default network scope (**All**).

Filtering the available network scopes

You can search for a specific network scope using the **Filter** field available in network scope list.

1. Click the **Scope** list.

2. Enter the filter criteria in the **Filter** field.

The **Network Scope** list will be updated automatically based on the search criteria.

Creating a customized network scope

You can create a user-defined network scope from any objects in your AOR. You can create network scopes based on fabrics, products, product groups, or ports.

1. Click the **Scope** list.
2. Click **Manage Network Scopes**.
3. The **Edit Scopes** dialog box displays a list of existing user-defined network scopes in the **Network Scopes** list.

FIGURE 99 Edit Scopes dialog box

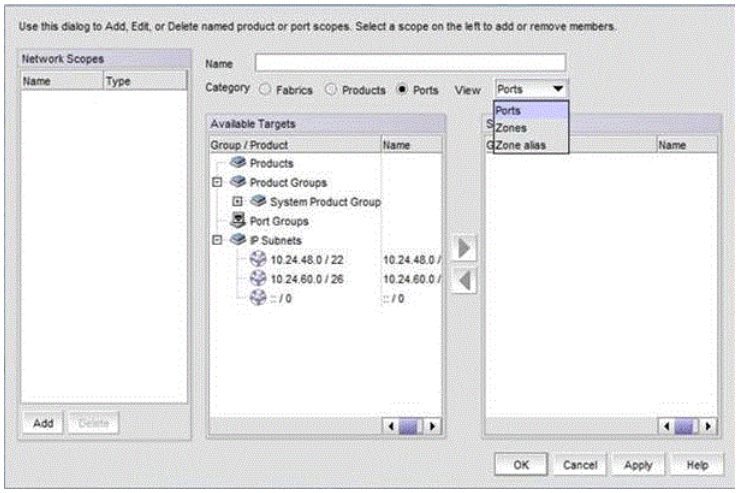

4. Click **Add**.

   A new network scope displays in the **Network Scopes** list.

5. Enter a name for the scope in the **Name** field.

6. Select one of the following **Category** options**:**

   - **Fabrics** — Select to create your network from one or more fabrics.
   - **Products** — Select to create your network from one or more products or product groups.
   - **Ports** — Select to create your network from one or more ports or port groups or zones or zone aliases.

7. Select one of the following in the **View** list under the **Ports** category:

   - **Ports (default)** — Select to display the existing network based on the Ports category with one or more ports or port groups in the **Available Targets** list.
   - **Zones** — Select to display the list of zones that belong to the defined zone database along with its members in the Fabric > Zone(s) > Zone members tree structure. If a zone has a zone alias as a member, the structure is Fabric > Zone(s) > Zone alias > Zone members under the alias. The current active zone configuration of the fabric is highlighted with a green square icon.
   - **Zone alias** — Select to display the list of zone aliases of the fabric along with its members in the **Fabric > Zone alias(es) > Zone members** tree structure. The tree contains all the zone aliases configured in the fabric irrespective of the zone.

     When **Zones** or **Zone alias** is selected, you can move only the entire zone to the RHS and not the individual members of the zone separately. This helps in the dynamic update to the zone in the network scope.

     You can add, remove, or rename a zone or alias member or zone alias through the **Zoning** dialog box or through the CLI. When a zone or zone alias is renamed, it is considered removed and all the members associated with the renamed zone or zone alias are removed from the network scope.

     If you try to change from one category (**Ports, Zones,** or **Zone alias**) to another after the selection to the **Selected Targets** list, a warning message displays. Click **Yes** to change the category and continue with the selection or click **No** to retain the previous category.

8. Select one or more of the objects you want to include in the network from the **Available Targets** list and click the right arrow button.

   The objects display in the **Selected Targets** list. To remove an object from the **Selected Targets** list, select it and click the left arrow button.

9. Click **OK** to save your changes and close the **Edit Scopes** dialog box.

Editing a user-defined network scope

You can edit any user-defined network scope.

1. Click the **Scope** list.
2. Click **Manage Network Scopes**.

   The **Edit Scopes** dialog box displays a list of existing user-defined network scopes in the **Network Scopes** list.

3. Select the network scope you want to edit in the **Network Scopes** list.
4. Change the name for the scope in the **Name** field, if needed.
5. To add objects, select one or more of the objects you want to include in the network from the **Available Targets** list and click the right arrow button.

   The objects display in the **Selected Targets** list.

6. To remove an object from the **Selected Targets** list, select it and click the left arrow button.
7. Click **OK** to save your changes and close the **Edit Scopes** dialog box.

Deleting a user-defined network scope

You can delete any user-defined network scope.

1. Click the **Scope** list.
2. Click **Manage Network Scopes**.

   The **Edit Scopes** dialog box displays a list of existing user-defined network scopes in the **Network Scopes** list.

3. Select the network you want to delete in the **Network Scopes** list.
4. Click **Delete**.
5. Click **OK** to save your changes and close the **Edit Scopes** dialog box.

Setting the time scope

Setting the time scope in the **Scope** list configures and displays data for all applicable widgets in the dashboard. The global duration allows you to select a specific time range as shown in the following list:

- **30 Minutes —** Displays data for 30 minutes.
- **1 Hour —** Displays data for 1 hour.
- **6 Hours —** Displays data for 6 hours.
- **12 Hours —** Displays data for 12 hours.

- **1 Day —** Displays data for 24 hours.
- **3 Days —** Displays data for 3 days.
- **1 Week —** Displays data for 1 week.
- **1 Month —** Displays data for 30 days.

The time scope allows you to select the end time and duration in the last 30 days from the current time.

To set the time scope, select one of the following options under **End time:**

- **Stay current** — Displays data for the current time based on the selected network scope and duration.
- **Historical** — Displays data for a specific date and time based on the selected network scope and duration. The **Historical** option displays a calender with a 30-day timeline. The end date in the calendar is the current date and the calendar will show the last 30 days from the current date. You select a specific date in the last 30 days from current date, for which you want to display data. When you select the **Historical** option, a message bar (Figure 100) displays.

  FIGURE 100 Historical mode message bar

  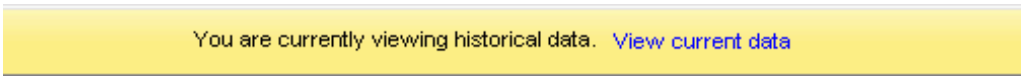

  You can click **View current data** on the message bar to view the current data. If you select the **Stay current** option, all the widgets will start refreshing data in the configured auto-refresh interval.

  Use the spin box to select a specific time for the selected date.

  NOTE
  If the start date is within 8 days of the current date, then the granularity of the spin box will be 5 minutes. If the start date is beyond 8 days (8 to 30), the granularity will be 30 minutes.

  The value in the spin box is incremented or decremented with the granularity based on the start date.

  You can also enter the time manually in the spin box. If the selected date is 8 days or more before the current date, then the entered time will be rounded to the 30-minutes granularity.

  For example, if you enter 03:12 PM and the selected date falls beyond 8 days, the time will be rounded to 3:30 PM.

  If the selected time falls beyond the actual current time, then you cannot manually enter or select a time in the spin box.

Customizing the time scope

1. From the **Scope** list, select the **Historical** option under **End time.**

   The calendar displays the last 30 days from the current date.

2. Select a specific date.
3. Select or enter a specific time in the spin box.
4. Click **OK** and the selected time displays under **Time Scope**.

   The applicable widgets in the dashboard will show the data based on the time scope selected.

For example, if you select the date and time as September 9, 2013, 12:00 PM and the set duration is 30 minutes, all the applicable widgets in the dashboard will show data ranging from 11:30 AM to 12:00 PM of September 9, 2013.

NOTE
If you select the **Stay current** option and click **OK**, all the applicable widgets get refreshed and display data based on the current time and the selected **Duration**.

NOTE
The time scope does not affect Status widgets and Inventory widgets.

Dashboard playback

You can use dashboard control buttons (Pause, Rewind, and Forward) to view the available data of the dashboard and widgets in playback mode. Auto-refresh of data will not occur in playback mode.

- Pause button — Use the Pause button to pause function in playback mode.
- Rewind or Forward button — Use the Rewind or Forward button in playback mode to show the data automatically for every 20 seconds based on the selected time scope.

  If you click the Rewind or Forward button for the first time, data refresh will be in 1x speed with 5 minutes or 30 minutes granularity based on the time scope.

  For example, consider the current time as September 9, 2013, 7:00 AM. In the **Scope** list, if you set the duration as 30 minutes and the time scope as September 9, 2013, 6:15 AM, data will be displayed for all the applicable widgets based on the given time range (September 9, 2013, 5:45 AM to September 9, 2013, 6:15 AM).

  In this example, September 9, 2013, 5:45 AM will be the start time and the selected September 9, 2013, 6:15 AM will be the end time.

  If you click the Forward button, all the applicable widgets in the dashboard will show the data with the granularity of 5 minutes as follows:

  - First refresh will display data from September 9, 2013, 5:50 AM to September 9, 2013, 6:20 AM.
  - Second refresh will display data from September 9, 2013, 5:55 AM to September 9, 2013, 6:25 AM.
  - Third refresh will display data from September 9, 2013,6:00 AM to September 9, 2013, 6:30 AM.

  If the start date is within 8 days from the current date, then the granularity will be 5 minutes. If the start date is beyond 8 days (8 to 30 days), then the granularity will be 30 minutes.

  For example, consider the current time as September 9, 2013, 6:00 AM and the selected time is August 25, 2013, 6:00 AM.

  If you click the Forward button, all the applicable widgets in the dashboard will show the data with the granularity as of 30 minutes as follows:

  - First refresh will display data from August 25, 2013, 6:00 AM to August 25, 2013, 6:30 AM.
  - Second refresh will display data from August 25, 2013, 6:30 AM to August 25, 2013, 7:00 AM.
  - Third refresh will display data from August 25, 2013, 7:00 AM to August 25, 2013, 7:30 AM.

  If you click the Rewind or Forward button successively for the second, third, and fourth time, the data refresh will change to 2x, 3x, and 4x speed, respectively. On the fifth click, the data refresh returns to 1x speed. A label indicates the playback speed. A tooltip for the label will display the duration of the data refresh, based on granularity.

  For example, if you select the start time within 8 days and set 3x speed, the tooltip will show Speed (3x):15 minutes.

  If you select the start time beyond 8 days and set 3x speed, the tooltip will show Speed (3x):90 minutes.

Once the playback reaches the current time, it will change to the **Stay current** option and the widgets will start refreshing in regular intervals.

Default dashboards

The Management application provides preconfigured dashboards which provide high-level overview of the network, the current states of managed devices, and performance of devices, ports, and traffic on the network.

Campus dashboard

The Campus dashboard provides the following preconfigured status widgets and performance monitors:

- Top Product CPU Utilization monitor
- Status widget
- Events widget
- Top Product Memory Utilization monitor
- Top Port Utilization Percentage monitor
- Top Products with Unused Ports monitor
- IP Status widget
- Top Product Temperature monitor
- Port Traffic Distribution monitor
- Port Utilization Distribution monitor

Product Status and Traffic dashboard

The Product Status and Traffic dashboard provides the following preconfigured status widgets and performance monitors:

- SAN Inventory widget
- Status widget
- Events widget
- Custom Events widget
- COMPASS Drifts widget
- Out of Range Violations widget
- Top Product Memory Utilization monitor
- Top Product CPU Utilization monitor
- Top Products with Unused Ports monitor
- Top Port Utilization Percentage monitor (includes details for all ports, Initiator ports, ISL ports, and Target ports)
- Bottom Port Utilization Percentage monitor (includes details for all ports, Initiator ports, ISL ports, and Target ports)

SAN Port Health dashboard

The SAN Port Health dashboard provides the following preconfigured status widgets and performance monitors for the ISL, Host, and Target ports:

- Port Health Violations widget
- Bottlenecked Ports widget
- Top Port CRC Errors monitor
- Top Port Sync Losses monitor
- Top Port Link Failures monitor
- Top Port C3 Discards RX TO monitor
- Top Port Link Resets monitor
- Top Port Encode Error Out monitor

Storage Port Health dashboard

The Storage Port Health dashboard provides the following preconfigured performance monitors:

- FC Storage Port Widgets
  - Top Port Encode Error Out
  - Bottlenecked Ports
  - Top Port Link Failures
  - Top Port Sync Losses
  - Top Port C3 Discards RX TO
  - Top Port Link Resets
- IP and FC Storage Port Widgets
  - Port Health Violation
  - Top Port Utilization
  - Top Port CRC Errors

WAN Vision dashboard

The WAN Vision dashboard enables you to monitor and troubleshoot FCIP tunnel and circuit statistics. This dashboard is only supported on the 16 Gbps 24-FC port, 18 GbE port switch and 32 Gbps, Router Extension blade running Fabric OS 8.1.0 or later.

The WAN Vision dashboard provides the following preconfigured status widgets and performance monitors:

- Top Tunnel Utilization monitor
- Top Tunnel Dropped Packets monitor
- Top Circuit Utilization monitor
- Top Circuit FC Utilization monitor
- Top Circuit IP Extension Utilization monitor
- Top Circuit Jitter monitor
- Top Circuit RTT monitor
- Top Duplicate Acknowledge monitor
- Top Slow Start monitor
- Top Out of Order monitor
- Out of Range Violations widget

Status widgets

The Management application provides the following preconfigured status widgets:

- Bottlenecked Ports widget — Table view of bottlenecked ports and number of violations for each bottlenecked port in the SAN. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.
- Events widget — Bar chart view of events grouped by severity and range.
- Custom Events widget — Bar chart view of call home and special events grouped by severity and range.
- Host Adapter Inventory widget — Stacked bar chart view of Host Adapters grouped by selected category
- Out of Range Violations widget — Table view of all out-of-range threshold violations reported by your SAN devices
- Port Health Violations widget — Table view of out-of-range port health violations reported by your SAN devices. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.
- SAN Inventory widget — Stacked bar chart view of FC devices grouped by operational status and selected category
- SAN Status widget — Pie chart view of FC devices categorized by operational status
- Status widget — List view of various status attributes
- VM Alarms widget — Table view of alarms received from vCenter products

Bottlenecked Ports widget

The **Bottlenecked Ports** widget (Figure 101) displays the bottlenecked port violations for the specified fabric and time range in a table. There are four bottlenecked port widgets: All, ISL, Initiator, and Target.

FIGURE 101 Bottlenecked Ports widget

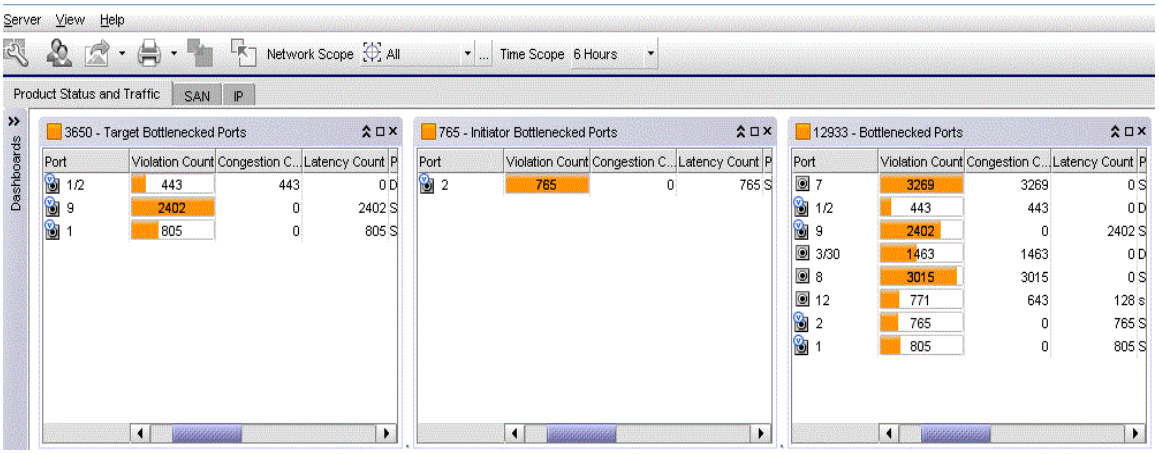

The **Bottlenecked Ports** widget includes the following data:

- Severity icon/violation count/widget title — The color of the worst severity and the total number of ports with congestion and latency violations are displayed before the widget title.
- **Port** — The port identifier, such as port name, number, address, WWN, user port number, or zone alias.

- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Violation Count —** The number of bottleneck violations for the port during the selected time range. This is based on bottleneck configuration. Each trap or alert sent by the switch and the Management application counts as one violation. For more information, refer to "Bottleneck detection" on page 988.
  - **Congestion —** The number of bottleneck violations caused due to congestion for the port during the selected time range.
  - **Latency —** The number of bottleneck violations caused due to latency for the port during the selected time range.

    NOTE
    The Bottleneck violation for the AN -1010 event is shown in the **Latency** count. The **Latency** count increases to 1 in all **Bottlenecked Ports** widgets.
- **Product —** The product label, such as product name, IP address, node WWN, domain ID, or zone alias.
- **Type** — The port type.
- **Storage Type** — The type of storage port (for example, iSCSI, NAS).
- **Attached Device (MACs)** — The MAC address of the attached device.
- **Identifier** — The port identifier, such as port name, number, address, WWN, user port number, or zone alias.
- **Port Number** — The port number.
- **State** — Whether the port is online or offline.
- **Status** — Whether the port is online or offline.

Customizing the Bottlenecked Ports widget

You can customize the widget to display data for a specific fabric and duration:

- To display data for a specific fabric or group, refer to "Creating a customized network scope" on page 210.
- To display data for a specific duration, refer to "Customizing the time scope" on page 213.

Accessing additional data from the Bottlenecked Ports widget

Right-click a row in the widget to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.

Bottleneck Graph dialog box

The **Bottleneck Graph** dialog box (Figure 102) displays the statistics for the selected ports based on the time period.

FIGURE 102 Bottleneck Graph dialog box

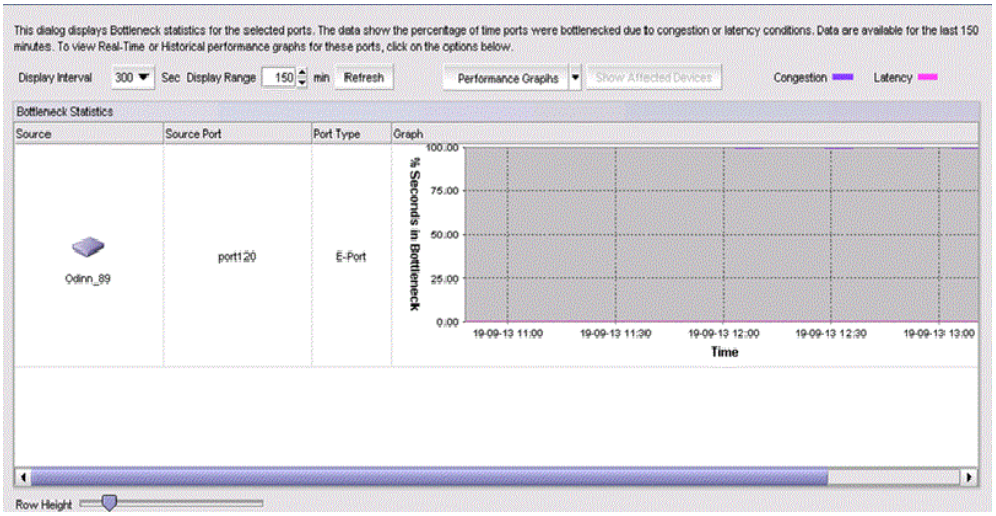

The **Bottleneck Graph** dialog box displays event information for a specific duration by selecting one of the following from the time period:

- If the dashboard time period is 30 minutes, then the **Display Range** is 30 minutes and the **Display Interval** is 60 seconds.
- If the dashboard time period is one hour, then the **Display Range** is 60 minutes and the **Display Interval** is 300 seconds.
- If the dashboard time period is greater than one hour, then the **Display Range** is 150 minutes and the **Display Interval** is 300 seconds.

  NOTE
  On launching the graph, for SIM ports that are running on Fabric OS 7.2.0 or earlier, an error message displays.

Events widget

The **Events** widget (Figure 103) displays the number of events by severity level for a specified network scope, specified time scope, and duration as a stacked bar graph.

FIGURE 103 Events widget

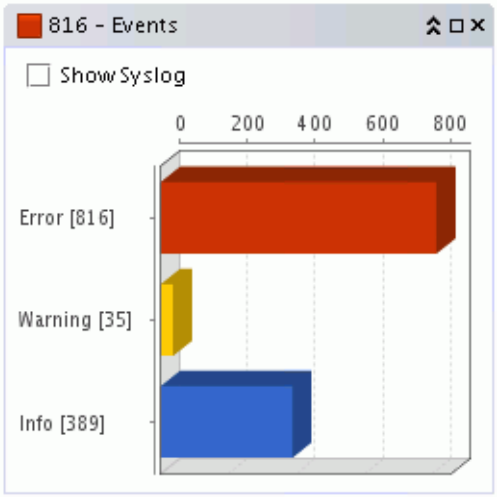

The **Events** widget includes the following data:

- Severity icon/event count/widget title — The color of the worst severity followed by the event count with that severity displays before the widget title.
- **Show Syslog** check box — Select to include Syslog information (default) on the Event Summary.
- Bar chart — The event severity using the color codes in Table 19.

TABLE 19 Event severity color codes

Color	Severity
Red ( )	Emergency
Brick Red( )	Alert
Brick Red ( )	Critical
Brick Red ( )	Error
Gold ( )	Warning
Grey ( )	Notice
Blue ( )	Info

- Network Scope — Select to display the events count of the products present in the selected network scope.
- Time Scope — Select to display the events count of the products present in the selected time range and duration.

NOTE
Application events excluding object level are displayed only if **All** is selected in the Network Scope.

The **Events** widget only includes events from products that are in your AOR.

The x-axis represents the number of occurrences of a particular event severity during the selected time period. If you pause on a bar, a tooltip shows the number of events with that severity level during the selected time period. Also, for each severity, the cumulative number of traps, application events, and security events is reported next to the horizontal bar. If Syslog messages are included, then they are included in the count. To conserve space, the number is shown as is or truncated to the nearest 1,000 ("K") or 1,000,000 ("M").

By default, Syslog events are included in the summary; however, because Syslog events occur at a much higher frequency than other events and therefore could skew the bars for the other events, you can exclude Syslog events. If they are excluded, they will not be displayed in the legend. Users' selections are persisted (per user per server).

Customizing the Events widget

You can customize the **Events** widget to display events for a specific network scope and time scope to display Syslog details.

- Include Syslog information (default) on the **Event Summary** pane by selecting the **Show Syslog** check box. To exclude Syslog information, clear the **Show Syslog** check box.
- To display data for a specific fabric or group, refer to "Setting the network scope" on page 210.

Accessing additional data from the Events widget

Double-click a bar in the **Events** widget to navigate to an event custom report (HTML) that displays the events corresponding to the network scope and time scope selected.

For information about report details, refer to "Fault Management" on page 1131.

Custom Events widget

The Custom Events widget (Figure 104) displays the number of Call Home events and Special events by severity level for a specified network scope, specified time scope, and duration as a stacked bar graph. An event is marked as a Call Home event if it is listed on the **Available Call Event Types** list of the **Call Home Event Filter** dialog box. The refresh time of the widget is 5 seconds.

FIGURE 104 Custom Events widget

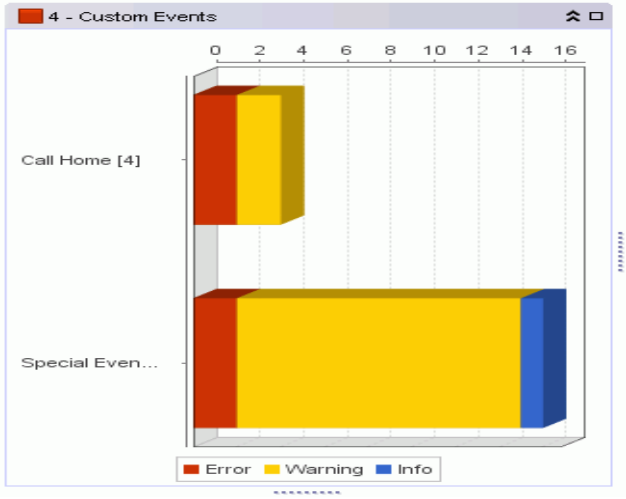

The **Custom Events** widget includes the following data:

- Severity icon/event count/widget title — The color of the worst severity followed by the event count with that severity displays before the widget title.
- Call Home Events bar chart — Displays the Call Home events count grouped by severity using the color codes in Table 19.
- Special Events bar chart — Displays the Special events count grouped by severity using the color codes in Table 19.
- Network Scope — Select to display the events count of the products present in the selected network scope.
- Time Scope — Select to display the events count of the products present in the selected time range and duration.

The **Custom Events** widget only includes events from products that are in your AOR.

The x-axis represents the number of occurrences of a particular event severity during the selected time period. If you pause on a bar, a tooltip shows the number of events with that severity level during the selected time period.

Accessing additional data from the Custom Events widget

Double-click a severity color on the Call Home event bar or Special event bar in the **Custom Events** widget to navigate to an event custom report (HTML) that displays the events corresponding to the network scope and time scope selected.

Double-click a color legend in the **Custom Events** widget to navigate to an event custom report (HTML) that displays both the Call Home events and Special events corresponding to the network scope and time scope selected.

For information about report details, refer to "Fault Management" on page 1131.

COMPASS Drifts widget

The **COMPASS Drifts** widget displays the number of configuration drifts and failed checks for a specified network scope and time scope as a stacked bar chart.

The **COMPASS Drifts** widget includes the following data:

- Widget title — The name of the widget.
- Widget summary — The product count for each status (configuration drifts and failed checks) displays underneath the widget title.
- Stacked bar chart — The number of configuration drifts and failed checks using the color codes. The bar chart displays each group as a separate color on the chart. Tooltips showing the number of configuration drifts and failed checks are shown when you pause on the bar.
- Color legend — Displays the color legend below the bar chart using the following color codes:
  - Red — Displays for drifts.
  - Blue — Displays for failed checks.
- Time Scope — The time scope.

Accessing additional data from the Drifts widget

Double-click a bar in the **COMPASS Drifts** widget to navigate to the **COMPASS Drifts Detailed Viiew** dialog box.

The COMPASS drift data includes the following:

- **Time** — The date and time that the drift occurred.
- **Product** — The name or IP address of the product on which the drift occurred.
- **Object Name** — The name of the product on which the drift occurred.
- **Template** — The name of the linked template.

- **Switch Setting** — The switch setting.
- **Template Setting** — The template setting.

Host Adapter Inventory widget

The **Host Adapter Inventory** widget (Figure 105) displays the host adapter products inventory as stacked bar graphs.

FIGURE 105 Host Adapter Inventory widget

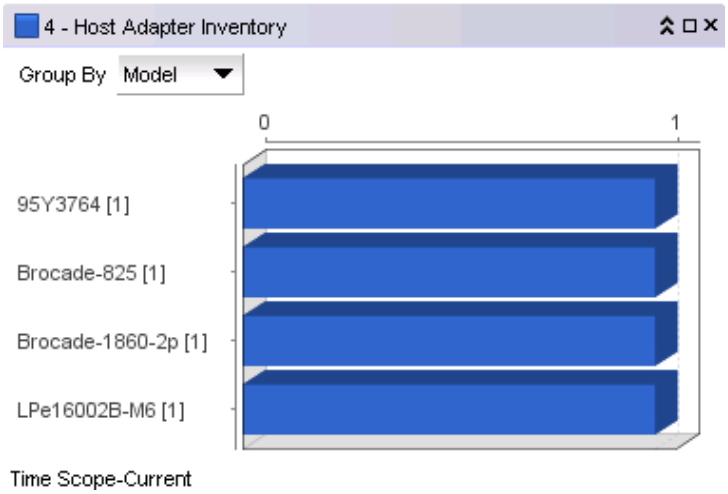

The **Host Adapter Inventory** widget includes the following data:

- Severity icon/Host product count/widget title — The color of the worst severity and the Host product count with that severity displays before the widget title.
- **Group By** list — Use to customize this widget to display a specific grouping. Options include: **Model** (default), **Location**, **Driver**, **BIOS**, and **OS Type**.
- Bar chart — Displays each group as a separate bar on the graph. Displays the current state of all Host products discovered for a group in various colors on each bar. Tooltips showing the number of devices in that state are shown when you pause on the bar.
- Time Scope — The time scope.

Customizing the Host Adapter Inventory widget

You can customize the **Host Adapter Inventory** widget to display product inventory for a specific grouping. The group type and number of products in the group displays to the left of the associated bar; for example, 2.3.0.005 [3], where 2.3.0.005 is the driver number and [3] is the number of products running that driver level.

- Change the grouping by selecting one of the following from the **Group By** list:
  - **Model** — Displays the Host product inventory by model.
  - **Location** — Displays the Host product inventory by physical location.
  - **Driver** — Displays the Host product inventory by driver.
  - **BIOS** — Displays the Host product inventory by BIOS (boot code image version).
  - **OS Type** — Displays the Host product inventory by operating system.

    NOTE
    The OS type is blank when the host has only Emulex host adapters with the Brocade drivers not installed.

- Zoom in on an area of the widget by dragging the mouse (upper left corner to lower right corner) to select one or more bars.

  NOTE
  If the ratio between the longest and shortest bar reaches 5000:1, you should maximize the widget prior to using zoom.

  To return the widget to its original state, reverse the selection (drag from lower right corner to upper left corner).

Accessing additional data from the Host Adapter Inventory widget

Double-click a bar in the **Host Adapter Inventory** widget to navigate to the **Host Adapter Inventory Report**.

SAN Inventory widget

The **SAN Inventory** widget (Figure 106) displays the SAN products inventory as stacked bar graphs.

FIGURE 106SAN Inventory widget

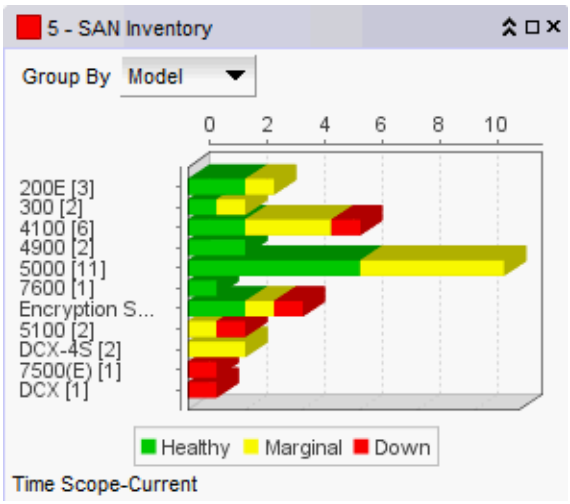

The **SAN Inventory** widget includes the following data:

- Severity icon/product count/widget title — The color of the worst severity followed by the number of products with that severity displays before the widget title.
- **Group By** list — Use to customize this widget to display a specific group of products. Options include: **Firmware**, **Model**, **Location**, and **Contact**.
- Bar chart — The product status as a percentage of the total number of products.

The bar chart displays each group as a separate bar on the graph. Displays the current state of all products discovered for a group in various colors on each bar. Tooltips showing the number of devices in that state are shown when you pause on the bar.

- Color legend — Displays the color legend below the bar chart using the following color codes:
  - Green — Healthy: Status obtained from the SAN switch based on Fabric Watch or Monitoring and Alerting Policy Suite (MAPS) thresholds configured on the switch.
  - Yellow — Marginal: Status obtained from the SAN switch based on Fabric Watch or MAPS thresholds configured on the switch.
  - Red — Down: Status obtained from the SAN switch based on Fabric Watch or MAPS thresholds configured on the switch.
  - Blue — Not Reachable: SAN switch is not reachable by HTTP.
  - Gray — Unknown: Temporary status that displays when switch asset collection is in progress. Once switch asset collection is complete, the current status is obtained from the switch.
- Time Scope — The time scope.

Customizing the SAN Inventory widget

You can customize the **SAN Inventory** widget to display the product inventory for a specific group. The group type and number of devices in the group displays to the left of the associated bar; for example, v7.0.0 [3], where v7.0.0 is the firmware number and [3] is the number of devices running that firmware level.

- Change the grouping by selecting one of the following from the **Group By** list:
  - **Firmware —** The product inventory by firmware release.
  - **Model** — The product inventory by model.
  - **Location —** The product inventory by physical location.
  - **Contact —** The product inventory by contact name.
- Zoom in on an area of the widget by dragging the mouse (upper left corner to lower right corner) to select one or more bars.

  NOTE
  If the ratio between the longest and shortest bar reaches 5000:1, you should maximize the widget prior to using zoom.

  To return the widget to its original state, reverse the selection (drag from lower right corner to upper left corner).

Accessing additional data from the SAN Inventory widget

Double-click a section in the **SAN Inventory** widget to navigate to the **SAN Products - *Status*** dialog box (where ***Status*** is the section of the widget you selected). For more information, refer to "Viewing additional SAN product data" on page 227.

NOTE
It takes a few moments to populate newly discovered products in the **SAN Products - *Status*** dialog box (where ***Status*** is the section of the widget you selected).

SAN Status widget

The **SAN Status** widget (Figure 107) displays the device status as a pie chart.

FIGURE 107SAN Status widget

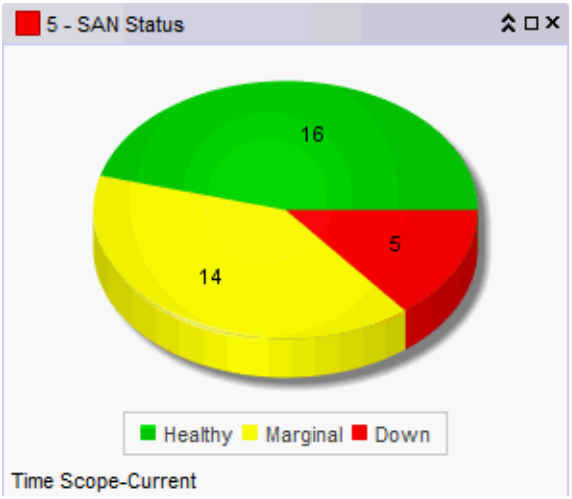

The **SAN Status** widget includes the following data:

- Severity icon/product count/widget title — The color of the worst status followed by the number of products with that status displays before the widget title.
- Pie chart — The device status as a percentage of the total number of devices.

  The pie chart displays the percentage in various colors on each slice. Tooltips showing the number of devices in that state are shown when you pause on the slice. When there is one status category with less than one percent of the total number of devices, the status widget displays the number of devices in each category on each slice.
- Color legend — Displays the color legend below the bar chart using the following color codes:
  - Green — Healthy
  - Yellow — Marginal
  - Red — Down
  - Blue — Not Reachable
  - Gray — Unknown
- Time Scope — The time scope.

Accessing additional data from the SAN Status widget

Double-click a section in the **SAN Status** widget to navigate to the **SAN Products - *Status*** dialog box (where ***Status*** is the section of the widget you selected). For more information, refer to "Viewing additional SAN product data" on page 227.

NOTE
It takes a few moments to populate newly discovered products in the **SAN Products - *Status*** dialog box (where ***Status*** is the section of the widget you selected).

Viewing additional SAN product data

1. Double-click a section in the **SAN Status** widget.

   The **SAN Products - *Status*** dialog box (where ***Status*** is the section of the widget you selected) displays with the following fields and components:

   - **Product —** The product name.
   - **Fabric** — The fabric associated with the product.
   - **Product Type —** The type of product.
   - **State —** The state for the product and the port.
   - **Status —** The status for the product and the port.
   - **Tag —** The tag number of the product.
   - **Serial # —** The serial number of the product.
   - **Model** — The model number of the product.
   - **Port Count —** The number of ports on the product.
   - **Firmware —** The firmware version of the product.
   - **Location** — The physical location of the product. This field is editable at the fabric level.
   - **Contact —** The name of the person or group you should contact about the product. This field is editable at the fabric level.

2. Right-click any row in the table to access the corresponding shortcut menu for the device. For more information about shortcut menus, refer to "SAN shortcut menus" on page 1304.
3. Click **Close**.

Status widget

The **Status** widget (Figure 108) displays the number of products managed and the number of events within the selected event time range.

FIGURE 108 Status widget

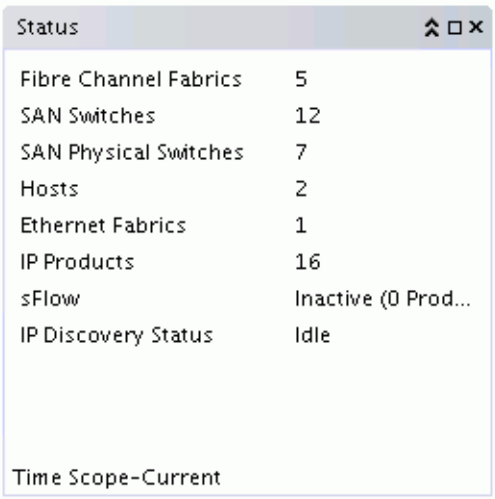

The **Status** widget displays the following items for each product license:

- Fibre Channel Fabrics — The number of managed fabrics.
- SAN Switches — The number of managed SAN switches.
- SAN Physical Switches — The number of discovered physical SAN switches.
- Hosts — The number of managed hosts.
- Time Scope — The time scope.

VM Alarms widget

NOTE
Enabling the **VM Alarms** widget requires discovery of vCenters.

The **VM Alarms** widget displays the vCenter alarms for the specified fabric and time range in a table.

The **VM Alarms** widget includes the following data:

- Severity icon/widget title — The worst severity of the data shown next to the widget title.
- **VM** — Virtual Machine name.
- **Host —** Host name.
- **Total** — Number of alarms triggered by the following violations: VM disk aborts, VM disk resets, VM disk usage (kbps), and VM total disk latency (ms).
  - **Latency —** Number of latency violations.
  - **Usage —** Number of usage violations.
  - **Aborts —** Number of abort violations.
  - **Resets —** Number of reset violations.

Customizing the VM Alarms widget

You can customize the **VM Alarms** widget to display data for a specific fabric and duration.

- To display data for a specific fabric or group, refer to "Creating a customized network scope" on page 210.
- To display data for a specific duration, refer to "Customizing the time scope" on page 213.

Accessing additional data from the VM Alarms widget

- Right-click a row in the widget to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "SAN shortcut menus" on page 1304.
- Double-click a row in the widget to navigate to the **VM Troubleshooting - *VM_Name*** (***Host_Name***) dialog box (where ***VM_Name*** (***Host_Name***) is the name of the virtual machine and associated host). For more information, refer to "Host Management" on page 477.

Monitoring and Alerting Policy Suite widgets

NOTE
MAPS is only supported on a licensed version of the Management application with SAN management.

NOTE
MAPS is only supported on FC devices running Fabric OS 7.2.0 or later with the Fabric Vision license.

NOTE
MAPS is not supported on DCB devices.

The Monitoring and Alerting Policy Suite (MAPS) is an optional storage area network (SAN) health monitor that allows you to enable each switch to constantly monitor its SAN fabric for potential faults and automatically alert you to problems long before they become costly failures.

The widget displays the number of MAPS threshold violations for all network objects (such as ports, trunks, switches, and circuits) for all MAPS-capable devices.

The widget also includes the Fabric Watch threshold violations for devices running Fabric OS 6.4.0 or later with the Fabric Watch license or FC devices running Fabric OS 7.2.0 or later with the Fabric Vision license, but not migrated to MAPS.

The MAPS widgets display on the main **Dashboard** tab. The Management application provides the following preconfigured MAPS widgets:

- Out of Range Violations widget — Table view of all out of range threshold violations reported by your SAN devices .
- Port Health Violations widget — Table view of out of range port health violations.

Out of Range Violations widget

The **Out of Range Violations** widget (Figure 109) displays the number of violations for each MAPS category, Fabric Watch category, and the number of network objects (such as ports, trunks, switches, and circuits) for SAN devices with the MAPS violation and Fabric Watch violation based on the selected fabric and a specified time range.

By default, this widget refreshes every minute. If any violations occur on fabrics in your area of responsibility (AOR) during the minute refresh time frame, the widget refreshes every 10 seconds. If you delete, discover, or unmonitor a device, the widget refreshes.

FIGURE 109 Out of Range Violations widget

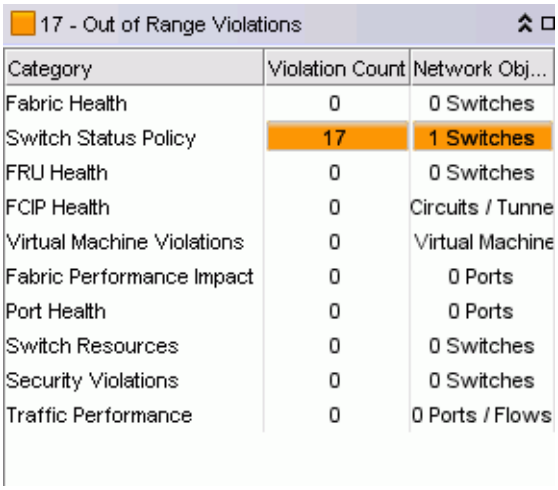

The **Out of Range Violations** widget includes the following fields and components:

- Severity icon/product count/widget title — The color of the worst severity and the number of products with that severity displays before the widget title.
- **Category** — A list of the MAPS and Fabric Watch dashboard categories. Always displays whether or not there is an associated violation. Categories include:
  - Fabric Health
  - FCIP Health
  - FRU Health
  - Port Health
  - Backend Port Health
  - Security Violations
  - Switch Resources
  - Switch Status Policy
  - Traffic Performance
  - Virtual Machine Violations
- **Violation Count** — The total number of MAPS and Fabric Watch rule violations for each category. Always displays whether or not there is a violation.
- **Network Object Count** — The number and network object type (such as switch, virtual machine, port, trunk, and so on) with a MAPS and Fabric Watch violation for each category. Always displays whether or not there is a violation.

  NOTE
  For FCIP Health, the Network Object Count is based on the number of VE_port and circuit combinations with a MAPS violation. For example, if switch A and switch B are connected through one circuit, and both switch A and switch B report a violation, the Network Object Count is 2, because the circuit on switch A is considered to be on a different network object than the circuit on switch B.

Customizing the Out of Range Violations widget

You can customize the widget to display violations for a specific fabric or group and time frame.

- To display data for a specific fabric or group, refer to "Creating a customized network scope" on page 210.
- To display data for a specific duration, refer to "Customizing the time scope" on page 213.
- Sort the contents by clicking the column header. Click the same column header again to reverse the sort order.

Accessing additional data from the widget

- Right-click any row and select **Violations** to navigate to the **Violations** dialog box.
- Double-click the **Port Health** category row (or right-click and select **Port Health Violations**) to navigate to the **Port Health Violations** widget. For more information, refer to "Port Health Violations widget" on page 231.
- Double-click the **Virtual Machine Violations** category row to navigate to the **VM Alarms** widget. For more information, refer to the "VM Alarms widget" on page 228.
- Double-click any category row, other than **Port Health** and **Virtual Machine Violations**, to navigate to the **Violations** dialog box.

Port Health Violations widget

The **Port Health Violations** widget (Figure 110) displays the number of violations for each product based on the selected fabric and a specified time range. There are four port health violation widgets: All, ISL, Initiator, and Target.

FIGURE 110 Port Health Violations widget

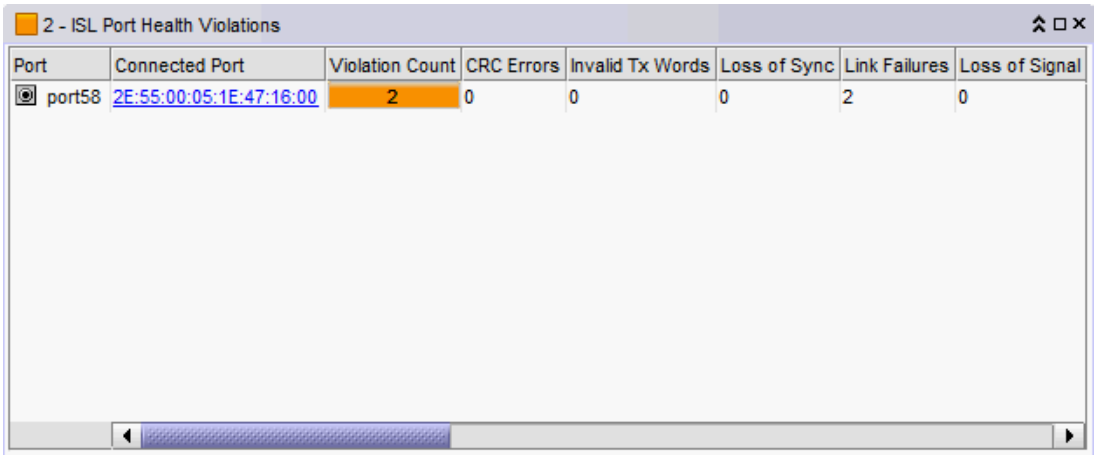

The **Port Health Violations** widget displays the following data for each product:

- Severity icon/port count/widget title — The color of the worst severity and the number of products with that severity displays before the widget title.
- **Product** — A product label such as product name, IP address, node WWN, domain ID, or zone alias.
- **Port** — A port identifier such as port name, number, address, WWN, user port number, or zone alias.

  NOTE
  All non-FC ports display either the MAC address or the port name instead of WWN.

- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Violation Count** — The number of MAPS and Fabric Watch rule violations for the port.
- **CRC Errors** — The number of times an invalid cyclic redundancy check error occurs on a port or a frame that computes to an invalid CRC.
- **Invalid Tx Words** — The number of times an invalid transmission word error occurs on a port.
- **Loss of sync** — The number of times a synchronization error occurs on the port.
- **Link Failures** — The number of times a link failure occurs on a port or sends or receives NOS.
- **Loss of Signal** — The number of times that a signal loss occurs in a port.
- **Protocol Errors** — The number of times a protocol error occurs on a port.
- **Link Reset** — The ports on which the number of link resets exceed the specified threshold value.
- **C3TXTO** — The number of Class 3 discards frames because of timeouts.
- **State changes** — The state of the port has changed for one of the following reasons:
  - The port has gone offline.
  - The port has come online.
  - The port is faulty.
- **SFP Current** — The amount of supplied current to the SFP transceiver.
- **SFP Receive Power** — The amount of incoming laser, in ?watts, to help determine if the SFP transceiver is in good working condition.
- **SFP Transmit Power** — The amount of outgoing laser, in ?watts. Use this to determine the condition of the SFP transceiver.
- **SFP Voltage** — The amount of voltage supplied to the SFP transceiver.
- **SFP Temperature** — The physical temperature of the SFP transceiver, in degrees Celsius.
- **SFP Power On Hours** — The number of hours the 16 Gbps SFP transceiver is powered on.
- **Storage Type** — The type of storage port (for example, iSCSI, NAS).
- **Attached Device (MACs)** — The MAC address of the attached device.

For SAN violations, the following categories display as blank.

- **Abnormal Frame Terminations** — The number of frames abnormally terminated.
- **Symbol Errors** — The number of undefined or invalid symbols received.
- **IFG (InterFrame Gap) Errors** — The interframe gap between successive frames that is violated.

Customizing the Port Health Violations widget

You can customize the widget to display violations for a specific fabric and time frame.

- To display data for a specific fabric or group, refer to "Creating a customized network scope" on page 210.
- To display data for a specific duration, refer to "Customizing the time scope" on page 213.
- Sort the contents by clicking the column header. Click the same column header again to reverse the sort order.

Accessing additional data from the widget

- Right-click a row in the widget to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "SAN shortcut menus" on page 1304.
- Right-click any row and select **Locate** to locate the particular device to which the port belongs in the **Network Objects** products list.
- Double-click a row to navigate to the **Violations** dialog box.

Performance monitors

The performance monitors provide a high-level overview of the performance on the network. This allows you to easily check the performance of devices, ports, and traffic on the network. The performance monitors also provide several features to help you quickly access performance metrics and reports.

The dashboards update every ten minutes regardless of the currently selected tab (SANor Dashboard) or the SAN size.

You can change the default size of the status widgets and performance monitors by placing the cursor on the divider until a double arrow displays. Click and drag the adjoining divider to resize the window.

The Top N, Bottom N, and Distribution monitors with the following measures will display data for only one discovered logical switch in the chassis:

- Memory utilization percentage
- CPU utilization percentage
- Temperature
- Fan speed
- System up time

The Management application provides the preconfigured performance monitors, refer to Table 20.

TABLE 20 Preconfigured performance monitors

Monitor title	Description	Data collectors
"Top FCoE Port Alignment Errors monitor"	Table view of the alignment errors measure	All SAN TE port collector
"Top Port PCS Block Errors monitor"	Table view of the PCS Block Error. Monitors based on the port types: All ports, Initiator Ports, ISL Ports, and Target Ports.	All SAN FC port collector
"Top Port C3 Discards monitor"	Table view of the C3 discards measure	All SAN FC port collector
"Top Port C3 Discards RX TO monitor"	Table view of the C3 discards RX TO measure. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.	All SAN FC port collector
"Top Port CRC Errors monitor"	Table view of the CRC errors measure. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.	All SAN FC port collector, All SAN TE port collector
"Top Port Discards monitor"	Table view of the discards measure	Port discard count collector
"Top Port Encode Error Out monitor"	Table view of the encode error out measure. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.	All SAN FC port collector

TABLE 20 Preconfigured performance monitors (Continued)

Monitor title	Description	Data collectors
"Top Port Link Failures monitor"	Table view of the top port link failures. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.	All SAN FC port collector
"Top Port Link Resets monitor"	Table view of the top port link resets. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.	All SAN FC port collector
Top Port Overflow Errors	Table view of the overflow errors measure	All SAN TE port collector
Top Port Receive EOF	Table view of the received end-of-frames measure	All SAN TE port collector
Top Port Runtime Errors	Table view of the runtime errors measure	All SAN TE port collector
Top IP Server/IP ISL/IP Storage Port Runtime Errors	Table view of the IP server or IP ISL or IP storage runtime errors measure	All IP
Top Port Sync Losses	Table view of the top port synchronization losses. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.	All SAN FC port collector
Top Port Too Long Errors	Table view of the too long errors measure	All SAN TE port collector
Top Port Traffic	Table view of the traffic measure	All SAN FCIP tunnel collector, All SAN FC port collector, port throughput collector, All SAN TE port collector
Top Port Underflow Errors	Table view of the underflow errors measure	All SAN TE port collector
Top Port Utilization Percentage	Table view of the port utilization percentage measure. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.	All SAN FCIP tunnel collector, All SAN FC port collector, port utilization collector, All SAN TE port collector
Bottom Port Utilization Percentage	Table view of the port utilization percentage measure. There are four versions of this monitor based on the type of port: All ports, initiator ports, ISL ports, and Target ports.	All SAN FCIP tunnel collector, All SAN FC port collector, port utilization collector, All SAN TE port collector
Top Product CPU Utilization	Table view of the CPU utilization percentage measure	All SAN products collector
Top Product Memory Utilization	Table view of the memory utilization percentage measure	All SAN products collector
Top Product Response Time	Table view of the response time measure	All SAN products collector
Top Product Temperature	Table view of the temperature measure	All SAN products collectorSystem temperature collector
Top Products with Unused Ports	Table view of the products with unused ports measure	All SAN Product collector, Ports Not in Use Collector
Invalid Transmission Widget	Table view of the number of invalid transmissions	All SAN FC Port collector

The preconfigured performance monitors can be turned off, hidden, and edited; however, you cannot delete the preconfigured monitors.

You can also create new performance monitors to display on the dashboard. For more information, refer to "User-defined performance monitors" on page 268.

Displaying performance monitors on the dashboard

1. From the **Dashboards** expand navigation bar, double-click the desired dashboard.

   The selected **Dashboard** displays.

2. Click the **Customize Dashboard** icon.

   The **Customize Dashboard** dialog box displays.

3. Select the check box in the **Display** column for each performance monitor you want to display on the **Dashboard**.
4. Click **OK**.

Top FCoE Port Alignment Errors monitor

The **Top FCoE Port Alignment Errors** performance monitor displays the top ports with alignment errors in a table.

The **Top FCoE Port Alignment Errors** performance monitor includes the following data:

- Threshold icon/object count/monitor title — The color associated with the threshold and number of objects within that threshold displays next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Alignment Errors** — The number (error count) of alignment errors for the duration specified in the monitor.
- **Alignment Errors/sec** — The number (error rate) of alignment errors per second for the duration specified in the monitor.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To edit a port performance monitor, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from top or bottom port monitors

Double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Historical Graphs/Tables** dialog box for the selected measures. For more information, refer to "Performance Data" on page 959.

Top Port C3 Discards monitor

The **Top Port C3 Discards** monitor (Figure 111) displays the top ports with Class 3 frames discarded in a table. There are four port widgets: All, ISL, Initiator, and Target.

FIGURE 111 Top Port C3 Discards monitor

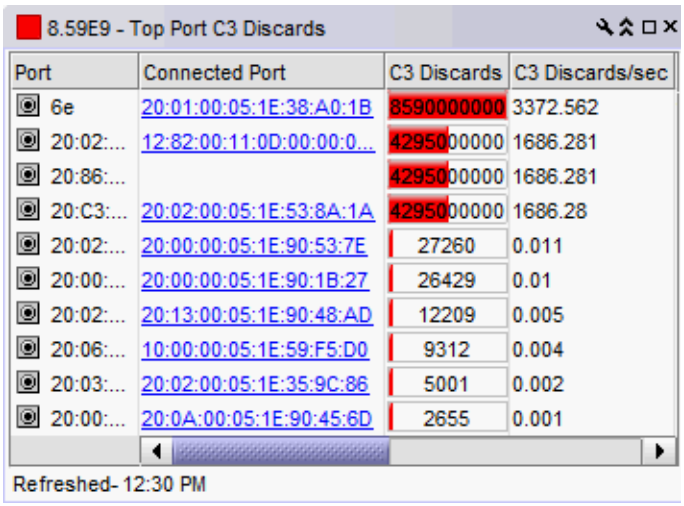

The **Top Port C3 Discards** monitor includes the following data:

- **Count/monitor title** — The count for the data based on the error count shown next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **C3 Discards/sec** — The number (error rate) of Class 3 discard errors per second for the duration specified in the monitor.
- **C3 Discards** — The number (error count) of Class 3 discard errors for the duration specified in the monitor.
- **C3 Discards TX TO** — The number of transmitted Class 3 frames discarded due to timeout.
- **C3 Discards RX TO** — The number of received Class 3 frames discarded due to timeout.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Port C3 Discards monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Port C3 Discards RX TO monitor

The **Top Port C3 Discards RX TO** monitor displays the top ports with receive Class 3 frames received at this port and discarded at the transmission port due to timeout in a table.

FIGURE 112 Top Port C3 Discards RX TO monitor

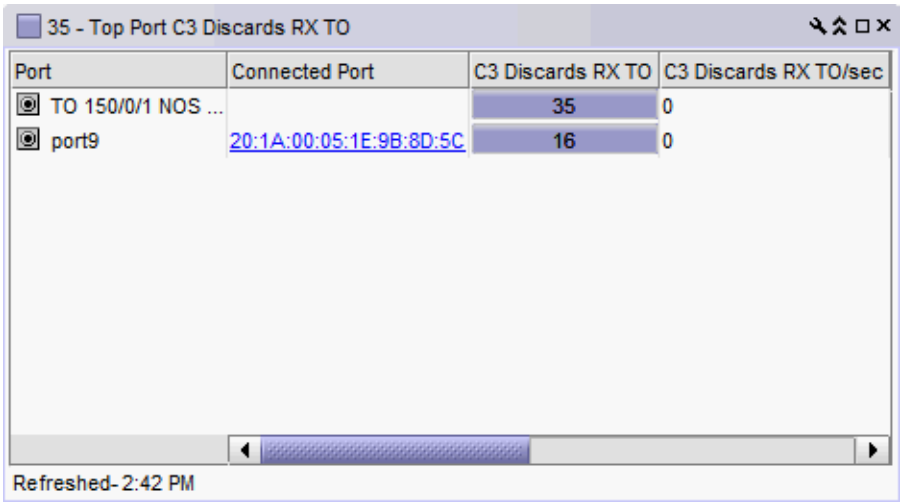

The **Top Port C3 Discards RX TO** monitor includes the following data:

- **Count/monitor title** — The count for the data based on the error count shown next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **C3 Discards RX TO/sec** — The number (error rate) of Class 3 frames received at this port and discarded at the transmission port due to timeout errors per second for the duration specified in the monitor.
- **C3 Discards RX TO** — The number (error count) of Class 3 frames received at this port and discarded at the transmission port due to timeout errors for the duration specified in the monitor.
- **C3 Discards** — The number of Class 3 frames discarded.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.

- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Port C3 Discards RX TO monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "SAN shortcut menus" on page 1304.
- In a Top N or Bottom N C3 Discards TX TO and C3 Discards RX TO monitors, right-click an FC-port row and select **Discarded Frames** to navigate to the **Discarded Frames** dialog box. For more information, refer to "Viewing discarded frames from a port" on page 441.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Port CRC Errors monitor

The **TopPort CRC Errors** monitor (Figure 113)contain cyclic redundancy check (CRC) errors in a table.

FIGURE 113 Top Port CRC Errors monitor

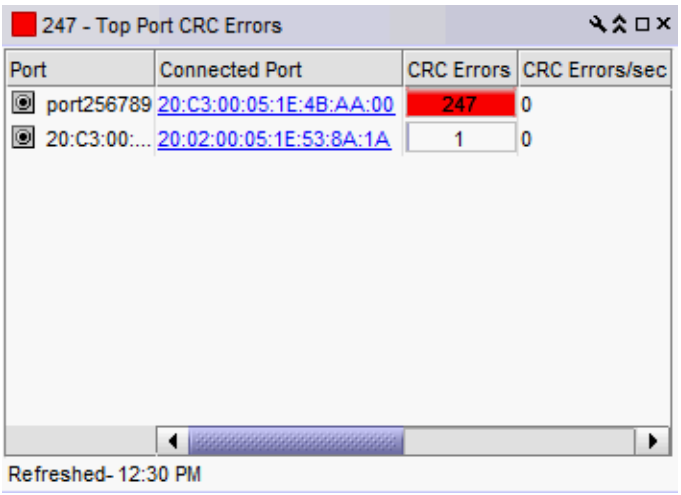

The **Top Port CRC Errors**monitor includes the following data:

- **Count/monitor title** — The count for the data based on the error count shown next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **CRC Errors/sec** — The number (error rate) of cyclic redundancy check (CRC) errors per second for the duration specified in the monitor.

- **CRC Errors** — The number (error count) of cyclic redundancy check (CRC) errors for the duration specified in the monitor.
- **Link Failures** — The number of link failures.
- **Sequence Errors** — The number of sequence errors.
- **Invalid Transmissions** — The number of invalid transmissions.
- **C3 Discards Tx TO** — The number of transmitted Class 3 frames discarded due to timeout.
- **PCS Block Errors** — The number of PCS block errors.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Port CRC Errors monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Port Encode Error Out monitor

The **Top Port Encode Error Out** monitor (Figure 114) displays the top ports with encoding errors outside of frames in a table.

FIGURE 114 Top Port Encode Error Out monitor

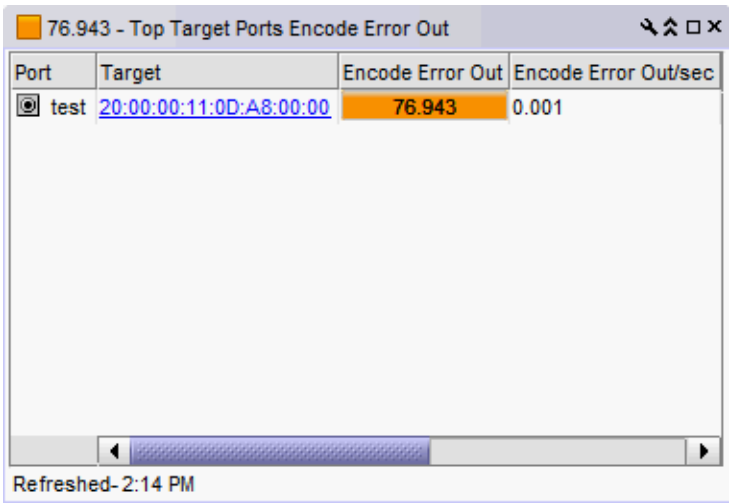

The **Top Port Encode Error Out** monitor includes the following data:

- **Count/monitor title** — The count for the data based on the error count shown next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Encode Error Out/sec** — The number (error rate) of encoding errors outside of frames per second for the duration specified in the monitor.
- **Encode Error Out** — The number (error count) of encoding errors outside of frames for the duration specified in the monitor.
- **CRC Errors** — The number of CRC errors.
- **Link Failures** — The number of link failures.
- **Invalid Transmissions** — The number of invalid transmissions.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Storage Type — The type of storage port (for example, iSCSI, NAS).**
- **Attached Device (MACs) — The MAC address of the attached device.**
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Port Encode Out Errors monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Port PCS Block Errors monitor

The Top Port PCS Block Errors monitor displays the top top ports with Physical Coding Sublayer (PCS) block errors outside of frames in a table.

NOTE
PCS block errors are only applicable on 10 and 16 Gbps ports.

The Top Port PCS Block Errors monitor includes the following data:

- Count/monitor title — The count for the data based on the error count shown next to the monitor title.
- Port — The port affected by this monitor.
- Connected_Port_Link

- (where Connected_Port_Link is Connected Port, Initiator, or Target) — Displays one of the following:
  - Connected Port — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - Initiator — The initiator port on the connected device. Click to launch the device properties dialog box.
  - Target — The target port on the connected device. Click to launch the device properties dialog box.
- PCS Block Errors/sec — The number (error rate) of PCS block errors outside of frames per second for the duration specified in the monitor.
- PCS Block Errors — The number (error count) of PCS block errors outside of frames for the duration specified in the monitor.
- Link Failures – The number of link failures.
- Invalid Transmissions – The number of invalid transmissions.
- Product — The product affected by this monitor.
- Type — The type of port (for example, U-Port).
- Identifier — The port identifier.
- Port Number — The port number.
- State — The port state (for example, Enabled).
- Status — The port status (for example, Up).
- Refreshed — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Top Port Link Failures monitor

The **Top Port Link Failures** monitor (Figure 115) displays the top ports with link failures in a table.

FIGURE 115 Top Port Link Failures monitor

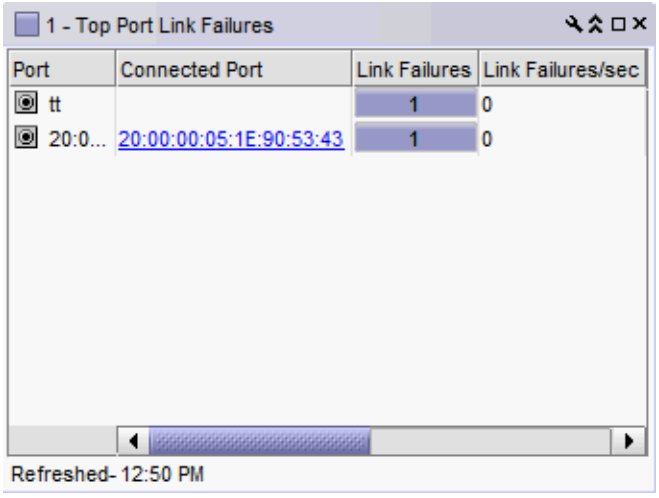

The **Top Port Link Failures** monitor includes the following data:

- **Count/monitor title** — The count for the data based on the error count shown next to the monitor title.
- **Port** — The port affected by this monitor.

- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Link Failures/sec** — The number (error rate) link failure errors per second for the duration specified in the monitor.
- **Link Failures** — The number (error count) of link failure errors.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Port Link Failures monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the SAN **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Port Link Resets monitor

The **Top Port Link Resets** monitor (Figure 116) displays the top ports with link resets in a table.

FIGURE 116 Top Port Link Resets monitor

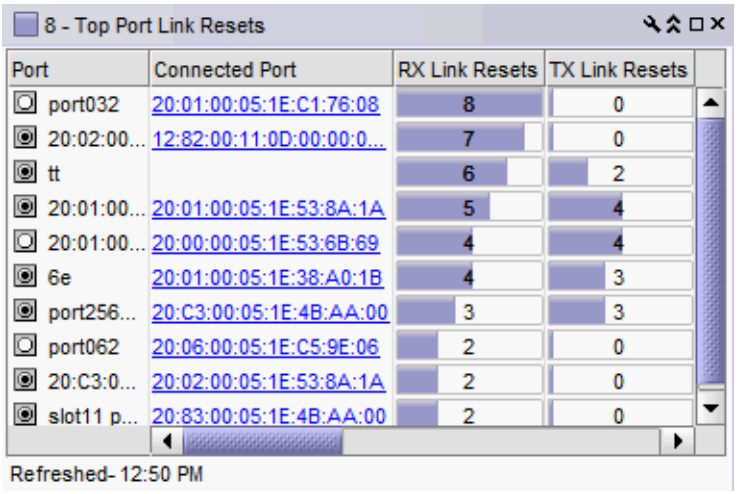

The **Top Port Link Resets** monitor includes the following data:

- **Count/monitor title** — The count for the data based on the error count shown next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **RX Link Resets /sec** — The number (error rate) receive link reset errors per second for the duration specified in the monitor.
- **RX Link Resets** — The number (error count) of receive link reset errors.
- **TX Link Resets/sec** — The number (error rate) of transmit link reset errors for the duration specified in the monitor.
- **TX Link Resets** — The number (error count) of transmit link reset errors.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Port Link Resets monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the SAN **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top FCoE Port Overflow Errors monitor

The **Top FCoE Port Overflow Errors** performance monitor (Figure 117) displays the top ports with overflow errors in a table.

FIGURE 117 Top FCoE Port Overflow Errors performance monitor

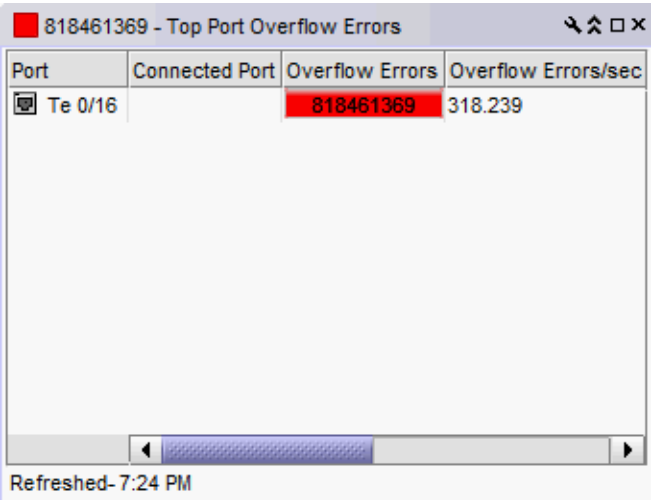

The Top **FCoE** Port Overflow Errors performance monitor includes the following data:

- Threshold icon/object count/monitor title — The color associated with the threshold and number of objects within that threshold displays next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Overflow Errors**— The number (error count) of overflow errors for the duration specified in the monitor.
- **Overflow Errors/sec** — The number (error rate) of overflow errors per second for the duration specified in the monitor.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To edit a port performance monitor, refer to "Editing a preconfigured performance monitor" on page 267.

Top FCoE Port Receive EOF monitor

The **Top FCoE Port Receive EOF** performance monitor displays the top ports with received end-of-frames in a table.

The **Top FCoE Port Receive EOF** performance monitor includes the following data:

- Threshold icon/object count/monitor title — The color associated with the threshold and number of objects within that threshold displays next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Receive EOF** — The number (count) of end of frames received.
- **Receive EOF/sec** — The number (rate) of end of frames received per second for the duration specified in the monitor.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To edit a port performance monitor, refer to "Editing a preconfigured performance monitor" on page 267.

Top FCoE Port Runtime Errors monitor

The **Top FCoE Port Runtime Errors** performance monitor displays the top ports with runtime errors in a table.

The **Top FCoE Port Runtime Errors** performance monitor includes the following data:

- Threshold icon/object count/monitor title — The color associated with the threshold and number of objects within that threshold displays next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Runtime Errors**— The number (error count) of runtime errors for the duration specified in the monitor.
- **Runtime Errors/sec** — The number (error rate) of runtime errors per second for the duration specified in the monitor.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).

- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To edit a port performance monitor, refer to "Editing a preconfigured performance monitor" on page 267.

Top Port Sync Losses monitor

The **Top Port Sync Losses** monitor (Figure 118) displays the top ports with synchronization failures in a table.

FIGURE 118 Top Port Sync Losses monitor

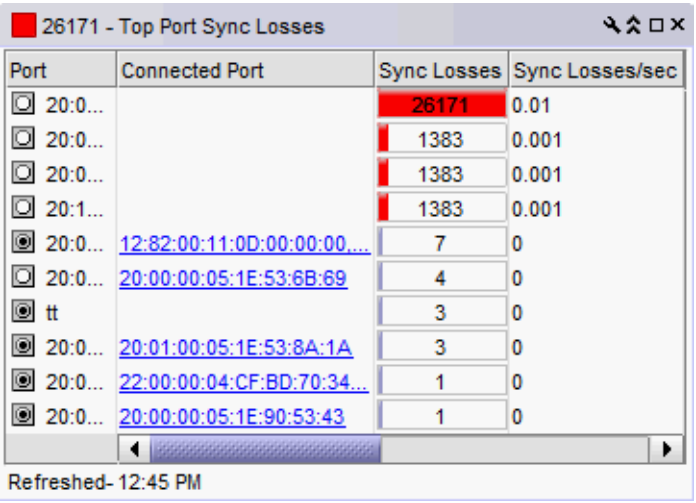

The **Top Port Sync Losses** monitor includes the following data:

- Severity icon/monitor title — The color of the worst severity of the data shown next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Sync Losses —** The number of synchronization failures for the port.
- **Sync Losses/sec —** The number of synchronization failures for the port per second.
- **CRC Errors** — The number of CRC errors.
- **Link Failures** — The number of link failures.
- **Invalid Transmissions** — The number of invalid transmissions.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Storage Type** — The type of storage port (for example, iSCSI, NAS).
- **Attached Device (MACs)** — The MAC address of the attached device.
- **Identifier** — The port identifier.

- **Port Number** — The port number.
- **State** — The port state (for example, Online).
- **Status** — The port status (for example, In_Sync, No_Sync).
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Port Link Resets monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Custom: Historical Performance Graphs** dialog box. For more information, refer to "Performance Data" on page 959.

Top FCoE Port Too Long Errors monitor

The **Top FCoE Port Too Long Errors** performance monitor displays the top ports with frames longer than the maximum frame size allowed errors in a table.

The **Top FCoE Port Too Long Errors** performance monitor includes the following data:

- Threshold icon/object count/monitor title — The color associated with the threshold and number of objects within that threshold displays next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Too Long Errors** — The number (error count) of frames longer than the maximum frame size allowed errors for the duration specified in the monitor.
- **Too Long Errors/sec** — The number (error rate) of frames longer than the maximum frame size allowed errors per second for the duration specified in the monitor.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To edit a port performance monitor, refer to "Editing a preconfigured performance monitor" on page 267.

Top Port Traffic monitor

The **Top Port Traffic** monitor (Figure 119) displays the top ports with receive and transmit traffic in a table.

FIGURE 119 Top Port Traffic monitor

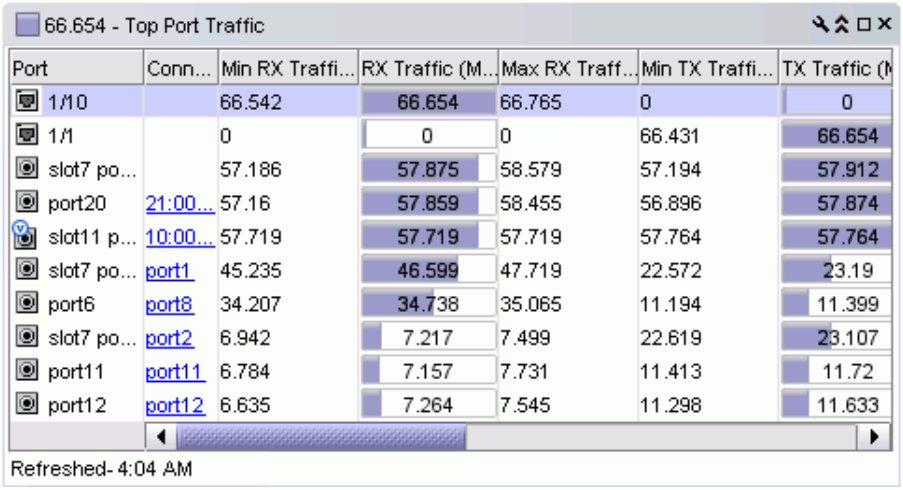

The **Top Port Traffic** monitor includes the following data:

- Severity icon/monitor title — Displays the worst severity of the data shown next to the monitor title.

  NOTE
  The **Top Port Traffic** widget displays the threshold colors based on the port speed. Click edit icon of the widget to customize the threshold values.

- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Min RX Traffic (MB/s)** — The minimum receive traffic in megabits per second.
- **RX Traffic (MB/s)** — The top receive traffic in megabits per second.
- **Max Traffic (MB/s)** — The maximum receive traffic in megabits per second.
- **Min TX Traffic (MB/s)** — The minimum transmit traffic in megabits per second.
- **TX Traffic (MB/s)** — The top transmit traffic in megabits per second.
- **Max TX Traffic (MB/s)** — The maximum transmit traffic in megabits per second.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).

- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Port Traffic monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top FCoE Port Underflow Errors monitor

The **Top FCoE Port Underflow Errors** performance monitor displays the top ports with underflow errors in a table.

The **Top FCoE Port Underflow Errors** performance monitor includes the following data:

- Threshold icon/object count/monitor title — The color associated with the threshold and number of objects within that threshold displays next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Underflow Errors**— The number (error count) of underflow errors for the duration specified in the monitor.
- **Underflow Errors/sec** — The number (error rate) of underflow errors per second for the duration specified in the monitor.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To edit a port performance monitor, refer to "Editing a preconfigured performance monitor" on page 267.

Top Port Utilization Percentage monitor

The **Top Port Utilization Percentage** monitor (Figure 120) displays the top port utilization percentages in a table.

FIGURE 120 Top Port Utilization Percentage monitor

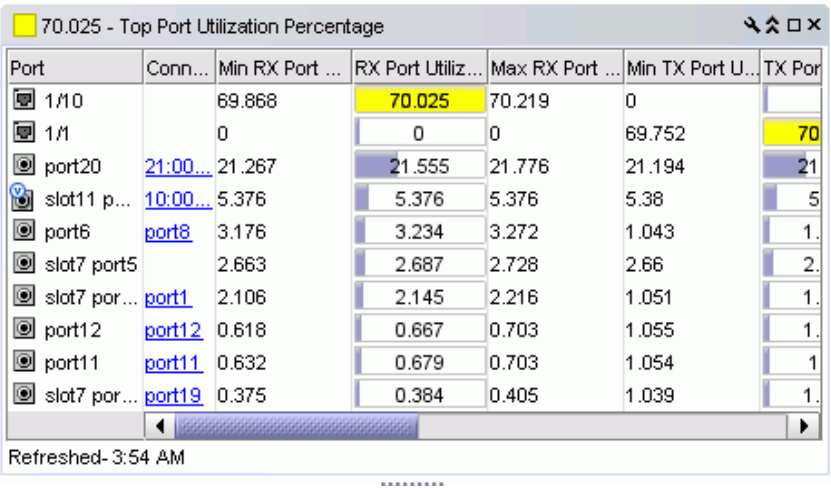

The **Top Port Utilization Percentage** monitor includes the following data:

- **Count/monitor title** — The count for the data shown next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Min RX Port Utilization Percentage** — The minimum receive port utilization percentages.
- **RX Port Utilization Percentage** — The top receive port utilization percentages.
- **Max RX Port Utilization Percentage** — The maximum receive port utilization percentages.
- **Min TX Port Utilization Percentage** — The minimum transmit port utilization percentages.
- **TX Port Utilization Percentage** — The top transmit port utilization percentages.
- **Max TX Port Utilization Percentage** — The maximum transmit port utilization percentages.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Port Utilization Percentage monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Bottom Port Utilization Percentage monitor

The **Bottom Port Utilization Percentage** monitor (Figure 121) displays the bottom port utilization percentages in a table.

FIGURE 121 Bottom Port Utilization Percentage monitor

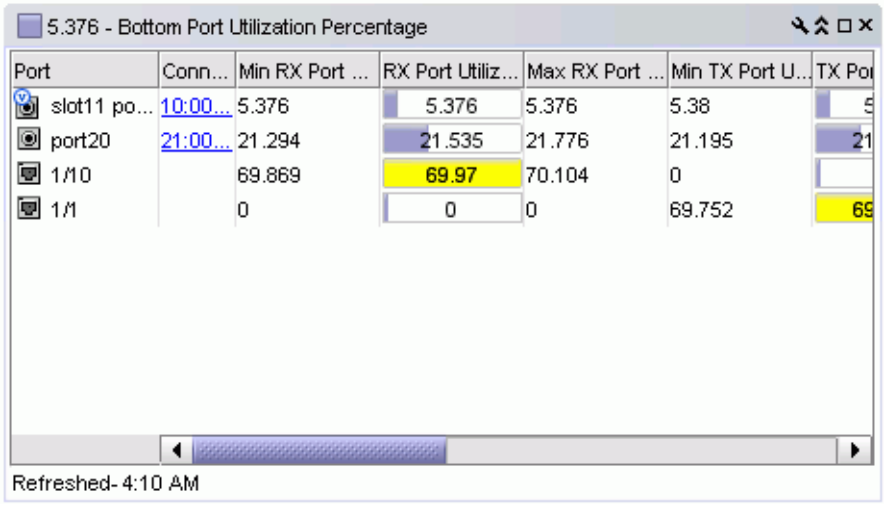

The **Bottom Port Utilization Percentage** monitor includes the following data:

- **Count/monitor title** — The count for the data shown next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- **Min RX Port Utilization Percentage** — The minimum receive port utilization percentages.
- **RX Port Utilization Percentage** — The bottom receive port utilization percentages.
- **Max RX Port Utilization Percentage** — The maximum receive port utilization percentages.
- **Min TX Port Utilization Percentage** — The minimum transmit port utilization percentages.
- **TX Port Utilization Percentage** — The bottom transmit port utilization percentages.
- **Max TX Port Utilization Percentage** — The maximum transmit port utilization percentages.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.

- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Bottom Port Utilization monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Product CPU Utilization monitor

The **Top Product CPU Utilization** monitor (Figure 122) displays the top product CPU utilization percentages in a table.

FIGURE 122 Top Product CPU Utilization monitor

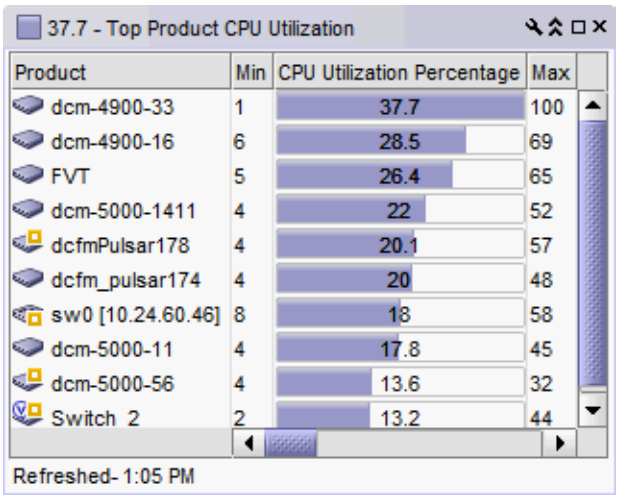

The **Top Product CPU Utilization** monitor includes the following data:

- **Count/monitor title** — The count for the data shown next to the monitor title.
- **Product** — The product affected by this monitor.
- **Min** — The minimum value of the measure in the specified time range.
- **CPU Utilization Percentage** — The CPU utilization percentages.
- **Max** — The maximum value of the measure in the specified time range.
- **Fabric** — The fabric to which the device belongs.
- **Product Type** — The type of product (for example, switch).
- **State** — The product state (for example, Offline).
- **Status** — The product status (for example, Reachable).

- **Tag** — The product tag.
- **Serial #** — The serial number of the product.
- **Model** — The product model.
- **Port Count** — The number of ports on the product.
- **Firmware** — The firmware level running on the product.
- **Location** — The location of the product.
- **Contact** — A contact name for the product.
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Product CPU Utilization monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Product Memory Utilization monitor

The **Top Product Memory Utilization** monitor (Figure 123) displays the top product memory utilization percentages in a table.

FIGURE 123 Top Product Memory Utilization monitor

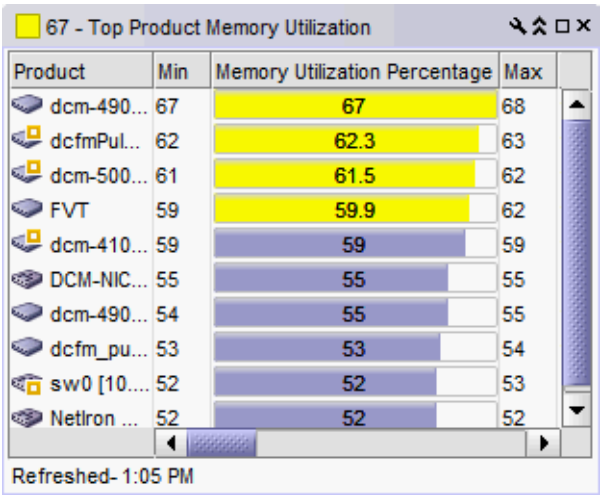

The **Top Product Memory Utilization** monitor includes the following data:

- **Count/monitor title** — The count for the data shown next to the monitor title.
- **Product** — The product affected by this monitor.
- **Min** — The minimum value of the measure in the specified time range.
- **Memory Utilization Percentage** — The top memory utilization percentages.
- **Max** — The maximum value of the measure in the specified time range.

- **Fabric** — The fabric to which the device belongs.
- **Product Type** — The type of product (for example, switch).
- **State** — The product state (for example, Offline).
- **Status** — The product status (for example, Reachable).
- **Tag** — The product tag.
- **Serial #** — The serial number of the product.
- **Model** — The product model.
- **Port Count** — The number of ports on the product.
- **Firmware** — The firmware level running on the product.
- **Location** — The location of the product.
- **Contact** — A contact name for the product.
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Product Memory Utilization monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Product Response Time monitor

The **Top Product Response Time** monitor (Figure 124) displays the top product response time in a table.

FIGURE 124 Top Product Response Time monitor

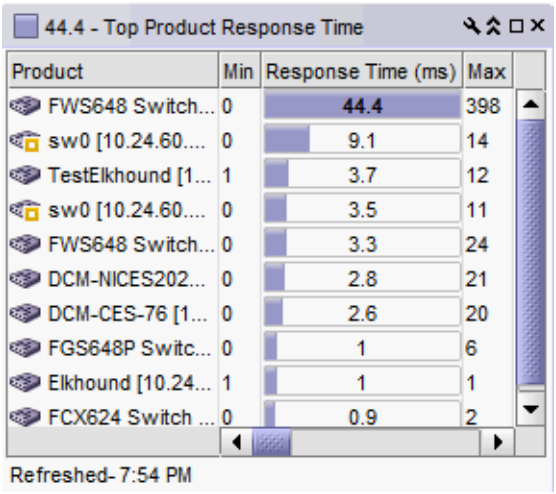

The **Top Product Response Time** monitor includes the following data:

- Severity icon/response time/monitor title — The worst severity of the data and the response time displays next to the monitor title.
- **Product** — The product affected by this monitor.
- **Min** — The minimum value of the measure in the specified time range.
- **Response Time (ms) —** The top response time in milliseconds.
- **Max** — The maximum value of the measure in the specified time range.
- **Fabric** — The fabric to which the device belongs.
- **Product Type** — The type of product (for example, switch).
- **State** — The product state (for example, Offline).
- **Status** — The product status (for example, Reachable).
- **Tag** — The product tag.
- **Serial # —** The serial number of the product.
- **Model** — The product model.
- **Port Count —** The number of ports on the product.
- **Firmware —** The firmware level running on the product.
- **Location —** The location of the product.
- **Contact —** A contact name for the product.
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Product Response Time monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Product Temperature monitor

The **Top Product Temperature** monitor (Figure 125) displays the top product temperature in a table.

FIGURE 125 Top Product Temperature monitor

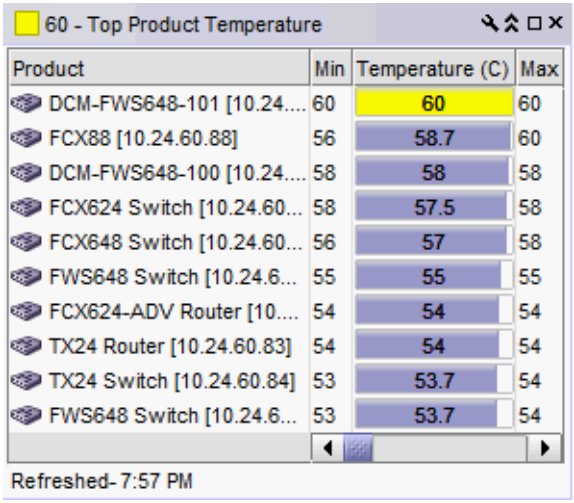

The **Top Product Temperature** monitor includes the following data:

- **Severity icon/temperature/monitor title** — The worst severity of the data and the temperature displays next to the monitor title.
- **Product** — The product affected by this monitor.
- **Min** — The minimum value of the measure in the specified time range.
- **Temperature** — The top temperatures.
- **Max** — The maximum value of the measure in the specified time range.
- **Fabric** — The fabric to which the device belongs.
- **Product Type** — The type of product (for example, switch).
- **State** — The product state (for example, Offline).
- **Status** — The product status (for example, Reachable).
- **Tag** — The product tag.
- **Serial #** — The serial number of the product.
- **Model** — The product model.
- **Port Count** — The number of ports on the product.
- **Firmware** — The firmware level running on the product.
- **Location** — The location of the product.
- **Contact** — A contact name for the product.
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Product Temperature monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Top Products with Unused Ports monitor

The **Top Products with Unused Ports** monitor displays the top products with ports not in use in a table.

The **Top Products with Unused Ports** monitor includes the following data:

- **Count/monitor title** — The count for the data shown next to the monitor title.
- **Product** — The product affected by this monitor.
- **Min** — The minimum value of the measure in the specified time range.
- **Ports Not In Use** — The number of ports not in use for the product.
- **Max** — The maximum value of the measure in the specified time range.
- **Fabric** — The fabric to which the device belongs.
- **Product Type** — The type of product (for example, switch).
- **State** — The product state (for example, Offline).
- **Status** — The product status (for example, Reachable).
- **Tag** — The product tag.
- **Serial # —** The serial number of the product.
- **Model** — The product model.
- **Port Count —** The number of ports on the product.
- **Firmware —** The firmware level running on the product.
- **Location —** The location of the product.
- **Contact —** A contact name for the product.
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Top Product with Unused Ports monitor

- Right-click a row in the monitor to access the shortcut menu available for the associated device. For more information about shortcut menus, refer to "Application Menus" on page 1295.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to "Performance Data" on page 959.

Port Traffic Distribution monitor

The **Port Traffic Distribution** monitor (Figure 119) displays the top ports with receive and transmit traffic in a table.

FIGURE 126 Port Traffic Distribution monitor

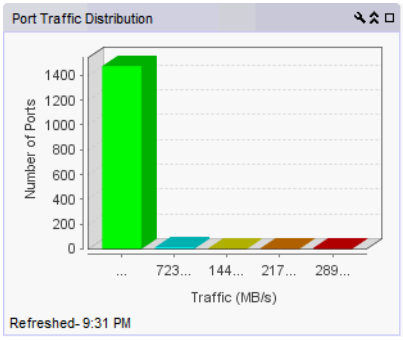

The **Port Traffic Distribution** monitor includes the following data:

- Monitor title — The monitor title.
- Y-axis — The y-axis displays the number of ports.
- X-axis — The x-axis displays the traffic distribution in MB/s. Pause on a bar in the graph to view the tooltop including the traffic range and the number of the ports in that range. Click a bar on the graph to launch the **Port Traffic Distribution Data Details.**
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Port Traffic Distribution monitor

1. Click a bar on the graph to launch the **Top Port Traffic Distribution Data Details**.
    - **Port** — The port affected by this monitor.
    - *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
        - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
        - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
        - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
    - **RX/TX** — Whether it is receive or transmit traffic.
    - **Traffic (MB/s)** — The traffic in megabits per second.
    - **Storage Type** — The storage type.
    - **Attached Device (MACs)** — The MAC address of the attached device.
    - **Product** — The product affected by this monitor.
    - **Type** — The type of port (for example, U-Port).
    - **Identifier** — The port identifier.
    - **Port Number** — The port number.
    - **State** — The port state (for example, Enabled).
    - **Status** — The port status (for example, Up).

2. Click **Close**.

Port Utilization Distribution monitor

The **Port Utilization Distribution** monitor (Figure 119) displays the top ports with receive and transmit traffic in a table.

FIGURE 127 Port Utilization Distribution monitor

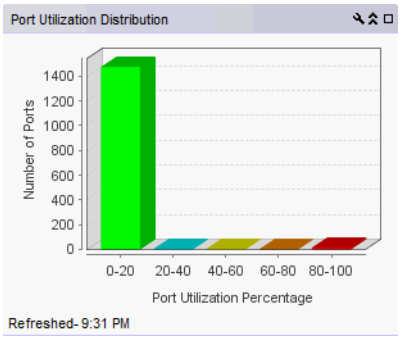

The **Port Utilization Distribution** monitor includes the following data:

- Monitor title — The monitor title.
- Y-axis — The y-axis displays the number of ports.
- X-axis — The x-axis displays the traffic utilization percentage range. Pause on a bar in the graph to view the tooltop including the port utilization percentage range and the number of the ports in that range. Click a bar on the graph to launch the **Port Utilization Distribution Data Details.**
- **Refreshed** — The time of the last update for the monitor.

To customize the monitor to display data by a selected time frame as well as customize the display options, refer to "Editing a preconfigured performance monitor" on page 267.

Accessing additional data from the Port Utilization Distribution monitor

1. Click a bar on the graph to launch the **Port Utilization Distribution Data Details**.
    - **Port** — The port affected by this monitor.
    - *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
        - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
        - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
        - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
    - **RX/TX** — Whether it is receive or transmit traffic.
    - **Port Utilization Percentage** — The port utilization percentage.
    - **Storage Type** — The storage type.
    - **Attached Device (MACs)** — The MAC address of the attached device.
    - **Product** — The product affected by this monitor.
    - **Type** — The type of port (for example, U-Port).
    - **Identifier** — The port identifier.

- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).

2. Click **Close**.

Top Tunnel Utilization monitor

The Top Tunnel Utilization monitor displays the top tunnel utilization data, which is transmitted and received data at the ends of the tunnel, in a table.

FIGURE 128 Top Tunnel Utilization monitor

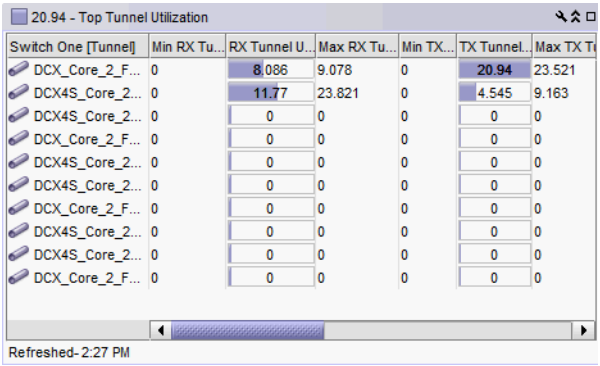

The Top Tunnel Utilization monitor includes the following data:

- Count/monitor title — The count for the data based on the utilization shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- Min RX Tunnel Utilization — The minimum receive traffic utilization percentage.
- RX Tunnel Utilization — The receive traffic utilization percentage.
- Max RX Tunnel Utilization — The maximum receive traffic utilization percentage.
- Min TX Tunnel Utilization — The minimum transmit traffic utilization percentage.
- TX Tunnel Utilization — The transmit traffic utilization percentage.
- Max TX Tunnel Utilization — The maximum transmit traffic utilization percentage.
- Switch Two — The destination switch IP address and FID of the FCIP tunnel.
- Total Circuits — The total number of circuits assigned to the tunnel.
- Operational Status — The operational status of the tunnel.
- Administrative Status — The administrative status of a tunnel (enabled or disabled).
- Description — A free text description of the tunnel.

Top Tunnel Dropped Packets monitor

The Top Tunnel Dropped Packets monitor displays packet loss across the FCIP tunnel per second.

FIGURE 129 Top Tunnel Dropped Packets monitor

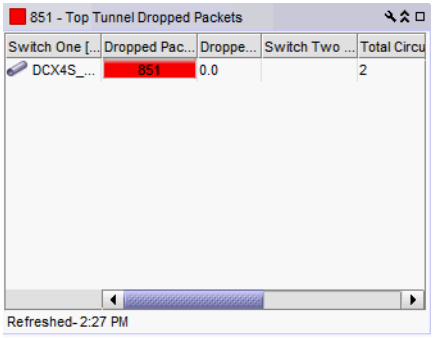

The Top Tunnel Dropped Packets monitor includes the following data:

- Count/monitor title — The count for the data based on the dropped packet count shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- Dropped Packets — The number of packets dropped across the FCIP tunnel.
- Dropped Packets/sec — The number of packets dropped across the FCIP tunnel per second.
- Switch Two [Tunnel] — The IP address and FID of the destination switch for the FCIP tunnel.
- Total Circuits — The total number of circuits assigned to the tunnel.
- Operational Status — The operational status of the tunnel.
- Administrative Status — The administrative status of a tunnel (enabled or disabled).
- Description — A free text description of the tunnel.

Top Circuit Utilization monitor

The Top Circuit Utilization monitor displays the overall circuit utilization for both FC and IP Extension traffic in a table.

FIGURE 130 Top Circuit Utilization monitor

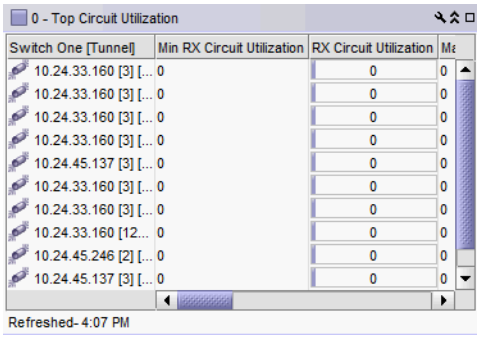

The Top Circuit Utilization monitor includes the following data:

- Count/monitor title — The count for the data based on the utilization shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- Min RX Circuit Utilization — The minimum receive traffic utilization percentage.
- RX Circuit Utilization — The receive traffic utilization percentage.
- Max RX Circuit Utilization — The maximum receive traffic utilization percentage.
- Min TX Circuit Utilization — The minimum transmit traffic utilization percentage.
- TX Circuit Utilization — The transmit traffic utilization percentage.
- Max TX Circuit Utilization — The maximum transmit traffic utilization percentage.
- Switch Two — The IP address and FID of the destination switch for the FCIP tunnel.
- Circuit ID — The identifier of the circuit.
- Circuit Status — The operational status of the circuit.
- Circuit Admin Status — The administrative status of a circuit (enabled or disabled).

Accessing additional data from the Top Circuit Utilization monitor

Double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Custom: Historical Performance Graph** dialog box for the selected measures. For more information, refer to "Generating and saving a historical performance graph" on page 972.

Top Circuit FC Utilization monitor

The Top Circuit FC Utilization monitor displays the top circuit utilization for FC traffic in a table.

FIGURE 131 Top Circuit FC Utilization monitor

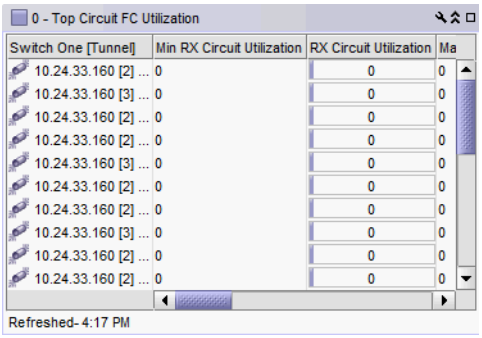

The Top Circuit FC Utilization monitor includes the following data:

- Count/monitor title - The count for the data based on the utilization shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- Min RX Circuit Utilization — The minimum FC receive traffic utilization percentage.
- RX Circuit Utilization — The FC receive traffic utilization percentage.
- Max RX Circuit Utilization — The maximum FC receive traffic utilization percentage.
- Min TX Circuit Utilization — The minimum FC transmit traffic utilization percentage.
- TX Circuit Utilization — The FC transmit traffic utilization percentage.

- Max TX Circuit Utilization — The maximum FC transmit traffic utilization percentage.
- Switch Two — The IP address and FID of the destination switch for the FCIP tunnel.
- Circuit ID — The identifier of the circuit.
- Circuit Status — The operational status of the circuit.
- Circuit Admin Status — The administrative status of a circuit (enabled or disabled).

Accessing additional data from the Top Circuit FC Utilization monitor

Double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Custom: Historical Performance Graph** dialog box for the selected measures. For more information, refer to "Generating and saving a historical performance graph" on page 972.

Top Circuit IP Extension Utilization monitor

The Top Circuit IP Extension Utilization monitor displays the top circuit utilization for IP Extension traffic in a table.

FIGURE 132 Top Circuit IP Extension Utilization monitor

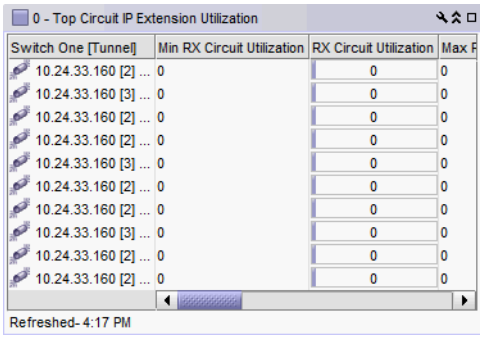

The Top Circuit IP Extension Utilization monitor includes the following data:

- Count/monitor title — The count for the data based on the utilization shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- Min RX Circuit Utilization — The minimum IP Extension receive traffic utilization percentage.
- RX Circuit Utilization — The IP Extension receive traffic utilization percentage.
- Max RX Circuit Utilization — The maximum IP Extension receive traffic utilization percentage.
- Min TX Circuit Utilization — The minimum IP Extension transmit traffic utilization percentage.
- TX Circuit Utilization — The IP Extension transmit traffic utilization percentage.
- Max TX Circuit Utilization — The maximum IP Extension traffic utilization percentage.
- Switch Two — The IP address and FID of the destination switch for the FCIP tunnel.
- Circuit ID — The identifier of the circuit.
- Circuit Status — The operational status of the circuit.
- Circuit Admin Status — The administrative status of a circuit (enabled or disabled).

Accessing additional data from the Top Circuit IP Extension Utilization monitor

Double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Custom: Historical Performance Graph** dialog box for the selected measures. For more information, refer to "Generating and saving a historical performance graph" on page 972.

Top Circuit Jitter monitor

The Top Circuit Jitter monitor displays the top circuit jitter variance, in milliseconds, in a table.

FIGURE 133 Top Circuit Jitter monitor

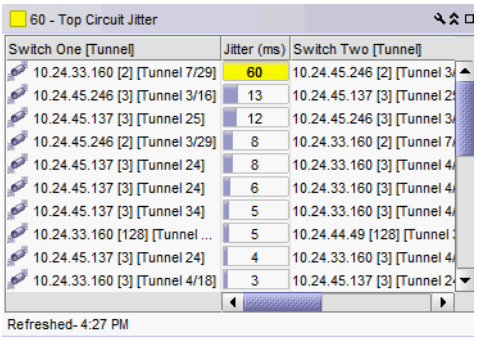

The Top Circuit Jitter monitor includes the following data:

- Count/monitor title — The count for the data based on the circuit jitter shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- Jitter (ms) — The variance, in milliseconds, in round-trip time of the circuit.
- Switch Two — The IP address and FID of the destination switch for the FCIP tunnel.
- Circuit ID — The identifier of the circuit.
- Circuit Status — The operational status of the circuit.
- Circuit Admin Status — The administrative status of a circuit (enabled or disabled).

Accessing additional data from the Top Circuit Jitter monitor

Double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Custom: Historical Performance Graph** dialog box for the selected measures. For more information, refer to "Generating and saving a historical performance graph" on page 972.

Top Circuit RTT monitor

The Top Circuit RTT monitor displays the top round-trip time of the circuit, in milliseconds, in a table.

FIGURE 134 Top Circuit RTT monitor

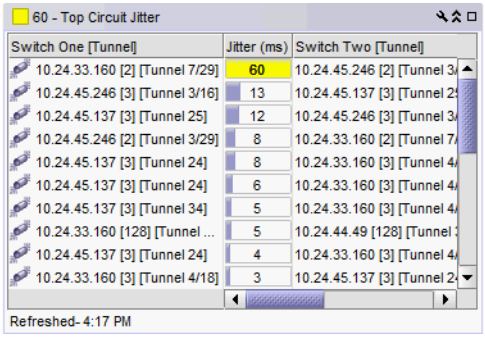

The Top Circuit RTT monitor includes the following data:

- Count/monitor title — The count for the data based on the round-trip time of the circuit shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- RTT (ms) — The round-trip time, in milliseconds, of the circuit.
- Switch Two — The IP address and FID of the destination switch for the FCIP tunnel.
- Circuit ID — The identifier of the circuit.
- Circuit Status — The operational status of the circuit.
- Circuit Admin Status — The administrative status of a circuit (enabled or disabled).

Accessing additional data from the Top Circuit RTT monitor

Double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Custom: Historical Performance Graph** dialog box for the selected measures. For more information, refer to "Generating and saving a historical performance graph" on page 972.

Top Duplicate Acknowledge monitor

The Top Duplicate Acknowledge monitor displays the top duplicate acknowledgment packets on a tunnel in a table.

FIGURE 135 Top Duplicate Acknowledge monitor

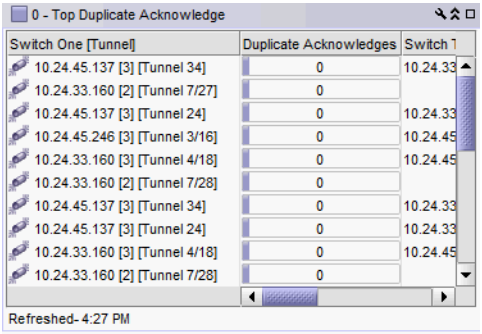

The Top Duplicate Acknowledge monitor includes the following data:

- Count/monitor title — The count for the data based on the number of duplicate acknowledgment packets on the tunnel shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- Duplicate Acknowledges — The number of duplicate acknowledgment packets.
- Switch Two — The IP address and FID of the destination switch for the FCIP tunnel.
- Circuit ID — The identifier of the circuit.
- Circuit Status — The operational status of the circuit.
- Circuit Admin Status — The administrative status of a circuit (enabled or disabled).

Accessing additional data from the Top Duplicate Acknowledge monitor

Double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Custom: Historical Performance Graph** dialog box for the selected measures. For more information, refer to "Generating and saving a historical performance graph" on page 972.

Top Slow Start monitor

The Top Slow Start monitor displays the top slow starts on a tunnel in a table.

FIGURE 136 Top Slow Start monitor

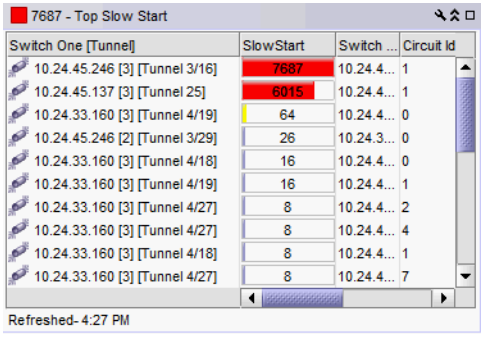

The Top Slow Start monitor includes the following data:

- Count/monitor title — The count for the data based on the number of slow starts on the tunnel shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- Slow Start — The number of slow starts.
- Switch Two — The IP address and FID of the destination switch for the FCIP tunnel.
- Circuit ID — The identifier of the circuit.
- Circuit Status — The operational status of the circuit.
- Circuit Admin Status — The administrative status of a circuit (enabled or disabled).

Accessing additional data from the Top Slow Start monitor

Double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Custom: Historical Performance Graph** dialog box for the selected measures. For more information, refer to "Generating and saving a historical performance graph" on page 972.

Top Out of Order monitor

The Top Out of Order monitor displays the top number of data packets delivered out of order on a tunnel in a table.

FIGURE 137 Top Out of Order monitor

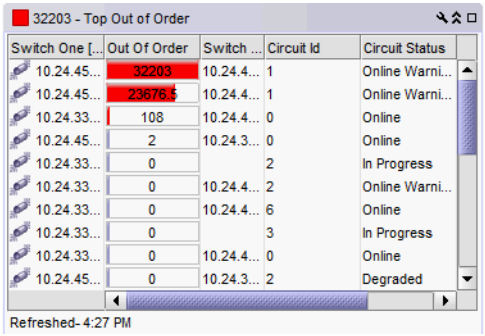

The Top Out of Order monitor includes the following data:

- Count/monitor title — The count for the data based on the number of data packets that was delivered out of order on the tunnel shown next to the monitor title.
- Switch One [Tunnel] — The IP address, FID, and tunnel ID of the source switch for the FCIP tunnel.
- Out of Order — The number of data packets that were delivered out of order.
- Switch Two — The IP address and FID of the destination switch for the FCIP tunnel.
- Circuit ID — The identifier of the circuit.
- Circuit Status — The operational status of the circuit.
- Circuit Admin Status — The administrative status of a circuit (enabled or disabled).

Editing a preconfigured performance monitor

You can customize the monitor to display data by a selected time frame as well as customize the display options.

1. Click the edit icon ( ) on the monitor.

   From the **Performance** tab of the **Customize Dashboard** dialog box, select the monitor you want to edit and click **Edit**.

2. Select the number of products to include in a selected measure by entering a number in the **For Top N, Bottom N Monitors, N=** field.

   Valid values are from 1 through 25. The default is 10.

3. Configure the monitor to show only values greater than or less than a specified value by completing the following steps.

   a. Select the **Show values** check box.

   b. Select **greater than** or **less than** from the list.

   c. Enter a value in the field.

4. Configure threshold numbers and associated colors by completing the following steps.

   You can define three threshold numbers in decreasing order and four threshold colors. The default values are as follows: 90 and above displays red; 75 and above displays orange; 60 and above displays yellow; and all others display blue.

a. Select the check box.
b. Enter a number in the field.
c. Click the color square to launch the **Color** dialog box.
    - To pick a color from a swatch, select the **Swatches** tab. Select a color from the display.
    - To specify a color based on hue, saturation, and brightness, click the **HSV** tab. Specify the hue (0 through 360 degrees), saturation (0 through 100%), value (0 through 100%), and transparency (0 through 100%).
    - To specify a color based on hue, saturation, and lightness, click the **HSL** tab. Specify the hue (0 through 360 degrees), saturation (0 through 100%), lightness (0 through 100%), and transparency (0 through 100%).
    - To specify a color based on values of red, green, and blue, click the **RGB** tab. Specify the values for red (0 through 255), green (0 through 255), blue (0 through 255), and alpha (0 through 255).
    - To specify a color based on values of cyan, magenta, yellow, and black, click the **CMYK** tab. Specify the values for cyan (0 through 255), magenta (0 through 255), yellow (0 through 255), black (0 through 255), and alpha (0 through 255).
    - To reset to the default color, click **Reset**.

5. Click **OK** to save your changes.

User-defined performance monitors

You can customize performance monitors specific to your needs. You can define up to 100 performance monitors; however, you can only display up to 30 performance monitors at a time.

Monitor types

You can create the following types of monitors:

- Top N (Products, Ports, and Traffic Flows monitors) — Displays the top number of products, ports, or traffic flows for the selected measure in a table.
- Bottom N (Products, Ports, and Traffic Flows monitors) Displays the bottom number of products, ports, or traffic flows for the selected measure in a table.
- Distribution (Products and Ports monitors) — Displays the number (distribution) of products or ports for each of the five percentage ranges defined for the selected measure in a bar graph
- Time Series (Products, Ports, and Traffic Flows monitors) — Displays the selected measures for products, ports, or traffic flows in a chart.
- Performance graph — Displays the configured performance graph on the dashboard.

Measures

Depending on the object (products, ports, traffic) you want to monitor, you can choose from the following measures:

- Product
    - Memory Utilization Percentage — The memory utilization percentage for the product.
    - CPU Utilization Percentage — The CPU utilization percentage for the product.
    - Temperature — The temperature in Celsius for the product.
    - Fan Speed — The fan speed in RPM for the product.
    - Response Time — The response time in seconds for the product.

- System Up Time — The system up time in days for the product.
- Ports Not In Use — The number of ports not in use for the product.
- Ping Packet Loss Percentage — The ping packet loss percentage for the product.
- AP Client Count — The number of AP clients for the product.

- Port
  - Common
    - Port Utilization Percentage — The memory utilization percentage.
    - Traffic — The traffic in mbps.
    - CRC Errors — The number of CRC errors.
  - FC
    - Link Resets — The number of link resets.
    - Signal Losses — The number of signal failures.
    - Sync Losses — The number of synchronization failures.
    - Link Failures — The number of link failures.
    - Sequence Errors — The number of sequence errors.
    - Invalid Transmissions — The number of invalid transmissions.
    - C3 Discards — The number of class 3 frames discarded.
    - C3 Discards TX TO — The number of transmitted class 3 frames discarded due to timeout.
    - C3 Discards RX TO — The number of received class 3 frames discarded due to timeout.
    - C3 Discards Unreachable — The number of class 3 frames discarded due to unreachable destination.
    - C3 Discards Other — The number of class 3 frames discarded due to other reasons.
    - Encode Error Out — The number of encode errors outside of the frame.
    - SFP Power — The SFP power in dbm.
    - SFP Voltage — The SFP voltage in mV.
    - SFP Current — The SFP current in mA.
    - SFP Temperature — The SFP temperature in Celsius.
    - Invalid Ordered Sets — The number of invalid ordered sets received at a port.
    - BB Credit Zero — The number of transitions in and out of the BB credit zero state.
    - Truncated Frames — The number of truncated frames received at a port.
  - FCIP
    - Cumulative Compression Ratio — The cumulative compression ratio for the FCIP tunnel.
    - Latency — The latency for the FCIP tunnel.
    - Dropped Packets — The number of dropped packets.
    - Link Retransmits — The number of retransmitted links.
    - Timeout Retransmits — The number of retransmits due to timeout.
    - Fast Retransmits — The number of fast retransmits triggered.
    - Duplicate Ack Received — The number of duplicate acknowledgments received.
    - Window Size RTT — The window size round trip time.
    - TCP Out of Order Segments — The number of segments received out of order.
    - Slow Start Status — The number of slow starts.
    - Current Compression Ratio — The current compression ratio for the FCIP tunnel.
    - RTT — The round trip time of the circuit.
    - Jitter — The circuit jitter variance.
    - Duplicate Acknowledges — The duplicate acknowledgment packets on a tunnel.
    - Slow Start — The slow starts on a tunnel.
    - Out of Order — The data packets delivered out of order on a tunnel.
  - IP
    - Errors — The number of errors.
    - Discards — The number of discarded frames.
    - Receive EOF — The number of end-of-frames received.

        - Underflow Errors — The number of underflow errors.
        - Overflow Errors — The number of overflow errors.
        - Alignment Errors — The number of alignment errors.
        - Runtime Errors — The number of run time errors.
        - Too Long Errors — The number of too long frame errors.
        -
    - Wireless
        - Dropped Events — The number of dropped events.
        - MAC Errors — The number of MAC errors.
        - Back Packets Received — The number of bad packets received.
        - Tx Errors — The number of transmit errors.
- Traffic flows
    - SCSI
        - Read Frame Count (frames) — The SCSI read command frame count as reported in the last data point received for the flow.
        - Write Frame Count (frames) — The SCSI write command frame count as reported in the last data point received for the flow.
        - Read Frame Rate (f/s) — The SCSI write frame rate per second as reported in the last data point received for the flow.
        - Write Frame Rate (f/s) — The SCSI write frame rate per second as reported in the last data point received for the flow.
        - Read Data (Bytes) — The SCSI read data in bytes as reported in the last data point received for the flow.
        - Write Data (Bytes) — The SCSI read data in bytes as reported in the last data point received for the flow.
        - Read Data Rate (Mbps) — The SCSI read frame in megabytes per second rate as reported by the last data point.
        - Write Data Rate (Mbps) — The SCSI write frame rate in megabytes per second as reported by the last data point.
    - Frame
        - Transmit Frame Count (frames) — The transmit frame count as reported in the last data point received for the flow.
        - Receive Frame Count (frames) — The received frame count as reported in the last data point received for the flow.
        - Transmit Frame Rate (f/s) — The transmit frame rate per second as reported in the last data point received for the flow.
        - Receive Frame Rate (f/s) — The received frame rate per second as reported in the last data point received for the flow.
        - Transmit Word Count (bytes) — The transmit word count in bytes as reported in the last data point received for the flow.
        - Receive Word Count (bytes) — The received word count in bytes as reported in the last data point received for the flow.
        - Transmit Throughput (Mbps) — The transmit throughput in megabytes per second as reported by the last data point.
        - Receive Throughput (Mbps) — The received throughput in megabytes per second as reported by the last data point.
        - Generator Transmit Frame Count (frames) — The transmit frame count as reported in the last data point received for the flow.
        - Generator Receive Frame Count (frames) — The received frame count as reported in the last data point received for the flow.
        - Mirrored Frames Count (frames) — The mirrored frame count as reported in the last data point received for the flow.
        - Mirrored Tx Frames (frames) —The mirrored transmit frame count as reported in the last data point received for the flow.
        - Mirrored Rx Frames (frames) —The mirrored received frame count as reported in the last data point received for the flow.

Top or bottom product performance monitors

The top or bottom product performance monitors (Figure 138) display the top or bottom number of products (for example, top 10 products) for the selected measure in a table.

FIGURE 138 Top or bottom product performance monitor example

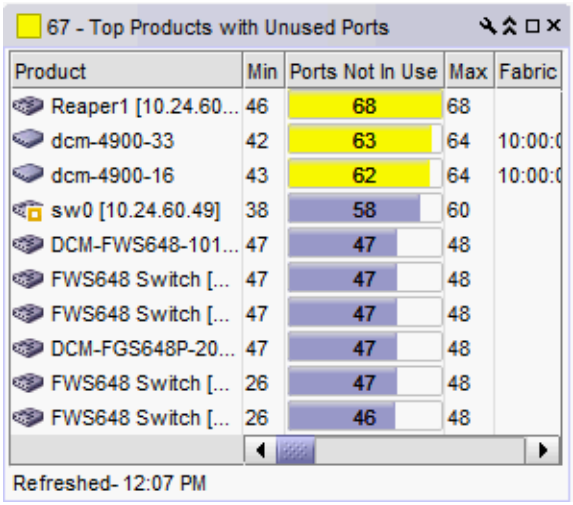

The top or bottom product performance monitor includes the following data:

- Threshold icon/object count/monitor title — The color associated with the threshold and number of objects within that threshold displays next to the monitor title.
- **Product** — The product affected by this monitor.
- **Min** — The minimum value of the measure in the specified time range.
- ***Measure_Type*** **—** The percentage bar of the selected measure.

  By default, products display sorted by the ***Measure_Type*** value (Top products sort from highest to lowest and bottom products sort lowest to highest). Click a column head to sort the columns by that value.
- **Max** — The maximum value of the measure in the specified time range.
- **Fabric** — The fabric to which the device belongs.
- **Product Type** — The type of product (for example, switch).
- **State** — The product state (for example, Offline).
- **Status** — The product status (for example, Reachable).
- **Tag** — The product tag.
- **Serial # —** The serial number of the product.
- **Model** — The product model.
- **Port Count —** The number of ports on the product.
- **Firmware —** The firmware level running on the product.
- **Location —** The location of the product.
- **Contact —** A contact name for the product.
- **Refreshed** — The time of the last update for the monitor.

To configure a product performance monitor, refer to "Configuring a user-defined product performance monitor" on page 275.

Accessing additional data from top or bottom product monitors

In a Top N or Bottom N monitor, double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Historical Graphs/Tables** dialog box for the selected measures. For more information, refer to "Performance Data" on page 959.

Top or bottom port performance monitors

The top or bottom port performance monitors (Figure 139) display the top or bottom number of ports (for example, bottom 10 ports) for the selected measure in a table.

FIGURE 139 Top or bottom port performance monitor example

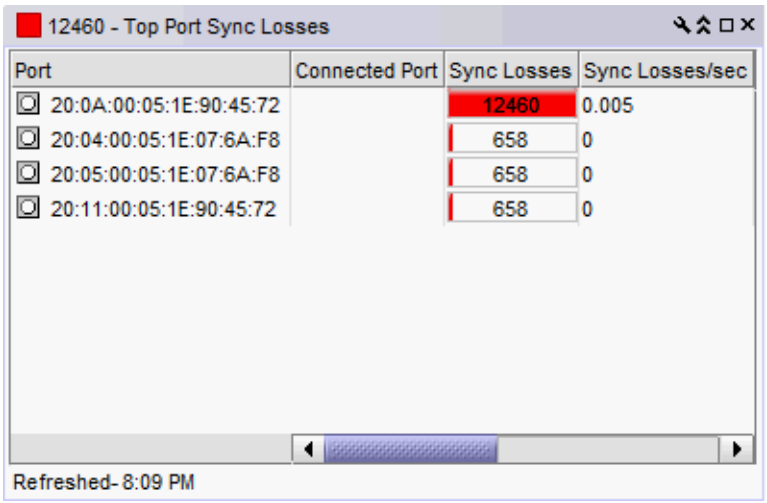

The top or bottom port performance monitor includes the following data:

- Threshold icon/object count/monitor title — The color associated with the threshold and number of objects within that threshold displays next to the monitor title.
- **Count/monitor title** — The count for the data based on the error count or error rate shown next to the monitor title.
- **Port** — The port affected by this monitor.
- *Connected_Port_Link* (where *Connected_Port_Link* is **Connected Port**, **Initiator**, or **Target**) — Displays one of the following:
  - **Connected Port** — The ISL or IFL port on the connected device. Click to launch the switch port properties dialog box.
  - **Initiator** — The initiator port on the connected device. Click to launch the device properties dialog box.
  - **Target** — The target port on the connected device. Click to launch the device properties dialog box.
- *Measure_Type* — The percentage bar of the selected measure. Depending on the selected measure, both the error rate (per second) and error count may display. For selected measures, more than one **Measure_Type** may display (for example RX and TX).

  By default, ports display sorted by the **Measure_Type** value (Top ports sort from highest to lowest and bottom ports sort lowest to highest). Click a column head to sort the columns by that value.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.

- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).
- **Refreshed** — The time of the last update for the monitor.

To configure a port performance monitor, refer to "Configuring a user-defined port performance monitor" on page 278.

Accessing additional data from top or bottom port monitors

- In a Top N or Bottom N monitor, double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Custom: Historical Performance Graph** dialog box for the selected measures. For more information, refer to "Generating and saving a historical performance graph" on page 972.

Distribution performance monitors

The distribution performance monitor (Figure 140) displays the distribution (number) of products or ports for each of the five percentage ranges defined for the selected measure in a bar graph.

FIGURE 140 Distribution performance monitor example

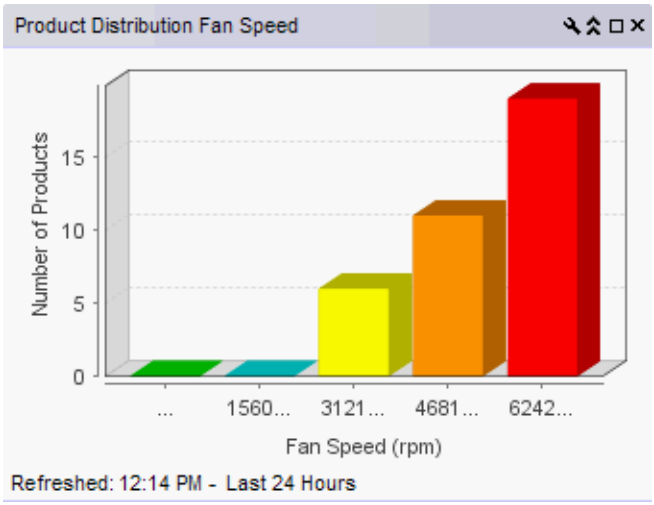

The distribution performance monitor includes the following data:

- Monitor title — The user-defined monitor title.
- **Number of Products/Ports** (y-axis) — The y-axis always displays a numbered range (zero to the maximum number of objects) for the products or ports affected by the selected measure.
- ***Measure_Type*** (x-axis) — The x-axis display depends on the ***Measure_Type*** you selected for this monitor. Each bar on the graph maps directly to one of the five percentage ranges defined for the monitor. ***Measure_Type*** includes the following measures:

TABLE 21 Product measures types

- Memory Utilization Percentage
- CPU Utilization Percentage
- Temperature (C)
- Fan Speed (rpm)
- Response Time (s)
- System Up Time (days)
- Ports Not In Use
- Ping Packet Loss Percentage
- AP Client Count

TABLE 22 Port measures types

Common

- Port Utilization Percentage
- Traffic
- CRC Errors

FC

- Link Resets
- Signal Losses
- Sync Losses
- Link Failures
- Sequence Errors
- Invalid Transmissions
- C3 Discards
- C3 Discards TX TO
- C3 Discards RX TO
- C3 Discards Unreachable
- C3 Discards Other
- Encode Error Out
- SFP Power
- SFP Voltage
- SFP Current
- SFP Temperature
- Invalid Ordered Sets
- BB Credit Zero
- Truncated Frames

FCIP

- Cumulative Compression Ratio
- Latency
- Dropped Packets
- Link Retransmits
- Timeout Retransmits
- Fast Retransmits
- Duplicate Ack Received
- Window Size RTT
- TCP Out of Order Segments
- Slow Start Status
- Current Compression Ratio
- RTT
- Jitter
- Duplicate Acknowledges
- Slow Start
- Out of Order

IP

- Errors
- Discards
- Receive EOF
- Underflow Errors
- Overflow Errors
- Alignment Errors
- Runtime Errors
- Too Long Errors

Wireless

- Dropped Events
- MAC Errors
- Back Packets Received
- Tx Errors

To configure a distribution performance monitor, refer to "Configuring a user-defined product performance monitor" on page 275 or "Configuring a user-defined port performance monitor" on page 278.

Accessing additional data from the Distribution monitors

- Place the cursor on a bar in the graph to display the number of products included in the count for the selected bar. For example, the tooltip "(Data Item 3, 22.6-33.8) = 6" means that there are six products within the third percentage range (displays the temperatures within the percentage range) for the selected measure (product temperature).
- Double-click a percentage range to navigate to the *Monitor_Title* **Distribution Data Details** dialog box. for more information, refer to "Viewing product distribution data details" on page 280 or "Viewing port distribution data details" on page 281.

Time series performance monitors

The time series performance monitors (Figure 141) display the selected measures in a chart.

FIGURE 141 Time series performance monitor example

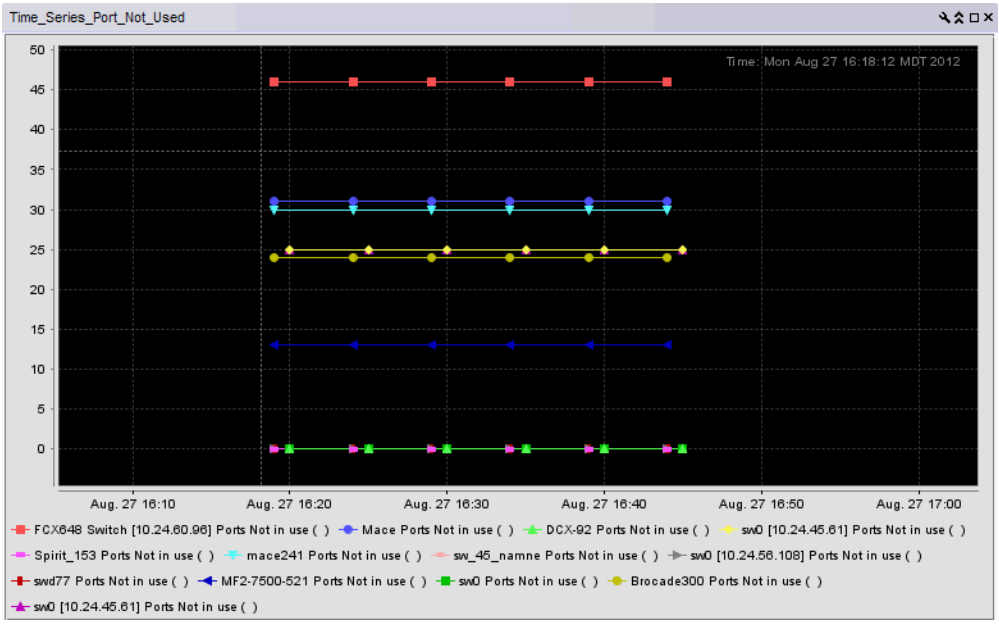

The time series performance monitor includes the following data:

- Monitor title — The user-defined monitor title.
- **Value** (y-axis) — The number of objects affected by this monitor.
- **Time** (x-axis) — The date and time the monitor collected the data.
- **Legend** (below the x-axis) — The line color and the associated data that each line represents.
- Network Scope — The network scope, such as Local or Published. Displays Local if you select the targets when creating the monitor. Displays Published if you select the **Use Network Scope** check box when creating the monitor.

Place the cursor on a data point in graph line to view details. Place the cursor on an Event icon to view the event details. Right-click the graph to access the graph shortcut menu (refer to "Configuring the performance graph" on page 1717).

To configure a time series performance monitor, refer to "Configuring a user-defined product performance monitor" on page 275 or "Configuring a user-defined port performance monitor" on page 278.

Configuring a user-defined product performance monitor

For creating a user-defined dashboard, refer to "Creating a user-defined dashboard" on page 204 and perform the following steps to configure a user-defined product performance monitor.

1. Click the **Customize Dashboard** icon.

   The **Customize Dashboard** dialog box displays.

2. Click the **Performance** tab.

3. Click **Add**.

   The **Add Performance Dashboard Monitor** dialog box displays.

4. Enter a unique title for the monitor.

   The title can be up to 256 characters in length.

5. Select the type of monitor you are creating from the **Monitor Type - Products** area:

   - **Top N** — Select to monitor the top N (number) products affected by the selected measure.
   - **Bottom N** — Select to monitor the bottom N (number) products affected by the selected measure.
   - **Distribution** — Select to monitor the selected measure for five defined distribution percentages.
   - **Time Series** — Select to monitor a selected measure for a range of time and specified target.

6. Select the product measure for the monitor in the **Measure** area:

   - **Memory Utilization Percentage**
   - **CPU Utilization Percentage**
   - **Temperature**
   - **Fan Speed**
   - **Response Time**
   - **System Up Time**
   - **Ports Not In Use**
   - **Ping Packet Loss Percentage**
   - **AP Client Count** (not available for Time Series monitors)

7. (Top N and Bottom N monitors only) Select the number products to include in a selected measure by entering a number in the **For Top N, Bottom N Monitors, N=** field.

   Valid values are from 1 through 25. The default is 10.

8. (Top N, Bottom N, and Distribution monitors only) Configure the monitor to show only values greater than or less than a specified value by completing the following steps.

   a. Select the **Show values** check box.

   b. Select **greater than** or **less than** from the list.

   c. Enter a value in the field.

9. (Top N, Bottom N, and Distribution monitors only) Configure threshold numbers and associated colors by completing the following steps.

   Depending on the monitor type you select, you can define up to four threshold numbers in increasing or decreasing order and up to five associated threshold colors.

   (Top N and Bottom N monitors only) The decreasing order defaults are as follows: 90 and above displays red, 75 and above displays orange, 60 and above displays yellow, and all others display blue. The maximum values allowed are -32,768 through 32,767 for SFP power and 0 through 32,767 for all other measures.

   (Distribution monitors only) The increasing order defaults are as follows: 0 through 20 displays green, 21 through 40 displays blue, 41 through 60 displays yellow, 61 through 80 displays orange, and 81 through 100 displays red.

a. (Top N and Bottom N monitors only) Select the check box.

b. Enter a number in the field.

c. Click the color square to launch the **Color** dialog box.

- To pick a color from a swatch, select the **Swatches** tab. Select a color from the display.
- To specify a color based on hue, saturation, and brightness, click the **HSV** tab. Specify the hue (0 through 360 degrees), saturation (0 through 100%), value (0 through 100%), and transparency (0 through 100%).
- To specify a color based on hue, saturation, and lightness, click the **HSL** tab. Specify the hue (0 through 360 degrees), saturation (0 through 100%), lightness (0 through 100%), and transparency (0 through 100%).
- To specify a color based on values of red, green, and blue, click the **RGB** tab. Specify the values for red (0 through 255), green (0 through 255), blue (0 through 255), and alpha (0 through 255).
- To specify a color based on values of cyan, magenta, yellow, and black, click the **CMYK** tab. Specify the values for cyan (0 through 255), magenta (0 through 255), yellow (0 through 255), black (0 through 255), and alpha (0 through 255).
- To reset to the default color, click **Reset**.

10. (Time series monitors only) Add targets for the monitor by clicking **Add** and completing steps in "Adding targets to a user-defined performance monitor" on page 277.

    Remove targets from the monitor by selecting one or more targets in the **Targets** list and clicking **Remove**.

11. Click **OK** on the **Add Performance Dashboard Monitor** dialog box.

    The **Customize Dashboard** dialog box displays with the new monitor in the **Performance Monitors** list.

12. Click **OK** on the **Customize Dashboard** dialog box.

    The new performance monitors display at the bottom of the dashboard.

Accessing additional data from user-defined product performance monitors

- In a Distribution monitor, double-click a percentage range to navigate to the *Measure_Type* **Distribution Data Details** dialog box. For more information, refer to "Viewing product distribution data details" on page 280 or "Viewing port distribution data details" on page 281.
- In a Top N or Bottom N product monitor, double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Historical Graphs/Tables** dialog box for the selected measures. For more information, refer to "Performance Data" on page 959.

Adding targets to a user-defined performance monitor

You can only add targets for Time Series monitors. For creating a user-defined dashboard, refer to "Creating a user-defined dashboard" on page 204 and perform the following steps to add targets to a user-defined product performance monitor.

1. Click the **Customize Dashboard** icon.

   The **Customize Dashboard** dialog box displays.

2. Click the **Performance** tab.

3. Click **Add**.

   The **Add Performance Dashboard Monitor** dialog box displays.

4. Select **Time Series** from the **Monitor Type - Product** or **Port** area.

5. Select the port measure for the monitor in the **Measure** area

6. Display data for a specific duration from the **Duration** options.
7. Click **Add** beneath the **Targets** table.

   The **Performance Dashboard Monitor Targets** dialog box displays.

   Depending on the type of measure you select, you can add SAN products/portsand FCIP tunnels to the list of targets.

   If you selected a product measure, continue with step 8.

   If you selected SAN port measure, continue with step 8.

   If you selected a FC IP port measure, go to step 11.
8. Click the **SAN** tab.
9. Select SAN targets from the **Available SAN Sources** list.
10. Click the right arrow button to move the targets to the **Selected Sources** list.
11. Select FCIP targets from the **Available** list.
12. Click the right arrow button to move the targets to the **Selected Sources** list.
13. Click **OK** on the **Performance Dashboard Monitor Targets** dialog box.

    The targets display in the **Targets** list of the **Add Performance Dashboard Monitor** dialog box.
14. Click **OK** on the **Add Performance Dashboard Monitor** dialog box.

    The **Customize Dashboard** dialog box displays with the new monitor in the **Performance Monitors** list.
15. Click **OK** on the **Customize Dashboard** dialog box.

    The performance monitors display at the bottom of the dashboard.

Configuring a user-defined port performance monitor

For creating a user-defined dashboard, refer to "Creating a user-defined dashboard" on page 204 and perform the following steps to configure a user-defined port performance monitor.

1. Click the **Customize Dashboard** icon.

   The **Customize Dashboard** dialog box displays.
2. Click the **Performance** tab.
3. Click **Add**.

   The **Add Performance Dashboard Monitor** dialog box displays.
4. Select the type of monitor you are creating from the **Monitor Type - Port** area:
   - **Top N** — Select to monitor the top N (number) ports affected by the selected measure.
   - **Bottom N** — Select to monitor the bottom N (number) ports affected by the selected measure.
   - **Distribution** — Select to monitor the selected measure for five defined distribution percentages.
   - **Time Series** — Select to monitor a selected measure for a range of time and specified targets.
5. Select the port measure for the monitor in the **Measure** area:

Common

- **Port Utilization Percentage**
- **Traffic**
- **CRC Errors**

FC

- **Link Resets**
- **Signal Losses**
- **Sync Losses**
- **Link Failures**
- **Sequence Errors**
- **Invalid Transmissions**
- **C3 Discards**
- **C3 Discards TX TO**
- **C3 Discards RX TO**
- **C3 Discards Unreachable**
- **C3 Discards Other**
- **Encode Error Out**
- **SFP Power**
- **SFP Voltage**
- **SFP Current**
- **SFP Temperature**
- **Invalid Ordered Sets**
- **BB Credit Zero**
- **Truncated Frames**

FCIP

- Cumulative Compression Ratio
- Latency
- Dropped Packets
- Link Retransmits
- Timeout Retransmits
- Fast Retransmits
- Duplicate Ack Received
- Window Size RTT
- TCP Out of Order Segments
- Slow Start Status
- Current Compression Ratio
- RTT
- Jitter
- Duplicate Acknowledges
- Slow Start
- Out of Order

IP

- Errors
- Discards
- Receive EOF
- Underflow Errors
- Overflow Errors
- Alignment Errors
- Runtime Errors
- Too Long Errors

Wireless

- Dropped Events
- MAC Errors
- Back Packets Received
- Tx Errors

6. (Top N and Bottom N monitors only) Select the number of ports to include in a selected measure by entering a number in the **For Top N, Bottom N Monitors, N=** text box.

   Valid values are from 1 through 25. The default is 10.

7. (Top N, Bottom N, and Distribution monitors only) Configure the monitor to show only values greater than or less than a specified value by completing the following steps.

   a. Select the **Show values** check box.

   b. Select **greater than** or **less than** from the list.

   c. Enter a value in the field.

8. (Top N, Bottom N, and Distribution monitors only) Configure threshold numbers and associated colors by completing the following steps.

   Depending on the monitor type you select, you can define up to four threshold numbers in increasing or decreasing order and up to five associated threshold colors.

   (Top N and Bottom N monitors only) The decreasing order defaults are as follows: 90 and above displays red, 75 and above displays orange, 60 and above displays yellow, and all others display blue. The maximum values allowed are -32,768 through 32,767 for SFP power and 0 through 32,767 for all other measures.

   (Distribution monitors only) The increasing order defaults are as follows: 0 through 20 displays green, 21 through 40 displays blue, 41 through 60 displays yellow, 61 through 80 displays orange, and 81 through 100 displays red.

a. (Top N and Bottom N monitors only) Select the check box.

b. Enter a number in the field.

c. Click the color square to launch the **Color** dialog box.

- To pick a color from a swatch, select the **Swatches** tab. Select a color from the display.
- To specify a color based on hue, saturation, and brightness, click the **HSV** tab. Specify the hue (0 through 360 degrees), saturation (0 through 100%), value (0 through 100%), and transparency (0 through 100%).
- To specify a color based on hue, saturation, and lightness, click the **HSL** tab. Specify the hue (0 through 360 degrees), saturation (0 through 100%), lightness (0 through 100%), and transparency (0 through 100%).
- To specify a color based on values of red, green, and blue, click the **RGB** tab. Specify the values for red (0 through 255), green (0 through 255), blue (0 through 255), and alpha (0 through 255).
- To specify a color based on values of cyan, magenta, yellow, and black, click the **CMYK** tab. Specify the values for cyan (0 through 255), magenta (0 through 255), yellow (0 through 255), black (0 through 255), and alpha (0 through 255).
- To reset to the default color, click **Reset**.

9. (Time series monitors only) Add targets for the monitor by clicking **Add** and completing the steps in "Adding targets to a user-defined performance monitor" on page 277.

   Remove targets from the monitor by selecting one or more targets in the **Targets** list and clicking **Remove**.

10. Click **OK** on the **Add Performance Dashboard Monitor** dialog box.

    The **Customize Dashboard** dialog box displays with the new monitor in the **Performance Monitors** list.

11. Click **OK** on the **Customize Dashboard** dialog box.

    The new performance monitors display at the bottom of the dashboard.

Accessing additional data from user-defined port performance monitors

- In a Distribution monitor, double-click a percentage range to navigate to the *Measure_Type* **Distribution Data Details** dialog box. For more information, refer to "Viewing product distribution data details" on page 280 or "Viewing port distribution data details" on page 281.
- In a Top N or Bottom N monitor, double-click a row or right-click a row and select **Show Graph/Table** to navigate to the **Custom: Historical Performance Graph** dialog box for the selected measures. For more information, refer to "Generating and saving a historical performance graph" on page 972.
- In a Top N or Bottom N C3 Discards TX TO and C3 Discards RX TO monitors, right-click an FC-port row (Fabric OS device running 7.1.0 or later) and select **Discarded Frames** to navigate to the **Discarded Frames** dialog box. For more information, refer to "Viewing discarded frames from a port" on page 441.

Viewing product distribution data details

Each bar on the product distribution graph maps directly to one of the five percentage ranges defined for the distribution performance monitor (refer to "Distribution performance monitors" on page 273).

1. Double-click a bar in the graph.

   The *Monitor_Title* **Data Details** dialog box displays.

2. Review the data.

   The product distribution data details include the following fields and components:

- **Product** — The name of the product affected by the selected measure.
- *Measure_Type* — This column depends on which measure you select for the monitor.
  - Memory Utilization Percentage — The memory utilization percentage for the product.
  - CPU Utilization Percentage — The CPU utilization percentage for the product.
  - Temperature — The temperature in Celsius for the product.
  - Fan Speed — The fan speed in RPM for the product.
  - Response Time — The response time in seconds for the product.
  - System Up Time — The system up time in days for the product.
  - Ports Not In Use — The number of ports not in use for the product.
  - Ping Packet Loss Percentage — The ping packet loss percentage for the product.
  - AP Client Count — The number of AP clients for the product.
- **Fabric** — The fabric to which the device belongs.
- **Product Type** — The type of product (for example, switch).
- **State** — The product state (for example, Offline).
- **Status** — The product status (for example, Reachable).
- **Tag** — The product tag.
- **Serial # —** The serial number of the product.
- **Model** — The product model.
- **Port Count —** The number of ports on the product.
- **Firmware —** The firmware level running on the product.
- **Location —** The location of the product.
- **Contact —** A contact name for the product.

3. Click **Close**.

Viewing port distribution data details

Each bar on the port distribution graph maps directly to one of the five percentage ranges defined for the distribution monitor (refer to "Distribution performance monitors" on page 273).

1. Double-click a bar in the graph.

   The *Monitor_Title* **Data Details** dialog box displays.

2. Review the data.

   The port distribution data details include the following fields and components:

   - **Port** — The port affected by the selected measure.
   - **TX/RX** — Whether the port is transmitting (TX) or receiving (RX) data. This column is not available for all measures.
   - *Measure_Type* — This column depends on which measure you select for the monitor.
     - Common
       - Port Utilization Percentage — The memory utilization percentage.
       - Traffic — The traffic in mbps.
       - CRC Errors — The number of CRC errors.

- FC
  - Link Resets — The number of link resets.
  - Signal Losses — The number of signal failures.
  - Sync Losses — The number of synchronization failures.
  - Link Failures — The number of link failures.
  - Sequence Errors — The number of sequence errors.
  - Invalid Transmissions — The number of invalid transmissions.
  - C3 Discards — The number of class 3 frames discarded.
  - C3 Discards TX TO — The number of transmitted class 3 frames discarded due to timeout.
  - C3 Discards RX TO — The number of received class 3 frames discarded due to timeout.
  - C3 Discards Unreachable — The number of class 3 frames discarded due to unreachable destination.
  - C3 Discards Other — The number of class 3 frames discarded due to other reasons.
  - Encode Error Out — The number of encode errors outside of the frame.
  - SFP Power — The SFP power in dbm.
  - SFP Voltage — The SFP voltage in mV.
  - SFP Current — The SFP current in mA.
  - SFP Temperature — The SFP temperature in Celsius.
  - Invalid Ordered Sets — The number of invalid ordered sets received at a port.
  - BB Credit Zero — The number of transition in and out of the BB credit zero state.
  - Truncated Frames — The number of truncated frames received at a port.
- FCIP
  - Cumulative Compression Ratio — The cumulative compression ratio for the FCIP tunnel.
  - Latency — The latency for the FCIP tunnel.
  - Dropped Packets — The number of dropped packets.
  - Link Retransmits — The number of retransmitted links.
  - Timeout Retransmits — The number of retransmits due to timeout.
  - Fast Retransmits — The number of fast retransmits triggered.
  - Duplicate Ack Received — The number of duplicate acknowledgments received.
  - Window Size RTT — The window size round trip time.
  - TCP Out of Order Segments — The number of segments received out of order.
  - Slow Start Status — The number of slow starts.
  - RTT — The round trip time of the circuit.
  - Jitter — The circuit jitter variance.
  - Duplicate Acknowledges — The duplicate acknowledgment packets on a tunnel.
  - Slow Start — The slow starts on a tunnel.
  - Out of Order — The data packets delivered out of order on a tunnel.
- IP
  - Errors — The number of errors.
  - Discards — The number of discarded frames.

        - Receive EOF — The number of end-of-frames received.
        - Underflow Errors — The number of underflow errors.
        - Overflow Errors — The number of overflow errors.
        - Alignment Errors — The number of alignment errors.
        - Runtime Errors — The number of run time errors.
        - Too Long Errors — The number of too long frame errors.
    - Wireless
        - Dropped Events — The number of dropped events.
        - MAC Errors — The number of MAC errors.
        - Back Packets Received — The number of bad packets received.
        - Tx Errors — The number of transmit errors.
- **Product** — The product affected by this monitor.
- **Type** — The type of port (for example, U-Port).
- **Identifier** — The port identifier.
- **Port Number** — The port number.
- **State** — The port state (for example, Enabled).
- **Status** — The port status (for example, Up).

3. Click **Close**.

Traffic flow dashboard monitors

NOTE
Traffic flow monitors are only supported on devices running Fabric OS 7.2 and later with the Fabric Vision license.

You can use the dashboard to monitor traffic flows. To monitor a flow, you must first create and activate the flow in Flow Vision (refer to "Flow Vision".

Traffic flow monitor types

You can create the following types of monitors for traffic flows:

- Top N — Displays the top number of traffic flows for the selected measure in a table.
- Bottom N — Displays the bottom number of traffic flows for the selected measure in a table.
- Time Series — Displays the selected measures for a range of time and specified targets for traffic flows in a chart.
- Performance graph — Displays the configured performance graph on the dashboard.

Traffic flow measures

You can use the following measures to create your traffic flow monitors:

- SCSI
  - Read Frame Count (frames) — The SCSI read command frame count as reported in the last data point received for the flow.
  - Write Frame Count (frames) — The SCSI write command frame count as reported in the last data point received for the flow.
  - Read Frame Rate (f/s) — The SCSI write frame rate per second as reported in the last data point received for the flow.
  - Write Frame Rate (f/s) — The SCSI write frame rate per second as reported in the last data point received for the flow.
  - Read Data (Bytes) — The SCSI read data in bytes as reported in the last data point received for the flow.
  - Write Data (Bytes) — The SCSI read data in bytes as reported in the last data point received for the flow.
  - Read Data Rate (Mbps) — The SCSI read frame in megabytes per second rate as reported by the last data point.
  - Write Data Rate (Mbps) — The SCSI write frame rate in megabytes per second as reported by the last data point.
- Frame
  - Transmit Frame Count (frames) — The transmit frame count as reported in the last data point received for the flow.
  - Receive Frame Count (frames) — The received frame count as reported in the last data point received for the flow.
  - Transmit Frame Rate (f/s) — The transmit frame rate per second as reported in the last data point received for the flow.
  - Receive Frame Rate (f/s) — The received frame rate per second as reported in the last data point received for the flow.
  - Transmit Word Count (bytes) — The transmit word count in bytes as reported in the last data point received for the flow.
  - Receive Word Count (bytes) — The received word count in bytes as reported in the last data point received for the flow.
  - Transmit Throughput (Mbps) — The transmit throughput in megabytes per second as reported by the last data point.
  - Receive Throughput (Mbps) — The received throughput in megabytes per second as reported by the last data point.
  - Generator Transmit Frame Count (frames) — The transmit frame count as reported in the last data point received for the flow.
  - Generator Receive Frame Count (frames) — The received frame count as reported in the last data point received for the flow.
  - Mirrored Frames Count (frames) — The mirrored frame count as reported in the last data point received for the flow.
  - Mirrored Tx Frames (frames) —The mirrored transmit frame count as reported in the last data point received for the flow.
  - Mirrored Rx Frames (frames) —The mirrored received frame count as reported in the last data point received for the flow.

Traffic flow performance graph monitor

The traffic flow performance monitors display (Figure 142) the selected measures in a chart.

FIGURE 142 Traffic flow performance graph monitor example

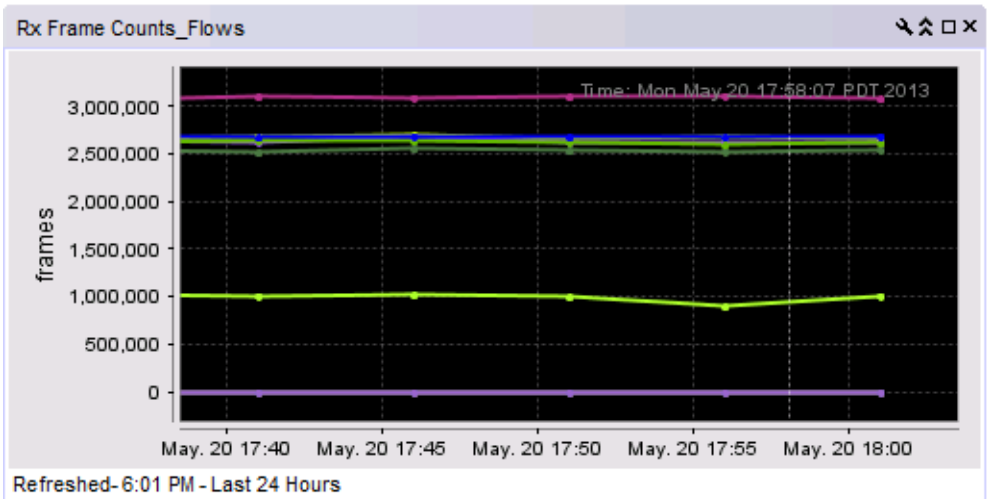

The traffic flows performance monitor includes the following data:

- Monitor title — The user-defined monitor title.
- **Value** (y-axis) — The number of objects affected by the selected measure.
- **Time** (x-axis) — The time the monitor collected the data.
- **Legend** (below the x-axis) — The line color and the associated data that each line represents.

Accessing additional data from traffic flows performance graph monitors

- Place the cursor on a data point in graph line to view details.
- Right-click the graph to access the graph shortcut menu (refer to "Configuring the performance graph" on page 1717).

Top or bottom traffic flow performance monitor

The top or bottom traffic flow performance monitors (Figure 143) top or bottom number of flows for the selected measure in a table.

FIGURE 143 Top traffic flow monitor example

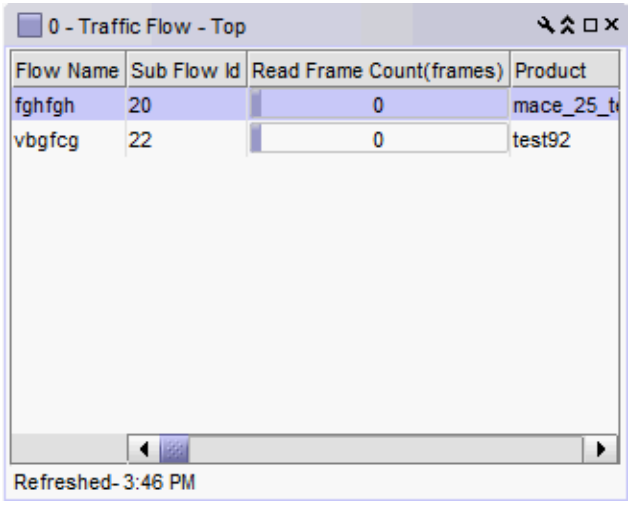

The top or bottom flow performance monitor includes the following data:

- Threshold icon/object count/monitor title — The color associated with the threshold and the number of objects within that threshold are displayed next to the monitor title.
- **Flow Name** — The name of the flow.
- **Sub Flow ID** — The sub-flow identifier.
- *Measure_Type* — The percentage bar of the selected measure. For a list of selected measures, refer to "Traffic flow measures" on page 284.

  By default, flows display sorted by the *Measure_Type* value (top flows sort from highest to lowest and bottom flows sort lowest to highest). Click a column head to sort the columns by that value.
- **Product** — The device name.
- **Source** — The source device identifier.
- **Destination** — The destination device identifier.
- **Feature** — The active feature for the sub-flow definition. Valid values include: Generator, Monitor, or Mirror.
- **Rx Port** — The receive (ingress) port.
- **Tx Port** — The transmit (egress) port.
- **LUN** — The LUN values defined in the flow.
- **Bi-direction** — Whether or not the flow is bi-directional. Valid values are Yes or No.
- **Flow Definition Persistence** — Whether or not to persist flow definition over device reboot.
- **SCSI Commands** — List of provisioned SCSI commands.
- **Size** — The size of the frame payload.
- **Pattern** — The pattern of the frame payload.

Accessing additional data from traffic flow performance monitors

- Right-click a row in the table to access the shortcut menu and select one of the following options:
  - **Show Graph/Table** — Launches the **Flow Graphing** dialog box with the selected measures (sub-flows) to be plotted.
  - **Locate** — Move the focus to the **SAN** tab with the associated switch highlighted.
  - **Monitor** — Launches the **Monitor - Flow Vision** dialog box with the selected sub-flows in the **Active Flows** list.
  - **Table** — Use to configure the table (refer to "Customizing application tables" on page 308).
- Right-click column head to configure the table (refer to "Customizing application tables" on page 308).

Time series traffic flow performance monitor

The time series traffic flow performance monitors display (Figure 144) the selected measure in a chart.

FIGURE 144 Traffic flow performance monitor example

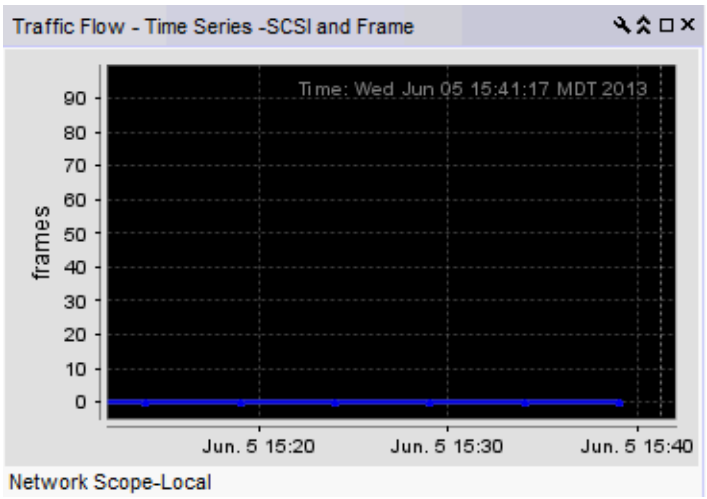

The time series traffic flow performance monitor includes the following data:

The time series performance monitor includes the following data:

- Monitor title — The user-defined monitor title.
- **Value** (y-axis) — The number of objects affected by this monitor.
- **Time** (x-axis) — The date and time the monitor collected the data.
- **Legend** (below the x-axis) — The line color and the associated data that each line represents.

Place the cursor on a data point in graph line to view details. Place the cursor on an Event icon to view the event details. Right-click the graph to access the graph shortcut menu (refer to "Configuring the performance graph" on page 1717).

To configure a time series performance monitor, refer to "Configuring a user-defined traffic flow performance monitor" on page 288.

Configuring a traffic flows monitor from a performance graph

1. Configure the performance graph.

   To configure a traffic flows performance graph, refer to //link to flow vision//.

2. Click **Save As Widget** to create a monitor of the graph data for the dashboard.

   The **Historical Chart Monitor** - *Date_Time* dialog box displays (where *Date_Time* is the date and time the monitor was created).

3. Modify the title, if necessary, and click **OK**.

4. Click **OK** on the message.

   The new monitor is added to the **Performance** tab of the **Customize Dashboard** dialog box.

5. From the dashboard, click the **Customize Dashboard** icon.

   The **Customize Dashboard** dialog box displays.

6. Click the **Performance** tab.

   The new performance monitor displays at the bottom of the Performance Monitors list.

7. Select the **Display** check box for the new monitor.

8. Click **OK** on the **Customize Dashboard** dialog box.

Configuring a user-defined traffic flow performance monitor

For creating a user-defined dashboard, refer to “Creating a user-defined dashboard” on page 346 and perform the following steps to configure a user-defined port performance monitor.

1. Click the **Customize Dashboard** icon.

   The **Customize Dashboard** dialog box displays.

2. Click the **Performance** tab.

3. Click **Add**.

   The **Add Performance Dashboard Monitor** dialog box displays.

4. Select the type of monitor you are creating from the **Monitor Type - Traffic Flows** area:

   - **Top N** — Select to monitor the top N (number) ports affected by the selected measure.
   - **Bottom N** — Select to monitor the bottom N (number) ports affected by the selected measure.
   - **Time Series** — Select to monitor one or more measures for a range of time and specified targets.

5. Select the traffic measure for the monitor in the **Measure** area:

For Time Series monitors, you can select more than one measure.

SCSI

- Read Frame Count (frames)
- Write Frame Count (frames)
- Read Frame Rate (f/s)
- Write Frame Rate (f/s)
- Read Data (Bytes)
- Write Data (Bytes)
- Read Data Rate (Mbps)
- Write Data Rate (Mbps)

Frame

- Transmit Frame Count (frames)
- Receive Frame Count (frames)
- Transmit Frame Rate (f/s)
- Receive Frame Rate (f/s)
- Transmit Word Count (bytes)
- Receive Word Count (bytes)
- Transmit Throughput (Mbps)
- Receive Throughput (Mbps)
- Generator Transmit Frame Count (frames)
- Generator Receive Frame Count (frames)
- Mirrored Frames Count (frames)
- Mirrored Tx Frames (frames)
- Mirrored Rx Frames (frames)

6. (Top N and Bottom N monitors only) Select the number of ports to include in a selected measure by entering a number in the **For Top N, Bottom N Monitors, N=** text box.

   Valid values are from 1 through 25. The default is 10.

7. (Top N and Bottom N monitors only) Configure the monitor to show only values greater than or less than a specified value by completing the following steps.

   a. Select the **Show values** check box.

   b. Select **greater than** or **less than** from the list.

   c. Enter a value in the field.

8. (Top N and Bottom N monitors only) Configure threshold numbers and associated colors by completing the following steps.

   Depending on the monitor type you select, you can define up to four threshold numbers in increasing or decreasing order and up to five associated threshold colors.

   (Top N and Bottom N monitors only) The decreasing order defaults are as follows: 90 and above displays red, 75 and above displays orange, 60 and above displays yellow, and all others display blue. The maximum values allowed are -32,768 through 32,767 for SFP power and 0 through 32,767 for all other measures.

   a. (Top N and Bottom N monitors only) Select the check box.

   b. Enter a number in the field.

   c. Click the color square to launch the **Color** dialog box.

      - To pick a color from a swatch, select the **Swatches** tab. Select a color from the display.
      - To specify a color based on hue, saturation, and brightness, click the **HSV** tab. Specify the hue (0 through 360 degrees), saturation (0 through 100%), value (0 through 100%), and transparency (0 through 100%).
      - To specify a color based on hue, saturation, and lightness, click the **HSL** tab. Specify the hue (0 through 360 degrees), saturation (0 through 100%), lightness (0 through 100%), and transparency (0 through 100%).
      - To specify a color based on values of red, green, and blue, click the **RGB** tab. Specify the values for red (0 through 255), green (0 through 255), blue (0 through 255), and alpha (0 through 255).
      - To specify a color based on values of cyan, magenta, yellow, and black, click the **CMYK** tab. Specify the values for cyan (0 through 255), magenta (0 through 255), yellow (0 through 255), black (0 through 255), and alpha (0 through 255).
      - To reset to the default color, click **Reset**.

9. (Time series monitors only) Add targets for the monitor by clicking **Add** and completing the steps in "Adding targets to a traffic flow performance monitor" on page 290.

   Remove targets from the monitor by selecting one or more targets in the **Targets** list and clicking **Remove**.

10. Click **OK** on the **Add Performance Dashboard Monitor** dialog box.

    The **Customize Dashboard** dialog box displays with the new monitor in the **Performance Monitors** list.

11. Click **OK** on the **Customize Dashboard** dialog box.

    The new performance monitors display at the bottom of the dashboard.

Adding targets to a traffic flow performance monitor

You can only add targets for Time Series monitors.

1. Select a user-defined dashboard and click the **Customize Dashboard** icon.

   The **Customize Dashboard** dialog box displays.

2. Click the **Performance** tab.
3. Click **Add**.

   The **Add Performance Dashboard Monitor** dialog box displays.

4. Select **Time Series** from the **Traffic Flows** area.
5. Select the one or more measures for the monitor in the **Measure** area
6. Click **Add** beneath the **Targets** table.

   The **Performance Dashboard Monitor Targets** dialog box displays.

7. Select a fabric from the **Fabric** list.

   Flows defined in the selected fabric display in the **Available Flow** list. Both the **Available Flow** list and the **Selected Flow** list contain the following information:

   - **Sub Flow ID** — The sub flow identifier.
   - **Flow Name** — The name of the flow.
   - **Switch IP Address** — The IP address of the target switch.
   - **Source** — The source device identifier.
   - **Destination** — The destination device identifier.
   - **Feature** — The active feature for the sub flow definition. Valid values include: Generator, Monitor, or Mirror.
   - **LUN** — The LUN values defined in the flow.
   - **Bi-direction** — Whether or not the flow is bi-directional. Valid values are Yes or No.

8. Select **Port Type** from the list. You can monitor flows for VE_Ports alone by selecting **VE Ports** as port type.
9. Select the flow targets from the **Available Flow** list and click the right arrow button to move the targets to the **Selected Flow** list.

   Remove targets from the monitor by selecting one or more targets in the **Selected Flow** list and clicking the left arrow button.

10. Click **OK** on the **Performance Dashboard Monitor Targets** dialog box.

    The targets display in the **Targets** list of the **Add Performance Dashboard Monitor** dialog box.

11. Click **OK** on the **Add Performance Dashboard Monitor** dialog box.

    The **Customize Dashboard** dialog box displays with the new monitor in the **Performance Monitors** list.

12. Click **OK** on the **Customize Dashboard** dialog box.

    The new performance monitors display at the bottom of the dashboard.

View Management

SAN tab overview

The **SAN** tab (Figure 145) displays the Product List, Topology Map, Master Log, Utilization Legend, and Minimap.

> NOTE
> When you launch the Management application or navigate to a new view, the **SAN** tab displays with a gray screen over the Product List and Topology Map while data is loading.

You can change the default size of the display by placing the cursor on the divider until a double arrow displays. Click and drag the adjoining divider to resize the window. You can also show or hide an area by clicking the left or right arrow on the divider.

> NOTE
> Some areas may be hidden by default. To view areas of the **SAN** tab, select **View > Show Panels > All Panels**, or press **F12**.

FIGURE 145 Main window - SAN tab

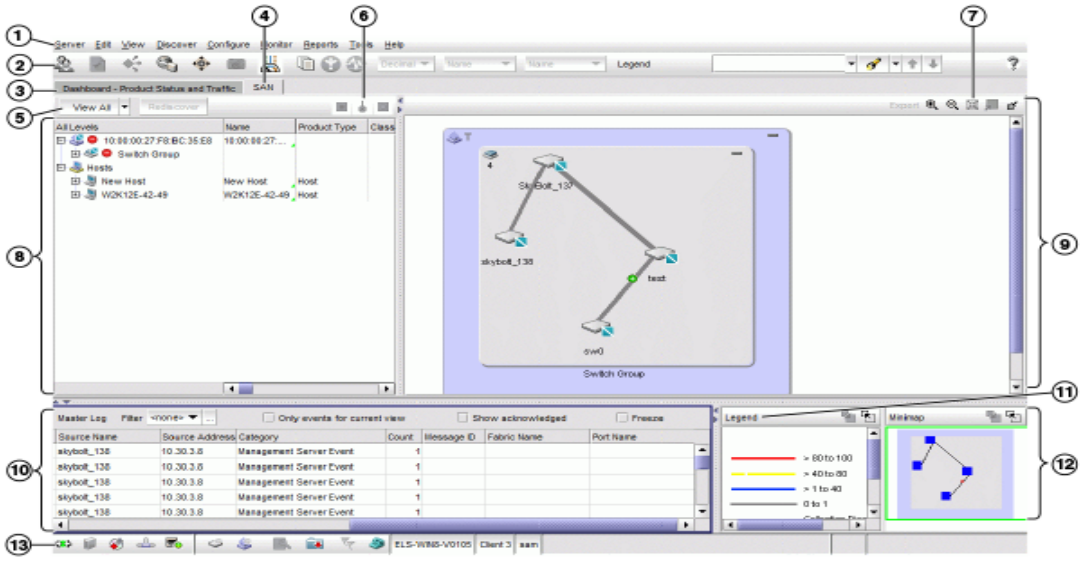

1. **Menu bar —** Lists commands you can perform on the **SAN** tab. Some menu items display as disabled unless you select the correct object from the product list or topology map. For a list of the many functions available on each menu, refer to "SAN main menus" on page 1295.
2. **SAN main toolbar —** Provides buttons that enable quick access to dialog boxes and functions. For a list of available commands, refer to "SAN main toolbar" on page 294.
3. **Dashboard tab —** Provides a high-level overview of the network managed by the Management application server. For more information, refer to "Dashboard Management" on page 199.
4. **SAN tab —** Displays the Master Log, Minimap, Connectivity Map (topology), and Product List.
5. **View All list —** Enables you to create, copy, or edit a view, select how to view the Product List (All Levels, Products and Ports, Products Only, or Ports Only), and select which view you want to display in the main window. For more information, refer to "View All list" on page 295. For step-by-step instruction about creating a view, refer to "Creating a customized view" on page 316.
6. **Port Display buttons** — Provides buttons that enable quick access to configuring how ports display. Not enabled until you discover a fabric or host. For more information, refer to "Port Display buttons" on page 296.
7. **Connectivity Map toolbar —** Provides tools for viewing the Connectivity Map as well as exporting the Connectivity Map as an image. Does not display until you discover a fabric. For more information, refer to "Connectivity Map toolbar" on page 296.
8. **Product List —** Lists the devices discovered in the Management application. For more detailed information, refer to "Product List" on page 297.
9. **Connectivity Map —** Displays the topology, including discovered and monitored devices and connections. For more information, refer to "Connectivity Map" on page 298.
10. **Master Log —** Displays all events that have occurred on the Management application. For more information, refer to "Master Log" on page 300.
11. **Utilization Legend —** (Trial and Licensed version only) Indicates the percentage ranges represented by the colored, dashed lines on the Connectivity Map. Only displays when you select **Monitor > Performance > View Utilization** or click the **Utilization** icon on the toolbar. For more information, refer to "Utilization Legend" on page 299.
12. **Minimap —** Displays a "bird's-eye" view of the entire topology. Does not display until you discover a fabric. For more information, refer to "Minimap" on page 301.
13. **Status bar —** Displays the connection, port, product, fabric, special event, call home, and backup status, as well as Server and User data. For more information, refer to "Status bar" on page 302.

SAN main toolbar

The toolbar is located beneath the Menu bar and provides icons to perform various functions.

FIGURE 146SAN main toolbar

The icons on your toolbar vary based on the licensed features on your system.

1. **Users —** Displays the **Users** dialog box. Use to configure users, user groups, and permissions.
2. **Properties —** Displays the **Properties** dialog box of the selected device or fabric. Use to view or edit device or fabric properties.
3. **Launch Element Manager —** Launches the Element Manager of the selected device. Use to configure a device through its Element Manager.

4. **Fabric discovery —** Displays the **Discover Fabrics** dialog box. Use to configure discovery.
5. **Zoning —** Displays the **Zoning** dialog box. Use to configure zoning.
6. **Track Fabric Changes —** Select to turn track fabric changes on or off for the selected device or group.
7. **View Utilization —** Displays or hides the utilization legend.
8. **View Report —** Displays the **View Reports** dialog box. Use to view available reports.
9. **Flow Vision** — Displays the **Flow Vision** dialog box. Use to configure Flow Vision.
10. **MAPS —** Displays the **MAPS** dialog box. Use to configure MAPS.
11. **Domain ID/Port # —** Use to set the domain ID or port number to display as decimal or hex in the Product List.
12. **Product Label —** Use to set the product label for the devices in the Connectivity Map and Product List.
13. **Port Label —** Use to set the port label for the devices in the Connectivity Map and Product List.
14. **Legend —** Use to view the topology legend. For more information, refer to "SAN product icons" on page 303.
15. **Product List Search —** Use to search for a device in the Product List. For detailed instructions, refer to "Search" on page 313.
16. **Help —** Displays the Online Help.

View All list

The **View All** list is located at the top-left side of the window and enables you to create, copy, or edit a view, select how to view the Product list (All Levels, Products and Ports, Products Only, or Ports Only), and select which view you want to display in the main window. The **View All** list does not display until you discover a fabric. To discover a fabric, refer to "Discovering fabrics" on page 35.

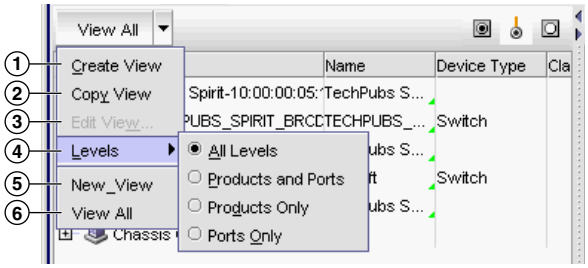

1. **Create View —** Select to create a new view.
2. **Copy View —** Select to copy an existing view.
3. **Edit View —** Select to edit an existing view.
4. **Levels —** Select the level at which you want to view the Product list, Options include: All Levels, Products and Ports, Products Only, or Ports Only.
5. ***View_Name* —** Any additional views that you create. Select which view you want to display in the main window.
6. **View All —** Select to display the default view of the main window.

Port Display buttons

The **Port Display** buttons are located at the top right of the Product List and enable you to configure how ports display. You have the option of viewing connected (or occupied) product ports, unoccupied product ports, or attached ports. Not enabled until you discover a fabric or host.

> NOTE
> Occupied/connected ports are those that originate from a device, such as a switch. Attached ports are ports of the target devices that are connected to the originating device.

FIGURE 147 Port Display buttons

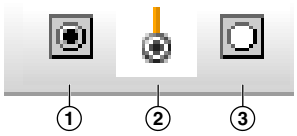

1. **Show/Hide Occupied Port** — Displays or hides the ports of the devices in the fabrics (present in the connectivity map) that are connected to other devices.
2. **Show/Hide Attached Port** — Displays or hides the attached ports of the target devices.
3. **Show/Hide Unoccupied Port** — Displays or hides the ports of the devices (shown in the connectivity map) that are not connected to any other device.

Connectivity Map toolbar

The Connectivity Map toolbar is located at the top right side of the **View** window and provides tools to export the topology, to zoom in and out of the Connectivity Map, collapse and expand groups, and fit the topology to the window. Not enabled until you discover a fabric.

FIGURE 148 The Connectivity Map toolbar

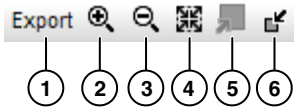

1. **Export —** Use to export the topology to a PNG file.
2. **Zoom In —** Use to zoom in on the Connectivity Map.
3. **Zoom Out —** Use to zoom out on the Connectivity Map.
4. **Fit in View —** Use to scale the map to fit within the Connectivity Map area.
5. **Expand —** Use to expand the map to show all ports in use on a device.
6. **Collapse —** Use to collapse the map to show only devices (hides ports).

Product List

The Product List, located on the **SAN** tab, displays an inventory of all discovered devices and ports. The Product List is a quick way to look up product and port information, including serial numbers and IP addresses.

To display the Product List, select **View > Show Panels > Product List** or press **F9**.

Note that the Product List can only display up to 9000 ports at a time. If you expand a fabric with enough ports to exceed the 9000 port limit, the Management application automatically collapses the least recently expanded fabrics until the port count is 9000 ports or less. Automatically collapsed fabrics do not display the + icon, which allows you to quickly determine which fabrics are manually or automatically collapsed. To expand an automatically collapsed fabric, right-click the fabric and select **Expand**. Also note that only manually collapsed fabrics expand when you use the **Expand All** command. **Expand All** does not affect automatically collapsed fabrics.

You can edit information in the Product List by double-clicking in a field marked with a green triangle. You can sort the Product List by clicking a column heading.

The following columns (presented here in alphabetical order) are included in the Product List.

- **Additional Port Info —** Displays additional port information.
- **All Levels —** Displays all discovered fabrics, groups, devices, and ports as both text and icons. Also, displays the status of the fabrics, groups, devices, and ports. For a list of icons that display in the **All Levels** column, refer to "Icon legend" on page 303.
- **Additional Port Info —** Displays additional information about the port.
- **Attached Port # —** Displays the number of the attached port.
- **BB Credit —** Displays the BB Credit of the port.
- **Class —** Displays the class value of the FICON device port.
- **Contact —** Displays the name of the person or group you should contact about the product. This field is editable at the fabric level.
- **Description —** Displays the description of the product. This field is editable at the fabric level.
- **Product Type —** Displays the type of product.
- **Domain ID —** Displays the Domain ID for the product in the format xx(yy), where xx is the normalized value and yy is the actual value on the wire.
- **FC Address —** Displays the Fibre Channel address of the port.
- **Firmware —** Displays the firmware version of the product.
- **IP Address —** Displays the IP address (IPv4 or IPv6 format) of the product.
- **Location —** Displays the physical location of the product. This field is editable at the fabric level.
- **Model —** Displays the model number of the product.
- **Name —** Displays the name of the product or port. This field is editable at the fabric, device, and port level.
- **Port # —** Displays the number of the port.
- **Port Count —** Displays the number of ports on the product.
- **Port Type —** Displays the type of port (for example, expansion port, node port, or NL_port).
- **Protocol —** Displays the protocol for the port.
- **Serial # —** Displays the serial number of the product.
- **Speed Configured (Gbps) —** Displays the actual speed of the port in Gigabits per second.
- **State —** Displays the state for the product and the port.
- **Status —** Displays the status for the product and the port.

- **Symbolic Name —** Displays the symbolic name for the port.
- **TAG —** Displays the tag number of the product.
- **Vendor —** Displays the name of the product's vendor.
- **WWN —** Displays the world wide name of the product or port.
- **Zone Alias —** Displays the zone alias of the product or port.
- **User-defined property labels** — Displays the user-defined property labels. You can create up to three user-defined property labels.

Product List functions

- Customize — Customize the Product list. For more information, refer to "Product List customization" on page 311.
- Sort — Click a column head to sort the list. Click a column head again to reverse the sort order.
- Two-way selection — Select a device in the Product List and that device is highlighted on the Topology Map and vice versa.
- Table shortcut menus — Right-click a column header in the Product List to view the menu. For a list of right-click menus, refer to "Customizing application tables" on page 308.

Connectivity Map

The Connectivity Map, which displays in the upper right area of the main widow, is a grouped map that shows physical and logical connectivity of SAN components, including discovered and monitored devices and connections. These components display as icons in the Connectivity Map. For a list of icons that display in the Connectivity Map, refer to the following tables:

- "SAN product icons" on page 303
- "SAN group icons" on page 304
- "Event icons" on page 306

FIGURE 149 Connectivity Map

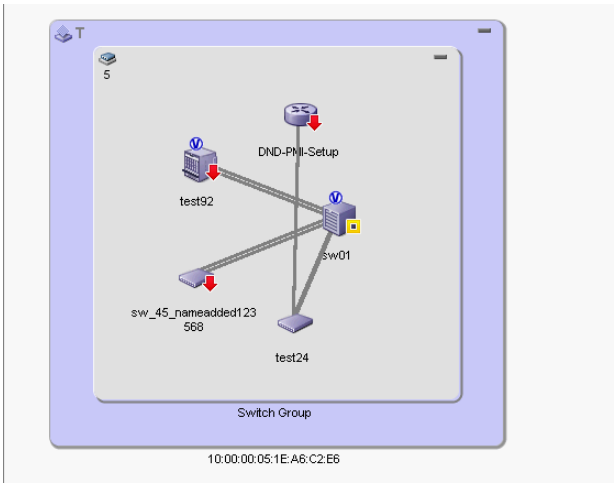

The Management application displays all discovered fabrics in the Connectivity Map by default. To display a discovered Host in the Connectivity Map, you must select the Host in the Product List. You can only view one Host and physical and logical connections at a time.

Connectivity Map functions

- Two-way selection — When you select an icon on the Topology Map, that device is highlighted in the Product List and vice versa.
- Device double-click — Double-click a device to launch Web Tools for the selected device.
- Zoom In/Zoom Out — Click the appropriate button to zoom in or out on the Topology Map.
- Tool tips — Mouse over a device or connection to view information.
- Right-click menus — Right-click a device to view the menu. For a list of right-click menus, refer to "SAN shortcut menus" on page 1304.

Utilization Legend

The Utilization Legend, which displays in the lower right corner of the main window, indicates the percentage ranges represented by the colored, dashed lines on the Connectivity Map. It only displays when you select **Monitor > Performance > View Utilization** or click the **Utilization** icon on the toolbar.

FIGURE 150 Utilization Legend

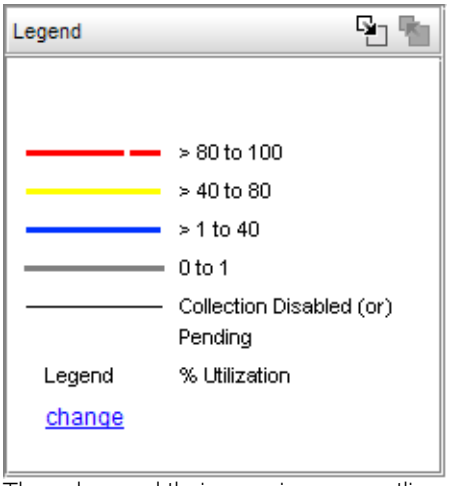

The colors and their meanings are outlined in the following table.

TABLE 23

Line Color	Utilization Defaults
Red line	80% to 100% utilization
Yellow line	40% to 80% utilization
Blue line	1% to 40% utilization
Gray line	0% to 1% utilization
Black line	Utilization disabled

For more information about the utilization legend, refer to "SAN connection utilization" on page 1000.

Master Log

The Master Log, which displays in the lower area of the main window, lists the events and alerts that have occurred on the SAN. If you do not see the Master Log, select **View > Show Panels > All Panels** or press **F5**.

The default order of the Master Log columns is 'Severity', 'Acknowledged', 'Last Event Server Time', and 'Description'. Which columns are displayed and in what order can be controlled through the "Customize Columns" dialog, as described in "Displaying columns" and in "Changing the order of columns". You can sort the Master Log by clicking a column heading. By default, the Master Log is sorted by the **Last Event Server Time** column. To filter information in the Master Log, refer to "Filtering events in the Master Log" on page 1195. To view event properties, refer to "Displaying event properties from the Master Log" on page 1193.

The following fields and columns are included in the Master Log:

- **Severity —** The severity of the event. When the same event (Warning or Error) occurs repeatedly, the Management application automatically eliminates the additional occurrences. For more information about events, refer to "Fault Management" on page 1131. For a list of the event icons, refer to "Event icons" on page 306.

  NOTE
  The mapsConfigRuleName and mapsQuietTimeExpirationTrap MAPS events only display the severity as Warning even when you customize the severity for the MAPS rule. You must use the MAPS Violations dialog boxes to verify the event severity.

- **Acknowledged —** Whether the event is acknowledged or not. Select the check box to acknowledge the event.
- **Source Name —** The product on which the event occurred.
- **Source Address —** The IP address (IPv4 or IPv6 format) of the product on which the event occurred.
- **Origin —** The event source type (for example trap, pseudo event, application, or syslog).
- **Category —** The type of event that occurred (for example, client/server communication events).
- **Description —** A description of the event.
- **Last Event Server Time —** The time and date the event last occurred on the server.
- **Count —** The number of times the event occurred.
- **Module Name —** The name of the module on which the event occurred.
- **Message ID —** The message ID of the event.
- **Product Address —** The IP address of the product on which the event originated.
- **Contributor —** The name of the contributor on which the event occurred.
- **Node WWN —** The world wide name of the node on which the event occurred.
- **Fabric Name —** The name of the fabric on which the event occurred.
- **Operational Status —** The operational status (such as, unknown, healthy, marginal, or down) of the product on which the event occurred.
- **First Event Product Time —** The time and date the event first occurred on the product.
- **Last Event Product Time —** The time and date the event last occurred on the product.
- **First Event Server Time —** The time and date the event first occurred on the server.
- **Audit —** The audit of the event.
- **Virtual Fabric ID —** The VFID of the product on which the event occurred.
- **Zone Alias** — Displays the zone alias of the product or port.

- **Remarks** — Displays the administrator name and the client IP address of the original initiator who triggered the switch support save.

Minimap

The **Minimap**, which displays in the lower right corner of the main window, is useful for getting a bird's-eye view of the topology, or to quickly jump to a specific place on the topology. To jump to a specific location on the topology, click that area on the Minimap. A close-up view of the selected location displays on the topology.

Use the Minimap to view the entire topology and to navigate more detailed map views. This feature is especially useful if you have a large topology. Does not display until you discover a device.

FIGURE 151SAN Minimap

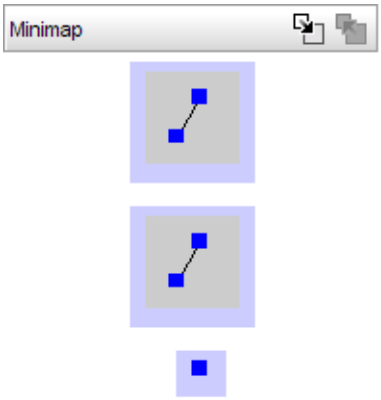

Anchoring or floating the Minimap

You can anchor or float the Minimap to customize your main window.

- To float the Minimap and view it in a separate window, click the **Detach** icon ( ) in the upper right corner of the Minimap.
- To anchor the Minimap and return the Minimap to its original location on the main window, do one of the following steps:
  - Click the **Attach** icon ( ) in the upper right corner of the Minimap.
  - Click the **Close** icon ( ) in the upper right corner of the Minimap.
  - Double-click the logo in the upper left corner of the Minimap.
  - Click the logo in the upper left corner of the Minimap and select **Close** (**ALT** + **F4**).

Resizing the Minimap

On an anchored Minimap, place the cursor on the left border of the Minimap until a double-pointed arrow displays. Click and drag the adjoining divider.

On a floating Minimap, place the cursor on a border of the Minimap until a double-pointed arrow displays. Click and drag to change the window size.

Status bar

The status bar displays at the bottom of the main window. The status bar provides a variety of information about the SAN and the application. The icons on the status bar change to reflect different information, such as the current status of products, fabrics, and backup.

FIGURE 152Status Bar

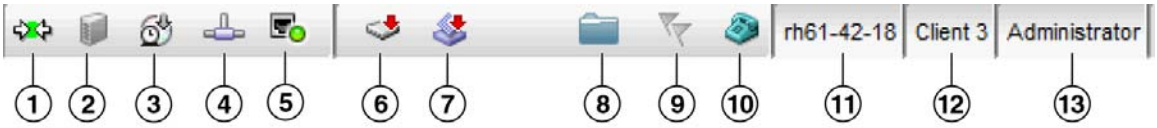

The icons on your status bar will vary based on the licensed features on your system.

1. **Connection Status —** Displays the Server-Client connection status. Also displays whether the client topology is in sync with the server. Resynchronize with the server by restarting the client.
2. **Server Status** — Displays the status of the server disk space (for example, low or sufficient).
3. **Server Backup Status —** Displays a backup status icon, which allows you to determine the current backup status. Right-click and select **Backup now** to begin back up immediately. Right-click and select **Configure backup** to launch the **Options** dialog box - **Server Backup** pane and configure backup. Let the pointer pause on the backup status icon to display the following information in a tooltip.
   - **Backup in Progress icon —** Backup started at hh:mm:ss, in progress... *XX* files in *Directory_Name* are backed up.
   - **Countdown to Next Scheduled Backup icon —** Waiting for next backup to start.
   - **Backup Disabled icon —** Backup is disabled.
   - **Backup Failed icon —** Backup failed at hh:mm:ss mm/dd/yyyy.
4. **Network Size Status —** Displays a memory allocation status icon, which allows you to determine the current network size status. Double-click the icon to launch the **Memory Allocation** pane of the **Options** dialog box. Let the pointer pause on the backup status icon to display the following information in a tooltip.
   - **Network size within limits icon —** Network size is within the recommended count.
   - **Ne work size exceeds limits icon —** Network size exceeds the recommended count.
5. **Server Port Status —** Displays port status for the following ports: CIM Indication for Event Handling, CIM Indication for HCM Proxy, FTP, SCP/SFTP, SNMP Trap, Syslog, , Web Server (HTTP), and Web Server (HTTPS). Click to launch the **Port Status** dialog box. For more information about port status, refer to "Viewing port status" on page 12.
6. **Product Status —** Displays the status of the most degraded device in the SAN. For example, if all devices are operational except one (which is degraded), the Product Status displays as degraded. Click this icon to open the **Product Status Log**.
7. **Fabric Status —** Displays the state of the fabric that is least operational, based on ISL status. The possible states are: operational, unknown, degraded or failed. Select a product or fabric from the Connectivity Map or Product List and click this icon to open the related **Fabric Log** (only available for persisted fabrics).
8. **Configuration Policy Manager Status —** Displays whether or not a policy manager has failed or partially failed. Click to launch the **Configuration Policy Manager** dialog box. For more information about configuration policy manager, refer to "Viewing configuration policy manager status" on page 1115.
9. **Special Events** — Displays whether or not a special event has been triggered. Click to launch the **Special Events** dialog box. For more information about special events, refer to "Creating an event action definition" on page 1154.
10. **Call-Home Status —** (Trial and Licensed version only) Displays a call home status icon when one or more product are discovered, which allows you to determine the current call home status. Click to launch the **Call Home Notification** dialog box. For more information about Call Home status and icons, refer to "Viewing Call Home status" on page 339.

11. **Server Name —** Displays the name of the Server to which you are connected. Click to launch the **Server Properties** dialog box. For more information, refer to "Viewing server properties" on page 11.
12. **Total Users —** Displays the number of clients logged into the server. Click to launch the **Active Sessions** dialog box. For more information, refer to "Viewing active sessions" on page 10.
13. **User's ID —** Displays the user ID of the logged in user. Click to launch the **User Profile** dialog box. For more information, refer to "User profiles" on page 153.
14. **Trial license** (Not shown) — Displays the trial expiration information to the right of the User's ID.

Icon legend

Various icons are used to illustrate devices and connections in a network. The following tables list icons that display on the Connectivity Map and Product List.

SAN product icons

The following table lists the manageable SAN product icons that display on the topology. Fabric OS manageable devices display with blue icons. Unmanageable devices display with gray icons. Some of the icons shown only display when certain features are licensed.

TABLE 24

Icon	Description	Icon	Description
	Fabric		Fabric OS Switch and Blade Switch
	Fabric OS Director (Vertical blades)		Fabric OS Director (Horizontal blades)
	Fabric OS 32 Gbps Backbone Chassis (Vertical blades)		Fabric OS 32 Gbps Backbone Chassis (Horizontal blades)
	Fabric OS DCB Switch		Fabric OS 64-port, 32 Gbps switch
	Fabric OS Router		Storage
	Fabric OS FC Switch in Access Gateway mode (single-fabric connected)		Fabric OS FC Switch in Access Gateway mode (multiple-fabric connected)
	Fabric OS DCB Switch in Access Gateway mode (single-fabric connected)		Fabric OS DCB Switch in Access Gateway mode (multiple-fabric connected)
	iSCSI Target		iSCSI Initiator

Host product icons

The following table lists the manageable Host product icons that display on the topology. Fabric OS manageable devices display with blue icons. Unmanageable devices display with gray icons. Some of the icons shown only display when certain features are licensed.

TABLE 25

Icon	Description	Icon	Description
	HBA		HBA Mezzanine Card
	CNA		CNA Mezzanine Card
	Unmanaged HBA		Any IO
	Host		Unmanaged Host
	VM Host		Virtual HBA
	Ethernet Cloud		Layer 2 Cloud

SAN group icons

The following table lists the manageable SAN product group icons that display on the topology.

TABLE 26

Icon	Description	Icon	Description
	Switch Group		Host Group
	Storage Group		Unknown Fabric Group
	Unmanaged Fabric Group		Chassis Group

Host group icons

The following table lists the manageable Host product group icons that display on the topology.

TABLE 27

Icon	Description	Icon	Description
	Host Group		

SAN port icons

The following table lists the SAN port icons that display in the Product List.

TABLE 28

Icon	Description
	Occupied FC Port
	Unoccupied FC Port
	Attached FC Port
	Trunk (port group)
	IP and 10 GE Port
	Attached IP and 10 GE Port
	Attached-to-Cloud 10 GE Port
	Virtual Port
	Virtual FCoE Port
	Attached FCoE Port
	Pre-boot Virtual Port
	Virtual Attached Port
	Mirror Port
	Bottleneck Port
	Analytics E_Port (AE_Port) or Analytics F_Port (AF_Port)

SAN product status icons

The following table lists the product status icons that display on the topology.

TABLE 29

Icon	Status
No icon	Healthy/Operational
	Attention
	Bottleneck

TABLE 29

Icon	Status
	Degraded/Marginal
	Device Added
	Device Removed/Missing
	Down/Failed
	Routed In
	Routed Out
	Unknown/Link Down
	Unreachable

Event icons

The following table lists the event icons that display on the topology and Master Log. For more information about events, refer to "Fault Management" on page 1131.

TABLE 30

Event Icon	Description
	Emergency
	Alert
	Critical
	Error
	Warning
	Notice
	Informational
	Debug

Customizing the main window

You can customize the main window to display only the data you need by displaying different levels of detail on the Connectivity Map (topology) or Product List.

Zooming in and out of the Connectivity Map

You can zoom in or out of the Connectivity Map to see products and ports.

Zooming in

To zoom in on the Connectivity Map, use one of the following methods:

- Click the zoom-in icon ( ) on the Connectivity Map toolbar.
- Press CTRL and the plus sign on the number pad on the keyboard.
- Use the **Zoom** dialog box.
    a. Select **View > Zoom**.
    The **Zoom** dialog box displays.

    FIGURE 153 Zoom dialog box

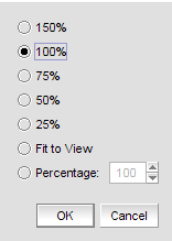

    b. Select a zoom percentage.
    c. Click **OK** to save your changes and close the **Zoom** dialog box.

Zooming out

To zoom out of the Connectivity Map, use one of the following methods:

- Click the zoom-out icon ( ) on the Connectivity Map toolbar.
- Press CTRL and the minus sign on the number pad on the keyboard.
- Use the **Zoom** dialog box.
    a. Select **View > Zoom**.
    The **Zoom** dialog box displays.
    b. Select a zoom percentage.
    c. Click **OK** to save your changes and close the **Zoom** dialog box.

Showing levels of detail on the Connectivity Map

You can configure different levels of detail on the Connectivity Map, making device management easier.

Viewing fabrics

To view only fabrics, without seeing groups, products, or ports, select **View > Show> Fabrics Only**.

Viewing groups

To view only groups and fabrics, without seeing products, or ports, select **View > Show> Groups Only**.

Viewing products

To view products, groups, and fabrics, select **View > Show> All Products**.

Viewing ports

To view all ports, select **View > Show> All Ports**.

Exporting the topology

You can save the topology to an image (PNG format).

1. Click **Export** in the toolbar.

   The **Export Topology To PNG File** dialog box displays.

2. Browse to the directory where you want to export the image.
3. Edit the name in the **File Name** field, if necessary.
4. Click **Save**.

   If the file name is a duplicate, a message displays. Click **Yes** to replace the image or click **No** to go back to the **Export Topology To PNG File** dialog box and change the file name.

   The **File Download** dialog box displays.

5. Click **Open** to view the image or click **Cancel** to close the dialog box.

Customizing application tables

You can customize any table in the Management application main interface (for example, the Master Log or the Product List) or in individual dialog boxes in the following ways:

- Display only specific columns
- Display columns in a specific order
- Resize the columns to fit the contents
- Sort the table by a specific column or multiple columns
- Copy information from the table to another application
- Export information from the table
- Search for information

- Expand the table to view all information
- Collapse the table

Displaying columns

To only display specific columns, complete the following steps.

1. Right-click anywhere in the table and select **Customize** or **Table > Customize**.

   The **Customize Columns** dialog box displays.

   FIGURE 154 Customize Columns dialog box

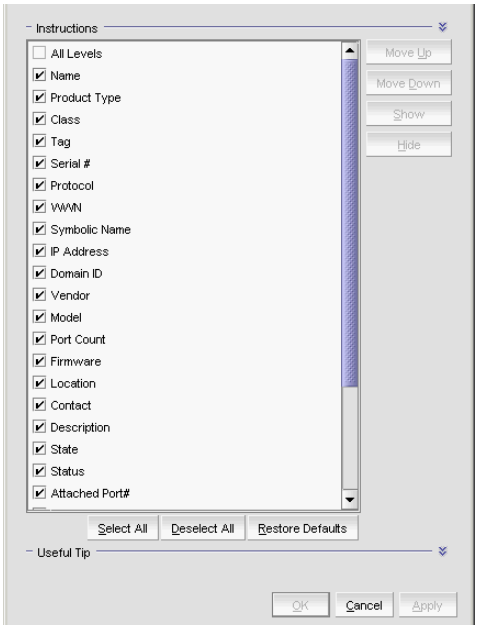

2. Choose from the following options:

   - Select the check box to display a column.
     OR
     Select the column name and click **Show**.
   - Clear the check box to hide a column.
     OR
     Select the column name and click **Hide**.
   - Click **Select All** to select all check boxes.
   - Click **Deselect All** to clear all check boxes.
   - Click **Restore Defaults** to restore the table to the original settings.

3. Click **OK**.

Changing the order of columns

To change the order in which columns display, choose from one of the following options.

Rearrange columns in a table by dragging and dropping the column to a new location.

OR

1. Right-click anywhere in the table and select **Customize** or **Table > Customize**.

   The **Customize Columns** dialog box displays.

2. Select the name of the column you want to move and use the **Move Up** button and **Move Down** button to move it to a new location.
3. Click **OK**.

Resizing the columns

You can resize a single column or all columns in the table.

To resize a single column, right-click the column header and select **Size Column to Fit** or **Table > Size Column to Fit**.

To resize all columns in the table, right-click anywhere in the table and select **Size All Columns to Fit** or **Table > Size All Columns to Fit**.

Sorting table information

To sort the table by a single column, click the column header.

To reverse the sort order, click the column header again.

To sort the table by multiple columns, complete the following steps.

1. Click the primary column header.
2. Press CTRL and click a secondary column header.

Copying table information

You can copy the entire table or a specific row to another application (such as Notepad, Excel, Word, and so on).

1. Choose from one of the following options:
   - Right-click anywhere in the table and select **Table > Copy Table**.
   - Select the table row that you want to export and select **Table > Copy Row.**
2. Open the application to which you want to copy the Product List information.
3. Select **Edit > Paste** (or press CTRL + V).
4. Save the file.

Exporting table information

You can export the entire table or a specific row to a text file.

1. Choose from one of the following options:
   - Right-click anywhere in the table and select **Table > Export Table**.
   - Select the table row that you want to export and select **Table > Export Row.**

The **Save table to a tab delimited file** dialog box displays.

2. Browse to the location where you want to save the file.
3. Enter the file name in the **File Name** field.
4. Click **Save**.

Searching for information in a table

You can search for information in the table by any of the values found in the table.

1. Right-click anywhere in the table and select **Table > Search**.

   The focus moves to the Search field.

   FIGURE 155 Search field

2. Enter all or part of the search text in the Search field and press **Enter**.

   The first instance is highlighted in the table.

3. Press **Enter** to go to the next instance of the search text.

Expanding and collapsing tables

You can expand a table to display all information or collapse it to show only the top level.

To expand the entire table, right-click anywhere in the table and select **Expand All** or **Table > Expand All**.

To collapse the entire table, right-click anywhere in the table and select **Collapse All** or **Table > Collapse All**.

Product List customization

NOTE
Properties customization requires read and write permissions to the Properties - Add / Delete Columns privilege.

You can customize the Product List by creating user-defined fabric, product, and port property labels. You can also edit or delete user-defined property labels, as needed.

You can create up to three user-defined property labels from the Product List for each of the following object types: fabric, product, and port properties. Product and fabric property labels created from the Product List display in the Product List and the **Properties** dialog box. Port property labels created from the Product List display in the Product List and the **Properties** dialog box. User-defined properties must be unique across all **Properties** dialog boxes and the Product List.

Property fields containing a green triangle ( ) in the lower right corner are editable. To edit a field with a green triangle ( ), click in the field and make your changes.

Adding a property label

You can create up to three user-defined fabric, product, and port property labels from the Product List. To add a new property label (column heading), complete the following steps.

1. Right-click any column heading on the Product List and select **Add Column**.

   The **Add Property** dialog box displays.

2. Enter a label and description for the property.

   The label must be unique and can be up to 30 characters.

   The description can be up to 126 characters.

3. Select the property type from the **Type** list.

   Options include: Fabric, Product, or Port.

4. Click **OK**.

   The new property displays in the last column of the Product List as well as the associated Properties dialog box based on the selected type.

   Property fields containing a green triangle ( ) in the lower right corner are editable. To edit a field with a green triangle ( ), click in the field and make your changes.

Editing a property label

You can only edit labels that you create on the Product List.

To edit a user-defined property label (column heading), complete the following steps.

1. Right-click the column heading on the Product List for the property you want to edit and select **Edit Column**.

   The **Edit Property** dialog box displays.

2. Change the label and description for the property, as needed.

   The label must be unique and can be up to 30 characters.

   The description can be up to 126 characters.

   You cannot change the property type.

3. Click **OK**.

   The property details are updated in the Product List as well as the Properties dialog box.

   Property fields containing a green triangle ( ) in the lower right corner are editable. To edit a field with a green triangle ( ), click in the field and make your changes.

Deleting a property label

You can only delete labels that you created on the Product List. To delete a label, complete the following steps.

1. Right-click the user-defined column heading on the Product List you want to delete and select **Delete Column**.
2. Click **Yes** on the confirmation message.

   The column you selected is deleted from the Product List as well as the Properties dialog box.

Search

You can search for a objects by text or regular expression.

- **Text** — Enter a text string in the search text box. This search is case sensitive.

  For example, if you are searching for a device in the Product List, you can enter the first five characters in a device name. All products in the Product List that contain the search text display highlighted.
- **Regular Expression** — Enter a Unicode regular expression in the search text box. (For hints, refer to "Regular Expressions" on page 1375.) All products in the Product List that contain the search text display highlighted. This search is case insensitive.

  For example, you might need to search ports. To search for a port using a Unicode regular expressions, enter "2/1|2/2|2/3". This search will find Ports 2/1, 2/2, and 2/3 on all devices.

The Search features contains a number of components. The following graphic illustrates the various areas, and descriptions of them are listed below.

1. Text field — Enter the text or unicode regular expression for which you want to search.
2. Search list — Select one of the following options:
   - **Text** option — Select this option if you entered a text string in the text field.
   - **Regular Expression** option — Select this option if you entered a unicode regular expression in the text field.
   - **Clear Search** command — Select this option to clear the search text field
   - **Help** command — Select this option to view help for this feature.
3. Search up button — Click to search upward in the list.
4. Search down button — Click to search downward in the list.

Searching for a device

You can search for a device by name, WWN, or device type. When searching in the Connectivity Map, make sure you search the right view (**View > Manage View > Display View >** *View_Name*) with the appropriate options of port display (**View > Port Display >** *Display_Option*) and connected end devices (**View > Port Display > Show All**) enabled.

NOTE
Search does not search automatically collapsed fabrics. You must expand the fabric (right-click and select **Expand**) and repeat the search.

To search for a device, complete the following steps.

1. Enter your search criteria in the search field.

   NOTE
   To search for a device, the device must be discovered and display in the topology.

2. Choose one of the following options:

   - Select **Text** from the search list and enter a text string in the search text box.

     This search is case sensitive.

   - Select **Regular Expression** from the search list and enter a Unicode regular expression in the search text box.

     This search is case insensitive

3. Press **Enter** or click the search icon.

   The search results display highlighted.

   If the search finds more than one match, a message displays, advising you to restrict the search by restricting the search by node (refer to "Restricting a search by node" on page 314) or by looking for exact matches (refer to "Searching for an exact match" on page 315).

Restricting a search by node

When a device is assigned to a product group, it may be listed in the Product node, as well as Product Groups node. Therefore the search results include the device under both the Product node and the Product Group node.

NOTE
Search does not search automatically collapsed fabrics. You must expand the fabric (right-click and select **Expand**) and repeat the search.

NOTE
To search for a device, the device must be discovered and display in the topology.

To restrict the search only to specific nodes, complete the following steps.

1. Select the Product node or Product Group node that you want to search.
2. Choose one of the following options:
   - Select **Text** from the search list.
   - Select **Regular Expression** from the search list.
3. Enter your search criteria in the search field.

- **Text** — Enter a text string in the search text box. This search is case sensitive.

  For example, you can enter the first five characters in a device name. All products in the Product List that contain the search text display highlighted.

- **Regular Expression** — Enter a Unicode regular expression in the search text box. (For hints, refer to "Regular Expressions" on page 1375.) All products in the Product List that contain the search text display highlighted. This search is case insensitive.

4. Press **Enter** or click the search icon.

   The search results display highlighted.

Searching for an exact match

NOTE
Search does not search automatically collapsed fabrics. You must expand the fabric (right-click and select **Expand**) and repeat the search.

To search for an exact match, complete the following steps.

1. Choose one of the following options:
   - Select **Text** from the search list.
   - Select **Regular Expression** from the search list.
2. Enter your search criteria in the search field.
   - **Text** — Enter a text string in the search text box. This search is case sensitive.

     For example, you can enter the first five characters in a device name. All products in the Product List that contain the search text display highlighted.

   - **Regular Expression** — Enter a Unicode regular expression in the search text box. (For hints, refer to "Regular Expressions" on page 1375.) All products in the Product List that contain the search text display highlighted. This search is case insensitive.
3. Press **Ctrl** and click the search icon.

   The search results display highlighted.

   Example

   If you search for IP address "192.1.1.101" and then press CTRL and click the search icon, the application only highlights "192.1.1.101". This search does not highlight "SI-101 [192.1.1.101]".

   If you search for port "1/2" and then press CTRL and click the search icon, the application only highlights port "1/2". This search does not highlight ports "1/2", "1/20", "1/21", "1/22", and so forth.

Clearing search results

To clear search results, select **Clear Search** from the search list.

SAN view management overview

You can customize the topology by creating views that include certain fabrics or devices and then switch between the views to see specific information about those fabrics or devices.

If you discover or import a network with more than approximately 2,000 devices, the devices display on the Product List, but not on the Topology Map. Instead, the Topology Map shows a message stating that the topology cannot be displayed. To resolve this issue, create a new view to filter the number of devices being discovered. Refer to "Creating a customized view" on page 316 for instructions.

Creating a customized view

You may want to customize the Product List and Connectivity Map to simplify management of large SANs by limiting the topology size or Product List columns.

For each customized view, you can specify the fabrics and hosts that display on the Connectivity Map, as well as the columns and device groupings that display on the Product List.

Customized view settings reside on the server. Only users with the same login to the same server can see and select the view settings. No individual user can have access to the views created by another user.

If you select a customized view and new devices are discovered, those new devices display in the customized view if they belong in that view category or fabric.

1. Select **View > Manage View > Create View**.

   The **Create View** dialog box displays.

   FIGURE 156 Create View dialog box - Fabrics tab

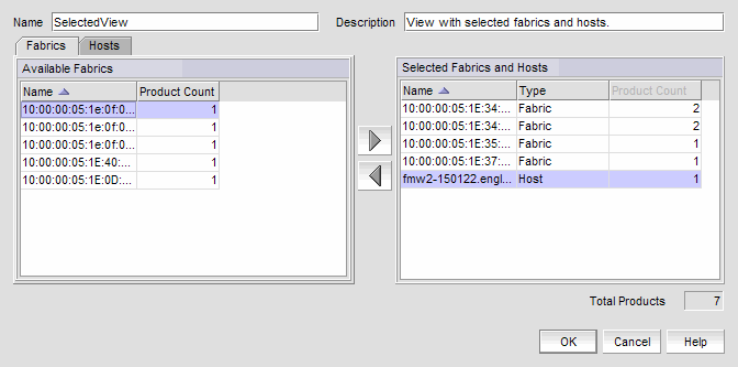

2. Enter a name (128-character maximum) in the **Name** field and a description (126-character maximum) in the **Description** field for the view.

   NOTE
   You cannot use the name "View" or "View All" in the **Name** field.

   NOTE
   You cannot use an existing name in the **Name** field.

3. Click the **Fabrics** tab.
4. In the **Available Fabrics** table, select the fabrics you want to include in the view and click the right arrow button to move your selections to the **Selected Fabrics and Hosts** table.

The **Available Fabrics** table displays the names and the number of products in the available fabrics. If this table is blank, it may be because all fabrics have been selected and are displayed in the **Selected Fabrics and Hosts** table.

To select more than one row, press CTRL and click individual rows. To select multiple sequential rows, press SHIFT and click on a sequence of rows.

5. Click the **Hosts** tab.

FIGURE 157 Create View dialog box - Hosts tab

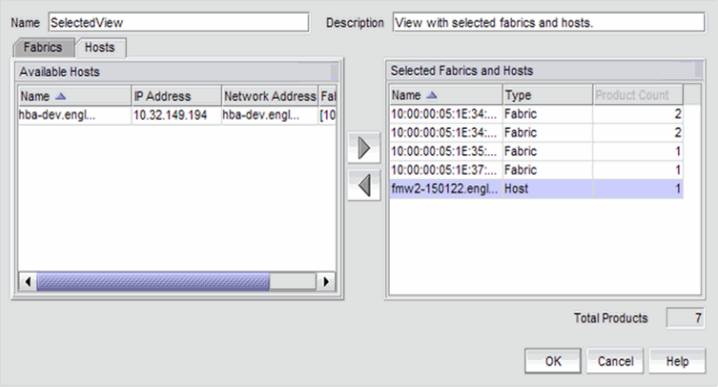

6. In the **Available Hosts** table, select the hosts you want to include in the view and click the right arrow button to move your selections to the **Selected Fabrics and Hosts** table.

The **Available Hosts** table displays the name, IP address, network address of the available hosts and the fabric in which the host is located. If this table is blank, it may be because all hosts have been selected and are displayed in the **Selected Fabrics and Hosts** table.

To select more than one row, press CTRL and click individual rows. To select multiple sequential rows, press SHIFT and click on a sequence of rows.

7. Confirm that all the fabrics and hosts you selected display in the **Selected Fabrics and Hosts** table.

The **Selected Fabrics and Hosts** table displays the name, type (host or fabric), number of products in the selected host or fabric.

8. Click **OK** to save the customized view and close the **Create View** dialog box.

The new view displays automatically in the main window of the Management application.

Editing a customized view

You can only edit customized views that you have created.

1. Select **View > Manage View > Edit View >** *View_Name*.

The **Edit View** dialog box displays.

FIGURE 158 Edit View dialog box - Fabrics tab

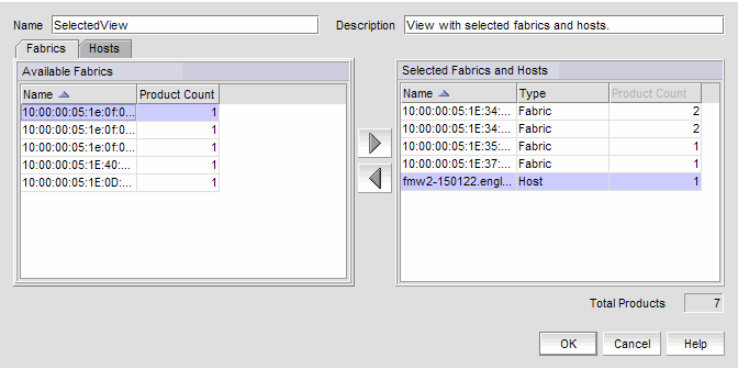

2. Click the **Fabrics** tab.
3. In the **Available Fabrics** table, select the fabrics you want to include in the view and use the right arrow button to move your selections to the **Selected Fabrics and Hosts** table.

   The **Available Fabrics** table displays the names and the number of products in the available fabrics. If this table is blank, it may be because all fabrics have been selected and are displayed in the **Selected Fabrics and Hosts** table.

   To select more than one row, press CTRL and click individual rows. To select multiple sequential rows, press SHIFT and click on a sequence of rows.
4. Click the **Hosts** tab.

FIGURE 159 Edit View dialog box - Hosts tab

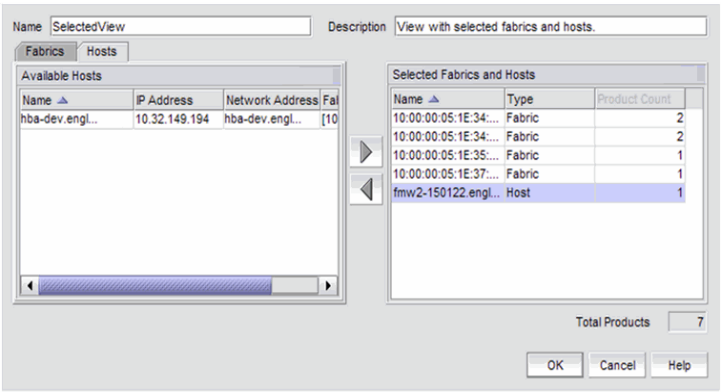

5. In the **Available Hosts** table, select the fabrics you want to include in the view and use the right arrow button to move your selections to the **Selected Fabrics and Hosts** table.

   The **Available Hosts** table displays the name, IP address, network address of the available hosts and the fabric in which the host is located. If this table is blank, it may be because all hosts have been selected and are displayed in the **Selected Fabrics and Hosts** table.

   To select more than one row, press CTRL and click individual rows. To select multiple sequential rows, press SHIFT and click on a sequence of rows.
6. To remove fabrics and hosts from a view, select the fabrics and hosts you want to remove in the **Selected Fabrics and Hosts** table and click the left arrow button.
7. Confirm that all the fabrics and hosts you selected display in the **Selected Fabrics and Hosts** table.

The **Selected Fabrics and Hosts** table displays the name, type (host or fabric), number of products in the selected host or fabric.

8. Click **OK** to save your changes and close the **Edit View** dialog box.
9. Verify your changes on the main window of the Management application.

Deleting a customized view

To delete a customized view, use the following procedure.

1. Select **View > Manage View > Delete View >** *View_Name*.
2. Click **Yes** on the message.

   If you delete the current view, the view changes to the default view (View All).

Copying a view

To copy a customized view, use the following procedure.

1. Use one of the following methods to open the **Copy View** dialog box:
   - Select **View > Manage View > Copy View >** *View_Name*.
   - Select **Copy View** from the **View All** list. The **View All** list does not display until you discover a fabric or host.

     The **Copy View** dialog box displays the name of the view you are copying.

FIGURE 160 Copy View dialog box

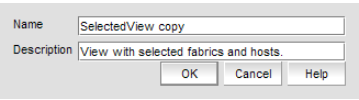

2. Enter a name (128-character maximum) in the **Name** field and a description (126-character maximum) in the **Description** field for the view.

   NOTE
   You cannot use the name “View” or “View All” in the **Name** field.

   NOTE
   You cannot use an existing name in the **Name** field.

3. In the **Available Fabrics** table, select the fabrics you want to include in the view and use the right arrow button to move your selections to the **Selected Fabrics and Hosts** table.

   The **Available Fabrics** table displays the names and the number of products in the available fabrics. If this table is blank, it may be because all fabrics have been selected and are displayed in the **Selected Fabrics and Hosts** table.

   To select more than one row, press CTRL and click individual rows. To select multiple sequential rows, press SHIFT and click on a sequence of rows.

4. In the **Available Hosts** table, select the fabrics you want to include in the view and use the right arrow button to move your selections to the **Selected Fabrics and Hosts** table.

The **Available Hosts** table displays the name, IP address, network address of the available hosts and the fabric in which the host is located. If this table is blank, it may be because all hosts have been selected and are displayed in the **Selected Fabrics and Hosts** table.

To select more than one row, press CTRL and click individual rows. To select multiple sequential rows, press SHIFT and click on a sequence of rows.

5. To remove fabrics and hosts from a view, select the fabrics and hosts you want to remove in the **Selected Fabrics and Hosts** table and click the left arrow button.
6. Confirm that all the fabrics and hosts you selected display in the **Selected Fabrics and Hosts** table.

   The **Selected Fabrics and Hosts** table displays the name, type (host or fabric), number of products in the selected host or fabric.
7. Click **OK** to save your changes and close the **Copy View** dialog box.

   NOTE
   When you open a new view, the **SAN** tab displays with a gray screen over the Product List and Topology Map while data is loading.
8. Verify that the copied view displays on the main window of the Management application.

SAN topology layout

You can customize various parts of the topology, including the layout of devices and connections and groups' background colors, to easily and quickly view and monitor devices in your SAN.
The following menu options are available on the **View** menu. Use these options to customize the topology layout.

- **Map Display.** Select to specify a new layout for the desktop icons, background color for groups, and line type for connections between icons.
- **Domain ID/Port #.** Select to set the display domain IDs and port numbers in decimal or hex format.
  - **Decimal.** Select to display all domain IDs and user and attached port numbers in decimal format.
  - **Hex.** Select to display all domain IDs and user and attached port numbers in hex format.
- **Product Label.** Select to configure which product labels display.

  NOTE
  Changes apply to all fabrics present in the topology when the **Product Label** option is selected.

  - **Name (Product).** Displays the product name as the product label.
  - **Node WWN.** Displays the world wide name as the product label.
  - **IP Address.** Displays the IP address as the product label.
  - **Domain ID.** Displays the domain ID as the product label.
  - **Zone Alias.** Displays the zone alias as the product label.
- **Port Label.** Select to configure which port labels display.

NOTE
Changes apply to the selected fabric or the fabric to which the selected item belongs.

- **Name.** Displays the name as the port label. If the port has not been given a name, the WWN of the port displays.
- **Port.** Displays the slot and port as the port label for a chassis switch and the port number for a switch.
- **Port Address.** Displays the port address as the port label.
- **Port WWN.** Displays the port world wide name as the port label.
- **User Port #.** Displays the user's port number as the port label.
- **Zone Alias.** Displays the zone alias as the port label.

- **Port Display.** Select to configure how ports display.
  - **Occupied Product Ports.** Select to display the ports of the devices in the fabrics (present in the Connectivity Map) that are connected to other devices.
  - **UnOccupied Product Ports.** Select to display the ports of the devices (shown in the Connectivity Map) that are not connected to any other device.
  - **Attached Ports.** Select to display the attached ports of the target devices.
  - **Switch to Switch Connections.** Select to display the switch-to-switch connections. Switch-to-switch connections only display when the **Attached Ports** option is also selected.

Customizing the layout of devices on the topology

You can customize the layout of devices by group type or for the entire Connectivity Map. Customizing the layout makes it easier to view the SAN and manage its devices. Group types include Fabric, Host, Storage, Router and Switch groups.

1. Right-click a group or the Connectivity Map and select **Map Display**.

   The **Map Display Properties** dialog box displays. The **Map Display Layout** list varies depending on what you selected (group type or Connectivity Map).

FIGURE 161 Map Display Properties dialog box

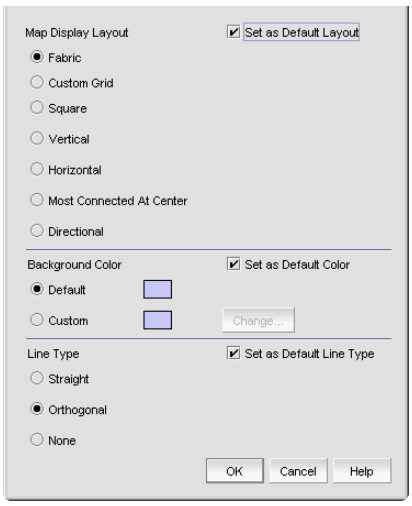

2. Select one of the following options from the **Map Display Layout** list:

   - **Free Form.** Select to display the devices in the default format for Switch Groups and Router Groups.
     When the **Free Form** map display layout is selected, the **View > Show Ports** menu command is unavailable.
   - **Fabric.** Only available for the group type "Fabric". Select to display the devices in the default format.

- **Custom Grid.** Select to be able to drag and drop product or group icons into a variable grid to reorganize the topology. The grid prevents icons from obscuring other icons. If enabled on a group, devices can only be moved within the group. If enabled on a fabric, groups can only be moved within the fabric. A device cannot be moved outside of its group.
- **Square.** Select to display the device icons in a square configuration. Default for Host and Storage groups.
- **Vertical.** Select to display the device icons vertically.
- **Horizontal.** Select to display the device icons horizontally.
- **Most Connected at Center.** Select to display the node that has the most connections at the center of the topology.
- **Directional.** Select to display the internal nodes in a position where they mirror the external groups to which they are connected.

3. Select the **Set as Default Layout** check box.
4. Click **OK** on the **Map Display Properties** dialog box to change the device layout on the topology.

Customizing the layout of connections on the topology

You can change the way inter-device connections display on the topology.

1. Right-click a group or the Connectivity Map and select **Map Display**.

   The **Map Display Properties** dialog box displays.

2. Select one of the following options from the **Line Type** list:
   - **Straight**. Select to display connections using straight lines.
   - **Orthogonal**. Select to display connections in orthogonal grid lines. Disabled if **Free Form** is selected in **Map Display Layout** area.
   - **None**. Select to hide the connections between devices.
3. Select the **Set as Default Line Type** check box.
4. Click **OK** on the **Map Display Properties** dialog box to change the line type on the topology.

Changing a group background color

You can customize the topology by changing the background color of a group.

1. Right-click a group or the Connectivity Map and select **Map Display**.

   The **Map Display Properties** dialog box displays.

2. Select the **Custom** option and click **Change**.

   The **Choose a background color** dialog box displays (Figure 162).

FIGURE 162 Choose a background color dialog box

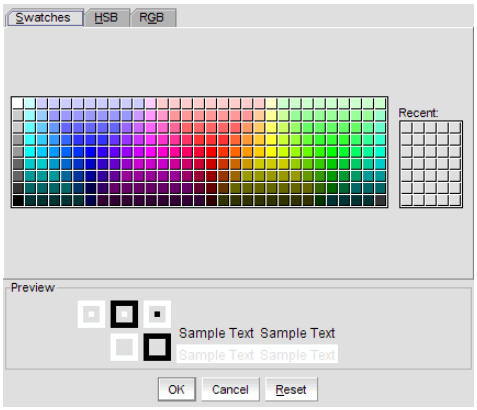

3. Select a color from the swatches tab and click **OK**.
   - To specify a color based on hue, saturation, and value, click the **HSV** tab. Specify the hue (0 to 359 degrees), saturation (0 to 100%), value (0 to 100%), and transparency (0 to 100%).
   - To specify a color based on hue, saturation, and lightness, click the **HSL** tab. Specify the hue (0 to 360 degrees), saturation (0 to 100%), lightness (0 to 100%), and transparency (0 to 100%).
   - To specify a color based on values of red, green, and blue, click the **RGB** tab. Specify the values for red (0 to 255), green (0 to 255), blue (0 to 255), and alpha (0 to 255) or enter a color code in the **Color Code** field.
   - To specify a color based on values of cyan, magenta, yellow, and black, click the **CMYK** tab. Specify the values for cyan (0 to 255), magenta (0 to 255), yellow (0 to 255), black (0 to 255), and alpha (0 to 255).
4. Select the **Set as Default Color** check box.
5. Click **OK** to change the background color on the topology.
6. Click **OK** on the **Map Display Properties** dialog box.

Reverting to the default background color

To revert back to the default background color, complete the following steps.

1. Right-click a group and select **Map Display**.

   The **Map Display Properties** dialog box displays.
2. Select the **Default** option.
3. Click **OK** on the **Map Display Properties** dialog box.

Changing the product label

To change the product label, complete the following steps.

1. Select a product in the Connectivity Map or Product List.
2. Select **View > Product Label**, and select one of the following options:
   - **Name (Product).** Displays the product name as the product label.
   - **WWN.** Displays the world wide name as the product label.
   - **IP Address.** Displays the IP address as the product label.

- **Domain ID.** Displays the domain ID as the product label.
- **Zone Alias.** Displays the zone alias as the product label.

Changes apply to all fabrics present in the topology when the **Product Label** option is selected.

Changing the port label

To change the port label, complete the following steps.

1. Select a port in the Connectivity Map or Product List.
2. Select **View > Port Label,** and select one of the following options:
   - **Name.** Displays the name as the port label.
   - **Port.** Displays the port number as the port label.
   - **Port Address.** Displays the port address as the port label.
   - **Port WWN.** Displays the port world wide name as the port label.
   - **User Port #.** Displays the user's port number as the port label.
   - **Zone Alias.** Displays the zone alias as the port label.

   All port labels within the fabric to which the selected item belongs change to the selected port label type.

Changing the port display

You have the option of viewing connected (or occupied) product ports, unoccupied product ports, or attached ports.

NOTE
Connected (or occupied) ports are those that originate from a device, such as a switch. Attached ports are ports of the target devices that are connected to the originating device.

1. Select **View > Port Display**, and select one of the following options:
   - **Occupied Product Ports**. Select to display the ports of the devices in the fabrics (present in the Connectivity Map) that are connected to other devices.
   - **Unoccupied Product Ports**. Select to display the ports of the devices (shown in the Connectivity Map) that are not connected to any other device.
   - **Attached Ports**. Select to display the attached ports of the target devices.
   - **Switch to Switch Connections**. Select to display the connections between devices. Switch-to-switch connections only display when the **Attached Ports** option is also selected.
2. Repeat step 1 to select more than one port display option.

Grouping on the topology

To simplify management, devices display in groups. Groups are shown with background shading and are labeled appropriately. You can expand and collapse groups to easily view a large topology.

Collapsing groups

To collapse a single group on the topology, choose one of the following options:

- Click the icon at the top right-hand corner of the group on the topology ( ).
- Double-click in the group, but not on a device.
- Right-click in a group, but not on a device, and select **Collapse** from the shortcut menu.

To collapse all groups on the topology by one level, click the **Collapse** button on the Connectivity Map toolbar ( ).

Expanding groups

To expand a group on the topology, do one of the following:

- Double-click the group icon.
- Right-click the group icon and select **Expand** from the shortcut menu.

To expand all groups on the topology by one level, click the **Expand** button on the Connectivity Map toolbar ( ).

Viewing connections

You can view the connections in a fabric using one of the following methods:

- Select a fabric and then select **View > Connected End Devices** and select **Include Virtual Devices**, **Hide All, Show All,** or **Custom**.
- Right-click the fabric and select **Connected End Devices > Include Virtual Devices**, **Hide All, Show All,** or **Custom** from the shortcut menu.

NOTE
Selecting **Hide All** disables the **Include Virtual Devices** option.

Configuring custom connections

NOTE
Active zones must be available on the fabric.

To create a display of the connected end devices participating in a single zone or group of zones, complete the following steps.

1. Select a fabric on the topology and select **View > Connected End Devices > Custom**.

   The **Connected End Devices - Custom display for** *Fabric* dialog box displays with a list of devices participating in a single zone or a group of zones in the **Zones in** *Fabric* list.

2. Select the zones you want to include in the connection in the **Zones in** *Fabric* list.
3. Select the application to which you want to add the selected zones in the **Application** list.
4. Click the right arrow button to move the zones to the **Selected Zones** list.

5. Click **Save**.

   The **Save Application** dialog box displays.

6. Enter a new name in the **Application Name** field.
7. Click **OK** on the **Save Application** dialog box.
8. Click **OK** on the **Connected End Devices - Custom display for** *Fabric* dialog box.

   The saved custom connection configuration displays in the **Connected End Devices** menu.

Deleting a custom connection configuration

NOTE
Active zones must be available on the fabric.

To delete a custom connection configuration, complete the following steps.

1. Select a fabric on the topology and select **View > Connected End Devices > Custom**.

   The **Connected End Devices - Custom display for** *Fabric* dialog box displays.

2. Select the configuration you want to delete in the **Application** list.
3. Click **Delete**.
4. Click **OK** on the confirmation message.
5. Click **OK** on the **Connected End Devices - Custom display for** *Fabric* dialog box.

Call Home

Call Home overview

NOTE
Call Home is supported on Windows systems for all e-mail Call Home centers and is supported on UNIX for the e-mail Call Home centers.

Call Home notification allows you to configure the Management application server to automatically send an e-mail alert or dial in to a support center to report system problems on specified devices (Fabric OS switches, routers, and directors). If you are upgrading from a previous release, all of your Call Home settings are preserved.

Call Home supports multiple Call Home centers which allows you to configure different devices to contact different Call Home centers. When you make any Call Home configuration changes or a Call Home event trigger occurs, the Management application generates an entry to the Master Log.

You can configure Call Home for the following Call Home centers,

- Brocade E-mail (Windows and UNIX)
- EMC (Windows and Linux)

NOTE
Beginning with Network Advisor 14.2.0 or later, the EMC Call Home center is available by default and the **Set inventory report interval** check box is clear by default.

- IBM E-mail (Windows and UNIX)
- NetApp E-mail (Windows and UNIX)
- Oracle E-mail (Windows and UNIX)

Call Home allows you to automate tasks that occur when the Call Home event trigger is fired. When a Call Home event trigger occurs, the Management application generates the following actions:

- Sends an e-mail alert to a specified recipient or dials in to a support center.
- Triggers supportSave on the switch (if supportSave is enabled on the switch) prior to sending an alert. The supportSave location is included in the alert.
- Adds an entry to the Master Log file and screen display.
- Generates an XML report (only available with EMC Call Home centers) with the product details, which is sent with the e-mail alert.

  NOTE
  The EMC Call Home center is not available by default; however, you can enable the EMC Call Home center. For step-by-step instructions, refer to "Enabling the EMC Call Home center" on page 336.

- Generates an HTML report for e-mail-based Call Home centers.

For more information about Call Home events, refer to "Call Home Event Tables" on page 1323. For more information about events, refer to "Fault Management" on page 1131.

Call Home allows you to perform the following tasks:

- Assign devices to and remove devices from the Call Home centers.
- Define filters from the list of events generated by Fabric OS devices.
- Edit and remove filters available in the Call Home Event Filters table.
- Apply filters to and remove filters from the devices individually or in groups.
- Edit individual Call Home center parameters to dial a specified phone number or e-mail a specific recipient.
- Enable and disable individual devices from contacting the assigned Call Home centers.
- Show or hide Call Home centers on the display.
- Enable and disable Call Home centers.

Viewing Call Home configurations

To view Call Home center configurations, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays (Figure 163).

FIGURE 163 Call Home dialog box

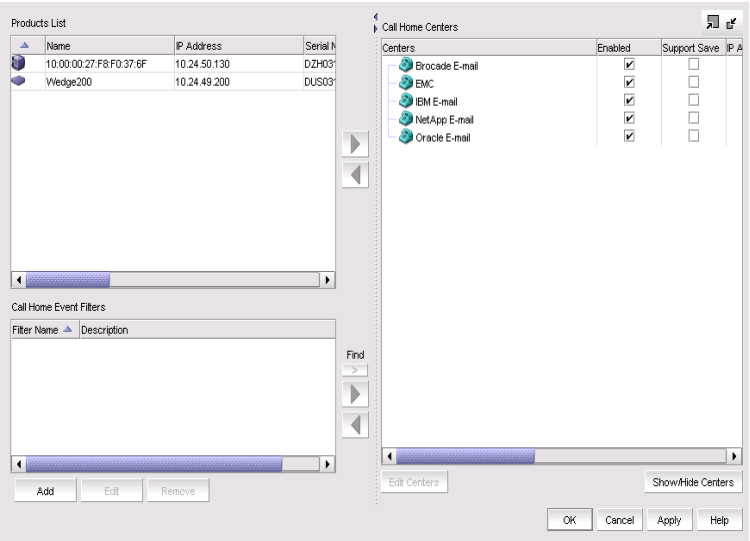

The **Call Home** dialog box contains the following fields and components:

- **Products List** — Displays all discovered products. The list allows for multiple selections and manual sorting of columns. This list displays the following information:
  - **Product Icon —** The status of the products' manageability.
  - **Name —** The name of the product.
  - **IP Address** — The IP address (IPv4 or IPv6 format) of the product.
  - **Serial No —** The serial number of the switch.
  - **Node WWN —** The node world wide name of the product.
  - **Fabric Name —** The name of the fabric.
  - **Vendor —** The vendor ID of the product.
  - **Call Home Status —** One of the following Call Home statuses for the product.
    - **Enabled** — The product is manageable and Call Home is enabled.
    - **Disabled** — The product is manageable and Call Home is disabled.
    - **Not Manageable** — The product is discovered but not manageable.
    - **Server Not Registered —** The server is not registered to receive Call Home events from the product.

    NOTE
    Call Home status only displays for Fabric OS products.
  - **DomainID** — The domain ID of the product.
  - **Product Type** — The type of product (switch, Layer 2 switch, router, or director).
- Right arrow buttons (top) — Click to assign the selected product to the selected Call Home center (refer to "Assigning a device to the Call Home center" on page 339). Disabled when no product is selected in the **Products List** or when more than one Call Home center is selected in the **Call Home Centers** list.
- Left arrow button (top) — Click to remove the selected product from the selected Call Home center (refer to "Removing a device from a Call Home center" on page 340). Disabled when no product or Call Home center is selected in the **Call Home Centers** list.

- **Call Home Event Filters** list — Displays all Call Home event filters. This list displays the following information:
  - **Filter Name** — The name of the event filter.
  - **Description** — The description of the event filter.
- **Add** button — Click to open the **Call Home Event Filter** dialog box and add an event filter (refer to "Defining an event filter" on page 341).
- **Edit** button — Click to open the **Call Home Event Filter** dialog box and edit an event filter (refer to "Defining an event filter" on page 341).
- **Remove** button — Click to remove the event filter (refer to "Removing an event filter from the Call Home Event Filters list" on page 344) from the **Call Home Event Filters** list.
- **Find** button (>) — Click to find all instances of the selected event filter in the **Call Home Centers** list.
- Right arrow button (bottom) — Click to assign the selected event filter (refer to "Assigning an event filter to a Call Home center" on page 342 or "Assigning an event filter to a device" on page 342) to the selected Call Home center or product. Disabled when no event filter is selected in the **Call Home Event Filters** list.
- Left arrow button (bottom) — Click to remove the selected event filter (refer to "Removing all event filters from a Call Home center" on page 343 or "Removing an event filter from a device" on page 344) from the selected Call Home center or product. Disabled when no event filter, product, or Call Home center is selected in the **Call Home Centers** list.
- **Call Home Centers** list — The Call Home centers, products assigned to the Call Home centers, and event filters assigned to the Call Home centers and products. This list displays the following information:
  - **Centers —** A tree with Call Home centers as the parent node, assigned products as subnodes, and event filters as the child node to the assigned products.
  - **Enabled** check box — Select the check box to enable the associated Call Home center or clear the check mark to disable the Call Home center. By default, all check boxes are selected during a fresh install.
  - **Support Save** check box — Select the check box to enable supportSave, which collects diagnostic information on Fabric OS switches.
  - **IP Address —** The IP address of the product.
  - **Serial No —** The serial number of the switch.
  - **Node WWN —** The node WWN of the product.
  - **Fabric Name —** The name of the fabric.
  - **Vendor —** The vendor of the product.
  - **Call Home Status —** One of the following Call Home statuses for the product:
    - **Enabled** — The product is manageable and Call Home is enabled.
    - **Disabled** — The product is manageable and Call Home is disabled.
    - **Not Manageable** — The product is discovered but not manageable.
    - **Server Not Registered —** The server is not registered to receive Call Home events from the product.

    NOTE
    Call Home status only displays for Fabric OS products.
  - **DomainID** — The domain ID of the product.
  - **Product Type** — The type of product (switch, Layer 2 switch, router, or director).
- **Edit Centers** button — Select a call home center in the **Centers** list anc click **Edit** to open the **Configure Call Home Center** dialog box and modify Call Home center information (refer to "Editing a Call Home center" on page 332).
- **Show/Hide Centers** button — Click to open the **Centers** dialog box and add or delete a Call Home center (refer to "Showing a Call Home center" on page 331 or "Hiding a Call Home center" on page 331).

2. Click **OK** to close the **Call Home** dialog box.

Showing a Call Home center

To show a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Click **Show/Hide Centers** (beneath the **Call Home Centers** list).

   The **Centers** dialog box displays with a predefined list of Call Home centers.

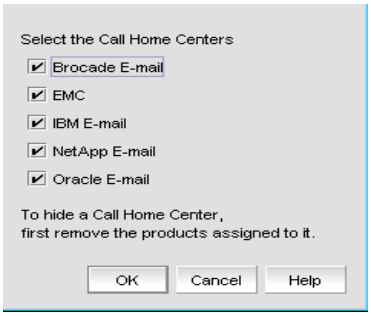

3. Select the check boxes of the Call Home centers you want to display.

   Clear the check box to hide the Call Home center.

4. Click **OK** on the **Centers** dialog box.

   The **Call Home** dialog box displays with the selected Call Home centers listed in the **Call Home Centers** list.

Hiding a Call Home center

> NOTE
> Before you can hide a Call Home center, you must remove all assigned products.

To hide a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Click **Show/Hide Centers** (beneath the **Call Home Centers** list).

   The **Centers** dialog box displays with a predefined list of Call Home centers.

3. Clear the check boxes of the Call Home centers you want to hide and click **OK**.

   The **Call Home** dialog box displays with only the selected Call Home centers listed in the **Call Home Centers** list.

Editing a Call Home center

To edit a Call Home center, select from the following procedures:

Editing an e-mail Call Home center

E-mail Call Home centers are available for Brocade, EMC, IBM, NetApp, and Oracle. To edit one of these Call Home centers, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the Call Home center you want to edit (**Brocade E-mail**, **IBM E-mail**, **NetApp E-mail**, or **Oracle E-mail**) in the **Call Home Centers** table.

3. Click **Edit Centers** (beneath the **Call Home Centers** table).

   The **Configure Call Home Center** dialog box displays (Figure 164).

   FIGURE 164 Configure Call Home Center dialog box (Brocade, EMC, IBM, NetApp, or Oracle E-mail option)

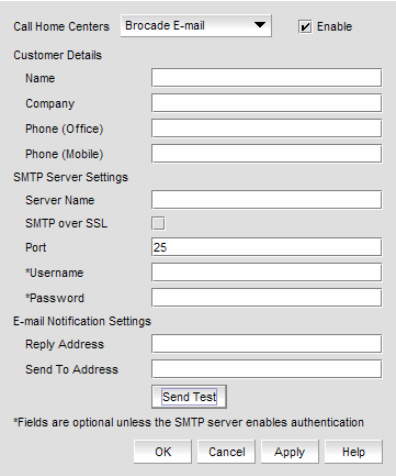

4. Make sure the Call Home center type you selected displays in the **Call Home Centers** list.

   If the Call Home center type is incorrect, select the correct type from the list.

5. Select the **Enable** check box to enable this Call Home center.

6. Enter your contact name in the **Customer Details - Name** field.

7. Enter your company name in the **Customer Details - Company** field.

8. Enter the phone number of the customer contact in the **Customer Details - Phone (Office)** field.

9. Enter the mobile phone number of the customer contact in the **Customer Details - Phone (Mobile)** field.

10. Enter the name of the e-mail server in the **SMTP Server Settings - Server Name** field.

11. Select the **SMTP over SSL** check box to enable secure communication between the SMTP server and the Management application.

12. Enter the port number of the server in the **SMTP Server Settings - Port** field.

    The default is 465 if SMTP over SSL is enabled; otherwise, the default is 25.

13. Enter a user name in the **SMTP Server Settings - Username** field.

    This is a required field when the SMTP server authentication is enabled.

14. Enter a password in the **SMTP Server Settings - Password** field.

    This is a required field when the SMTP server authentication is enabled.

15. Enter your e-mail address in the **E-mail Notification Settings - Reply Address** field.

    You can enter more than one e-mail address, separating each with a semi-colon. To send a text message or page by way of e-mail, use the following format: number@carrier.com (where number is your phone number and carrier.com is the SMS server; for example, 3035551212@txt.att.net (text message) or 3035551212@page.att.net (page)).

    NOTE
    Check with your carrier for the exact e-mail address format.

16. Enter an e-mail address in the **E-mail Notification Settings - Send To Address** field.

17. Click **Send Test** to test the mail server.

    The selected Call Home center must be enabled to test the mail server.

    A faked event is generated and sent to the selected Call Home center. You must contact the Call Home center to verify that the event was received and in the correct format. To see the content included in an e-mail message, refer to "Call Home alert e-mail messages" on page 333.

18. Click **OK** to close the "Test Event Sent" message.

19. Click **OK**.

    The **Call Home** dialog box displays with the Call Home center you edited highlighted in the **Call Home Centers** list.

20. Click **OK** to close the **Call Home** dialog box.

Call Home alert e-mail messages

When an event triggers a Call Home alert, an e-mail message is sent to the selected Call Home center. The e-mail message includes the following information:

- E-mail subject line — [*Severity* - *Event_Reason_Code* - *FRU_Code* or *Event_Type* - *Factory_Serial_Number*] Call Home Alert about product *IP_Address* with support save information

  A potential e-mail subject line is shown in the following example:

  [3 - 1427 - FW-1427 - AMH0344D006] Call Home Alert about product 172.26.24.85 with support save information

- E-mail content — Provides the following information about the triggered event:
  - Event Description — Details about the event that triggered the alert. Includes the following data:
    - Product WWN
    - Product IP address
    - Time
    - SupportSave location
  - Management Server Information — Details about the Management server. Includes the following data:
    - Server Name
    - Server IP
    - Server Version
  - Contact Information — Customer contact information. Includes the following data:
    - Customer Name
    - Contact Name
    - Phone 1
    - Phone 2
  - Source — Details about the product. Includes the following data:
    - Firmware Version
    - Supplier Serial number
    - Factory Serial number
    - IP Address
    - Model number
    - Type
    - Product Name
    - Product WWN
    - Ethernet IP
    - Ethernet IP Mask
    - FCIP
    - FCIP Mask
    - Product Type
    - Domain ID
    - Product Manufacturer
    - Product Type Number
    - Manufacturing Plant
    - Product Status
    - Status Reason
  - Event — Details about the triggered event. Includes the following data:
    - Event Time
    - Event Severity
    - Event Reason Code
    - FRU Code/Event Type
    - Event Description
  - Event Data — Information about the triggered event. Includes the following data:
    - Event level
    - Event number
    - Event count
    - Event time
    - Event Message Id
    - Event Description

- Last 30 Events on the Product (Brocade E-mail and NetApp E-mail only) — Table with the last 30 product and product status events. The first event is always the event that triggered the e-mail alert. Includes the following data for each event:
  - Event level
  - Event number
  - Count
  - Time
  - Message ID
  - Description

Editing the EMC Call Home center

NOTE
The EMC Call Home center is not available by default; however, you can enable the EMC Call Home center. For step-by-step instructions, refer to "Enabling the EMC Call Home center" on page 336.

To edit an EMC Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the **EMC** Call Home center you want to edit in the **Call Home Centers** list.

3. Click **Edit Centers** (beneath the **Call Home Centers** list).

   The **Configure Call Home Center** dialog box displays (Figure 165).

   FIGURE 165 Configure Call Home Center dialog box (EMC option)

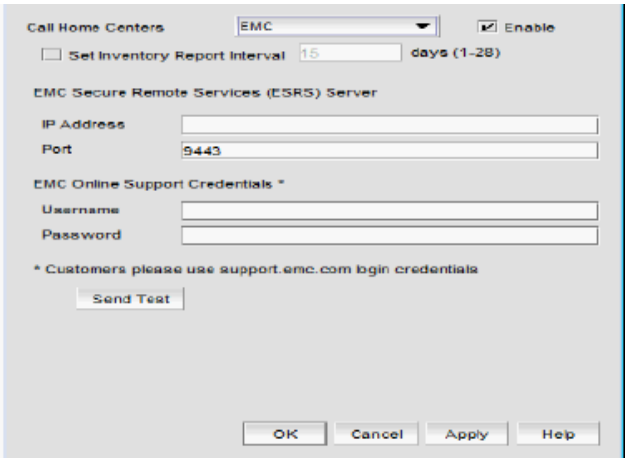

4. Make sure the **EMC** Call Home center type displays in the **Call Home Centers** list.

   If the Call Home center type is incorrect, select the correct type from the list.

5. Select the **Enable** check box to enable this Call Home center.

6. Set the time interval at which to check the Call Home center by selecting the **Set inventory interval at ___ days (1-28)** check box and entering the interval in the field.

7. Enter the IP address and port in the **EMC Secure Remote Services Server Centre - IP Address#** and **Port #** fields.

8. Enter the username and password in the **EMC Online Support Credentials - Username# and Password #** fields.

   The customers must use support.emc.com login credentials.

9. Click **Send Test** to test the mail server.

   The selected Call Home center must be enabled to test the mail server.

   A faked event is generated and sent to the selected Call Home center. You must contact the Call Home center to verify that the event was received and in the correct format. To see the content included in an e-mail message, refer to "Call Home alert e-mail messages" on page 333.

10. Click **OK**.

    The **Call Home** dialog box displays with the Call Home center you edited highlighted in the **Call Home Centers** list.

11. Click **OK** to close the **Call Home** dialog box.

    NOTE
    The user has to lookup in master log to confirm the successful addition of a device to the EMC Call Home center.

    NOTE
    After migration, the devices will not be automatically moved under **Call Home Centers** list. Only after the successful configuration of the **EMC Secure Remote Services Server Centre** details in the **Configure Call Home Center** dialog box, the device will be auto-assigned to the Call Home center.

Enabling the EMC Call Home center

To enable the EMC Call Home center, complete the following steps.

- On Windows systems, complete the following steps.
  a. Open a command prompt and navigate to the *Install_Home*\utilities directory.
  b. Enable the EMC Call Home center by typing `enableEMCcallhomeinNonCMCNE.bat dbusername dbpassword enable` and pressing **Enter**.

     For example, `enableEMCcallhomeinNonCMCNE.bat dcmadmin passw0rd enable`.

     Disable the EMC Call Home center by typing `enableEMCcallhomeinNonCMCNE.bat dbusername dbpassword disable` and pressing **Enter**.
- On UNIX systems, complete the following steps.
  a. Open a terminal and navigate to the *Install_Home*\utilities directory.
  b. Enable the EMC Call Home center by typing `enableEMCcallhomeinNonCMCNE.sh` *dbusername dbpassword* `enable|disable` and pressing **Enter**.

     For example, `enableEMCcallhomeinNonCMCNE.sh dcmadmin passw0rd enable`.

     Disable the EMC Call Home center by typing `enableEMCcallhomeinNonCMCNE.sh` *dbusername dbpassword* `disable` and pressing **Enter**.

Enabling a Call Home center

To enable a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the **Enable** check box of the Call Home center you want to enable in the **Call Home Centers** list.
3. Click **OK** to close the **Call Home** dialog box.

Enabling supportSave

NOTE
SupportSave is only supported on products running Fabric OS 7.0 or later or Network OS 2.1.X or later.

NOTE
Beginning with Fabric OS 8.0.1 and later, SupportSave is enabled by default for EMC Call Home centers.
The EMC Call Home center is not available by default; however, you can enable the EMC Call Home center. For step-by-step instructions, refer to "Enabling the EMC Call Home center" on page 336.

When you enable supportSave through the Call Home center, all Call Home events trigger the supportSave operation and the supportSave stored location on the FTP server is transmitted with the Call Home event.

To enable a supportSave for a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the **Support Save** check box of the Call Home center or device for which you want to enable supportSave in the **Call Home Centers** list.
3. Click **OK** to close the **Call Home** dialog box.

Testing the Call Home center connection

Once you add and enable a Call Home center, you should verify that Call Home is functional.

To verify Call Home center functionality, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.
2. Click **Edit Centers** (beneath the **Call Home Centers** list).

   The **Configure Call Home Center** dialog box displays.

3. Select the Call Home center you want to check in the **Call Home Centers** list.
4. Make sure that the **Enabled** check box is selected.

   NOTE
   You must configure the Call Home center before you test the connection. To configure a Call Home center, refer to "Editing a Call Home center" on page 332.

5. Click **Send Test**.

   A faked event is generated and sent to the selected Call Home center. You must contact the Call Home center to verify that the event was received and in the correct format.

   NOTE
   The **Sent Test** function is not supported on the EMC Call Home center.
   The EMC Call Home center is not available by default; however, you can enable the EMC Call Home center. For step-by-step instructions, refer to "Enabling the EMC Call Home center" on page 336.

6. Click **OK** to close the "Test Event Sent" message.
7. Click **OK** to close the **Configure Call Home Center** dialog box.
8. Click **OK** to close the **Call Home** dialog box.

Disabling a Call Home center

When a Call Home center is disabled, no devices can send Call Home events to the Call Home center. However, the devices and event filters assigned to the disabled Call Home center are not removed. You can still perform the following actions on a disabled Call Home center:

- Edit Call Home center configuration.
- Add devices and event filters to the Call Home center.

To disable a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Clear the **Enable** check box of the Call Home center you want to disable in the **Call Home Centers** list.

   The selected Call Home center and its devices and event filters become unavailable. However, the Call Home center is not disabled until you save your changes. When a device is assigned to the Call Home center, a confirmation message displays.

3. Click **OK** to confirm.
4. Click **OK** to close the **Call Home** dialog box.

Viewing Call Home status

You can view Call Home status from the main Management application window or from the **Call Home Notification** dialog box.

The Management application enables you to view the Call Home status at a glance by providing a Call Home status icon on the status bar. Table 31 illustrates and describes the icons that indicate the current status of the Call Home function.

TABLE 31 Call Home icons

Icon	Description
	Normal — Displays when Call Home is enabled on all devices and no filters are applied.
	Degraded — Displays when Call Home is enabled on all devices and at least one filter is active.
	Disabled — Displays when any of the following conditions are met: • At least one device's Call Home is disabled. • At least one non-manageable device. • At least one device does not have the Management server registered as a trap recipient.

To view more detail regarding Call Home status, click the **Call Home** icon. The **Call Home Notification** dialog box displays the following information for the list of devices that have assigned filters or Call Home disabled:

- **Product** — The name of the device. Click to go to the device in the topology.
- **IP Address** — The IP address (IPv4 or IPv6 format) of the device.
- **Status** — The status of the device. The possible status options include:
  - **Enabled** — The device is manageable, Call Home is enabled, and a filter is applied.
  - **Disabled** — Call Home is disabled on at least one device or Call Home is disabled from the **Call Home** dialog box.
  - **Not Manageable** — Manageability is lost.
  - **Server Not Registered** — The server is not registered to receive Call Home events from this device.

    NOTE
    Call Home status only displays for Fabric OS products.
- **Filter** — The name of the active event filter assigned to the device.
- **Call Home** button — Click to launch the **Call Home** dialog box, where you can configure Call Home centers.

Assigning a device to the Call Home center

Discovered devices (switches, routers, and directors) are not assigned to a corresponding Call Home center automatically. You must manually assign each device to a Call Home center before you use Call Home.

To assign a device or multiple devices to a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the devices you want to assign to a Call Home center in the **Products List**.

3. Select the Call Home center to which you want to assign the devices in the **Call Home Centers** list.

   You can only assign a device to one Call Home center at a time.

4. Click the right arrow button.

   The selected devices display beneath the selected Call Home center. Devices assigned to a Call Home center do not display in the **Products List**.

5. Click **OK** to close the **Call Home** dialog box.

Removing a device from a Call Home center

To remove a device or multiple devices from a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the Call Home center from which you want to remove devices in the **Call Home Centers** list.
3. Select the devices you want to remove from the selected Call Home center.
4. Click the left arrow button.

   A confirmation message displays.

5. Click **OK**.

   The selected devices are removed from the Call Home center and display in the **Products List**.

6. Click **OK** to close the **Call Home** dialog box.

Removing all devices and filters from a Call Home center

To remove all devices and filters from a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the Call Home center from which you want to remove devices and filters in the **Call Home Centers** list.
3. Click the left arrow button.

   A confirmation message displays.

4. Click **OK**.

   All devices assigned to the selected Call Home center display in the **Products List**. Any assigned filters are also removed.

5. Click **OK** to close the **Call Home** dialog box.

Defining an event filter

To define an event filter, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Click **Add** beneath the **Call Home Event Filter** list.

   The **Call Home Event Filter** dialog box displays.

3. Enter a name for the filter in the **Name** field.
4. Enter a name for the description in the **Description** field.
5. Select the check box for the events you want to include in the filter in the **Available Call Home Event Types** list.

   To exclude the event, clear the check box. By default, all check boxes are selected during a new installation. Click **Select All** to select all event types in the list or select **Unselect All** to clear the selected event types in the list. For more information about Call Home events, refer to "Call Home Event Tables".

   The **Available Call Home Event Types** list displays the following information:

   - **Description —** The description of the event.
   - **Type — The type of firmware for the selected event.**
   - **FRU Code/Event Type —** The field-replaceable unit (FRU) code and event type for the event.
   - **Severity —** The severity of the event.
   - **Event Reason Code —** The event reason code of the event.

6. Click **OK** on the **Call Home Event Filter** dialog box.

   The event filter name and the description are displayed in the **Call Home** dialog box.

   To assign event filters to a Call Home center or a device, refer to "Assigning an event filter to a Call Home center" on page 342 or "Assigning an event filter to a device" on page 342.

   NOTE
   All events are sent to call home, if the events are not defined or the filter is not created.

7. Click **OK** to close the **Call Home** dialog box.

NOTE
Fabric Watch is not supported in Fabric OS 7.4.0 and later. The Fabric Watch events listed in the **Call Home Event Filter** dialog box are not applicable for switches running Fabric OS 7.4.0 and later.

Call Home for virtual switches

For virtual switches, there are two types of Call Home events:

- FRU-based Call Home events, which are triggered at the chassis level
- Port-based Call Home events, which are triggered for each virtual switch

NOTE
FW-1444 will be generated only when the RASLog is configured as an action for the FRU class. FW-1444 and FW-1447 events are configured at the logical switch level and Call Home is triggered accordingly. Therefore, the Management application sends a Call Home alert at the logical switch level and not at the chassis level.

Assigning an event filter to a Call Home center

Event filters allow Call Home center users to log in to a Management server and assign specific event filters to the devices. This limits the number of unnecessary or "acknowledge" events and improves the performance and effectiveness of the Call Home center.

You can only select one event filter at a time; however, you can assign the same event filter to multiple devices or Call Home centers. When you assign an event filter to a Call Home center, the event filter is assigned to all devices in the Call Home center. For more information about Call Home events, refer to "Call Home Event Tables".

NOTE
You cannot assign an event filter to a Call Home center that does not contain devices.

To assign an event filter to a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the event filters you want to assign in the **Call Home Event Filters** list.
3. Select the Call Home centers to which you want to assign the event filters in the **Call Home Centers** list.
4. Click the right arrow button.

   The selected event filters are assigned to the selected Call Home centers.

5. Click **OK** to close the **Call Home** dialog box.

Assigning an event filter to a device

To assign an event filter to a device, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the event filter you want to assign in the **Call Home Event Filters** list.

   For more information about Call Home events, refer to "Call Home Event Tables".

3. Select one or more devices to which you want to assign the event filter in the **Call Home Centers** list.

4. Click the right arrow button.

   The selected event filter is assigned to the selected devices. The event filter displays beneath the specified device or all of the devices under the specified Call Home center.

5. Click **OK** to close the **Call Home** dialog box.

Overwriting an assigned event filter

A device can only have one event filter at a time; therefore, when a new filter is applied to a device that already has a filter, you must confirm the new filter assignment.

To overwrite an event filter, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the event filter you want to apply in the **Call Home Event Filters** list.

   For more information about Call Home events, refer to "Call Home Event Tables".

3. Select the devices to which you want to apply the event filter in the **Call Home Centers** list.
4. Click the right arrow button.

   For existing event filters, a confirmation messages displays.

5. Click **Yes**.

   The selected event filter is applied to the selected devices. The event filter displays beneath the specified device or all of the devices under the specified Call Home center.

6. Click **OK** to close the **Call Home** dialog box.

Removing all event filters from a Call Home center

To remove all event filters from a Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Choose one of the following options in the **Call Home Centers** list:

   - Right-click a Call Home center and select **Remove Filters**.
   - Select a Call Home center and click the left arrow button.

     All event filters assigned to the Call Home center are removed.

3. Click **OK** to close the **Call Home** dialog box.

Removing an event filter from a device

To remove an event filter from a device, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Choose one of the following options in the **Call Home Centers** list:

   - Right-click a device to which the event filter is assigned and select **Remove Filter**.
   - Select an event filter assigned to a device and click the left arrow button. Press **CTRL** and click to select multiple event filters assigned to multiple devices.

   All event filters assigned to the device are removed.

3. Click **OK** to close the **Call Home** dialog box.

Removing an event filter from the Call Home Event Filters list

To remove an event filter from the Call Home Event Filters list, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the event filter you want to remove in the **Call Home Event Filters** list.

3. Click **Remove**.

   - If the event filter is not assigned to any devices, a confirmation message displays asking if you want to remove the event filter. Click **Yes**.
   - If the event filter is assigned to any devices, a confirmation message displays informing you that removing this event filter will remove it from all associated devices. Click **Yes**.

   The event filter is removed from any associated devices and the **Call Home Event Filters** list.

   To determine to which devices the event filter is assigned, select the event filter and then click the **Find** button (>).

4. Click **OK** to close the **Call Home** dialog box.

Searching for an assigned event filter

To find all devices to which an event filter is assigned, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

   The **Call Home** dialog box displays.

2. Select the event filter you want to find in the **Call Home Event Filters** list.

3. Click the **Find** button (**>**).

   All instances of the event filter are highlighted in the **Call Home Centers** list.

   If the selected event filter is not assigned to any devices in the **Call Home Centers** list, a not found message displays.

Third-Party tools

About Third-party tools

NOTE
Installing tools is only available with the Trial and Licensed versions.

You can open other software products (such as, Firefox, Windows Explorer, Web Tools, Element Managers, FCR Configuration, HCM Agent, and so on) you frequently use from the **Tools** menu or shortcut menus.

You can add third-party tools to the **Tools** menu or shortcut menus to open other software products you frequently use.

Starting third-party tools from the application

You can open third-party tools from the **Tools** menu or a device's shortcut menu. Remember that you cannot open a tool that is not installed on your computer. You must install the tool on your computer and add the tool to the **Tools** menu or the device's shortcut menu.

NOTE
Installing tools is only available with the Trial and Licensed versions.

To open an application, complete the following steps.

1. Select the device.
2. Use one of the following techniques:
   - Select **Tools > Product Menu >** *Tool_Name*.
   - Select **Tools >** *Tool_Name*.
   - Right-click the device, and select the tool from the menu.

     If the third-party tool is a web-based application, you must enter the IP address of the applications server as a parameter to be able to open the application. For step-by-step instructions about entering the IP address of the server, refer to "Entering the server IP address of a tool" on page 355.

Launching a Telnet session

You can use Telnet to log in and issue command line-based commands to a device.

NOTE
The device must have a valid IP address. If the device does not have a valid IP address, the Telnet selection will not be available on the **Tools** menu or the shortcut menu. You must right-click the device icon, select **Properties,** and enter the device's IP address before you can open a Telnet session.

Launching a Telnet session from the SAN tab

To launch a Telnet session, complete the following steps.

On the Connectivity Map, right-click a device and select **Telnet** or **Telnet through Server**.

NOTE
Telnet through Server is only supported on Windows systems.

OR

1. Select the switch to which you want to connect.
2. Select **Tools > Product Menu > Telnet**.

   The Telnet session window displays.

   NOTE
   On Linux systems, you must use Ctrl + Backspace to delete text in the Telnet session window.

Launching Element Manager

Element Managers are used to manage Fibre Channel switches and directors. You can open a device's Element Manager directly from the application.

To launch a device's Element Manager, complete the following steps.

On the Connectivity Map, double-click the device you want to manage.

The Element Manager displays.

OR

On the Connectivity Map, right-click the device you want to manage and select **Element Manager > Hardware**.

The Element Manager displays.

OR

1. Select a device.
2. Select **Configure > Element Manager > Hardware**.

   The Element Manager displays.

OR

1. Select a device.
2. Click the Element Manager icon on the toolbar.

   The Element Manager displays.

Launching Web Tools

Use Web Tools to enable and manage Fabric OS access gateway, switches, and directors. You can open Web Tools directly from the application. For more information about Web Tools, refer to the ***Web Tools Administrator's Guide.*** For more information about Fabric OS access gateway, switches, and directors, refer to the documentation for the specific device.

To launch Web Tools, complete the following steps.

> NOTE
> Web Tools requires Oracle JRE. For the current supported JRE version for Web Tools, refer to the Release Notes. For the website listing patch information, go to *http://www.oracle.com/technetwork/java/javase/downloads/index.html*.

> NOTE
> You must have Element Manager - Product Administration privileges for the selected device to launch Web Tools. If you do not have Element Manager - Product Administration privileges, you must enter those credentials to launch Web Tools. For more information about privileges, refer to "User Privileges" on page 1333.

On the Connectivity Map, double-click the Fabric OS device you want to manage.

Web Tools displays.

OR

On the Connectivity Map, right-click the Fabric OS device you want to manage and select **Element Manager > Hardware**.

Web Tools displays.

OR

1. Select a Fabric OS device.
2. Select **Configure > Element Manager > Hardware**.

   Web Tools displays.

OR

1. Select a Fabric OS device.
2. Click the Element Manager icon on the toolbar.

   Web Tools displays.

   NOTE
   When you close the Management application client, any Web Tools instance launched from the clients closes as well.

Launching FCR Configuration

Use FCR Configuration to launch the FC Routing module, which enables you to share devices between fabrics without merging the fabrics. You can open the FC Routing module directly from the Management application. For more information about FC Routing, refer to the *Web Tools Administrator's Guide*.

The FCR Configuration option is available only for the following devices with Fabric OS 7.0 or later:

- Fabric OS extension switch
- Fabric OS Directors configured with an extension blade
- Fabric OS 1U, 8 Gbps 40-port FC Switch (with Integrated Routing license)
- Fabric OS 2U, 8 Gbps 80-port FC Switch (with Integrated Routing license)
- Fabric OS directors configured with an FC 8 Gbps 16-port Blade (with Integrated Routing license)
- Fabric OS directors configured with an FC 8 Gbps 32-port Blade (with Integrated Routing license)
- Fabric OS directors configured with an FC 8 Gbps 48-port Blade (with Integrated Routing license)

  Note that on the FC 8 Gbps 48-port Blade, the Shared Area ports, for example, 16-47, cannot be configured as EX_Ports

  NOTE
  You must have Element Manager - Product Administration privileges for the selected device to launch Web Tools. If you do not have Element Manager - Product Administration privileges, you must enter those credentials to launch Web Tools. For more information about privileges, refer to "User Privileges" on page 1333.

On the Connectivity Map, right-click the Fabric OS device you want to configure and select **Element Manager > Router Admin**.

OR

1. Select a Fabric OS device.

2. Select **Configure > Element Manager > Router Admin**.

   The FC Routing module displays.

   NOTE
   When you close the Management application client, any Web Tools instance launched from the clients closes as well.

Launching Name Server

Use Name Server to view entries in the Simple Name Server database. You can open the Name Server module directly from the Management application. For more information about Name Server, refer to the ***Web Tools Administrator's Guide***.

NOTE
You must have Element Manager - Product Administration privileges for the selected device to launch Web Tools. If you do not have Element Manager - Product Administration privileges, you will need to enter those credentials to launch Web Tools. For more information about privileges, refer to "User Privileges" on page 1333.

On the Connectivity Map, right-click the Fabric OS device you want to configure and select **Element Manager > Name Server**.

The Name Server module displays.

OR

1. Select a Fabric OS device.
2. Select **Configure > Element Manager > Name Server**.

   The Name Server module displays.

   NOTE
   When you close the Management application client, any Web Tools instance launched from the clients closes as well.

Launching HCM Agent

Use Fabric OS HCM Agent to enable and manage Fabric OS HBAs. You can open HCM Agent directly from the application. For more information about HCM Agent, refer to the ***HCM Agent Administrator's Guide.*** For more information about Fabric OS HBAs, refer to the documentation for the specific device.

To launch a Fabric OS HBA's Element Manager, complete the following steps.

NOTE
You must have Element Manager - Product Administration privileges for the selected device to launch HCM Agent. If you do not have Element Manager - Product Administration privileges, you will need to enter those credentials to launch HCM Agent. For more information about privileges, refer to "User Privileges" on page 1333.

On the Connectivity Map, double-click the Fabric OS HBA or CNA device you want to manage.

HCM Agent displays.

OR

On the Connectivity Map, right-click the Fabric OS HBA or CNA device you want to manage and select **Element Manager**.

HCM Agent displays.

OR

1. Select a Fabric OS HBA or CNA.
2. Select **Configure > Element Manager > HCM**.

   HCM Agent displays.

Launching Fabric Watch

> NOTE
> Fabric Watch is not supported on Fabric OS 7.4.0 and during firmware download from Fabric OS version 7.3 to 7.4. You must enable MAPS to migrate the existing Fabric Watch policies and monitor the switches. While monitoring 7.3 switches, the 7.4 switches present will be filtered out.

Use Fabric Watch as a health monitor that allows you to enable each switch to constantly monitor its SAN fabric for potential faults and automatically alerts you to problems long before they become costly failures. For more information about Fabric Watch, refer to the *Fabric Watch* ***Administrator's Guide.*** For more information about Fabro OS access gateway, switches, and directors, refer to the documentation for the specific device.

To launch Fabric Watch, complete the following steps.

> NOTE
> You must have Fabric Watch privileges for the selected device to launch Fabric Watch. If you do not have Fabric Watch privileges, you must enter those credentials to launch Fabric Watch. For more information about privileges, refer to "User Privileges" on page 1333.

> NOTE
> You must have the Fabric Watch license for the selected device.

On the Connectivity Map, right-click the Fabric OS device you want to monitor and select **Fabric Watch > Configure**.

Fabric Watch displays.

OR

1. Select a Fabric OS device.
2. Select **Monitor > Fabric Watch > Configure**.

   Fabric Watch displays.

Single sign-on support for IBM

NOTE
Single sign-on is not supported with IBM Tivoli Storage Productivity Center version 5.1.1 and later.

The Management application supports single sign-on (SSO) for IBM products such as IBM Tivoli Storage Productivity Center (TPC) or IBM Systems Director. Although SSO is not required, it creates a more seamless experience between the Management application server and IBM TPC or IBM Systems Director. There are several functions within the IBM TPC that launch the Management application client. If SSO is not enabled, each time the Management application client is launched, you must verify your Management application credentials. By enabling SSO, the Management application can authenticate against IBM TPC and launch the specified dialog box directly. This reduces the number of authentication steps required by you.

To configure the Management application to support SSO, complete the following steps.

1. Create the trust store on the IBM product.

   The trust store is used to establish SSL communication between the Management application and the IBM product for authentication. For instructions, refer to the IBM Systems Director or TPC documentation about configuring users.

2. Configure the Management application by completing the following steps.

   a. Copy the trust store to the Management application directory (***Install_Home***\bin\tpc).

      The Management application directory is located in ***Install_Home***\bin\tpc (Windows systems) or ***Install_Home***/bin/tpc (UNIX systems).

      The trust store is located where you specified in step 1.

   b. Open a **Command Prompt** window.

   c. Type **cd *Install_Home*\bin\tpc** and press **Enter** to go to the tpc directory.

   d. Type **tpcssosetup.bat** (Windows systems) or **sh tpcssosetup** (UNIX systems) with the following parameters:

      ```
      IP of the host where IBM product is running as the 1st parameter,
      The port number as the 2nd parameter, the default is 16311,
      The trust store name as the 3rd parameter,
      The password for the trust store as the 4th parameter,
      Basic authentication user name, this is a user in the LDAP server where IBM product authenticate with,
      as the 5th parameter, and basic authentication user's password the 6th parameter
      ```

      Example (Windows systems)

      ```
      tpcssosetup 10.32.1.1 16311 ibm_higgins_sso_10.32.1.1.jks password tipadmin super123
      ```

      Example (UNIX systems)

      ```
      sh tpcssosetup 10.32.1.1 16311 ibm_higgins_sso_10.32.1.1.jks password
      tipadmin super123
      ```

   e. Press **Enter** to configure single sign-on for the Management application.

3. Create a new user account in the Management application, including user name, password, and resource group.

   This account must match the IBM Systems Director or TPC user account. To create a user account, refer to "Creating a new user account" on page 138.

4. Make sure any switches you need to manage are discovered by the Management application. Add the switch or fabric into the Management applciation by selecting **Discovery > Setup > Add Fabric**.

   To discover a switch or fabric, refer to "Discovering fabrics" on page 35.

5. Restart the Management application.

Launch in context support for IBM

This Management application supports launch in context (LIC) for IBM products such as IBM Tivoli Storage Productivity Center (TPC) or IBM Systems Director. The Management application includes a package to deploy and remove the LIC menus for IBM TPC on Windows systems.

1. Copy tpc_*Application_Name*_ldf.zip to any directory on the TPC host.

   This procedure uses the *Install_Home*\conf\tpc\Win32 directory as an example.

2. Unzip the file and choose one of the following options:

   - To deploy the package, complete the following steps.
     a. Open the *Install_Home*\conf\tpc directory.
     b. Select **Start > Programs > Accessories > Command Prompt**.

     The **Command Prompt** window displays.

     c. Type **cd** *Install_Home*\conf\tpc and press **Enter** to go to the tpc directory.
     d. Type **tpc***Application_Name***ldfdeployer.bat** with the following the parameters and press **Enter** to to deploy the package.

     ```
     TIP install directory, no space, as the 1st parameter,
     Application_Name server domain as the 2nd parameter,
     Application_Name server name as the 3rd parameter, and
     Application_Name server port number, default 80, as the 4th parameter
     ```

     Example of deployment parameters

     ```
     tpcldfdeployer C:\Progra~1\IBM\tivoli\tip brocade.com myhost.engliab 80
     ```

   - To undeploy the package, complete the following steps.
     a. Open the *Install_Home*\conf\tpc directory.
     b. Select **Start > Programs > Accessories > Command Prompt**.

     The **Command Prompt** window displays.

     c. Type **cd** *Install_Home*\conf\tpc and press **Enter** to go to the tpc directory.
     d. Type **tpc***Application_Name***ldfundeployer.bat** with the first parameter and **Enter** to remove the package.

     First parameter is as follows:

     ```
     TIP install directory, no space, as the 1st parameter,
     ```

     Example

     ```
     tpcApplication_Nameldfundeployer  C:\Progra~1\IBM\tivoli\tip
     ```

3. Open the WSADMIN for TIP on the TPC server (C:\Program Files\IBM\tivoli\tip\bin\wsadmin.bat).
4. Type `$AdminTask modifyESSWSFedConfiguration {-domain ".domainname.com" -secure false}` and press **Enter**.

   NOTE
   The dot (.) in front of domainname is mandatory.

5. Restart the TCP data server for the menu to display.

Available LIC points

NOTE
Launch in Context (LIC) requires Oracle JRE. For the current supported JRE version for LIC, refer to the Release Notes. For the website listing patch information, go to *http://www.oracle.com/technetwork/java/javase/downloads/index.html*.

NOTE
LIC requires a Trial or Licensed version.

LIC enables you to launch the following dialog boxes:

- **Audit Log** dialog box
- **Bottleneck Detection** dialog box
- **DCB Configuration** dialog box
- *DCB_Name* **Edit Switch** dialog box, **QoS** tab
- **Configure Names** dialog box
- **Create View** dialog box
- **Device Connectivity Troubleshooting** dialog box
- **E-mail Event Notification Setup** dialog box
- **Encryption Center** dialog box
- **Event Log** dialog box
- **Fabric Binding** dialog box
- **Fabric Device Sharing Diagnosis** dialog box
- *Fabric_Name* **Historical Performance Graph** dialog box
- **FCIP Tunnels** dialog box
- **FCoE Configuration** dialog box
- **FICON Log** dialog box
- **Firmware Management** dialog box
- **Logical Switches** dialog box
- Main Interface
- **Port Fencing** dialog box
- **Product Status Log** dialog box
- **Real Time Port Picker** dialog box
- **Router Configuration - Connect Edge Fabric** *Fabric_Name* dialog box
- **Save Switch Configuration** dialog box
- **Security Log** dialog box
- **Set End-to-End Monitors** dialog box
- **Set Threshold Policies** dialog box
- **SMIA Configuration Tool** dialog box
- **Switch Configuration Repository** dialog box

- **Syslog Log** dialog box
- **Syslog Forwarding** dialog box
- **Technical Support Data** dialog box
- **Trace Route** dialog box
- **View Reports** dialog box (**Fabric Ports Report**)
- **View Reports** dialog box (**Historical Performance Report**)
- **VLAN Configuration** dialog box
- **Zoning** dialog box

Adding a tool

You can specify third-party tools so they appear on the **Setup Tools** dialog box. From there, you can add them to the **Tools** menu and then open the tools directly from the Management application.

To add a tool, complete the following steps.

1. Select **Tools > Setup**.

   The **Setup Tools** dialog box displays.

2. Click the **Tools Menu** tab.
3. Click **Define**.

   The **Define Tools** dialog box displays (Figure 166).

   FIGURE 166 Define Tools dialog box

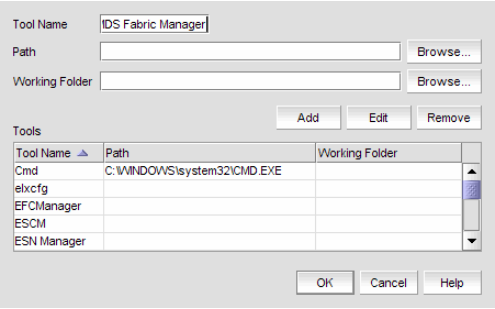

4. Type the tool's name in the **Tool Name** field as you want it to appear on the **Tools** menu.
5. Type or browse to the path of the executable file in the **Path** field.
6. Type or browse to the path of the folder that you want to set as your working folder in the **Working Folder** field.
7. Click **Add** to add the tool.

   The **Setup Tools** dialog box displays with the new tool added to the **Tools Menu Item** table.

   NOTE
   You must click **Add** before clicking **OK**; otherwise, your changes will be lost.

8. Click **OK** to save your work and close the **Define Tools** dialog box.

   To add this tool to the **Tools** menu, refer to "Adding an option to the Tools menu" on page 355.

9. Click **OK** to save your work and close the **Setup Tools** dialog box.

Entering the server IP address of a tool

If the third-party tool is a web-based application, you must enter the IP address of the applications server as a parameter to be able to open the application.

To enter the server IP address, complete the following steps.

1. Select **Tools > Setup**.

   The **Setup Tools** dialog box displays.

2. Click the **Tools Menu** tab.

   The **Tool Menu Items** table displays all configured tools, including the tool name as it displays on the **Tools** menu, parameters, and keystroke shortcuts.

3. Select the tool you want to edit in the **Tool Menu Items** table.

   The settings for the selected tool display in the fields at the top of the dialog box.

4. Edit the IP address of the server (for example, `http://`*IP_Address* or `http://`*IP_Address*`:`*Port_Number*) in the **Parameters** field.

5. Click **Edit**.

   NOTE
   You must click **Edit** before clicking **OK**; otherwise, your changes will be lost.

6. Click **OK** to save your work and close the **Setup Tools** dialog box.

Adding an option to the Tools menu

You can add third-party tools to the **Tools** menu which enables you to launch tools directly from the application.

To add an option to the **Tools** menu, complete the following steps.

1. Select **Tools > Setup**.

   The **Setup Tools** dialog box displays.

2. Click the **Tools Menu** tab.

   The **Tool Menu Items** table displays all configured tools, including the tool name as it displays on the **Tools** menu, parameters, and keystroke shortcuts (Figure 167).

FIGURE 167 Setup Tools dialog box (Tools Menu tab)

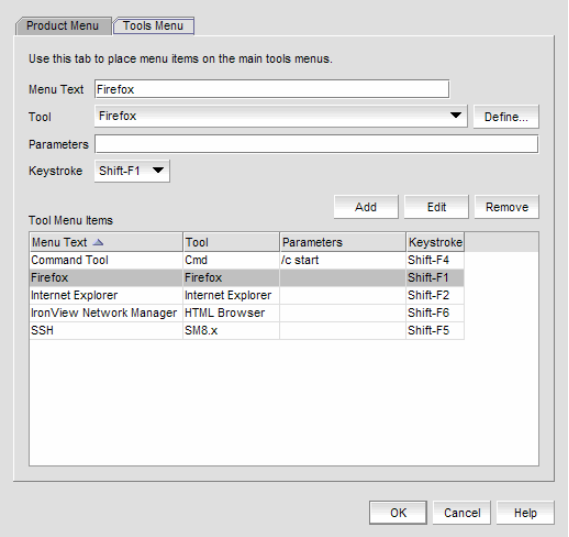

3. Type a label for the option as you want it to appear on the **Tools** menu in the **Menu Text** field.
4. Select the application from the **Tool** list, or click **Define** if you want to specify a new tool.

   To specify a new tool, refer to "Adding a tool" on page 354.
5. (*Optional*) Enter parameters, such as a URL, in the **Parameters** field.
6. (*Optional*) Select a keyboard shortcut in the **Keystroke** list.

   NOTE
   You cannot assign the same keyboard shortcut to two different tools.
7. Click **Add**.

   The new tool displays in the **Tool Menu Items** table.

   NOTE
   You must click **Add** before clicking **OK**; otherwise, the new menu option is not created.
8. Click **OK** to save your work and close the **Setup Tools** dialog box.

   The tool you configured now displays on the **Tools** menu.

Changing an option on the Tools menu

You can edit parameters for third-party tools that display on the **Tools** menu.

To edit an option on the **Tools** menu, complete the following steps.

1. Select **Tools > Setup**.

   The **Setup Tools** dialog box displays.
2. Click the **Tools Menu** tab.

The **Tool Menu Items** table displays all configured tools, including the tool name as it displays on the **Tools** menu, parameters, and keystroke shortcuts.

3. Select the tool you want to edit in the **Tool Menu Items** table.

   The settings for the selected tool display in the fields at the top of the dialog box.

4. Edit the label for the option as you want it to appear on the **Tools** menu in the **Menu Text** field.
5. Select the application from the **Tool** list.
6. Edit the parameters, such as a URL, in the **Parameters** field.
7. Select a new keyboard shortcut in the **Keystroke** list.
8. Click **Edit**.

   NOTE
   You must click **Edit** before clicking **OK**; otherwise, your changes will be lost.

9. Click **OK** to save your work and close the **Setup Tools** dialog box.

Removing an option from the Tools menu

You can remove a tool from the third-party tool list.

To remove an option on the **Tools** menu, complete the following steps.

1. Select **Tools > Setup**.

   The **Setup Tools** dialog box displays.

2. Click the **Tools Menu** tab.
3. Select the row of the tool you want to remove in the **Tools Menu Items** table.
4. Click **Remove**.

   If the tool is not being utilized, no confirmation message displays.

5. Click **Update** to remove the tool.
6. Click **OK** to save your work and close the **Setup Tools** dialog box.

Adding an option to a device's shortcut menu

You can add an option to a device's shortcut menu.

To add an option to the device's shortcut menu, complete the following steps.

1. Select **Tools > Setup**.

   The **Setup Tools** dialog box displays.

2. Click the **Product Menu** tab.

   The **Product Popup Menu Items** table displays all configured shortcut menu options.

3. Type or select the text in the **Menu Text** list as you want it to appear on the menu.
4. Choose one of the following options:
   - To display the menu option only for devices that meet the conditions listed, select the **Match Conditions** option.
   - To display the menu option on the shortcut menus for all devices, select the **All** option.

   If you select **All**, skip to step 8. Otherwise, continue to step 5.
5. Select the appropriate type in the **Condition 1 Property** name list.
6. Enter the appropriate value for the selected property in the **Condition 1 Value** field.
7. (*Optional*) Select the **Condition 2 Property** type and enter the **Value** for that property type (Condition 1 and Condition 2 must be true) to define a second condition to be simultaneously true.

   NOTE
   To set up a condition where Condition 1 or Condition 2 must be true, define two menu items, one for each condition.
8. Select the tool that you want to launch from the **Tool** list, or click **Define** to add a tool.

   To specify a new tool, refer to "Adding a tool" on page 354.
9. Select the **Append device ID** check box to specify the parameter used when opening the tool.
   - To specify that the device's IP address should be used when opening the tool, select the **IP Address** option.
   - To specify that the device's Node WWN should be used when opening the tool, select the **Node WWN** option.
10. Click **Add** to add the new menu item.

    It displays in the **Product Popup Menu Items** table.

    NOTE
    You must click **Add** before clicking **OK**; otherwise, your changes will be lost.
11. Click **OK** to save your work and close the **Setup Tools** dialog box.

Changing an option on a device's shortcut menu

You can change the parameters for a tool that displays on a device's shortcut menu.

To edit an option on the device's shortcut menu, complete the following steps.

1. Select **Tools > Setup**.

   The **Setup Tools** dialog box displays.
2. Click the **Product Menu** tab.

   The **Product Popup Menu Items** table displays all configured shortcut menu options.
3. Select the menu item you want to change in the **Product Popup Menu Items** table.

   The settings for the selected menu item display in the fields at the top of the dialog box.
4. Edit or select the text in the **Menu Text** list as you want it to appear on the menu.

5. Choose one of the following options:
   - To display the menu option only for devices that meet the conditions listed, select the **Match Conditions** option.
   - To display the menu option on the shortcut menus for all devices, select the **All** option.

     If you select **All**, skip to step 8. Otherwise, continue to step 5.
6. Change the type in the **Condition 1 Property** name list.
7. Change the value for the selected property in the **Condition 1 Value** field.
8. (*Optional*) Change the **Condition 2 Property** type or edit the **Value** for that property type (Condition 1 and Condition 2 must be true) to edit a second condition to be simultaneously true.

   NOTE
   To set up a condition where Condition 1 or Condition 2 must be true, define two menu items, one for each condition.
9. Select the tool from the **Tool** list that you want to launch, or click **Define** to add a tool.

   To specify a new tool, refer to "Adding a tool" on page 354.
10. Select the **Append device ID** check box to specify the parameter used when opening the tool.
    - To specify that the device's IP address should be used when opening the tool, select the **IP Address** option.
    - To specify that the device's Node WWN should be used when opening the tool, select the **Node WWN** option.
11. Click **Edit**.

    NOTE
    You must click **Edit** before clicking **OK**; otherwise, your changes will be lost.
12. Click **OK** to save your work and close the **Setup Tools** dialog box.

Removing an option from a device's shortcut menu

You can remove a tool that displays on a device's shortcut menu.

To remove an option on the device's shortcut menu, complete the following steps.

1. Select **Tools > Setup**.

   The **Setup Tools** dialog box displays.
2. Click the **Product Menu** tab.

   The **Product Popup Menu Items** table displays all configured menu options.
3. Select the menu item you want to remove in the **Product Popup Menu Items** table.
4. Click **Remove**.
5. Click **OK** to save your work and close the **Setup Tools** dialog box.

Microsoft System Center Operations Manager plug-in

NOTE
The System Center Operations Manager (SCOM) plug-in is only supported on Windows.

NOTE
The SCOM plug-in is only available on Professional Plus and Enterprise versions.

NOTE
You must have SCOM Management privileges to access the **Plug-in for SCOM** dialog box. For more information about privileges, refer to "User Privileges" on page 1333.

The Microsoft System Center Operations Manager (SCOM) plug-in allows fabric inventory information collected by the Management application to be displayed on the Microsoft SCOM console. The SCOM plug-in uses the SCOM SDK services to extend the SCOM console and present fabric inventory information. The SCOM plug-in serves dynamic HTML pages to the SCOM console.

The SCOM console displays the following information:

- Fabric and switch details
- End-to-end monitor statistics
- Events from the Management application (refer to "Configuring event forwarding to the SCOM console" on page 362)

The SCOM plug-in is supported on the following configurations:

- SCOM 2007 R2 or SCOM 2012
- Professional Plus and Enterprise Trial and Licensed version 11.0.0 and later

SCOM plug-in requirements

- Make sure you import the Management application management pack (*Management_Application_Name*.FabricView.xml) to the SCOM Server prior to registering the SCOM plug-in. The management pack is located in the following directory: *Install_Home*\scom.
- Make sure the Management application server host is managed by the SCOM Server in agent managed mode.
- Make sure the host name starts with an alphabet.
- Make sure the SCOM HealthService agent is running on the Management application server.
- Make sure you install the SCOM Console 2007 R2 software on the Management application server.
- (*Optional*) Enable SSL on the *Management_Application_Name* to use HTTPS Communication between SCOM Console and the Management application.
- Make sure that the fabric or switch is managed by the the Management application to view fabric and switch details.
- Make sure to enable performance monitoring at the SAN or fabric level to collect end-to-end monitor statistics. Refer to "SAN end-to-end monitoring" on page 979.

Registering a SCOM server

To register the SCOM server, complete the following steps.

1. Select **Tools > Plug-in for SCOM**.

   The **Plug-in for SCOM** dialog box displays.

2. Click **Add**.

   The **Add SCOM Server** dialog box displays.

3. Enter an IP address or fully qualified domain name for the SCOM host in the **Host** field.

   The Management application accepts IP addresses in IPv4 and IPv6 formats. The IPv4 format is valid when the operating system has IPv4 mode only or dual stack mode. The IPv6 format is valid when the operating system has IPv6 mode only or dual stack mode.

4. Enter the domain name in the **Domain** field.
5. Enter your user ID and password.
6. Click **OK**.
7. Click **Close**.

Editing a SCOM server

To edit the SCOM server, complete the following steps.

1. Select **Tools > Plug-in for SCOM**.

   The **Plug-in for SCOM** dialog box displays.

2. Select the server you want to edit and click **Edit**.

   The **Edit SCOM Server** dialog box displays. The **Host** field is not editable in the **Edit SCOM Server** dialog box.

3. Edit the domain name in the **Domain** field.
4. Enter your user ID and password.
5. Click **OK**.
6. Click **Close**.

Removing a SCOM server

To remove the SCOM server, complete the following steps.

1. Select **Tools > Plug-in for SCOM**.

   The **Plug-in for SCOM** dialog box displays.

2. Select the SCOM server you want to delete in the **SCOM Servers** table.
3. Click **Remove**.
4. Click **OK** on the confirmation message.
5. Click **Close**.

Configuring event forwarding to the SCOM console

You can configure the Management application to forward CallHome events as well as events based on severity to the SCOM comsole. You can also configure to include any events notes with the events.

To configure what events to forward to the SCOM console, complete the following steps.

1. Open the scom configuration file (*Install_Home*/conf/scom.conf) a text editor (such as Notepad).
2. Enable the Management application to forward CallHome events by changing the file as follows:

   `scom.forward.events.callhome=TRUE`

   By default, CallHome event forwarding to the SCOM console is enabled. To disable CallHome event forwarding, set scom.forward.events.callhome to FALSE.

3. Define the level of event severity that you want by defining the ***event_severity.***

   `scom.forward.events.severity=`***event_severity,*** where events with the configured ***event_severity*** and higher severity levels are forwarded.

   By default, the Error severity is defined. Therefore, events with Critical, Alert, and Emergency severity levels are forwarded. To forward all events, set ***event_severity*** to UNKNOWN, the lowest severity level. To forward no events, set ***event_severity*** to NONE.

   The valid event severities in increasing order of severity are as follows:

   NONE

   UNKNOWN

   INFO

   DEBUG

   NOTICE

   WARNING

   ERROR

   CRITICAL

   ALERT

   EMERGENCY

4. Include event notes, if any, with the forwarded events by changing the file as follows:

   `scom.forward.events.notes.include=true`

   By default, including notes with the fowarding events is enabled. To not include event notes, set `scom.forward.events.notes.include` to false.

5. Save and close the file.

Server Management Console

Server Management Console overview

NOTE
The Server Management Console (SMC) requires Oracle JRE. For the current supported JRE version for the SMC, refer to the Release Notes. For the website listing patch information, go to *http://www.oracle.com/technetwork/java/javase/downloads/index.html*.

The SMC is an automatically installed, stand-alone application for managing the Management application server. You can perform the following tasks using the SMC:

- From the Services tab, you can start, stop, refresh, and restart services on the server.
- From the Ports tab, you can view the Management application server or web server port number.
- From the AAA Settings tab (Enterprise Licensed version only), you can configure an authentication server (LDAP or Radius server), and establish authentication policies.
- From the Restore tab, you can restore server application data.
- From the Technical Support Information tab, you can collect information for technical support.
- From the HCM Upgrade tab, you can upgrade the Management application to use a new version of Host Connectivity Manager (HCM).
- From the SMI Agent Configuration Tool, you can configure the SMI Agent settings, such as security, CIMOM, and certificate management as well as launch Management application dialog boxes.

Launching the SMC on Windows

Open the **Server Management Console** from the **Start** menu on the Management application server.

You can also drag the SMC icon onto your desktop as a short cut.

Launching the SMC on Linux

NOTE
The Server Management Console is a graphical user interface and should be launched from the XConsole on Linux systems.

Perform the following steps to launch the Server Management Console on Linux systems.

1. On the Management application server, go to the following directory:

   Install_Directory/bin

2. Type the following at the command line:

```
./smc
OR
sh smc
```

Services tab

You must be logged in at the administrator (Windows systems) or root (UNIX systems) level to stop, start, and restart the Management application services. Stopping and restarting the Management application services causes clients connected to the server to lose connection, and they must re-log in to the server.

Monitoring and managing Management application services

To monitor the status of the Management application services, complete the following steps.

1. Launch the Server Management Console.
2. Click the **Services** tab (Figure 168).

FIGURE 168 Services tab

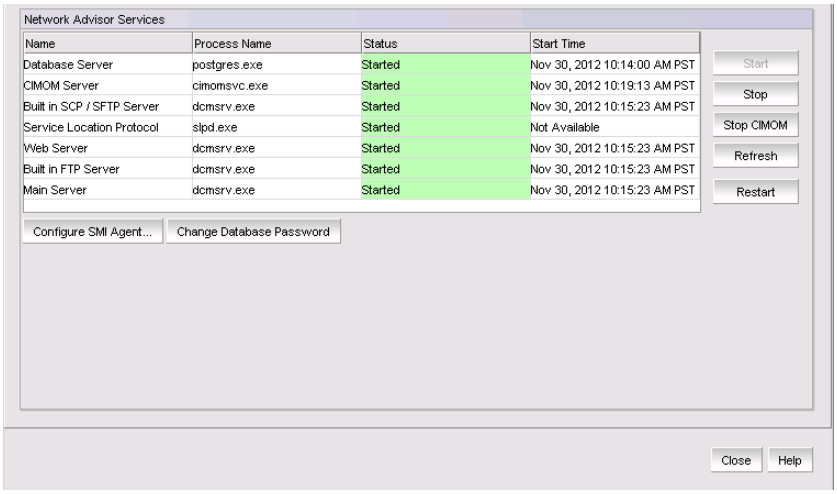

3. Review the following information for each available service.
   - **Name** — The name of the server; for example, FTP Server or Database Server.
   - **Process Name** — The name of the process; for example, postgres.exe (Database Server).

- **Status** — The status of the service; for example, started or stopped.
- **Start Time** — The date and time the service started. The Start Time for Service Location Protocol displays as 'Not Available'.

4. Click **Close** to close the Server Management Console.

Refreshing the server status

To refresh the server status for each of the Management application services, complete the following steps.

1. Launch the Server Management Console.
2. Click the **Services** tab.
3. Click **Refresh** to update the table with the latest status of the services in case the services were stopped or restarted outside of the Server Management Console.
4. Click **Close** to close the Server Management Console.

Stopping all services

To stop all services, complete the following steps.

1. Launch the Server Management Console.
2. Click the **Services** tab.
3. Click **Stop** to stop all services.

   Note that clicking **Restart** stops and then restarts all services.

4. Click **Close** to close the Server Management Console.

Stopping the CIMOM services

To stop the CIMOM (Common Information Model Object Manager) services, complete the following steps.

1. Launch the Server Management Console.
2. Click the **Services** tab.
3. Click **Stop CIMOM**.
4. Click **Close** to close the Server Management Console.

Starting all services

NOTE
The **Start** button restarts running services in addition to starting stopped services which causes client-server disconnect.

To start all services, complete the following steps.

1. Launch the Server Management Console.
2. Click the **Services** tab.

3. Click **Start** to start all services.

   NOTE
   If the server is configured to use an external FTP server, the Server Management Console does not attempt to start the built-in FTP service.

4. Click **Close** to close the Server Management Console.

Restarting all services

To stop and restart all services, complete the following steps.

1. Launch the Server Management Console.
2. Click the **Services** tab.
3. Click **Restart** to stop then restart all services.

   NOTE
   If the server is configured to use an external FTP server, the Server Management Console does not attempt to start the built-in FTP service.

4. Click **Close** to close the Server Management Console.

Changing the database password

Requires User Management read and write privilege.

1. Launch the Server Management Console.
2. Click the **Services** tab.
3. Click **Change Database Password**.

   The authentication **Login** dialog box displays.

4. Enter your Management application user name and password.
5. Click **OK**.

   The **Database Password** dialog box displays.

6. Select the database user name for which you want to change the password in the **User Name** field.

   Options include dcmadmin and dcmuser.

   Changing the dcmadmin password requires all Management application services, except for the database server, to be stopped and then re-started.

   Changing the dcmuser password requires all ODBC remote client sessions to be restarted.

7. Enter your current password in the **Old Password** field.
8. Enter you new password in the **New Password** and **Confirm New Password** fields.
9. Click **OK**.
10. Click **Yes** on the warning message.

Ports tab

Use the **Ports** tab of the Server Management Console to view the Management application server and Web server port numbers. The default Web Server port number is 80 (HTTP) or 443 (HTTPS). The Management application server default port number is 24600.

Viewing server port numbers

To view the Management application server or web server port number, complete the following steps.

1. Choose one of the following options:
   - For Windows systems, open the **Server Management Console** from the **Start** menu on the Management application server.
   - For Linux systems, on the Management application server, go to the *Install_Directory*/bin directory and type the following at the command line:

     ```
     ./smc
     OR
     sh smc
     ```
2. Click the **Ports** tab.
3. Review the following information for each available service.
   - *Management_Application_Name* **Server Port** text box — The Management application Server Port number. The default is 24600.
   - **Web Server Port # (HTTPS)** text box — The Web Server Port number for HTTPS. The default is 443.

     You can configure the server port settings from the **Options** dialog box (**Server Port** pane). For instructions, refer to "Configuring the server port" on page 127.

     You can also configure the server port settings from the configuration wizard. For instructions, refer to "Launching the Configuration Wizard" on page 6.
4. Click **Close** to close the Server Management Console.

AAA Settings tab

Authentication enables you to configure an authentication server and establish authentication policies. You can configure the Management application to authenticate users against the local database (Management application server), an external server (RADIUS, LDAP, CAC or TACACS+), or a switch. Authentication is configured to the local database by default. When you use an external server, the Management application sends the login information to the external server to make sure the name and password are valid.

If you configure primary authentication to an external or switch authentication, you can also configure secondary authentication to the local server. When you log in to the Management application, if the primary server is unavailable, the Management application attempts with the next configured primary server. If all primary servers are unavailable, then the Management application falls back to the secondary authentication. Fall back can occur when the server is unavailable, authentication fails, or the user is not found.

Configuring Radius server authentication

If you are using a Radius server for authentication, make the following preparations first:

- Make sure that the server you want to use is on the network that the Management application manages.
- Make sure that the Radius server and its user accounts have been properly configured (refer to "Radius server configuration" on page 380). For example, you must define the Management application client, users, and dictionary on the Radius server.

- Make sure that the external server and its user accounts have been properly configured. For example, you must define roles and areas of responsibility (AOR) in the external server to match the Management application roles and AOR.
- Select an **Authentication Type** (you will be prompted to provide a type in the **Add or Edit Radius Server** dialog box). The **Authentication Type** is the authentication policy you choose for handling authentication. The options are PAP and CHAP.
  - PAP, password protected protocol, is based on password verification. Passwords are not encrypted, and are not secure from eavesdroppers during transmission.
  - CHAP, challenge handshake protocol, uses a three-way handshake method of verification based on a shared secret. If you are using CHAP, have the shared secret available to you. You will need to type it in as a configuration parameter.
- Know the Shared Secret.
- Have the IP address of the server available.

  NOTE
  The Management server and the RADIUS server must be running the same Internet Protocol (IPv4 or IPv6). If the Management server and the RADIUS server are running different protocols, communication fails.

- Know the TCP port you are using and make sure it is open in the firewall. For Radius servers, ports 1812 or 1813 (actually UDP ports) are commonly used. Some older Radius server use 1645 or 1646 instead of 1812 and 1813; check with the Radius server vendor if you are not sure which port to specify.
- Know how long you want to wait between attempts to reach the server if it is busy. This is expressed as a timeout value (default is 3 seconds) in seconds. Values are between 1 and 15.
- Determine how many attempts (default is 3 times) to make to reach the server before stopping and assuming it is unreachable. Values are between 1 and 5.
- If possible, establish an active connection with the Radius server before configuration. This enables you to test the connection as part of the configuration procedure.

1. Select the **AAA Settings** tab (Figure 169).

FIGURE 169 AAA Settings tab

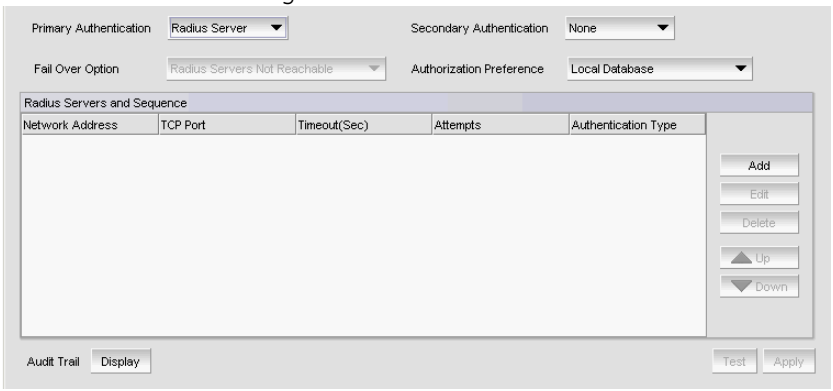

2. Select **Radius Server** from the **Primary Authentication** list.
3. Add or edit a Radius server by referring to "Configuring a Radius server" on page 369.
4. Rearrange the Radius servers in the table by selecting a server and click the **Up** or **Down** button to move it.
5. Delete a Radius server by selecting the server and click **Delete**.
6. Test the established active connection with the Radius server by clicking **Test**.

   Test attempts to contact the Radius server by issuing a **ping** command.

7. Set secondary authentication by selecting one of the following options from the **Secondary Authentication** list:
   - **Local Database**
   - **None**
8. Set the fall back condition to secondary authentication by selecting one of the following options from the **Fail Over Option** list:
   - **Radius Servers Not Reachable**
   - **Radius Authentication Failed**
9. Set the authorization preference by selecting one of the following options from the **Authorization Preference** list:
   - **Local Database**
   - **Primary Authentication Server**
10. Click **Apply** to save the configuration.

    To display the authentication audit trail, refer to "Displaying the client authentication audit trail" on page 380.
11. Click **Close** to close the Server Management Console.

    Confirm authentication and authorization by logging into the Management application server (refer to "Logging in to a server from the server machine" on page 2).

Configuring a Radius server

To add or edit a Radius server, complete the following steps.

1. Choose one of the following options from the **AAA Settings** tab:
   - Click **Add**.
   - Select an existing Radius server and click **Edit**.

     The **Add or Edit Radius Server** dialog box displays (Figure 170).

FIGURE 170 Add or Edit Radius Server

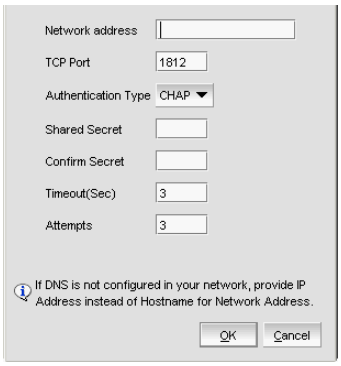

2. Enter the radius server's IP address in the **IP Address** field.
3. Enter the TCP port, if necessary, used by the Radius server in the **TCP Port** field.

   Default is 1812.
4. Select the authentication policy (PAP or CHAP) from the **Authentication Type** field.

   Default is CHAP.

5. Enter the shared secret in the **Shared Secret** and **Confirm Secret** fields.
6. Enter the timeout timer value (in seconds) that specifies the amount of time to wait between retries when the server is busy in the **Timeout (Sec)** field.

   Default is 3 seconds.
7. Enter the number of attempts to be made to reach a server before assuming it is unreachable in the **Attempts** field.

   Default is 3 attempts.
8. Click **OK** to return to the **AAA Settings** tab.

   The **Radius Servers and Sequence** table displays the following information:

   - **Network Address** — The network address of the Radius server.
   - **Authentication Type** — The authentication type (such as, CHAP).
   - **TCP Port** — The TCP port number of the Radius server.
   - **TimeOut (Sec)** — The timeout value in seconds specified when sending an authentication request to the server. Default is 3.
   - **Attempts** — The number of attempts made to reach a server before determining it is unreachable. Default is 3.

Configuring LDAP server authentication

NOTE
You cannot configure multiple Active Directory groups (domains) for the LDAP server.

NOTE
You cannot enter *Domain\User_Name* in the Management application dialog box for LDAP server authentication.

NOTE
By default, LDAP server is configured with *defaultNamingContext* attribute. The server cannot fetch data if the attribute is not available. Please contact Technical support or Network Administrator to add *defaultNamingContext* attribute manually.

If you configure the external LDAP server as the primary authentication server, make the following preparations first:

- Make sure that the external LDAP server and its user accounts have been properly configured (refer to "LDAP server configuration" on page 382). For example, you must define roles and areas of responsibility (AOR) in the external server to match the Management application roles and AOR.
- Make sure to configure the custom attributes "NmRoles" and "NmAors" on the LDAP server (refer to "Configuring roles and AORs on the external LDAP server" on page 383). NmRoles defines the Management application user roles (such as Host Administrator, Network Administrator, Operator, Report User Group, SAN System Administrator, Security Administrator, Security Officer, and Zone Administrator). NmAors defines the areas of responsibility (such as All Fabrics or All Hosts).

If you are using an LDAP server for authentication, make the following preparations first:

- Make sure that the LDAP server you want to use is on the network that the Management application manages.
- Have the IP address of the server available.
- Know the TCP port you are using. The LDAP server uses Transport Layer Security (TLS). LDAP over TLS generally uses port 389. If security is enabled the port number is 636. Check with the LDAP server administrator if you are not sure which port to specify.
- Know how long you want to wait between attempts (default is 3 seconds) to reach the server if it is busy. This is expressed as a timeout value in seconds. Values are between 1 and 15.

- Determine how many attempts (default is 3 times) to make to reach the server before stopping and assuming it is unreachable. Values are between 1 and 5.

  NOTE
  If the LDAP server's IP address is entered in the Management application, the LDAP server's hostname (if any) must still be known to the Management application host OS. The Management application server must be using a DNS server that knows the LDAP server's hostname, or you must manually add the LDAP server's hostname to the local hosts file (for Linux the file is located in /etc/hosts and for Windows the file is located in C:\Windows\System32\drivers\etc\hosts for Windows).

To configure an LDAP server for authentication, complete the following steps.

1. Select the **AAA Settings** tab.
2. Select **LDAP Server** from the **Primary Authentication** list.

   FIGURE 171 AAA Settings tab - LDAP server

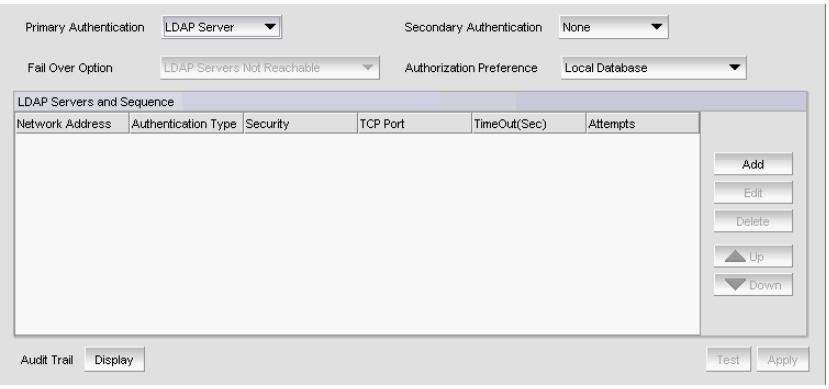

3. Add or edit an LDAP server by referring to "Configuring an LDAP server" on page 372.

   The **LDAP Servers and Sequence** table displays the following information:

   - **Network Address —** The network address of the LDAP server.
   - **Authentication Type —** The authentication type (such as CHAP).
   - **Security —** Whether or not security is enabled.
   - **TCP Port —** The TCP port number of the LDAP server.
   - **TimeOut (Sec) —** The timeout value in seconds specified when sending an authentication request to the server. Default is 3.
   - **Attempts —** The number of attempts made to reach a server before determining it is unreachable. Default is 3.

     NOTE
     You can add any number of LDAP servers to the configuration.

4. Rearrange the LDAP servers in the table by selecting a server and clicking the **Up** or **Down** button to move it.
5. Delete a LDAP server by selecting the server and clicking **Delete**.
6. Test the established active connection with the LDAP server by clicking **Test**.

   The **Test Authentication** dialog box displays.

7. Enter your user name and password and click **OK**.

   Test attempts to contact the LDAP server by issuing a **ping** command and verifies the following:

- Verifies connections to the LDAP Server
- Verifies authentication with the LDAP Server
- Verifies user privileges on the Local database

8. Set secondary authentication by selecting one of the following options from the **Secondary Authentication** list:
   - **Local Database**
   - **None**
9. Set the fall back condition to secondary authentication by selecting one of the following options from the **Switch to secondary authentication when** list:
   - **LDAP Servers Not Reachable**
   - **LDAP Authentication Failed**
   - **User Not Found in LDAP**
10. Set the authorization preference by selecting one of the following options from the **Authorization Preference** list:
    - **Local Database**
      - Use the LDAP server for authentication and the Management application local database for authorization.
      - The user name in the local database must match the LDAP user name (password does not need to match) and must have the appropriate roles and AORs. If the Management application user name and LDAP user name do not match, create the user and assign the respective roles and AORs (refer to "User Account Management" on page 135).
    - **Primary Authentication Server**
      - Use the LDAP server for authentication and authorization.
      - In the LDAP server, create new custom attributes (NmRoles & NmAors) in the AD server and assign the appropriate Roles and AORs (refer to "Configuring roles and AORs on the external LDAP server" on page 383).
        If this user already exists in the local database, the roles and AORs are overwritten with the new roles and AORs configured in the LDAP Server.
    - **LDAP Authorization**
      - Use to assign roles and AORs to user groups and not to individual users.
      - When roles and AORs are assigned to a group, all AD users in the group can obtain the roles and AORS assigned to the group. To assign roles and AORs to an AD Group, refer "Assigning roles and AORs to an AD group" on page 386.
        You do not need to create users in the local database.
11. Click **Apply** to save the configuration.

    To display the authentication audit trail, refer to "Displaying the client authentication audit trail" on page 380.
12. Click **Close** to close the Server Management Console.

    Confirm authentication and authorization by logging into the Management application server (refer to "Logging in to a server from the server machine" on page 2).

Configuring an LDAP server

To add or edit a LDAP server, complete the following steps.

1. Select the **AAA Settings** tab.
2. Select **LDAP Server** from the **Primary Authentication** list.
3. Choose one of the following options:
   - Click **Add**.

- Select an existing LDAP server and click **Edit**.

  The **Add or Edit LDAP Server** dialog box displays (Figure 172).

FIGURE 172 Add or Edit LDAP server

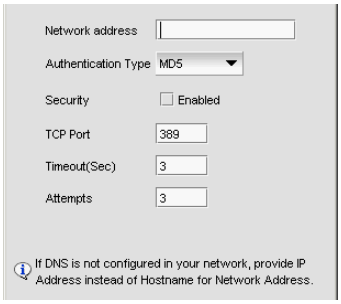

4. Enter the LDAP server's hostname in the **Network address** field.

   If DNS is not configured in your network, provide an IP address instead of the hostname.

5. Enable security by selecting the **Security Enabled** check box.

   When you enable security, the TCP port number automatically changes to port 636 and you must enable certificate services on the LDAP server.

6. Enter the TCP port used by the LDAP server in the **TCP Port** field.

   Default is 389 if security is not enabled. Default is 636 if security is enabled.

7. Enter the timeout timer value (in seconds) that specifies the amount of time to wait between retries when the server is busy in the **Timeout (Sec)** field.

   Default is 3 seconds.

8. Enter the number of attempts to be made to reach a server before assuming it is unreachable in the **Attempts** field.

   Default is 3 attempts.

9. Click **OK** to return to step 4 on the **AAA Settings** tab.

Configuring TACACS+ server authentication

If you are using a TACACS+ server for authentication, make the following preparations first:

- Make sure that the server you want to use is on the network that the Management application manages.
- Make sure that the external server and its user accounts have been properly configured. For example, you must define roles and areas of responsibility (AOR) in the external server to match the Management application roles and AOR.

To configure TACACS+ server authentication, complete the following steps.

1. Select the **AAA Settings** tab.
2. For **Primary Authentication**, select **TACACS+ Server**.

FIGURE 173 AAA Settings tab - TACACS+ server

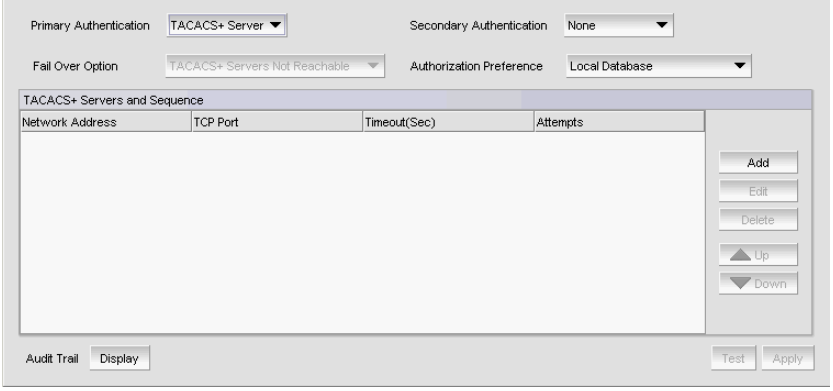

3. Add or edit a TACACS+ server by referring to "Configuring a TACACS+ server" on page 375.
4. Rearrange the TACACS+ servers in the table by selecting a server and click the **Up** or **Down** button to move it.
5. Delete a TACACS+ server by selecting the server and click **Delete**.
6. Test the established active connection with the TACACS+ server by clicking **Test**.

   The **Test Authentication** dialog box displays.
7. Enter your user ID and password and click **Test**.

   Test verifies your user ID and password for the local database and verifies user privileges on the Management application server.
8. Set secondary authentication by selecting one of the following options from the **Secondary Authentication** list:
   - **Local Database**
   - **None**
9. Set the fall back condition to secondary authentication by selecting one of the following options from the **Fail Over Option** list:
   - TACACS+ Server **Not Reachable**
   - TACACS+ Server **Authentication Failed**
10. Set the authorization preference by selecting one of the following options from the **Authorization Preference** list:
    - Local Database
    - Primary **Authentication Server**
11. Click **Apply** to save the configuration.

    To display the authentication audit trail, refer to "Displaying the client authentication audit trail" on page 380.
12. Click **Close** to close the Server Management Console.

    Confirm authentication and authorization by logging into the Management application server (refer to "Logging in to a server from the server machine" on page 2).

Configuring a TACACS+ server

To add or edit a TACACS+ server, complete the following steps.

1. Choose one of the following options from the **AAA Settings** tab:
   - Click **Add**.
   - Select an existing TACACS+ server and click **Edit**.

     The **Add or Edit TACACS+ Server** dialog box displays (Figure 172).

   FIGURE 174 Add or Edit TACACS+ Server

   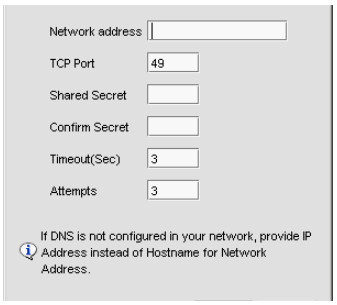

2. Enter the TACACS+ server's hostname in the **Network Address** field.

   If DNS is not configured in your network, provide an IP address instead of the hostname.

3. Enter the TCP port used by the TACACS+ server in the **TCP Port** field.

   Default is 49.

4. Enter the shared secret in the **Shared Secret** and **Confirm Secret** fields.

5. Enter the timeout timer value (in seconds) that specifies the amount of time to wait between retries when the server is busy in the **Timeout (Sec)** field.

   Default is 3 seconds.

6. Enter the number of attempts to be made to reach a server before assuming it is unreachable in the **Attempts** field.

   Default is 3 attempts.

7. Click **OK** to return to the **AAA Settings** tab.

   The **Radius Servers and Sequence** table displays the following information:
   - **Network Address —** The network address of the TACACS+ server.
   - **TCP Port —** The TCP port number of the LDAP server.
   - **TimeOut (Sec) —** The timeout value in seconds specified when sending an authentication request to the server. Default is 3.
   - **Attempts —** The number of attempts made to reach a server before determining it is unreachable. Default is 3.

Configuring Common Access Card authentication

NOTE
Common Access Card (CAC) authentication does not support SMI Agent and launch-in-context dialog boxes.

NOTE
CAC authentication is only supported with Active Client Library version 6.1 and 6.2.

NOTE
CAC authentication is only supported on Windows systems.

Common Access Card (CAC) authentication requires the following preparations:

- Make sure to connect the CAC reader to the Management application client workstation.
- Make sure to obtain and install the active client library on the client workstation. The active client library is not shipped with the Management application.
- Make sure to log in to the Management application client using a smartcard.
- Make sure that the Active Directory (AD) server you want to use is on the network that the Management application manages.
- Make sure that the Management application server and client system clocks are synchronized even if they are in different time zones.
- Make sure that the AD server you want to use is connected to the Management application client.
- Make sure you have the username and password of the Management application service account configured on the AD server to which the client is connected. It is recommended that you create and use the following name for this account: NetworkMangementSVC.

  NOTE
  If there are Management application clients from different domains, then each client's AD server must be configured with same user account and Kerberos Service Principal Name (SPN)

- Make sure you have the Kerberos SPN that is configured on the Key Distribution Center (KDC) of the AD server and map it to the Management application server account. It is recommended that you create and use the following name for this account: NetworkMangementSPN.

  If you need to add a Kerberos SPN to the KDC of the AD server, use the following command on the Management application client or the AD server to which the client is connected:

  `setspn -S` <*SPN*>/<*Management application server host name with domain name*><*AD server user account*>

  For example: setspn -S NetworkManagementSPN/DCM-VNext-65.JCB.com NetworkManagementSvc

  NOTE
  If there are multiple Management application servers, then a Kerberos Service Principal Name must be added for each server.

To configure CAC authentication, complete the following steps.

1. Select the **AAA Settings** tab.
2. Select **CAC** from the **Primary Authentication** list.

FIGURE 175 AAA Settings tab - CAC server

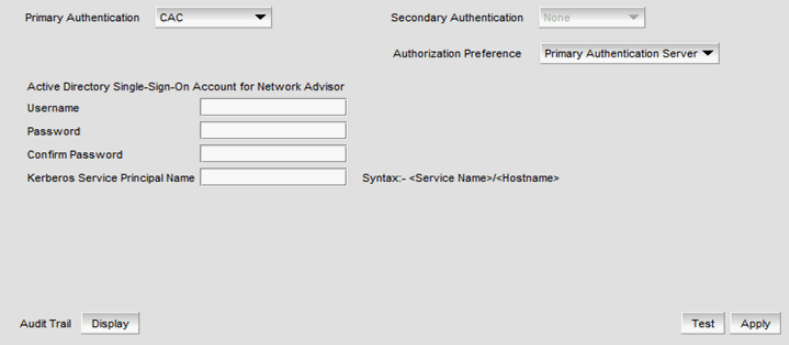

3. Set the authorization preference by selecting one of the following options from the **Authorization Preference** list:

   - **Local Database —** Uses the AD server for authentication and the Management application local database for authorization.
   - **Primary Authentication Server —** Uses the AD server for authentication and authorization.

   If you select **Primary Authentication Server** or **LDAP Authorization**, CAC authentication uses the same AD servers for authentication and authorization.

4. Enter the username for the Management application service account configured on the AD server in the **Username** field.
5. Enter the password for the Management application service account configured on the AD server in the **Password** and **Confirm Password** fields.
6. Enter the Kerberos SPN in the **Kerberos Service Principal Name** field.

   The SPN name uses the following syntax: <*Service_Name*>/<*Hostname*>, where hostname is the Management application server's host name with domain name. For example: NetworkManagementSPN/DCM-VNext-65.JCB.COM

7. Test the established active connection with the server by clicking **Test**.

   The **Test Authentication** dialog box displays. Test performs the following functions and verifications:

   - Obtains the Kerberos Ticket Granting Ticket (TGT) of the currently logged in user from Windows cached credentials.
   - Sends the TGT to the AD server to which the Management application server is connected and requests the session ticket for the SPN configured on AD server.

     Kerberos encryptsthe session ticket with the credentials of the AD server user account mapped to this SPN.
   - Logs on to the AD of the Management application server using the AD server single-sign-on (SSO) service account.
   - Verifies the service ticket by decrypting it using AD server SSO service account credentials.

8. Click **Apply** to save the configuration.

   To display the authentication audit trail, refer to "Displaying the client authentication audit trail" on page 380.

9. Click **Close** to close the Server Management Console.

Configuring switch authentication

Switch authentication enables you to authenticate a user account against the switch database and the Management application server. You can configure up to three switches and specify the fall back order if one or more of the switches is not available.

> NOTE
> Switch authentication is only supported on Fabric OS devices.

To configure switch authentication, complete the following steps.

1. Select the **AAA Settings** tab.
2. For **Primary Authentication**, select **Switch**.
3. Click **Add**.
4. Enter the switch IP address and click **OK**.

   You can add up to three switches.
5. Select a switch and click the **Up** or **Down** button to set the fall back order.
6. Select a switch and click **Delete** to remove a switch from the list.
7. Set secondary authentication by selecting one of the following options from the **Secondary Authentication** list:
   - **Local Database**
   - **None**
8. Click **Test**.

   The **Test Authentication** dialog box displays.
9. Enter your user ID and password and click **Test**.

   Test verifies your user ID and password on the switch and verifies user privileges on the Management application server.
10. Click **Apply** to save the configuration.

    To display the authentication audit trail, refer to "Displaying the client authentication audit trail" on page 380.
11. Click **Close** to close the Server Management Console.

Configuring Windows authentication

Windows authentication enables you to authenticate a user account against the Windows user accounts and the Management application server when running on Windows hosts.

The following list details the supported Windows authentication types and the associated platforms:

- NT domain authentication — supported on Windows XP/2003/2008 platforms only
- Windows Workgroup authentication — supported on Windows XP/2003/2008 platforms only
- Windows local user accounts — supported on Windows XP/2003/2008 platforms only.

To configure Windows authentication, complete the following steps.

1. Select the **AAA Settings** tab.
2. For **Primary Authentication**, select **Windows Domain**.

3. Enter the domain name in the **Windows Domain Name** field.
4. Set secondary authentication by selecting one of the following options from the **Secondary Authentication** list:
   - **Local Database**
   - **None**
5. Click **Test**.

   The **Test Authentication** dialog box displays.

1. In the **User ID** field, choose one of the following options:
   - To authenticate a user account against the current domain, enter your user name.
   - To authenticate a user account against a different domain, enter *Domain\User_Name*.
2. Enter your password in the **Password** field and click **OK**.

   Test verifies your user ID and password on the Windows domain and verifies user privileges on the Management application server.
3. Click **Apply** to save the configuration.

   To display the authentication audit trail, refer to "Displaying the client authentication audit trail" on page 380.
4. Click **Close** to close the Server Management Console.

Configuring local database authentication

Local database authentication enables you to authenticate a user account against the local database and the Management application server.

To configure local database authentication, complete the following steps.

1. Select the **AAA Settings** tab.
2. For **Primary Authentication**, select **Local Database**.
3. Click **Test**.

   The **Test Authentication** dialog box displays.
4. Enter your user ID and password and click **Test**.

   Test verifies your user ID and password for the local database and verifies user privileges on the Management application server.
5. Click **Apply** to save the configuration.

   To display the authentication audit trail, refer to "Displaying the client authentication audit trail" on page 380.
6. Click **Close** to close the Server Management Console.

Displaying the client authentication audit trail

All responses to authentication requests coming from clients are logged to an audit trail log file. This file is automatically backed up on the first day of every month.

1. Select the **AAA Settings** tab.
2. Click **Display** next to **Authentication Audit Trail**.

   The **Login** dialog box displays.
3. Enter your username and password in the appropriate fields and click **OK**.

   The defaults are Administrator and password, respectively.

   The **Authentication Audit Trail** log displays.

   The audit trail shows user names that have attempted to log in to the Management application, and changes to user authentication.
4. Click the **Client to Server Authentication** tab to view the client to server authentication status.
5. Click the **Authentication Settings Changes** tab to view the previous authentication changes.

Radius server configuration

NOTE
You must configure an Radius server as the primary authentication server (refer to "Configuring Radius server authentication" on page 367).

Depending on the Radius server you install, the configuration and dictionary files may have a different name than in the following procedures. If you are using a Radius server for authentication, complete the following procedures on the Radius server:

1. "Configuring Management application data on the Radius server" on page 380.
2. "Configuring user authorization for the Radius server" on page 381
3. "Configuring the dictionary file for the Radius server" on page 381

Configuring Management application data on the Radius server

The client configuration file contains the IP address, secret, and localhost name for the Management application server.

1. Open the client configuration file (such as clients.conf) a text editor (such as Notepad).
2. Enter the Management server data as follows:

```
client ip_address{
    secret      = user-defined_secret
    shortname   = localhost_name
}
```

For example:

```
client 172.26.3.76 {
        secret     = password
        shortname  = GVM1 server
}
```

3. Save and close the file.

Configuring user authorization for the Radius server

The user configuration file contains the individual user profiles.

1. Open the user configuration file (such as users.conf) a text editor (such as Notepad).
2. Enter the user data as follows:

   user_name User-Password = "*password*"
   NM-Roles-AORs-List = "nmRoles=*management_roles*; nmAORs=*management_AORs*"

   where *management_roles* is one or more of the following roles (separated by commas):

   Host Administrator, Network Administrator, Operator, Report User Group, SAN System Administrator, Security Administrator, Security Officer, and Zone Administrator

   and *management_AORs* is one or more of the following AORs (separated by commas):

   All Fabrics or All Hosts

   For example:

   ```
   jsmith        User-Password = "password"
   NM-Roles-AORs-List = "nmRoles=Host Administrator,Network Administrator,Operator,Report User Group,SAN
   System Administrator,Security Administrator,Security Officer,Zone Administrator; nmAORs=All FabricsAll
   Hosts"
   ```

3. Enter the following to make the default authentication type PAP:

   ```
   DEFAULT              Auth-Type = PAP
   ```

4. Save and close the file.

Configuring the dictionary file for the Radius server

The dictionary file defines the symbolic names for Radius attributes and values.

1. Copy the Management application dictionary file (dictionary.NM_AAA_dictionary) located in the *Install_Home*/docs/Auth directory to the Radius server dictionary directory.
1. Open the dictionary configuration file (dictionary.NM_AAA_dictionary) a text editor (such as Notepad).

   The dictionary file contains the following information:

   VENDOR *vendor_name* *vendor_id_number*

   BEGIN-VENDOR *vendor_name*

   ```
   ATTRIBUTE     NM-Roles-AORs-List     1     string


   END-VENDOR
   ```
   vendor_name

2. Change the attribute to use the sequence number 9 as follows.

   ```
   ATTRIBUTE     NM-Roles-AORs-List     9     string
   ```

3. Save and close the file.
4. Open the Radius server dictionary file in a text editor (such as Notepad).

5. Enter the following to add the Management application dictionary file to the Radius server dictionary file:

   ```
   $INCLUDE dictonary.NM_AAA_dictionary
   ```

6. Save and close the file.

LDAP server configuration

NOTE
You must have User Management Read and Write privileges to map roles and AORs to Active Directory (AD) groups.

NOTE
You must configure a Lightweight Directory Access Protocol (LDAP) server as the primary authentication server and set Authentication Server Groups as the authorization preference (refer to "Configuring LDAP server authentication" on page 370).

Authentication Server Groups enable you to configure user access rights to AD groups (including users, contacts, computers, and other AD groups) by assigning roles and AORs to groups in the Management application. LDAP provides user authentication and authorization using the AD service in conjunction with LDAP on the switch.

1. "Creating an AD user account" on page 382 (LDAP server)
2. "Assigning an AD user to an AD group" on page 383 (LDAP server)
3. "Defining user accounts on the external LDAP server" on page 383 (LDAP server)
4. "Assigning roles and AORs to an AD group" on page 386 (Management application server)

Creating an AD user account

To create a new user account in Active Directory Users and Computers, complete the following steps. For more information, click **F1** for help or refer to www.microsoft.com.

1. Open the Active Directory Users and Computers console.

   For example, on Windows 2008-R2, select **Start > Administrative Tools > Active Directory Users and Computers**.

   The **Active Directory Users and Computers** dialog box displays.

2. Right-click the **Users** folder and select **New > User**.

   The **New Object - User** dialog box displays.

3. Enter a name in the **First name** field.

   It is recommended that you use similar names for the **First name** and **User logon name** fields.

4. Enter a name in the **Full name** field
5. Enter a log on name in the **User logon name** field.
6. Click **Next**.
7. Select the **Password Never Expires** option and click **Next**.
8. Click **Finish**.

   The new user displays in the **Users** pane.

9. Right-click the new user in the **Users** pane and select **Reset Password**.
10. Assign a new password with at least one special character and one number and click **OK**..
11. Close the **Active Directory Users and Computers** dialog box.

Assigning an AD user to an AD group

To assign a new group in Active Directory Users and Computers, complete the following steps. For more information, click **F1** for help or refer to www.microsoft.com

1. Open the Active Directory Users and Computers console.

   For example, on Windows 2008-R2, select **Start > Administrative Tools > Active Directory Users and Computers**.

   The **Active Directory Users and Computers** dialog box displays.

2. Right-click the user in the **Users** pane and select **Add to a Group**.

   The **Select Group** dialog box displays.

3. Enter the group name in the **Enter the object name to select** text box and click **Check Names**.
4. Click **OK**.

Defining user accounts on the external LDAP server

If you configure the external LDAP server as the primary authentication server in the server management console, you must define roles and AORs in the external LDAP server to match the Management application roles and AORs.

Configuring roles and AORs on the external LDAP server

Open the Management console on the Active Directory installed server and complete the following steps.

1. Select **Start > Run**.
2. Type **mmc** and press **Enter**.
3. Select **File > Add/Remove Snap-in**.

   The **Add/Remove Snap-in** dialog box displays.

4. Click **Add**.

   The **Add Standalone Snap-in** dialog box displays.

5. Select **Active Directory Schema** from the **Available standalone snap-ins** list and click **Add**.

   If **Active Directory Schema** does not display the **Available standalone snap-ins** list, you must configure it on the LDAP server (refer to "Configuring the Active Directory Schema on the LDAP server" on page 384).

6. Click **Close** on the **Add Standalone Snap-in** dialog box.
7. Click **OK** on the **Add/Remove Snap-in** dialog box.
8. Right-click the **Attributes** folder (Console Root/Active Directory Schema/Attributes) and select **New > Attribute**.

   The **Create a New Attribute** dialog box displays.

9. Create the NmAors attribute by completing the following steps.
    a. Enter NmAors in the **Common Name** field.
    b. Enter NmAors in the **LDAP Display Name** field.
    c. Enter a unique object identifier in the **Unique x500 Object ID** field.
    d. Enter a description of the attribute in the **Description** field.
    e. Select **Case Insensitive String** in the **Syntax** list.
    f. Click **OK**.
10. Right-click the **Attributes** folder (Console Root/Active Directory Schema/Attributes) and select **New > Attribute**.
11. Create the NmRoles attribute by completing the following steps.
    a. Enter NmRoles in the **Common Name** field.
    b. Enter NmRoles in the **LDAP Display Name** field.
    c. Enter a unique object identifier in the **Unique x500 Object ID** field.
    d. Enter a description of the attribute in the **Description** field.
    e. Select **Case Insensitive String** in the **Syntax** list.
    f. Click **OK**.
12. Expand the **Classes** folder (Console Root/Active Directory Schema) and right-click **user** and select **Properties**.

    The **user Properties** dialog box displays.
13. Click the **Attributes** tab.
14. Click the **Add** button.

    The **Select Schema Object** dialog box displays.
15. Select **NmAors** and click **OK**.
16. Click the **Add** button on the **user Properties** dialog box.
17. Select **NmRoles** and click **OK**.
18. Click **OK** on the **user Properties** dialog box.
19. Close the Management console.
20. Restart the AD server.

    After you restart the AD server, go to "Configuring authorization details on the external LDAP server" on page 385.

Configuring the Active Directory Schema on the LDAP server

1. Select **Start > Run**.
2. Type **regsvr32 schmmgmt.dll** and press **Enter**.

    Make sure that the following message displays: Dll register Server in schmmgmt.dll succeeded.
3. Select **Start > Run**.
4. Type **mmc** and press **Enter**.

5. Select **File > Add/Remove Snap-in**.

   The **Add/Remove Snap-in** dialog box displays.

6. Click **Add**.

   The **Add Standalone Snap-in** dialog box displays.

7. Select **Active Directory Schema** from the **Available standalone snap-ins** list and click **Add**.
8. Click **Close** on the **Add Standalone Snap-in** dialog box.

   The Active Directory Schema displays in the **Add/Remove Snap-in** dialog box.

9. Click **OK** on the **Add/Remove Snap-in** dialog box.
10. Save this schema by selecting **File > Save As**.
11. Browse to the following location: C:\Windows\System32.
12. Change the extention to **.msc** and click **Save**.
13. Select **Start > Run**.
14. Type ***schema_name*.msc** and press **Enter**.

    Make sure that the schema that you just saved launches successfully. To configure roles and AORS for this schema, go to step 8 of "Configuring roles and AORs on the external LDAP server" on page 383.

Configuring authorization details on the external LDAP server

Open the **ADSI Edit** dialog box on the Active Directory installed server.

1. Select **Start > Run**.
2. Type **adsiedit.msc** and press **Enter**.

   The **ADSI Edit** dialog box displays.

3. Right-click **CN=*User_Name*** in the **CN=Users** directory and select **Properties**.

   Where ***User_Name*** is the name of the user you created in "Creating an AD user account" on page 382.

   The **CN=*User_Name* Properties** dialog box displays.

4. Select **NmAors** in the **Attributes** list and click **Edit**.

   The **String Attribute Editor** dialog box displays.

5. Enter the areas of responsibility (such as, All Fabricsand All Hosts) in the **Value** field and click **OK**..
6. Select **NmRoles** in the **Attributes** list and click **Edit**.
7. Enter the Management application user roles (such as Host Administrator, Network Administrator, Operator, Report User Group, SAN System Administrator, Security Administrator, Security Officer, and Zone Administrator) in the **Value** field and click **OK**.
8. Close the **ADSI Edit** dialog box.

Assigning roles and AORs to an AD group

Using Authentication Server Groups, you assign users to groups within the Authentication Server Groups server, and assign roles and AORs to the groups within the Management application.

To assign roles and AORs to an AD group, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click the **Authentication Server Groups** tab.
3. Select the roles and AORs you want to assign to the AD group in the **Available Roles / AORs** table.

   Select multiple roles and AORs by holding down the CTRL key and clicking more than one role and AOR.

4. Select the AD group to which you want to assign the selected roles and AORs in the **Active Directory Groups** table.

   If the AD group you want does not display in the table, refer to "Loading an AD group" on page 386.

5. Click the right arrow button.

   The selected roles and AORs are moved to the **Active Directory Groups** table.

6. Click **Apply** to save your work

   When you assign roles and AORs to an AD group and save the configurations, when you reopen the **Users** dialog box and select the **Authentication Server Groups** tab, only the configured AD group is available.

Removing roles and AORs from an AD group

To remove roles and AORs from an AD group, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click the **Authentication Server Groups** tab.
3. Select the roles and AORs you want to remove in the **Active Directory Groups** table.

   Select multiple roles and AORs by holding down the CTRL key and clicking more than one role and AOR.

4. Click the left arrow button.

   The selected roles and AORs are moved to the **Available Roles / AORs** table.

5. Click **OK** to save your work.

Loading an AD group

To load an AD group, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Click the **Authentication Server Groups** tab.

3. Click **Fetch**.

   The **Fetch AD Group** dialog box displays.

   FIGURE 176 Fetch AD Group dialog box

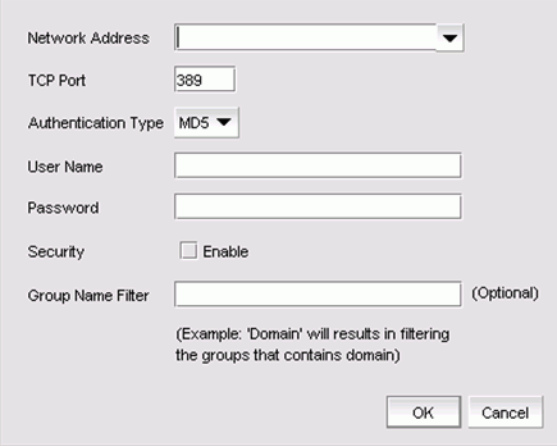

4. Select the LDAP server network address from the **Network Address** list.
5. Enter the TCP port number in the **TCP Port** field, if necessary.

   The default TCP port number is 389 if security is not enabled. The default TCP port number is 636 if security is enabled.

6. Select the authentication protocol **MD5** from the **Authentication Type** list.
7. Enter your LDAP server user login name in the **User Name** field.
8. Enter your LDAP server user login password in the **Password** field.
9. Select the **Security Enable** check box to enable the security channel between the Management application server and the LDAP server.

   When you enable security, the TCP port number automatically changes to port 636 and you must enable certificate services on the LDAP server.

10. (Optional) Enter the group name in the **Group Name Filter** field**.**

    You can specify the group name in the following formats:

    - *User*, *Domain* - Will fetch the group name that contains the user or the operator.
    - *User**, *Domain* - Will fetch the group name that starts with the user and contains the operator.
    - *User*, **Domain* - Will fetch the group name that starts with the user and ends with the operator.
    - **User**, *Domain* - Will fetch the group name that contains the user or the operator.

11. Click **OK**.

    The **Active Directory Groups** table displays with all AD groups available in the specified LDAP server, as well as any AD groups already mapped in the Management application server (local database).

    To assign or remove roles and AORs, refer to "Assigning roles and AORs to an AD group" on page 386 or "Removing roles and AORs from an AD group" on page 386.

12. Click **Close** to close the **Users** dialog box.

Deleting an AD group

Deleting an AD group deletes the roles and AORs assigned to the group and removes the group from the **Active Directory Groups** table.

To delete an AD group, complete the following steps.

1. Select one or more AD groups that you want to delete from the **Active Directory Groups** table.
2. Click **Delete**.
3. Click **Yes** on the confirmation message.
4. Click **OK** on the deletion successful message.
5. Click **OK** to save your work.

Restore tab

The **Restore** tab enables you to restore the application data files used by the Management application server.

Restoring the database

To restore application data files, you must know the path to the backup files. This path is configured from the **Server > Options** dialog box. For more information about backup, refer to "Server Data backup" on page 71.

NOTE
You cannot restore data from a earlier or later configuration (Trial or Licensed version) of the Management application.

NOTE
From the **Restore** tab, you can also restore server supportsave files. Restoring client supportsave files is not supported.

To restore the application data files, complete the following steps.

1. Click the **Services** tab.
2. Stop all services.
3. Click the **Restore** tab (Figure 177).

FIGURE 177 Restore tab

4. Click **Browse** to select the path (defined in the **Output Directory** field on the **Options** dialog box - **Backup** pane) to the database backup location.
5. Click **Restore**.

   Properties such as edition, package, version number, build number, OEM, application/release name, and server architecture are validated before the backup or server supportsave restoration. Upon successful validation, a message displays the status of the restore operation or an error message is displayed. Click **OK** to close the message and the SMC. For the restored data to take effect, relaunch the Configuration Wizard using the instructions in "Launching the Configuration Wizard" on page 6.

Technical Support Information tab

The **Technical Support Information** tab of the SMC allows you to capture technical support information for the Management application as well as the configuration files for all switches in discovered fabrics. This information is saved in a *zip* file in a location that you specify.

Capturing technical support information

To capture technical support information, complete the following steps.

1. Select the **Technical Support Information** tab.

FIGURE 178 Technical Support Information tab

2. Select the **Include database** check box to capture database server supportsave files and choose one of the following options:

   - Select **Partial** to exclude historical data and events from the database capture.
   - Select **Full** to include historical data and events in the database capture.

   NOTE
   It is recommended that you only capture the full database.

   NOTE
   You should only capture the full database when you need to debug Historical Performance Management or Historical Events issues.

   NOTE
   The backup data process performance is optimized when you backup large database. By using parallel processing capabilities provided by PostgreSQL the backup process is optimized and reduced by 50 percentage of the time consumed. For example, a backup of 4GB of data approximately takes 25 to 30 minutes.

3. Enter the path where you want to save the support data and a name for the support save file in the **Output Path** field.

   For example, *Full_Path**Support_Save_File_Name*.zip. You can also browse to the location you want to save the support data and append the file name to the path when you return to the **Technical Support Information** tab.

   If you do not specify an output path, the Management application automatically saves the data to the *Install_Home*/support directory. The default name of the Server Support Save is DCM-SS-*Time_Stamp*.

   NOTE
   For Linux systems, you cannot have blank spaces in the output path (target directory). If the output path contains blank spaces, the supportShow files are not complete.

4. Click **Capture**.

   A confirmation message displays when the capture is complete.

5. Click **OK**.

HCM Upgrade tab

The **HCM Upgrade** tab enables you to upgrade the Management application to include a new version of HCM.

Upgrading HCM on the Management server

To upgrade HCM, complete the following steps.

1. Select the **HCM Upgrade** tab.

   FIGURE 179 HCM Upgrade tab

   

2. Click **Browse** to select the HCM installation folder location (for example, C:\Program Files\BROCADE\Adapter on Windows systems and /opt/brocade/adapter on Linux systems).
3. Click **Upgrade**.
4. Click **Close**.

SMI Agent Configuration Tool

The **SMIA Configuration Tool** enables you to configure SMI (Storage Management Initiative) Agent settings, such as security, CIMOM, and certificate management. This tool is automatically installed with the Management application as part of the Server Management Console. This **SMIA Configuration Tool** consists of the following tabs:

- Home tab — enables you to access Management application features such as, fabric and host discovery, role-based access control, application configuration and display options, server properties, as well as the application name, build, and copyright.
- Authentication tab — enables you to configure mutual authentication for Client, CIMMOM server, and Indication using a secure protocol.
- CIMOM tab — enables you to configure the CIMOM server port, the CIMOM Bind Network Address, and the CIMOM log.
- Certificate Management tab — enables you to import Client and Indication certificates, export Server certificates, as well as view and delete current certificates.
- Summary tab — enables you to view the CIMOM server configuration and current configuration.

Launching the SMIA configuration tool on Windows

NOTE
All Management application services must be running before you can log into the **SMIA Configuration Tool**. To start the Management application services, click **Start** on the **Server Management Console dialog box.**

1. Launch the **Server Management Console** from the **Start** menu on the Management application server.

   You can also drag the SMC icon onto your desktop as a short cut.

2. Click **Configure SMI Agent on the Server Management Console dialog box.**

   The **Log In** dialog box displays.

   FIGURE 180 Log In dialog box

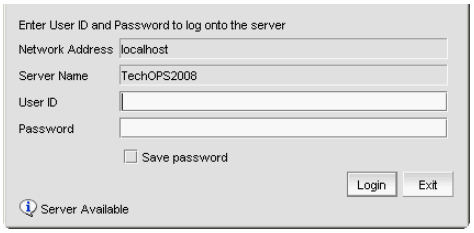

3. Enter your username and password in the appropriate fields.

   The defaults are Administrator and password, respectively. If you migrated from a previous release, your username and password do not change.

4. Select or clear the **Save password** check box to choose whether you want the application to remember your password the next time you log in.

5. Click **Login**.

   The **SMIA Configuration Tool** dialog box displays.

FIGURE 181 SMIA Configuration Tool dialog box

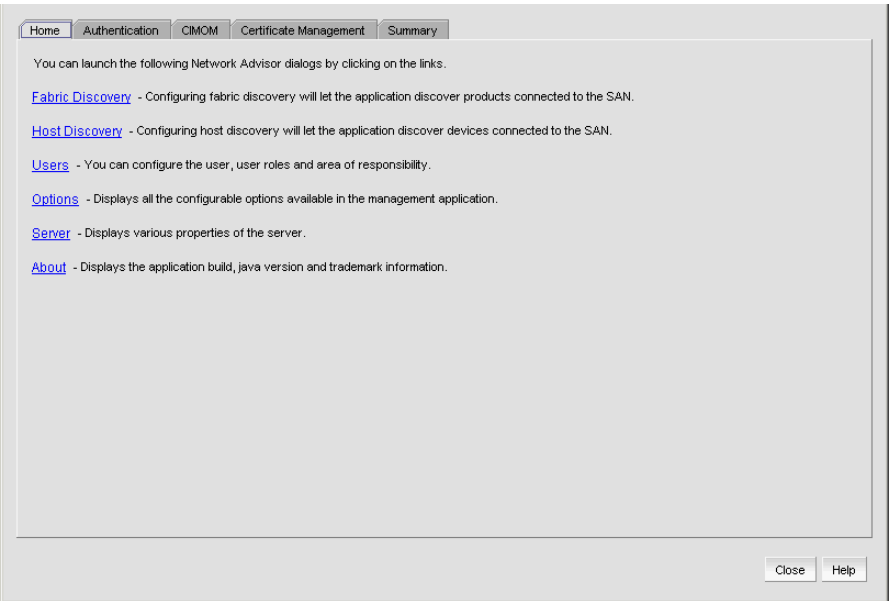

Launching the SMIA configuration tool on Unix

NOTE
All Management application services must be running before you can log into the **SMIA Configuration Tool**. To start the Management application services, click **Start** on the **Server Management Console dialog box.**

Perform the following steps to launch the Server Management Console on Unix systems.

1. On the Management application server, go to the following directory:

   Install_Directory/bin

2. Type the following at the command line:

   ```
   ./smc
   OR
   sh smc
   ```

3. Click **Configure SMI Agent on the Server Management Console dialog box.**

   The **Login** dialog box displays.

4. Enter your username and password in the appropriate fields and click **OK**.

   The defaults are Administrator and password, respectively. If you migrated from a previous release, your username and password do not change.

   The **SMIA Configuration Tool** dialog box displays.

Launching a remote SMIA configuration tool

To launch a remote SMIA configuration tool, complete the following steps.

1. Open a web browser and enter the IP address of the Management application server in the **Address** bar.

   If the web server port number does not use the default (443 if is SSL Enabled; otherwise, the default is 80), you must enter the web server port number in addition to the IP address. For example, *IP_Address*:*Web_Server_Port_Number*.

   The Management application web start screen displays.

2. Click the SMIA configuration tool application web start link.

   The **Log In** dialog box displays.

3. Enter your user name and password.

   The defaults are **Administrator** and **password**, respectively.

   NOTE
   Do not enter *Domain**User_Name* in the **User ID** field for LDAP server authentication.

4. Select or clear the **Save password** check box to choose whether you want the application to remember your password the next time you log in.

5. Click **Login**.

   The **SMIA Configuration Tool** dialog box displays

Service Location Protocol (SLP) support

The Management application SMI Agent uses Service Location Protocol (SLP) to allow applications to discover the existence, location, and configuration of WBEM services in enterprise networks.

You do not need a WBEM client to use SLP discovery to find a WBEM Server; that is, SLP discovery might already know about the location and capabilities of the WBEM Server to which it wants to send its requests. In such environments, you do not need to start the SLP component of the Management application SMI Agent.

However, in a dynamically changing enterprise network environment, many WBEM clients might choose to use SLP discovery to find the location and capabilities of other WBEM Servers. In such environments, start the SLP component of the Management application SMI Agent to allow advertisement of its existence, location, and capabilities.

SLP installation is optional and you can configure it during Management application configuration. Once installed, SLP starts whenever the Management application SMI Agent starts.

SLP support includes the following components:

- slpd script starts the slpd platform
- slpd program acts as a Service Agent (SA). A different slpd binary executable file exists for UNIX and Windows systems.
- slptool script starts the slptool platform-specific program
- slptool program can be used to verify whether SLP is operating properly or not. A different slptool exists for UNIX and Windows.

  By default, the Management application SMI Agent is configured to advertise itself as a Service Agent (SA). The advertised SLP template shows its location (IP address) and the WBEM Services it supports. The default advertised WBEM services show the Management application SMI Agent:

- accepts WBEM requests over HTTP without SSL on TCP port 5988
- accepts WBEM requests over HTTPS using SSL on TCP port 5989

slptool commands

Use the following slptool commands to verify whether the SLP is operating properly.

- slptool findsrvs service:service-agent

  Use this command to verify that the Management application SMI Agent SLP service is properly running as a Service Agent (SA).

  Example output: service:service-agent://127.0.0.1,65535

- slptool findsrvs service:wbem

  Use this command to verify that the Management application SMI Agent SLP service is properly advertising its WBEM services.

  Example outputs:

  service:wbem:https://10.0.1.3:5989,65535

  service:wbem:http://10.0.1.3:5988,65535

  This output shows the functionalities of the Management application SMI Agent:

  - accepts WBEM requests over HTTP using SSL on TCP port 5989
  - accepts WBEM requests over HTTP without SSL on TCP port 5988

- slptool findattrs service:wbem:https://*IP_Address*:*Port*

  NOTE
  Where *IP_Address*:*Port* is the IP address and port number that display when you use the slptool findsrvs service:wbem command.

  Use this command to verify that Management application SMI Agent SLP service is properly advertising its WBEM SLP template over the HTTP protocol.

  Example output:

  ```
  Install_Home\cimom\bin>slptool findattrs service:wbem:http://10.24.35.61:5988
  (template-type=wbem),(template-version=1.0),(template-description=This template describes the attributes
  used for advertising WBEM Servers), (template-url-syntax=http://10.24.35.61:5988),(service-hi-name=WBEM
  Solutions J WBEM Server),(service-hi-description=WBEM Solutions J WBEM Server),
  (service-id=WBEMSolutions:f1f65c3b-27f1-4b70-9ced-e412e93a8d5e),(CommunicationMechanism=CIM-XML),(OtherCo
  mmunicationMechanismDescription =null), (InteropSchemaNamespace=interop),(ProtocolVersion=1.2),
  (FunctionalProfilesSupported=Basic Read,Basic Write,Schema Manipulation, Instance
  Manipulation,Association Traversal,Query Execution,Qualifier
  Declaration,Indications),(FunctionalProfileDescriptions=null),(MultipleOperationsSupported=true),(Authent
  icationMechanismsSupported=Basic),(AuthenticationMechanismDescriptions=null),(Namespace=root/brocade1,int
  erop),(Classinfo=0,0),(RegisteredProfilesSupported=SNIA:SMI-S,DMTF:Profile Registration,SNIA:FC
  HBA,DMTF:LaunchInContext,SNIA:Fan,SNIA:Fabric,SNIA:Switch,DMTF:Role Based Authorization,SNIA:Power
  Supply,SNIA:Sensors,SNIA:Server)
  ```

- slptool findattrs service:wbem:http://*IP_Address*:*Port*

  NOTE
  Where *IP_Address*:*Port* is the IP address and port number that display when you use the slptool findsrvs service:wbem command.

Use this command to verify that the Management application SMI Agent SLP service is properly advertising its WBEM SLP template over the HTTPS protocol.

Example output:

```
Install_Home\cimom\bin>slptool findattrs service:wbem:
https://10.24.35.61:5989(template-type=wbem),(template-version=1.0),(template-description=This template
describes the attributes used for advertising WBEM
Servers),(template-url-syntax=https://10.24.35.61:5989),(service-hi-name=WBEM Solutions J WBEM
Server),(service-hi-description=WBEM Solutions J WBEM
Server),(service-id=WBEMSolutions:f1f65c3b-27f1-4b70-9ced-e412e93a8d5e),(CommunicationMechanism=CIM-XML),
(OtherCommunicationMechanismDescription
=null),(InteropSchemaNamespace=interop),(ProtocolVersion=1.2),(FunctionalProfilesSupported=Basic
Read,Basic Write,Schema Manipulation,Instance Manipulation,Association Traversal,Query
Execution,Qualifier Declaration,
Indications),(FunctionalProfileDescriptions=null),
(MultipleOperationsSupported=true),(AuthenticationMechanismsSupported=Basic),(AuthenticationMechanismDesc
riptions=null),(Namespace=root/brocade1,interop),(Classinfo=0,0),(RegisteredProfilesSupported=SNIA:SMI-S,
DMTF:Profile Registration,SNIA:FC HBA,DMTF:LaunchInContext,SNIA:Fan,SNIA:Fabric, SNIA:Switch,DMTF:Role
Based Authorization,SNIA:Power Supply,SNIA:Sensors,
SNIA:Server)
```

SLP on UNIX systems

This section describes how to verify the SLP daemon on UNIX systems.

SLP file locations on UNIX systems

- SLP log — *Install_Home*/cimom /cfg/slp.log
- SLP daemon — *Install_Home*/cimom /cfg/slp.conf

  You can reconfigure the SLP daemon by modifying this file.
- SLP register — *Install_Home*/cimom /cfg/slp.reg

  You can statically register an application that does not dynamically register with SLP using SLPAPIs by modifying this file. For more information about these files, read the comments contained in them, or refer to http://www.openslp.org/doc/html/UsersGuide/index.html.

Verifying SLP service installation and operation on UNIX systems

1. Open a command window.
2. Type % su root and press **Enter** to become the root user.
3. Type # *Install_Home*/cimom/bin/slptool findsrvs service:service-agent and press **Enter** to verify the SLP service is running as a Service Agent (SA).
4. Type # *Install_Home*/cimom/bin/slptool findsrvs service:wbem and press **Enter** to verify the SLP service is advertising its WBEM services.
5. Choose one of the following options to verify the SLP service is advertising the WBEM SLP template over its configured client protocol adapters.
   - Type # *Install_Home*/cimom /bin/slptool findattrs service:wbem:http://*IP_Address:Port* and press **Enter**.
   - Type # *Install_Home*/cimom /bin/slptool findattrs service:wbem:https://*IP_Address:Port* and press **Enter**.

     NOTE
     Where *IP_Address:Port* is the IP address and port number that display when you use the slptool findsrvs service:wbem command.

SLP on Windows systems

This section describes how to verify the SLP daemon on Windows systems.

SLP file locations on Windows systems

- SLP log — *Install_Home*\cimom \cfg\slp.log
- SLP daemon — *Install_Home*\cimom\cfg\slp.conf

  You can reconfigure the SLP daemon by modifying this file.
- SLP register — *Install_Home*\cimom\cfg\slp.reg

  You can statically register an application that does not dynamically register with SLP using SLPAPIs by modifying this file. For more information about these files, read the comments contained in them, or refer to http://www.openslp.org/doc/html/UsersGuide/index.html.

Verifying SLP service installation and operation on Windows systems

1. Launch the Server Management Console from the **Start** menu.
2. Click **Start** to start the SLP service.
3. Open a command window.
4. Type cd c:*Install_Home*\cimom \bin and press **Enter** to change to the directory where slpd.bat is located.
5. Type > slptool findsrvs service:service-agent and press **Enter** to verify the SLP service is running as a Service Agent.
6. Type > slptool findsrvs service:wbem and press **Enter** to verify the SLP service is advertising its WBEM services.
7. Choose one of the following options to verify the SLP service is advertising the WBEM SLP template over its configured client protocol adapters.
   - Type > slptool findattrs service:wbem:http://*IP_Address:Port* and press **Enter**.
   - Type > slptool findattrs service:wbem:https://*IP_Address:Port* and press **Enter**.

   NOTE
   Where *IP_Address:Port* is the IP address and port number that display when you use the slptool findsrvs service:wbem command.

Home tab

The **Home** tab of the **SMIA Configuration Tool** enables you to access the following Management application features or information:

- **Fabric Discovery** — enables you to view discovered fabrics, discover new fabrics, as well as edit the default SNMP configuration. For step-by-step instructions, refer to "Discovery" on page 33.
- **Host Discovery** — enables you to view discovered hosts, discover new hosts, as well as edit the default SNMP configuration. For step-by-step instructions, refer to "Host discovery" on page 51.
- **Users** — enables you to create or delete Management application users with System Administrator privileges. For step-by-step instructions, refer to "User accounts" on page 138.
- **Options** — enables you to configure the Management application settings. For step-by-step instructions, refer to "Application Configuration" on page 69.
- **Server** — enables you to view server properties. For step-by-step instructions, refer to "Viewing server properties" on page 11.
- **About** — enables you to display information about the Management application, including the build number, Java version, and trademark information.

- **Upgrade** button (Trial version only) — enables you to upgrade from managing 2560 switch ports to 9000 switch ports. For step-by-step instructions, refer to "Upgrading the Management application" on page 43.

Accessing Management application features

To access Management application features such as, fabric and host discovery, role-based access control, application configuration and display options, server properties, as well as the application name, build, and copyright, complete the following steps.

1. Click the **Home** tab, if necessary.
2. Select from the following to access the feature or dialog box.
   - **Fabric Discovery**
   - **Host Discovery**
   - **Users**
   - **Options**
   - **Server**
   - **About**
   - **Upgrade** (Trial version only)
3. Click **Close** to close the **SMIA Configuration Tool** dialog box.

Authentication tab

NOTE
You must have User Management Read and Write privileges to make changes on the CIMOM tab. For more information about privileges, refer to "User Privileges" on page 1333.

The **Authentication** tab enables you to configure mutual authentication for Client and Indication using a secure protocol.

Enabling or disabling CIM client and indication mutual authentication

When you enable client mutual authentication, all CIM client and indication requests to the SMI Agent must pass credentials (KeyStore and TrustStore) to validate the requests. The KeyStore file provides the credentials and the TrustStore file verifies the credentials. When you enable indication mutual authentication, both the CIM client and the CIMOM server maintain the TrustStore files.

The CIM client KeyStore file sends credentials to be validated by the CIMOM server TrustStrore file for any communication from the CIM client to the CIMOM server and the CIMOM server KeyStore file sends credentials to be validated by the CIM client TrustStrore file for any communication from the CIMOM server to the CIM client

To enable or disable CIM client and indication mutual authentication, complete the following steps.

1. Click the **Authentication** tab.

FIGURE 182 Authentication tab

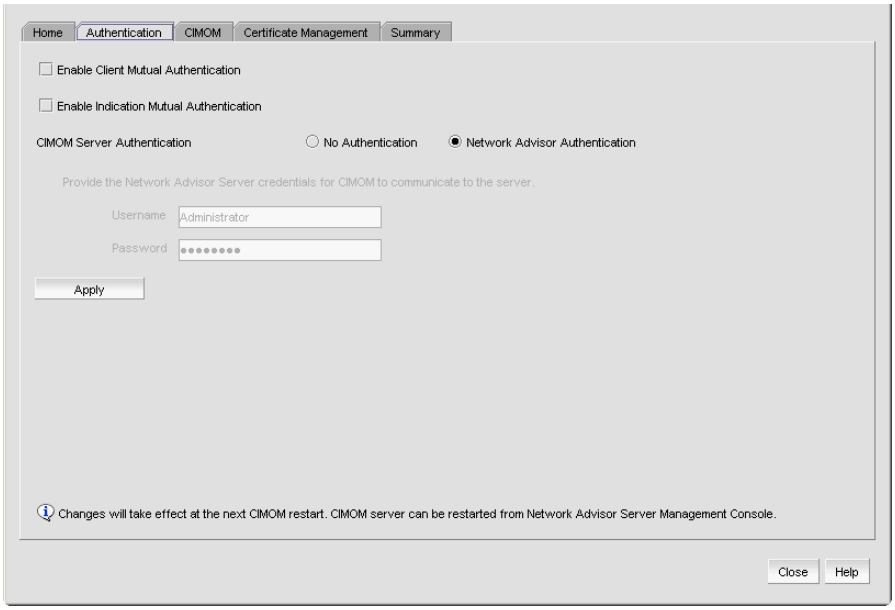

2. Select the **Enable Client Mutual Authentication** check box, as needed.

   If the check box is checked, CIM client mutual authentication is enabled. If the check box is clear (default), client mutual authentication is disabled.

3. Select the **Enable Indication Mutual Authentication** check box, as needed.

   If the check box is checked, indication mutual authentication is enabled. If the check box is clear (default), indication mutual authentication is disabled.

4. Click **Apply**.

   NOTE
   Changes on this tab take effect after the next CIMOM server restart.

   NOTE
   You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).

5. Click **Close** to close the **SMIA Configuration Tool** dialog box.

Configuring CIMOM server authentication

CIMOM server authentication is the authentication mechanism between the CIM client and the CIMOM Server. You can configure the CIMOM server to allow the CIM client to query the CIMOM server without providing credentials; however, the CIMOM server requires the Management application credentials to connect to the Management application server to retrieve the required data. Therefore, if you select no authentication, you must provide Management application credentials to retrieve data from the Management application server.

To configure CIMOM server authentication, complete the following steps.

1. Click the **Authentication** tab.
2. Choose from one of the following options:
   - Select **No Authentication** to allow the CIM client to query the CIMOM server without providing credentials; however, note that the CIMOM server requires the Management application credentials to connect to the Management application server to retrieve the required data. To provide Management application credentials, complete the following steps.
     a. Enter the Management application user name in the **Username** field.
     b. Enter the Management application user password in the **Password** field.
   - Select *Management_Application* **Authentication** to allow the CIM client to query the CIMOM server and the Management application server using the credentials configured on the **Users** tab.
3. Click **Apply**.

   NOTE
   Changes on this tab take effect after the next CIMOM server restart.

   NOTE
   You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).
4. Click **Close** to close the **SMIA Configuration Tool** dialog box.

CIMOM tab

NOTE
You must have SAN - SMI Operation Read and Write privileges to view or make changes on the CIMOM tab. For more information about privileges, refer to "User Privileges" on page 1333.

The **CIMOM** tab enables you to configure the CIMOM server port, the CIMOM Bind Network Address, and the CIMOM log.

Configuring the SMI Agent port number

To configure the SMI Agent port number, complete the following steps.

1. Click the **CIMOM** tab.

FIGURE 183 CIMOM tab

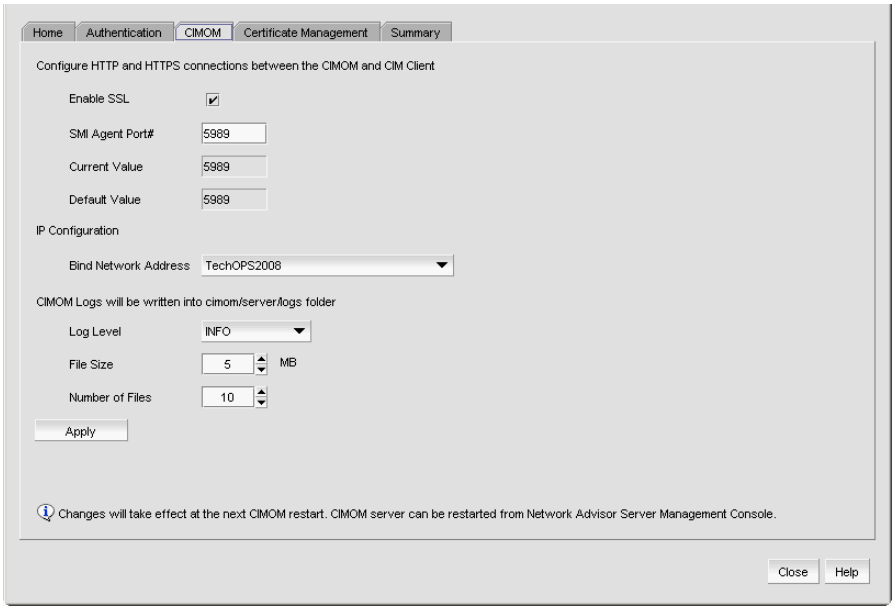

2. Select or clear the **Enable SSL** check box, to enable or disable SSL for the SMI Agent.

   NOTE
   Disabling SSL will disable Indication and Client Mutual Authentication.

   If the check box is checked (default), SSL is enabled. If the check box is clear, SSL is disabled.

3. Enter the SMI Agent port number in the **SMI Agent Port #** field.

   This port number must be within the range of 1 through 65535. Defaults are 5989 with SSL enabled and 5988 with SSL disabled.

4. Click **Apply**.

   NOTE
   Changes on this tab take effect after the next CIMOM server restart.

   NOTE
   You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).

   If you disabled SSL, a confirmation message displays. Click **Yes** to continue.

5. Click **Close** to close the **SMIA Configuration Tool** dialog box.

Configuring the CIMOM Bind Network Address

NOTE
You must have SAN - SMI Operation Read and Write privileges to view or make changes on the CIMOM tab. For more information about privileges, refer to "User Privileges" on page 1333.

To configure the network bind address, complete the following steps.

1. Click the **CIMOM** tab.
2. Select a network address from the **IP Configuration Bind Network Address** list to which you want to bind the CIMOM server.

   The default network address is the host system name.
3. Click **Apply**.

   NOTE
   Changes on this tab take effect after the next CIMOM server restart.

   NOTE
   You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).
4. Click **Close** to close the **SMIA Configuration Tool** dialog box.

Configuring the CIMOM log

NOTE
You must have SAN - SMI Operation Read and Write privileges to view or make changes on the **CIMOM** tab. For more information about privileges, refer to "User Privileges" on page 1333.

To configure the CIMOM log, complete the following steps.

1. Click the **CIMOM** tab.
2. Select a log category from the **Log Level** list to start logging support data for the server.

   Options include the following:

   - Off — select to turn off logging support data.
   - Severe — select to only log support data that indicates serious failures which prevent normal program operation.
   - Warning — select to only log support data that indicates a potential problem.
   - Info (default) — select to only log support data for informational messages.
   - Config — select to only log support data for static configuration messages used to assist in debugging problems associated with particular configurations.
   - Fine — select to only log message data used to provide trace information.
   - Finer — select to only log message data used to provide detailed trace information.
   - Finest — select to only log message data used to provide highly detailed trace information.
   - All — select to log support data for all messages.

3. Click **Apply**.

   NOTE
   Changes on this tab take effect after the next CIMOM server restart.

   NOTE
   You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).

4. Click **Close** to close the **SMIA Configuration Tool** dialog box.

Certificate Management tab

NOTE
You must have SMI Operation Read and Write privileges to view or make changes on the **Certificate Management** tab. For more information about privileges, refer to "User Privileges" on page 1333.

The **Certificate Management** tab enables you to manage your CIM client and Indication authentication certificates. Using this tab, you can perform the following operations:

- "Importing a certificate"
- "Viewing a certificate"
- "Exporting a certificate"
- "Deleting a certificate"

Importing a certificate

To import a certificate, complete the following steps.

1. Click the **Certificate Management** tab.

   FIGURE 184 Certificate Management tab

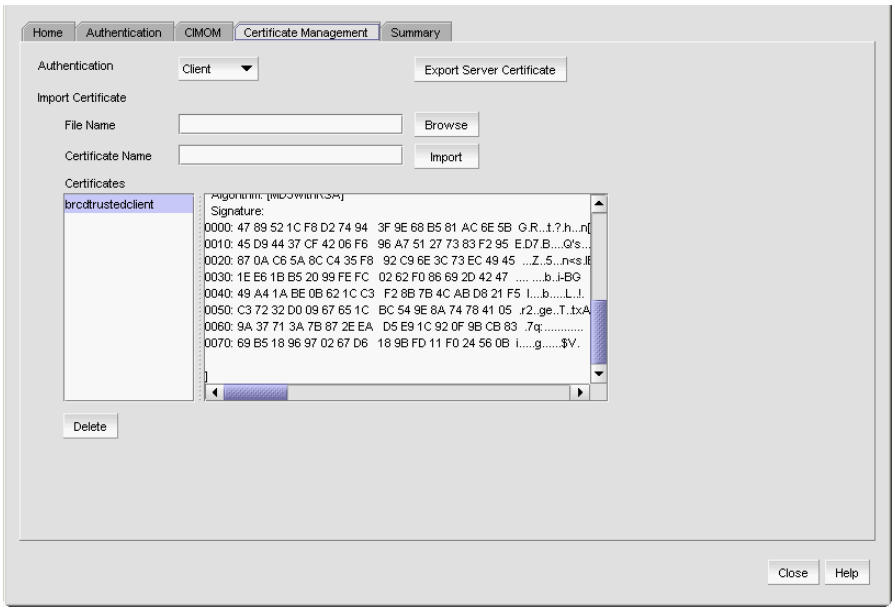

2. Select the **Client** or **Indication** from the **Authentication** list.

   The appropriate certificates display in the **Certificates** list.

3. Enter the full path or browse to the certificate you want to import (for example, on Windows the path is C:\Certificates\cimom-indication-auth2.cer and on Linux the path is opt/Certificates/cimom-indication-auth2.cer).

   You can only import certificate files with the CER extension (.cer).

4. Enter a name for the certificate in the **Certificate Name** field.

5. Click **Import**.

   The new certificate displays in the **Certificates** list and text box.

   If the certificate location is not valid, an error message displays. Click **OK** to close the message and reenter the full path to the certificate location.

   If you did not enter a certificate name, an error message displays. Click **OK** to close the message and enter a name for the certificate.

   If the certificate file is empty or corrupted, an error message displays. Click **OK** to close the message.

6. Click **Close** to close the **SMIA Configuration Tool** dialog box.

Viewing a certificate

NOTE
You must have SMI Operation Read and Write privileges to view the **Certificate Management** tab. For more information about privileges, refer to "User Privileges" on page 1333.

To view a certificate, complete the following steps.

1. Select **Client** or **Indication** from the **Authentication** list.

   The appropriate certificates display in the **Certificates** list.

2. Select the certificate you want to view in the **Certificates** list.

   The certificate details display in the **Certificates** text box.

3. Click **Close** to close the **SMIA Configuration Tool** dialog box.

Exporting a certificate

NOTE
You must have SMI Operation Read and Write privileges to view or make changes to the **Certificate Management** tab. For more information about privileges, refer to "User Privileges" on page 1333.

To export a certificate, complete the following steps.

1. Click the **Certificate Management** tab.

2. Select **Client** or **Indication** from the **Authentication** list.

   The appropriate certificates display in the **Certificates** list.

3. Select the certificate you want to export in the **Certificates** list.

4. Click **Export Server Certificate.**

   The **Save As** dialog box displays.

5. Browse to the directory where you want to export the certificate.
6. Edit the certificate name in the **File Name** field, if necessary.
7. Click **Save**.
8. Click **Close** to close the **SMIA Configuration Tool** dialog box.

Deleting a certificate

> NOTE
> You must have SMI Operation Read and Write privileges to view or make changes to the **Certificate Management** tab. For more information about privileges, refer to "User Privileges" on page 1333.

To delete a certificate, complete the following steps.

1. Click the **Certificate Management** tab.
2. Select **Client** or **Indication** from the **Authentication** list.

   The appropriate certificates display in the **Certificates** list.

3. Select the certificate you want to delete in the **Certificates** list.
4. Click **Delete**.
5. Click **Yes** on the confirmation message.

   The selected certificate is removed from the **Certificates** list.

6. Click **Close** to close the **SMIA Configuration Tool** dialog box.

Summary tab

The **Summary** tab enables you to view summary information about the Server configuration and the current configuration.

Viewing the configuration summary

To view summary information about the Server configuration and the current configuration, complete the following steps.

> NOTE
> Server configuration changes in the **Summary** tab only take effect after the CIMOM restart.

> NOTE
> You can only restart the server using the Server Management Console (**Start > Programs >** *Management_Application_Name* **12.X.X > Server Management Console**).

1. Click the **Summary** tab.

FIGURE 185 Summary tab

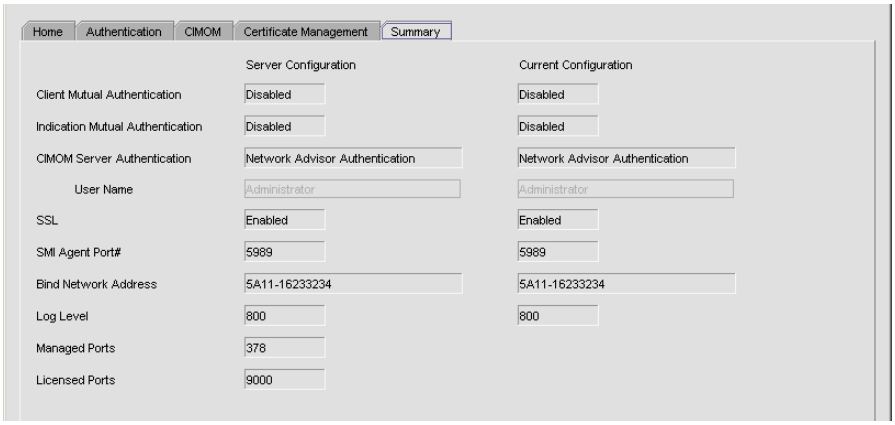

2. Review the summary.

NOTE
When the CIMOM server is stopped, the server configuration information does not display on the **Summary** tab.

The following information is included in the summary.

TABLE 32

Field/Component	Description
Client Mutual Authentication	Displays whether or not the client mutual authentication is enabled or disabled for the Server Configuration and the Current Configuration.
Indication Mutual Authentication	Displays whether or not the indication mutual authentication is enabled or disabled for the Server Configuration and the Current Configuration.
CIMOM Server Authentication	Displays whether or not the CIMOM server authentication is enabled or disabled for the Server Configuration and the Current Configuration.
User Name	Displays the user name for the Server Configuration and the Current Configuration. Only enabled if **CIMOM Server Authentication** is No Authentication.
SSL	Displays whether or not the SSL is enabled or disabled for the Server Configuration and the Current Configuration.
SMI Agent Port #	Displays the SMI Agent port number for the Server Configuration and the Current Configuration.
Bind Network Address	Displays the Bind Network address for the Server Configuration and the Current Configuration.
Log Level	Displays the log level for the Server Configuration and the Current Configuration. Options include the following: • 10000 — Off • 1000 — Severe • 900 — Warning • 800 — Info (default) • 700 — Config • 500 — Fine • 400 — Finer • 300 — Finest • 0 — All
Managed Ports	Displays the number of managed ports. For more information about managed port count rules, refer to "Managed count" on page 41.
Licensed Ports	Displays the number of licensed ports.

3. Click **Close** to close the **SMIA Configuration Tool** dialog box.

SAN Device Configuration

Configuration file management

(Professional only) Configuration files are stored as a flat file in the user chosen directory. For Windows platforms, the default prompted location is <*User Dir*>\Documents. For example, C:\Users\<*User_Name*>\Documents.

Professional only allows you to back up the configuration file manager and save switch configuration. For complete feature support, you must upgrade to Enterprise Edition.

(Trial and Licensed version) Configuration files are stored in an Postgress database on the Management application server. You can save entire configurations of switch configuration files and use them to ensure consistent switch settings in your fabric, propagate configuration settings to additional switches in the fabric, and troubleshoot the switches.

For more information about the database fields, refer to "Database Fields" on page 1397.

Saving switch configurations

NOTE
Saving switch configurations is supported only on Fabric OS switches.

NOTE
Saving switch configurations requires a Trial or Licensed version.

You can save a switch configuration to the repository using one of the following methods:

- On demand (**Configure > Configuration File > Backup Now**) (refer to "Saving switch configurations on demand or manually")
- Defining a schedule (**Configuration File > Schedule Backup**) (refer to "Scheduling switch configuration backup" on page 412)
- Defining adaptive backup (Discovery or Event-triggered) (refer to "Adaptive backup" on page 410)

Saving switch configurations on demand or manually

Configuration files are uploaded from the selected switches and stored in individual files only for the Professional edition. Files are named with the convention *cfg_switch_name_Date_Time*.

Use this procedure to immediately save switch configurations to the repository.

1. Select **Configure > Configuration File > Backup Now**.

   The **Save Switch Configurations** dialog box displays.

FIGURE 186 Save Switch Configurations dialog box

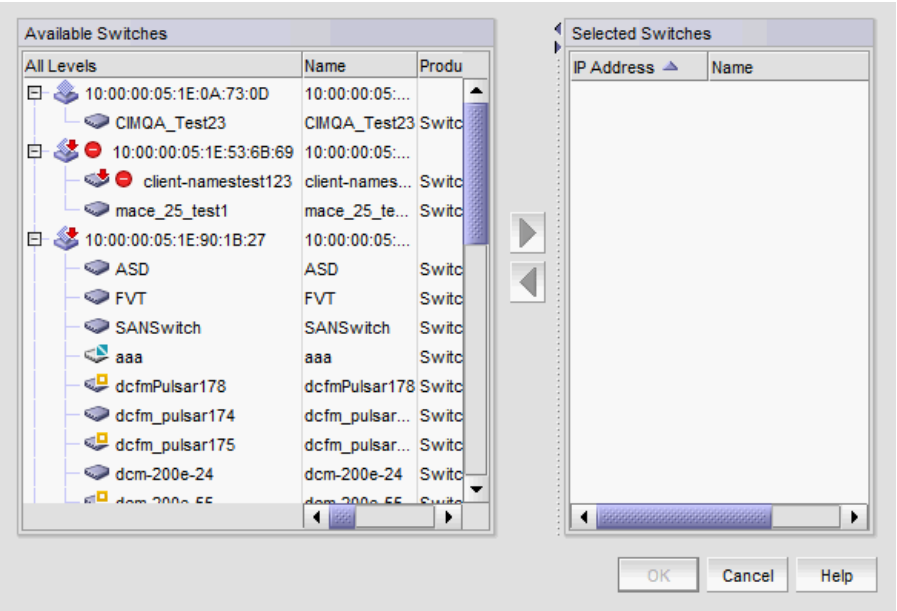

2. Select the switches for which you want to save configuration files from the **Available Switches** list.
3. Click the right arrow to move the selected switches to the **Selected Switches** list.
4. Click **OK**.

   Configuration files from the selected switches are saved to the repository.
5. (Professional only) Browse to the location where you want to save the switch configuration.
6. (Professional only) Click **Save Configuration**.

   Configuration files from the selected switches are saved to the selected location. You can use this file to restore the saved configuration through the device's Element Manager.

Adaptive backup

Adaptive backup is triggered based on fabric discovery and when configuration change events is received from a switch.

Discovery backup

Switch or fabric discovery automatically triggers discovery backup for all switches in the fabric which have the correct user credentials.

To discover a switch, refer to "Discovery" on page 33.

Discovery configuration files display in the **Configuration File Manager** dialog box with the **Backup Type** as **Discovery**.

Event -triggered backup

Event triggered backup is triggered when the switch undergoes configuration changes and on reception of audit events in the master log. The frequency of backup collection is every 15 minutes for all network sizes. The configuration backup will have the **Backup Type** as **Event Triggered**.

Triggering backup configuration on events

To trigger the backup configuration on events, complete the following steps.

1. Click the **SAN** tab.
2. Select **Configure** > **Configuration File** > **Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.
3. Click the **Switch Configurations** tab.

   The **Switch Configurations** tab of the **Configuration File Manager** dialog box displays.
4. Select the **Trigger Backup on Events** check box to collect backup configurations based on the events triggered.

Stopping collection of event-triggered backup configuration

To stop the triggered backup configuration on events, complete the following steps.

1. Click the **SAN** tab.
2. Select **Configure** > **Configuration File** > **Configuration File Manager.**

   The **Configuration File Manager** dialog box displays.
3. Click the **Switch Configurations** tab.

   The **Switch Configurations** tab of the **Configuration File Manager** dialog box displays.
4. Clear the **Trigger Backup on Events** check box to stop collecting backup configurations triggered by events.

Restoring a switch configuration for a selected device

The **Restore Switch Configuration** dialog box enables you to download a previously saved switch configuration to a selected device.

To restore a switch configuration, complete the following steps.

1. Right-click a device in the Product List or the Connectivity Map, and select **Configuration File > Restore**.

   The **Restore Switch Configuration** dialog box displays.

   FIGURE 187 Restore Switch Configuration dialog box

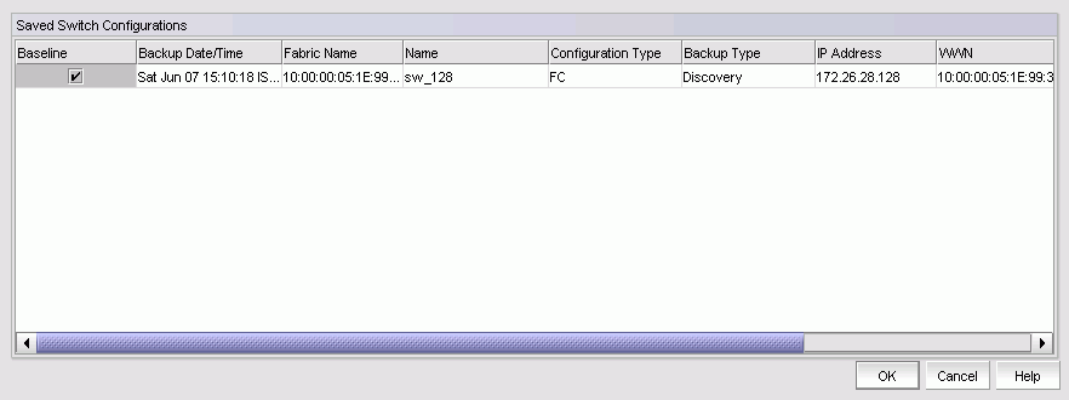

2. Select the switch configuration you want to download from the **Saved Switch Configurations** table.

3. Click **OK**.

   The configuration is downloaded to the device. If necessary, the restoration process prompts you to disable and reboot the device before the configuration begins. This lets you determine whether the configuration backup should be performed immediately or at a later time.

   When you restore a switch configuration on a Virtual Fabrics-configured chassis, the configuration data for the logical switches is downloaded to the switch as configured in the file. When you restore a switch configuration on a logical switch, only the selected logical switch configuration data is downloaded to the switch.

Scheduling switch configuration backup

NOTE
The Enhanced Group Management (EGM) license must be activated on a switch to perform this procedure and to use the supportSave module.

You can schedule a backup of one or more switch configurations. If a periodic backup is scheduled at the SAN level, that backup will apply to all switches from all fabrics discovered. Any new fabrics being discovered are automatically added to the list of fabrics to be backed up.

NOTE
If a backup is scheduled for more than one fabric and some of the fabrics contain common members, the backup will include the unique switch configuration values obtained from the fabrics.

Use this procedure to create a scheduled backup of switch configurations to the repository. To save switch configurations to the repository on demand, refer to "Saving switch configurations" on page 409.

The configuration files are stored in the Management application database.

1. Select **Configure > Configuration File > Schedule Backup.**

   The **Schedule Backup of Switch Configurations** dialog box displays.

FIGURE 188 Schedule Backup of Switch Configurations dialog box

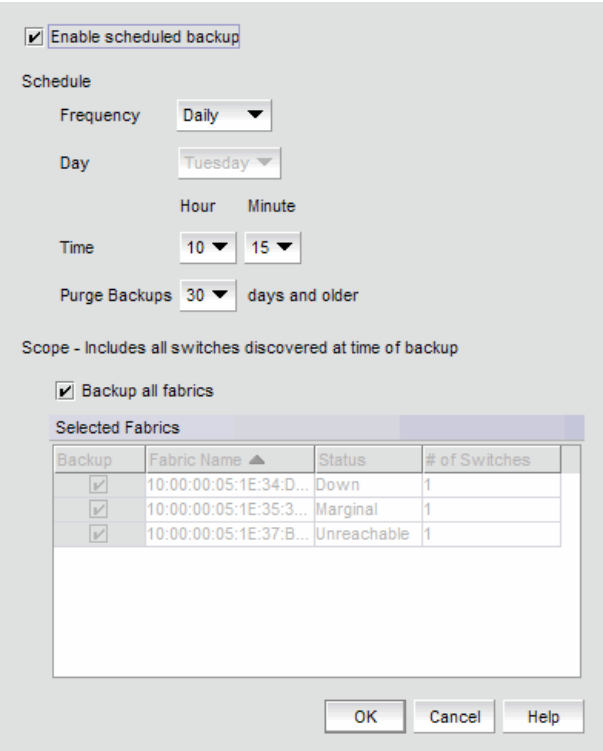

2. Select the **Enable scheduled backup** check box.
3. Set the **Schedule** parameters. These include the following:
   - The desired **Frequency** for backup operations (daily, weekly, monthly).
   - The **Day** you want backup to run.

     If **Frequency** is **Daily**, the **Day** list is unavailable.

     If **Frequency** is **Weekly**, choices are days of the week (Sunday through Saturday).

     If **Frequency** is **Monthly**, choices are days of the month (1 through 31).
   - The **Time** (hour, minute) you want backup to run.
   - The maximum age allowed before you **Purge Backups**.

     The number of days (7 through 90) before you purge backups should be at least one day more than the selected backup frequency. The default is 30 days.

     The backup purge thread runs every day at 12:30 PM and deletes all backup configurations that exceed the maximum age allowed.

     NOTE
     If you disable scheduled backup, purging backups is not disabled and continues to run at the selected frequency.
4. Choose one of the following options to determine the scope of the backup:
   - Select the **Backup all fabrics** check box, if necessary, to back up all switch configurations of discovered switches in all fabrics
   - Clear the **Backup all fabrics** check box and select the specific fabric **Backup** check boxes in the **Selected Fabrics** table to back up individual fabrics.

The **Selected Fabrics** table includes the following information:

- **Fabric Name** — The world wide name of the fabric selected for backup configuration.
- **Status** — The status of the fabric selected for backup configuration; for example, unknown or marginal.
- **# of Switches** — The number of switches that are configured on the fabric selected for backup configuration.

If any switches do not have the EGM license, a messages displays. Click **OK** to enable backup on the switches with the EGM license.

5. Click **OK**.

   Click **OK** on the confirmation message.

Viewing switch configurations

The **Switch Configurations** tab of the **Configuration File Manager** dialog box allows you to display the configuration for each switch or fabric, including the name of the switch, the firmware version, and the name of the switch configuration type.

To view configurations for a switch or fabric, complete the following steps.

1. Click the **SAN** tab.
2. Select **Configure** > **Configuration File** > **Configuration File Manager**.
3. Click the **Switch Configurations** tab.

   The **Switch Configurations** tab of the **Configuration File Manager** dialog box displays, as shown in Figure 189.

FIGURE 189 Switch Configurations tab

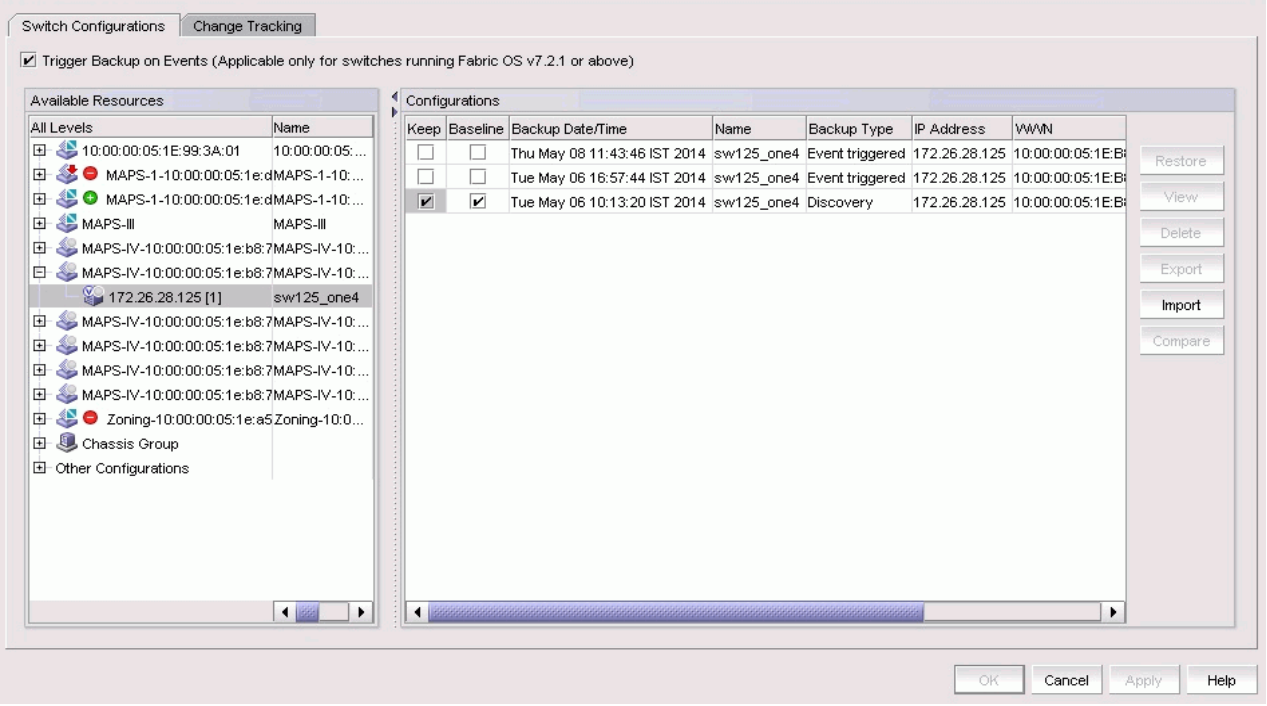

The **Switch Configurations** tab contains the following fields and components:

- **Trigger Backup on Events** check box — Select to collect backup configurations triggered by events (refer to "Event -triggered backup"). Clear the check box to not collect backup configurations triggered by events. **Trigger Backup on Events** is only supported on Fabric OS devices running 7.2.1 or later.
- **Available Resources** table — Displays the list of discovered fabrics and switches in the device tree view. Select one or more switches or fabrics to display the backup files in the **Configurations** table.

  The following are the configurations available in **Other Configurations:**

  - **Imported** — Displays the user-imported configurations.
  - **Unmanaged\Unmonitored Switch Configurations** — Displays the configurations with deleted and unmonitored switches.
- **Configurations** table — Displays the collected configuration backup files for the selected switches or fabrics. The **Configurations** table includes the following information:
  - **Keep** check box — Select to keep the associated configuration past the defined age limit. The configuration will be kept until it is manually deleted, or until the **Keep** check box is cleared to enable the age limit again**.**
  - **Baseline** check box — Select to compare all additional product configuration backup files to the baseline configuration for deviation. You can use the **Change Tracking** tab to compare the latest backup configuration file with the configuration that is designated as the baseline.

    NOTE
    For the selected baseline configuration file, **Keep** is marked as the default.
- **Backup Date/Time** — The date and time the last backup occurred. This is the backup that will be restored.
- **Name** — The name of the switch that will be restored.
- **Configuration Type** — The type of configuration for the switch (FC, DCB-running, or DCB-startup).
- **Backup Type** — The type of backup used to obtain the configuration files from the device. Backup options include the following types:
  - Discovery — Occurs when the discovery backup is obtained after the discovery process.
  - Event Triggered — Occurs when a trap is generated by the device during a configuration change.
  - Manual — Occurs when collecting the backup file for the switches or fabrics using Save Configuration dialog box.
  - Scheduled — Occurs when backup is obtained at the scheduled time.
- **IP Address** — The IP address of the switch that will be restored.
- **WWN** — The world wide name of the switch that will be restored.
- **Comments** — Comments regarding the switch.

4. Click the following buttons:

- **Restore** button — Select one or more configuration files from the **Configurations** table and click to restore that configuration. To restore a configuration, refer to "Restoring a configuration from the repository" on page 416.
- **View** button — Select a row in the **Configurations** table and click to display the contents of the selected configuration. Refer to "Viewing configuration file content" on page 417.
- **Delete** button — Select one or more configurations from the **Configurations** table and click to manually delete the configuration from the repository of the management server. Refer to "Deleting a configuration" on page 419.
- **Export** button — Select a configuration and click to launch the **Export Configuration** dialog box, which allows you to export the configurations to a text file. Refer to "Exporting a configuration" on page 419 for more information.
- **Import** button — Select a configuration and click to launch the **Import Configuration** dialog box, which allows you to import the configurations to a text file. Refer to "Importing a configuration" on page 420 for more information.

- **Compare** button — Select two configurations (same product or two different products) and click to launch the **Compare** dialog box with the differences between the two configurations highlighted. Refer to "Comparing switch configurations" on page 420 for more information.

  NOTE
  You cannot delete the baseline or the latest configuration.

5. Click **Cancel** to close.
6. Click **OK** to confirm.

Restoring a configuration from the repository

If you delete a fabric or switch from discovery, the configuration remains in the repository until you delete it manually. Stored configurations are linked to the switch WWN; therefore, if the IP address or switch name is changed and then rediscovered, the **Configuration File Manager** dialog box displays the new switch name and IP address for the old configuration.

NOTE
This feature requires a Trial or Licensed version.

1. Select **Configure > Configuration File > Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.

FIGURE 190 Saved Switch Configurations dialog box

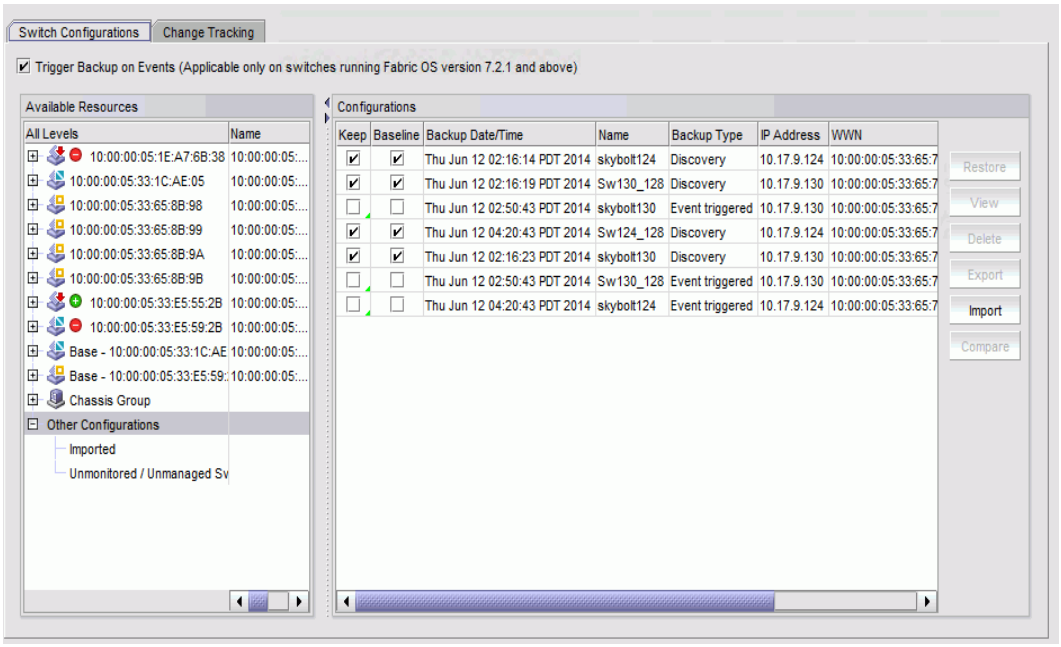

The **Saved Switch Configurations** table displays the following information:

- **Keep** check box — Select to keep the associated configuration past the defined age limit. The configuration will be kept until it is manually deleted, or until the **Keep** check box is cleared to enable the age limit again.
- **Baseline** check box — Select to compare all additional product configuration backup files to the baseline configuration for deviation. You can view configuration deviation status on the Status bar of the Management application window.
- **Backup Date/Time** — The date and time the last backup occurred. This is the backup that will be restored.

- **Name —** The name of the switch that will be restored.
- **Configuration Type —** The type of configuration for the switch (FC, DCB-running, or DCB-startup).
- **IP Address —** The IP address of the switch that will be restored.
- **WWN —** The world wide name of the switch that will be restored.
- **Backup Type** — The type of backup used to obtain the configuration files from the device. Backup options include the following types:
  - Discovery — The discovery backup is obtained after the discovery process.
  - Event Triggered — Occurs when a trap is generated by the device during a configuration change.
  - Manual — Occurs when you launch the **Save Switch Configuration** dialog box and save the switch configuration manually.
  - Scheduled — Occurs when backup is obtained at the scheduled time.

2. Select the configuration you want to restore, and click **Restore**.

   The configuration is downloaded to the device. If necessary, the restoration process prompts you to disable and reboot the device before the configuration begins. This lets you determine whether the configuration backup should be performed immediately or at a later time.

   If you confirm the restoration, the entire configuration is restored; you cannot perform selective download for specific configuration sections.

   You can also perform the following functions from the **Configuration File Manager** dialog box:

   - "Comparing switch configurations" on page 420
   - "Viewing configuration file content" on page 417
   - "Deleting a configuration" on page 419
   - "Exporting a configuration" on page 419
   - "Importing a configuration" on page 420

Viewing configuration file content

NOTE
This feature requires a Trial or Licensed version.

You can view switch configuration file content in a text file.

1. Right-click a device in the Product List or the Connectivity Map, and select **Configuration File > Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.

2. Click the **Switch Configurations** tab.

   The selected product in the Connectivity Map or the Product List is marked as default, or you can select the switch or fabric from the device tree.

3. Select a file in the **Configurations** table.

4. Click **View**.

   The configuration details display, Details include the backed-up switch, including boot parameters, licensing information, and configuration information. If you want to save the contents as a text file, click **Copy to Clipboard**, paste the copy into a text editor (such as Notepad), and save the file.

FIGURE 191 Configuration file content

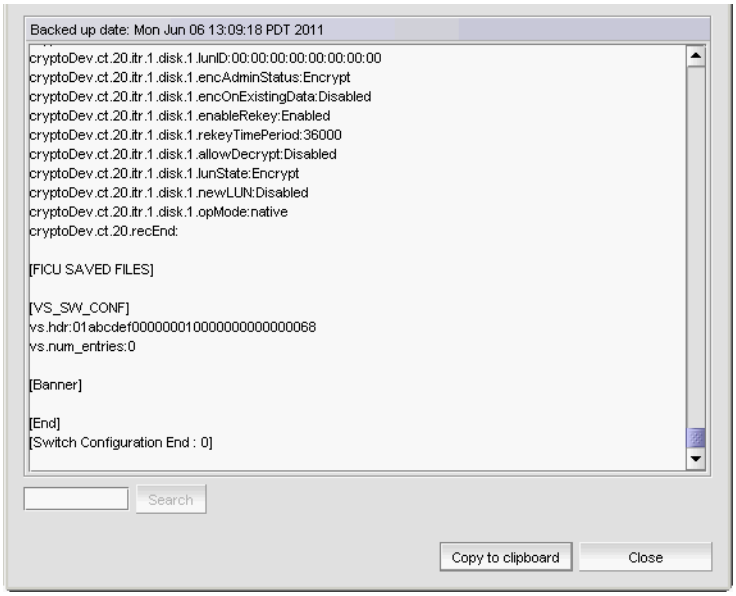

5. Click **Close** to close the dialog box.

Searching the configuration file content

NOTE
This feature requires a Trial or Licensed version.

To search the configuration file content, complete the following steps.

1. Right-click a device in the Product List or the Connectivity Map, and select **Configuration File > Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.

2. Click the **Switch Configuration** tab.

   The default product is selected in the Connectivity Map or Product List, or select the switch or fabric from the device tree.

3. Select a file in the **Configurations** table.

4. Click **View**.

   The configuration details display.

5. Enter the information you want to search for in the **Search** field and click **Search**.

   The text string you are searching for is highlighted in the dialog box. Continue clicking **Search** to scroll through the contents until you find the information you need. If the search item is not found, a “not found” message displays. Click **OK** to close the message.

FIGURE 192 Searching Configuration file content

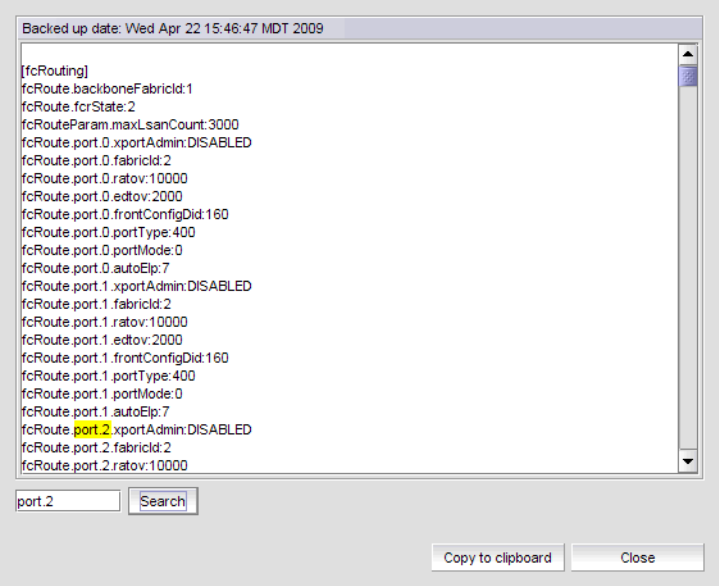

6. Click **Close** to close the dialog box.

Deleting a configuration

NOTE
This feature requires a Trial or Licensed version.

NOTE
Baseline configurations and the latest configurations will not be deleted.

1. Right-click a device in the Product List or the Connectivity Map, and select **Configuration File > Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.

2. Click the **Switch Configurations** tab.

   The default product is selected in the Connectivity Map or the Product List or you can select the switch or fabric from the device tree.

3. Select the configuration file you want to delete, and click **Delete**.

Exporting a configuration

NOTE
This feature requires a Trial or Licensed version.

1. Right-click a device in the Product List or the Connectivity Map, and select **Configuration File > Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.

2. Select the configuration you want to export, and click **Export**.

   The file appropriate to your operating system displays.

3. Click the file to select the location into which you want to export the configuration.
4. Click **Export**.

   The configuration is automatically named (*cfg_Switch_name_Date_and_Time)* and exported to the location you selected.

Importing a configuration

NOTE
This feature requires a Trial or Licensed version.

1. Right-click a device in the Product List or the Connectivity Map, and select **Configuration File > Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.

2. Click **Import**.

   The file appropriate to your operating system displays.

3. Click the file to select the file from which you want to import the configuration, and click **Import**.

   The imported configuration will be saved under **Available Resources** > **Other Configuration** > **Imported Configurations.**

Comparing switch configurations

The **Compare** dialog box allows you to display the contents of two configurations side-by-side. To compare two configurations, perform the following steps.

1. Click the **SAN** tab.
2. Select **Configure** > **Configuration** > **Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.

3. Click the **Switch Configurations** tab.
4. Select a switch to view the configurations.
5. Select two configurations and click **Compare**.

   The **Compare** dialog box displays, as shown in Figure 193.

FIGURE 193 Compare dialog box

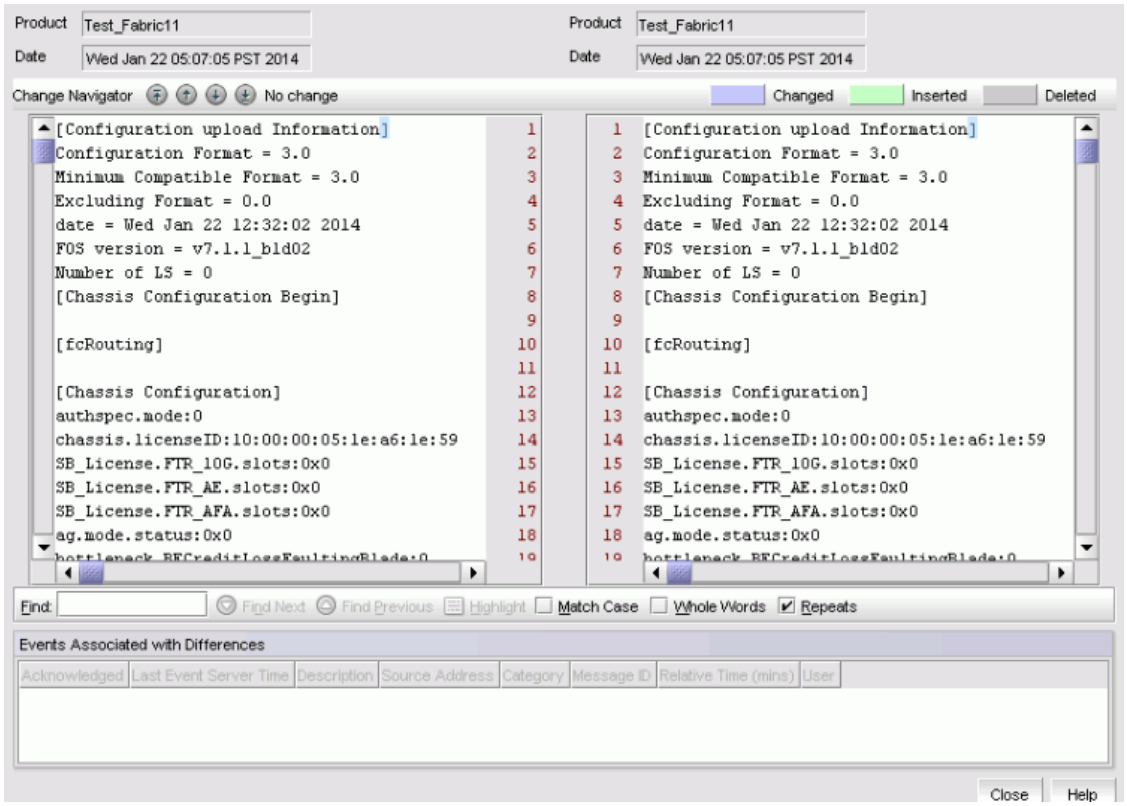

The **Compare** dialog box displays the following information:

- **Product** — The name of the switch.
- **Date** — Displays the date the switch configuration was taken.
- **Change Navigator** buttons/legend — The **Change Navigator** buttons and legends are enabled when there is at least one change between two compared files.
  - Go to first change button ( ) — Click to move to the first change.
  - Go to previous change button ( ) — Click to move to the previous change.
  - Go to next change button ( ) — Click to move to the next change.
  - Go to last change button ( ) — Click to move to the last change.
  - Number of changes label — Indicates the number of changes. If there are no differences, displays “No change”.
  - Differences legend — Displays the color legend for differences:
    - Changed status displays in blue.
    - Inserted status displays in green.
    - Deleted status displays in gray.
- **Phrase not found** icon — Displays when the search text string is not found.
- Configuration contents areas — Displays the contents of the selected configurations.
- **Find** — Enter a text string and take one of the following actions:
  - Click **Find Next** — Searches for the next matching string in the configuration.
  - Click **Find Previous** — Searches for the previous matching string in the configuration.
- **Highlight** button — Click to highlight the text string.
- **Match Case** check box — Click to render the search case-sensitive.

- **Repeats** check box — Click to continue the search at the top when the bottom is reached.
- **Whole Words** check box — Click to continue the search. Displays the combination of **Highlight**, **Match Case**, and **Repeats** searches.
- The **Events Associated with Differences** table is only available when you select two configuration backup files for the same product. The table list events (up to 100) associated with the configurations. Right-click an event and select **properties** to view the **Event Properties** dialog box (refer to "Displaying event properties from the Master Log" on page 2006).

  NOTE
  The **Events Associated with Differences** table is blank for configuration files triggered on a Fabric OS DCB device.

- **Acknowledge** check box — Select to acknowledge the event and remove it from the Master Log. The event is not removed from the **Events Associated with Differences** table.
- **Source Address** — IP address of the switch on which a change occurred.
- **Category** — Audit log event category. Options include product audit event and product event.
- **Description** — Description of the event.
- **Last Event Server Time** — Time and date the event last occurred on the server.
- **Message ID** — Message ID of the event.
- **Relative Time (mins)** — Relative time from the selected backup time to the time the event occured.
- **User** — Name of the user responsible for triggering the event.

6. Click **Close**.

Keeping a copy past the defined age limit

NOTE
This feature requires a Trial or Licensed version.

1. Select **Configure > Configuration File > Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.

2. Select the check box under **Keep** for the configuration you want to preserve.

   The configuration will be kept until it is manually deleted, or until the **Keep** check box is cleared to enable the age limit again. The file marked as **Baseline** will be marked as **Keep.**

3. Click **Apply** to save changes.

Tracking changes from the baseline configuration

Use change tracking to compare the latest backup configuration file with the configuration that is designated as the baseline.

NOTE
The **Change Tracking** tab is not displayed for imported, unmanaged, or unmonitored switches.

1. Click the **SAN** tab.
2. Select **Configure > Configuration File > Configuration File Manager**.

   The **Configuration File Manager** dialog box displays.

3. Click the **Change Tracking** tab.

   The **Change Tracking** tab of the **Configuration File Manager** dialog box displays, as shown in Figure 194 on page 423

   FIGURE 194 Change Tracking tab

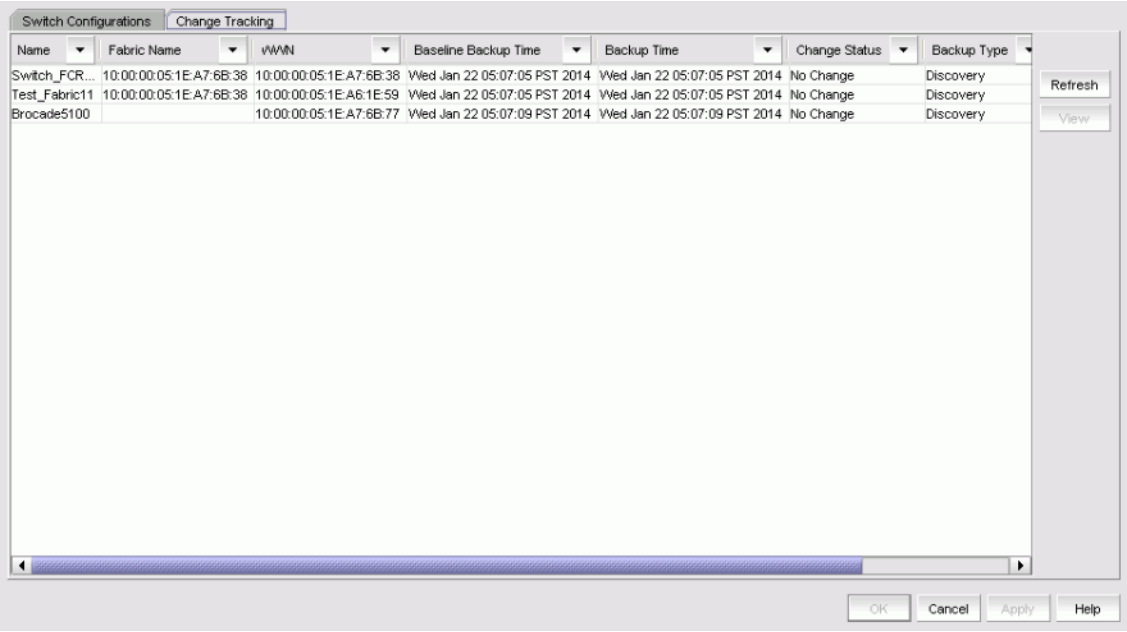

   The **Change Tracking** tab displays the following information:

   - **Name —** The switch name.
   - **Fabric Name —** The name of the fabric.
   - **WWN —** The world wide name of the switch selected to be the destination switch.
   - **Baseline Backup Time** — The date and time when the baseline configuration for the device was copied into the repository of the management server.
   - **Backup Time —** The time when the last backup attempt occurred for the selected device.
   - **Change Status** — The change status of the latest device backup.
   - **Backup Type** — Backup options include the following types:
     - Discovery **—** The discovery backup is obtained after the discovery process.
     - Event Triggered— Occurs when a trap is generated by the device during a configuration change.
     - Manual — Occurs when you click the **Save Running/Startup Configuration** button.
     - Scheduled — Occurs when backup is obtained at the scheduled time.

4. Perform one of the following actions:

   - Click the **Refresh** button to update the network and retrieve the latest data from the database.
   - Click the **View** button to compare the contents of the latest configuration and the baseline configuration. You can view only one configuration at a time.

5. Click **OK** to save the configuration.

6. Click **Cancel** to exit the dialog box.

Replicating configurations

NOTE
This feature requires a Trial or Licensed version.

You can replicate a switch SNMP configuration, the Fabric Watch configuration, Trace Destination configuration, or the entire configuration.

1. Select **Configure > Configuration File > Replicate > Configuration**.

   A wizard is launched to guide you through the process. The first step of the wizard, **Overview**, displays. There are seven steps in the Replicate Switch Configuration:

- **Overview**, which describes the wizard.
- **Configuration Type**, which allows you to select the type of configuration you wish to replicate. For more information about the fields and components of this step, refer to Table 33 on page 424.
- **Source Location**, which allows you to select the location of the configuration you wish to replicate. For more information about the fields and components of this step, refer to Table 34 on page 425.
- **Source Configuration**, which allows you to select the source switch to replicate. For more information about the fields and components of this step, refer to Table 35 on page 425.
- **Destination Switches**, which allows you to select the destination switch. For more information about the fields and components of this step, refer to Table 36 on page 426.
- **Validation**, which lists the configuration settings that you can validate before you replicate. For more information about the fields and components of this step, refer to Table 37 on page 427.
- **Summary**, which lists the replication settings that successfully ran on all the selected destination switches. For more information about the fields and components of this step, refer to Table 38 on page 427.

2. To proceed to the next step in the wizard, click **Next**. To return to the previous screen, click **Previous**.

TABLE 33 Step 2. Configuration Type

Field/Component	Description
All FC option	Replicates the entire configuration, including security settings. **Warning**: This is a disruptive operation and selected destination switches will be disabled prior to downloading the configuration.
Partial FC option	Replicates a part of the FC configuration. Select one of the following options: • **Fabric Watch** option — Lists switches with Fabric Watch configurations that you can replicate. • **SNMP** option — Lists switches with SNMP configurations that you can replicate. **Include system group configuration** check box — Select to include the SNMP system group configuration in the replication. • **Trace Destinations** option — Lists switches with trace destination configurations that you can replicate. **Warning**: This is a disruptive operation and selected destination switches will be disabled prior to downloading the configuration.
All DCB option	Replicates the entire DCB startup configuration.

TABLE 34 Step 3. Source Location

Field/Component	Description
Configuration File Manager option	Select to replicate the entire Configuration File Manager to the destination switches.
Configuration from the switch option	Select to assign a designated switch to the destination switch.
File in text format option	Select to choose a valid configuration file from the local file system by either typing in the complete path of the file in the text box or selecting the file using the **Browse** option on the **Source Configuration** screen.

TABLE 35 Step 4. Source Configuration

Field/Component	Description
Saved Switch Configuration table (**Configuration File Manager** only)	Lists the information related to the saved switch, if you selected **Configuration File Manager** on the **Source Location** screen.
Backup Date/Time (**Configuration File Manager** only)	The date and time the last backup occurred on the switch.
Fabric Name	The name of the fabric that is associated with the selected available switch.
Name	The name of the source switch to be replicated.
Configuration Type	The type of configuration.
IP Address	The IP address of the source switch to be replicated.
WWN	The world wide name of the source switch to be replicated.
Name	The name of the selected switch.
Comments	Comments regarding the switch.
Available Switches table (**Configuration from the switch** only)	Lists the information related to the available switches, if you selected **Configuration from the switch** on the **Source Location** screen.
All Levels	A list of all switches.
Additional Port Info	Additional information about the port.
Attached Port #	The number of the attached port.
BB Credit	The BB Credit of the port.
Class	The class value of the FICON device port.
Contact	The primary contact at the customer site.
Description	A description of the customer site.
Domain ID	The switch port's top-level addressing hierarchy of the domain.
FC Address	The Fibre Channel address of the port.
Firmware	The firmware version.
IP Address	The IP address of the switch.
Location	The customer site location.
Model	The name and model number of the hardware.
Name	The name of the switch.
Port #	The number of the port.
Port Count	The total number of ports.
Port Type	The type of port (for example, expansion port, node port, or NL_port).
Product Type	The type of product.
Protocol	The protocol for the port.

TABLE 35 Step 4. Source Configuration (Continued)

Field/Component	Description
Serial #	The serial number of the switch.
Speed Configured (Gbps)	The actual speed of the port in Gigabits per second.
State	The port state, for example, online or offline.
Status	The operational status of the port.; for example, unknown or marginal.
Symbolic Name	The symbolic name for the port.
Tag	The tag number of the port
Vendor	The hardware vendor's name.
WWN	The world wide name of the source switch to be replicated.
Zone Alias	The zone alias.
Configuration File field and **Browse** button (**File in text format** only)	Select a valid configuration file from the local file system by either typing in the complete path of the file in the text box or selecting the file using the **Browse** button.

TABLE 36 Step 5. Destination Switches

Field/Component	Description
Available Switches table	Lists the available switches you can select to be applied to the selected switches table.
All Levels	A list of all switches.
Additional Port Info	Additional information about the port.
Attached Port #	The number of the attached port.
BB Credit	The BB Credit of the port.
Class	The class value of the FICON device port.
Contact	The primary contact at the customer site.
Description	A description of the customer site.
Domain ID	The switch port's top-level addressing hierarchy of the domain.
FC Address	The Fibre Channel address of the port.
Firmware	The firmware version.
IP Address	The IP address of the switch.
Location	The customer site location.
Model	The name and model number of the hardware.
Name	The name of the switch.
Port #	The number of the port.
Port Count	The total number of ports.
Port Type	The type of port (for example, expansion port, node port, or NL_port).
Product Type	The type of product.
Protocol	The protocol for the port.
Serial #	The serial number of the switch.
Speed Configured (Gbps)	The actual speed of the port in Gigabits per second.
State	The port state, for example, online or offline.
Status	The operational status of the port.; for example, unknown or marginal.

TABLE 36 Step 5. Destination Switches (Continued)

Field/Component	Description
Symbolic Name	The symbolic name for the port.
Tag	The tag number of the port
Vendor	The hardware vendor's name.
WWN	The world wide name of the source switch to be replicated.
Zone Alias	The zone alias.
Right and left arrow buttons	Click to move the switches back and forth between the **Available Switches** table and the **Selected Switches** table.
Selected Switches table	Lists the switches selected as the destination switches.
Switch Name	The name of the switch selected to be the destination switch.
IP	The IP address of the switch selected to be the destination switch.
WWN	The world wide name of the switch selected to be the destination switch.
Current Firmware	The current firmware.
Status	The status of the switch .

TABLE 37 Step 6. Validation

Field/Component	Description
Validation Settings table	The replication settings that have been configured in previous steps; for example, the configuration type, source configuration, and destination settings. Click **Finish** to approve the settings.
Disable Destination Switch check box	Select to disable the destination switch during replication.

TABLE 38 Step 7. Summary

Field/Component	Description
Summary table	The replication settings that have been successfully applied to the selected destination switches; for example, the configuration type, source configuration, and destination settings. Click **Close** to close the dialog box.

Replicating security configurations

NOTE
This feature requires a Trial or Licensed version.

You can replicate an AD/LDAP Server, DCC, IP, RADIUS Server, or SCC security policy.

1. Select **Configure > Configuration File > Replicate > Security**.

   A wizard is launched to guide you through the process. The first step of the wizard, **Overview**, displays. There are seven steps in the **Replicate Switch Security Policy Configuration** wizard:

   - **Overview**, which describes the wizard.
   - **Configuration Type**, which allows you to select the type of configuration you wish to replicate. For more information about the fields and components of this step, refer to Table 39 on page 428.
   - **Select Source Switch**, which allows you to select the source device of the security policy configuration you wish to replicate. For more information about the fields and components of this step, refer to Table 40 on page 428.
   - **Select Destination Switches**, which allows you to select the destination devices. Only devices that can accept the selected security policy configuration display. For more information about the fields and components of this step, refer to Table 41 on page 429.

- **Validation**, which lists the configuration settings that you can validate before you replicate. For more information about the fields and components of this step, refer to Table 42 on page 429.
- **Summary**, which lists the replication settings that successfully ran on all the selected destination switches. For more information about the fields and components of this step, refer to Table 43 on page 429.

2. To proceed to the next step in the wizard, click **Next**. To return to the previous screen, click **Previous**.

TABLE 39 Step 2. Configuration Type

Field/Component	Description
AD/LDAP Server option	Select to replicate the Active Directory/Lightweight Directory Access Protocol (AD/LDAP) Server security policy. If both the source and destination devices are running Fabric OS 7.1 or later, also replicates the LDAP Role mapping configuration.
DCC Policy option	Select to replicate the Device Connection Control (DCC) security policy.
IP Policy option	Select to replicate the Internet Protocol (IP) Filter security policy.
RADIUS Server option	Select to replicate the Remote Authentication Dial-In User Service (RADIUS) Server security policy.
SCC Policy option	Select to replicate the Switch Connections Control (SCC) security policy.

TABLE 40 Step 3. Select Source Switch

Field/Component	Description
Available Switches table	Lists the devices from which you can select to replicate a security policy
Fabric Name	The name of the fabric that is associated with the selected available switch.
Switch Name	The name of the source switch to be replicated.
Switch IP Address	The IP address of the source switch to be replicated.
Switch WWN	The world wide name of the source switch to be replicated.
Name	The name of the selected switch.
Device Type	The type of device port.
Tag	The tag number of the port
Serial #	The serial number of the switch.
WWN	The switch port's world wide name.
IP Address	The switch port's IP address.
Domain ID	The switch port's top-level addressing hierarchy of the domain.
Vendor	The hardware vendor's name.
Model	The name and model number of the hardware.
Port Count	The total number of ports.
Firmware	The firmware version.
Location	The customer site location.
Contact	The primary contact at the customer site.
Description	A description of the customer site.
State	The port state, for example, online or offline.
Status	The operational status of the port.; for example, unknown or marginal.

TABLE 41 Step 4. Select Destination Switches

Field/Component	Description
Available Switches table	Lists the available switches you can select to be applied to the selected switches table.
Name	The name of the available switch.
Device Type	The type of device port.
Tag	The tag number of the port.
Serial #	The serial number of the switch.
WWN	The switch port's world wide name.
IP Address	The switch port's IP address.
Domain ID	The switch port's top-level addressing hierarchy of the domain.
Vendor	The hardware vendor's name.
Model	The name and model number of the hardware.
Port Count	The total number of ports.
Firmware	The firmware version.
Location	The customer site location.
Contact	The primary contact at the customer site.
Description	A description of the customer site.
State	The port state, for example, online or offline.
Status	The operational status of the port; for example, unknown or marginal.
Right and left arrow buttons	Click to move the switches back and forth between the **Available Switches** table and the **Selected Switches** table.
Selected Switches table	Lists the switches selected as the destination switches.
Switch Name	The name of the switch selected to be the destination switch.
IP	The IP address of the switch selected to be the destination switch.
WWN	The world wide name of the switch selected to be the destination switch.
Current Firmware Status	The status of the current firmware.

TABLE 42 Step 5. Validation

Field/Component	Description
Validation Settings table	The replication settings that have been configured in previous steps; for example, the configuration type, source configuration, and destination settings. Click **Finish** to approve the settings.
Disable Destination Switch check box	Select to disable the destination switch during replication.

TABLE 43 Step 6. Summary

Field/Component	Description
Summary table	The replication settings that have been successfully applied to the selected destination switches; for example, the configuration type, source configuration, and destination settings. Click **Close** to close the dialog box.

Enhanced group management

Use Enhanced Group Management (EGM), a separate licensed feature, to control access to specific features on Fabric OS devices. The features affected include the following:

- Firmware Download - enables you to perform group firmware download.

  For specific instructions for firmware download, refer to "Firmware management" on page 430.

- Security - enables you to perform Group Security Policy Replication.

  For specific instructions for security, refer to "Configuration file management" on page 409.

- Configuration Management - enables you to perform Group Configuration Upload and Replication.

  For specific instructions for configuration management, refer to "Tracking changes from the baseline configuration" on page 422.

Firmware management

A firmware file repository (Windows systems only) is maintained on the server in the following location: C:\Program Files*Install_Directory*\data\ftproot\Firmware\Switches\7.0\n.n.n\n.n.n

The firmware repository is used by the internal FTP, SCP, or SFTP server that is delivered with the Management application software, and may be used by an external FTP server if it is installed on the same platform as the Management application software. The repository is not available to FTP servers on external platforms.

NOTE
The repository is not available on external SCP or SFTP servers installed on the same platform as the Management application software.

NOTE
Non-disruptive firmware download (HCL) is not supported when downgrading from Fabric OS version 7.0 to 6.4. You must remove all non-default logical switches and disable Virtual Fabrics before downgrading.

NOTE
You cannot use Fabric OS firmware download with command line options in the Management application.

NOTE
You cannot configure firmware management through the Management application for Emulex adapters.

Downloading firmware

NOTE
You cannot use Fabric OS firmware download with command line options in the Management application.

NOTE
Beginning with Fabric OS 7.4.0 Fabric Watch support is deprecated and displays "Fabric Watch is not supported on Fabric OS v7.4 instead use MAPS to monitor the switches warning message, during the firmware upgrade from 7.3.0 to 7.4.0.

NOTE
Beginning with Fabric OS 7.4.0 IP Extension mode is supported on the 16 Gbps 24-FC port, 18 GbE port switches. During firmware download, if IP Extension mode is supported on the 16 Gbps 24-FC port, 18 GbE port switches, then the IP traffic through FCIP will be disruptive and a warning message "IP extension mode is enabled" displays.

NOTE
For Peer zone enabled switches, during firmware downgrade from Fabric OS 7.4 to 7.3, a confirmation message "Peer zones will be considered as normal zones after firmware downgrade on the following switches" displays.

You can download firmware using the **Firmware Management** dialog box.

1. Select **Configure > Firmware Management**.

   The **Firmware Management** dialog box displays.

2. Select the **Download** tab (Figure 195).

   FIGURE 195 Download tab

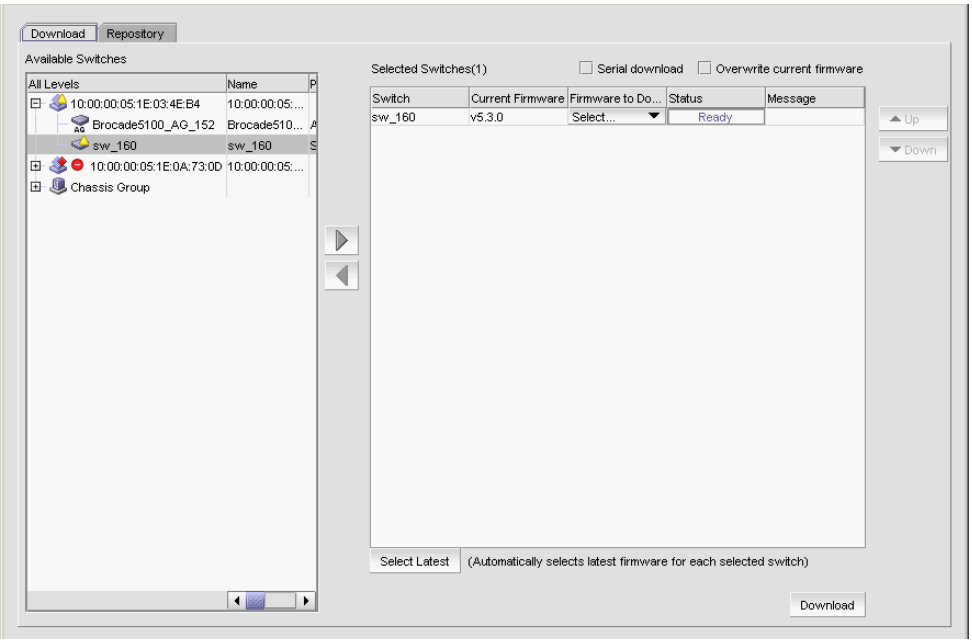

3. Select one or more switches from the **Available Switches** table.

   The **Available Switches** table lists the switches that are available for firmware download.

4. Click the right arrow to move the switches to the **Selected Switches** table.

   If you select any switches that do not support firmware download, a message displays. Click **OK** on the message.

   The switches that support firmware download display in the **Selected Switches** table. The current version displays in the **Current Firmware** column.

5. (Built-in FTP, SCP, or SFTP server) If you have your FTP, SCP, or SFTP server configured to use the built-in FTP, SCP, or SFTP server, select a specific version from the **Firmware to Download** column, or use **Select Latest** to automatically select the latest version. Go to step 8.

   If you have your FTP server configured to use an external FTP server, the **Firmware to Download** column is empty.

6. (External FTP, SCP, or SFTP server) If you configured an external FTP, SCP, or SFTP server, choose from one of the following options in the **External FTP/SCP/SFTP Server** area:

   - Select the **FTP server** option to download from the external FTP server and configure the following on the FTP server:

   a. Create a user name and password.

   b. Select the **Shared folders** link and set the firmware location as the home directory and select all check boxes under the **Files** and **Directories** attributes. Continue with step 7.

   - Select the **SCP Server** option to download from the external SCP server. Continue with step 7.

     NOTE
     The Management application only supports WinSSHD as the third-party Windows external SCP server. Firmware upgrade and downgrade through WinSSHD is only supported on devices running Fabric OS 7.0 or later.

   - Select the **SFTP Server** option to download from the external SFTP server. Continue with step 7.

     NOTE
     The Management application only supports WinSSHD as the third-party Windows external SFTP server. Firmware upgrade and downgrade through WinSSHD is only supported on devices running Fabric OS 7.0 or later.

7. (External FTP, SCP, or SFTP server) If you configured an external server, enter the path to the firmware directory in the **Firmware Directory** field.

   A confirmation message displays. Click **Yes** on the confirmation message.

   This field does not display if the external server is installed on the same machine as the Management application and occupies port 21.

8. To download the firmware to the selected switches one at a time, select the **Serial download** check box.

   Use the **Up** and **Down** buttons to determine the order in which the firmware is downloaded to the switches. If firmware download fails on one switch, all other switches in the queue will be skipped.

   If the **Serial download** check box is cleared, the download occurs in parallel on the switches (up to 20 at a time).

9. To overwrite the current firmware, even if the selected version is the same as the version currently running on the switch, click the **Overwrite current firmwares** check box.

10. Click **Download**.

    While the firmware is downloaded to the device, the **Status** column displays the current download status. Once firmware download is complete, the **Message** column displays whether the download was a success or failure.

    NOTE
    Firmware Download using the FTP is not supported if the **Enforce Secure Config Upload/Download** attribute is set as **Yes** in the FOS Switches. Incase the user downloads the firmware using the FTP, an error message displays as "**Secure Config Upload / Download** is enabled on the switch. Please select either SCP or SFTP for Firmware Download in **Server Options** dialog.

Firmware download support for HCL enabled Fabric OS 16 Gbps 24-FC port, 18 GbE port switches

In the Management application software, when the HCL enabled Fabric OS 16 Gbps 24-FC port, 18 GbE port switches with high availability (HA) configured FCIP tunnels are selected for Firmware Download, the firmware download will be initiated in a serial manner even though the user has not selected the **Serial Download** check box.

- If HCL is not enabled and HA is not configured for the FCIP tunnels, then the switches are considered as any other FOS switch and the firmware download happens in a regular way.
- If Parallel download is selected for a group of switches with HCL enabled Fabric OS 16 Gbps 24-FC port, 18 GbE port switch with HA configured FCIP tunnels present in them, the following warning message displays:

  "Firmware download will be serial on the following HCL supported switches since one or more FCIP tunnels between them are HA configured. Do you want to continue?".

  Name1:IP_address1; Name2:IP_address2; Name3:IP_Address3;..."
- When the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch is selected for Firmware Download, only if all the HA tunnels configured are online, the device is considered as HCL capable. If FCIP tunnels are either not configured or offline, then traffic disruption may occur during firmware download and the following warning messages displays:

One or more tunnels are not HA configured or offline and might result in traffic disruption

Displaying the firmware repository

The firmware repository is available on the **Firmware Management** dialog box. The Management application supports .zip and .gz compression file types for firmware files.

Initially, the firmware repository is configured to use the built-in FTP, SCP, or SFTP server. To use an external FTP server, refer to "Configuring an external FTP, SCP, or SFTP server" on page 125.

NOTE
The repository is not available on external SCP or SFTP servers installed on the same platform as the Management application software.

1. Select **Configure > Firmware Management**.

   The **Firmware Management** dialog box displays.

2. Select the **Repository** tab (Figure 196).

   Initially, the repository is empty. You must import firmware files into the repository (refer to "Importing a firmware file" on page 434). Imported firmware files are then displayed under **Firmware Repository**.

FIGURE 196 Repository tab

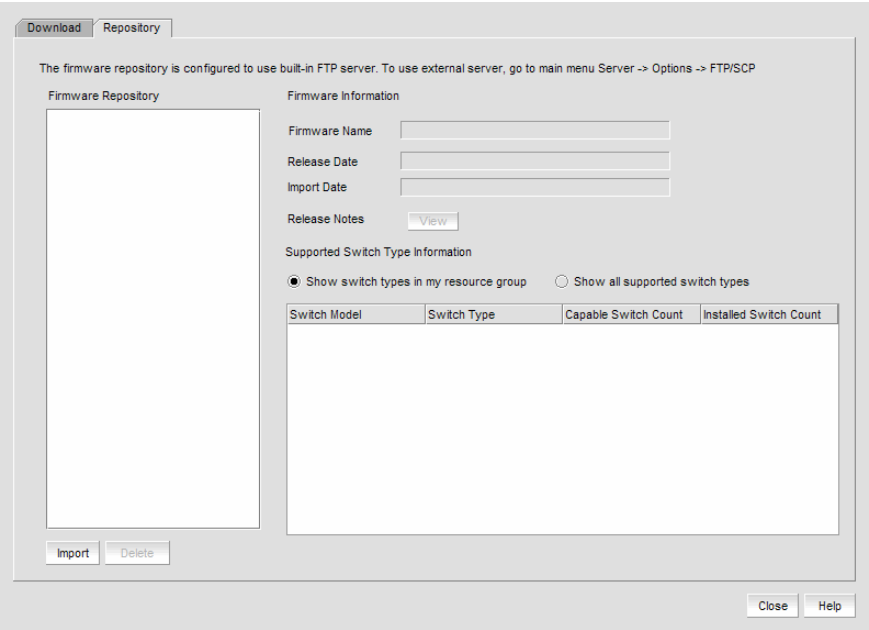

3. View information about a specific firmware file by selecting the firmware file in the **Firmware Repository**.

   The following information displays.

   - **Firmware Name —** Lists the version of the current installed firmware.
   - **Release Date —** Lists the date and time the firmware was released.
   - **Import Date —** Lists the date and time the firmware was imported.
   - **Release Notes View** button **—** Click to view the release notes, if imported, which contain information about downloading firmware.

     For internal built-in FTP, SCP, or SFTP servers or external SCP or SFTP servers running on the same system as the Management application, if there is a space in the release note file name, you will not be able to view the release notes.
   - **Supported Switch Type Information** table — Shows the switch type, capable switch count, and number of installed switches. You can choose one of two switch groups:
     - Show switch types in my resource group.
     - Show all supported switch types.

4. Click **Import** to launch the **Import Firmware from File** dialog box, which enables you to browse to the firmware location for importing. Refer to "Importing a firmware file" on page 434.
5. Click **Delete** to delete firmware files from the firmware repository. Refer to "Deleting a firmware file" on page 436.
6. Click **Close** to close the **Firmware Management** dialog box.

Importing a firmware file

You can import firmware files, release notes, and MD5 checksum files into the firmware repository.

1. Select **Configure > Firmware Management**.

   The **Firmware Management** dialog box displays.

2. Select the **Repository** tab (Figure 196).
3. Click **Import**.

   The **Import Firmware from File** dialog box displays (Figure 197).

   FIGURE 197 IImport Firmware from File dialog box

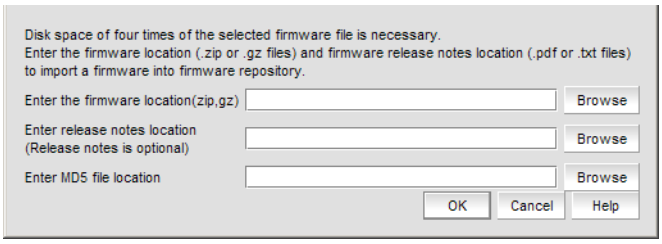

4. Enter or browse to the location of the firmware file.

   NOTE
   Firmware file import requires disk space that is four times the size of the selected file.

   The Management application supports .zip and .gz compression file types for firmware files.

5. (Optional) Enter or browse to the location of the release notes.

   The Management application supports .pdf and .txt file types for release notes.

   For internal built-in FTP, SCP, or SFTP servers or external SCP or SFTP servers running on the same system as the Management application, if there is a space in the release note file name, you can import the file. However, you will not be able to view the release notes.

6. Enter or browse to the location of the MD5 file (.md5 file type).

   If the MD5 checksum file is located in the same directory as the firmware file and has the same file name (with the md5 extension), this field is auto-populated.

   The MD5 checksum file can be obtained from the Fabric OS product download site in the same location as the firmware file. The MD5 checksum file cannot be downloaded directly from the site; however, you can open the file, copy and paste the contents into a new file, and save the file with the md5 extension in the same directory as the firmware file.

   The MD5 checksum file validates the firmware file twice; first when the firmware is downloaded to the client and again when the file is copied from the client to the server's repository.

   If you configure the Management application to enforce the MD5 checksum file import ("Enforcing MD5 file during import" on page 100), this field is not optional.

7. Click **OK**.

   You return to the **Repository** tab. The file is listed in the Firmware Repository when the import is complete and successful.

8. Click **Close** to close the **Firmware Management** dialog box.

Deleting a firmware file

Firmware files can be deleted from the Firmware Repository.

1. Select **Configure > Firmware Management**.

   The **Firmware Management** dialog box displays.

2. Select the **Repository** tab.
3. Select one or more firmware files from the Firmware Repository for deletion.
4. Click **Delete**.

   A confirmation dialog displays. Click **Yes** to confirm. The firmware file is deleted from the repository.

Switch password management

Switch password management enables you to change or reset the switch password for one or more users across multiple switches.

NOTE
You can change the switch password for root and factory users only by using the **Change Password** button because the current password is mandatory.

NOTE
If you change the switch password for one Fabric ID (FID) user name, the switch password changes for all FIDs that have the same user name.

NOTE
You should change the switch password before the expiriration date; however, if the switch password expires, you must provide valid credentials in the **Discovery** dialog box.

NOTE
You must change switch password using **Discovery Edit** dialog to avoid any failure in configuration when a switch is discovered and **passwdconfig** command **--hash** value is modified in the Command Line Interface (CLI). A master log event is triggered when the **passwdconfig --hash** command is modified in CLI.

Changing the switch password

To change the switch password, complete the following steps.

1. Select **Configure** > **Switch Password.**

   The **Manage Switch Password** dialog box displays, as shown in Figure 198.

FIGURE 198 Manage Switch Password dialog box

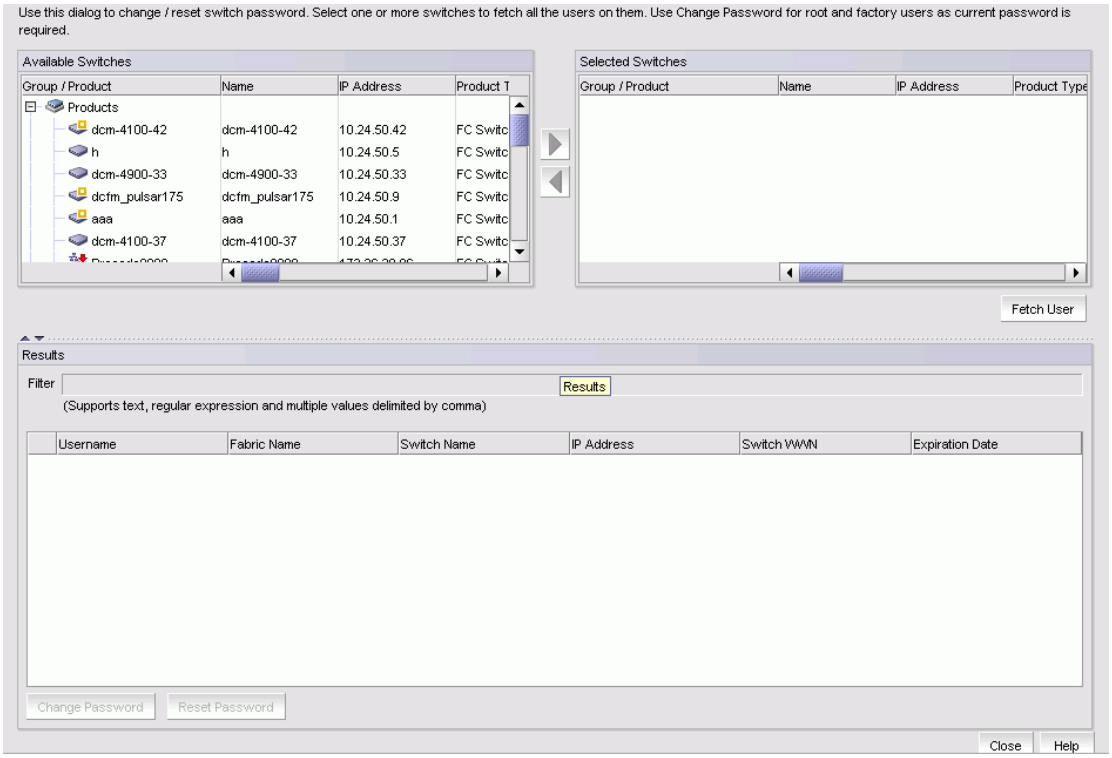

The **Manage Switch Password** dialog box includes the following components:

- **Available Switches** table — Displays the switches available in the current view of the application.
- **Selected Switches** table — Displays the selected switches.
- **Results** table — Displays the users associated with the selected switches.

2. Select the switches for which you want to change the switch password from the **Available Switches** table.

   NOTE
   You cannot change the switch password for unmanaged switches.

3. Click the right arrow to move the selected switches to the **Selected Switches** table.

   If you select an unmanaged switch, an error messages displays.

   Click **OK** to close the error message.

4. Click **Fetch User**.

   The **Filter** field in the **Results** table of the **Manage Switch Password** dialog box supports text, regular expressions, multiple values delimited with commas, and wildcards.

   The users for the selected switches display in the **Results** table (as shown in Figure 199).

FIGURE 199 Manage Switch Password dialog box - Results table

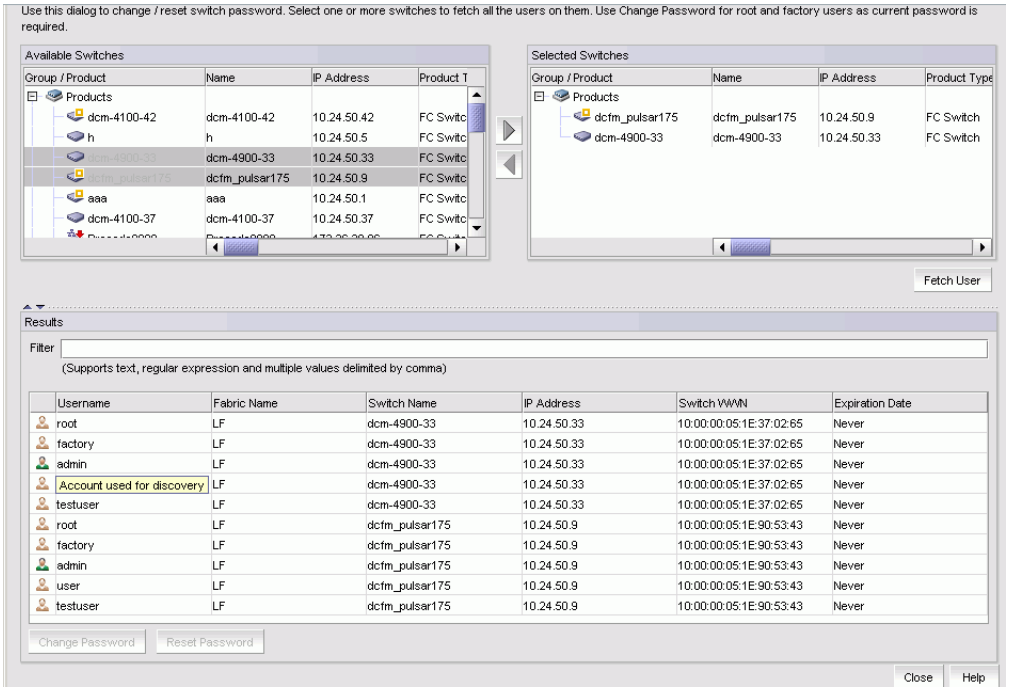

5. Select one or more users for whom you want to change the switch password from the **Results** table and click **Change Password**.

   The **Change Password** dialog box displays (as shown in Figure 200)..

FIGURE 200 Change Password dialog box

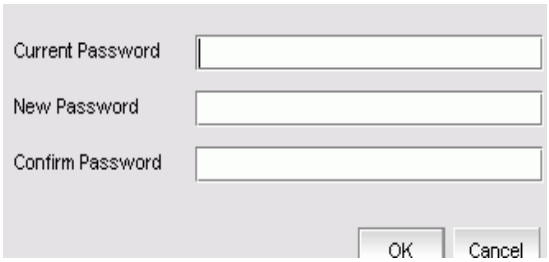

6. Enter the current password in the **Current Password** field.
7. Enter the new password in the **New Password** and **Confirm Password** fields.

   NOTE
   Passwords must be from 8 through 40 characters long and cannot contain a colon (:).

8. Click **OK**.

   The **Change Password Summary** dialog box displays.

   NOTE
   If the password change is successful for the “admin” (account used for discovery), the password is updated in the database.

Resetting the switch password

To reset the switch password, complete the following steps.

1. Select **Configure** > **Switch Password.**

   The **Manage Switch Password** dialog box is displayed (as shown in Figure 198).

2. Select the switches for which you want to reset the password from the **Available Switches** table.

   NOTE
   You cannot change the switch password for unmanaged switches.

   NOTE
   You cannot reset the password for root and factory users.

3. Click the right arrow to move the selected switches to the **Selected Switches** table.

   If you select an unmanaged switch, an error messages displays.

   Click **OK** to close the error message.

4. Click **Fetch User**.

   Th **Filter** field in the **Results** table of the **Manage Switch Password** dialog box supports text, regular expressions, multiple values delimited with commas, and wildcards.

   The users for the selected switches are listed in the **Results** table.

5. Select one or more users for whom you want to reset the switch password from the **Results** table and click **Reset Password**.

   The **Reset Password** dialog box displays (as shown in Figure 201).

   FIGURE 201 Reset Password dialog box

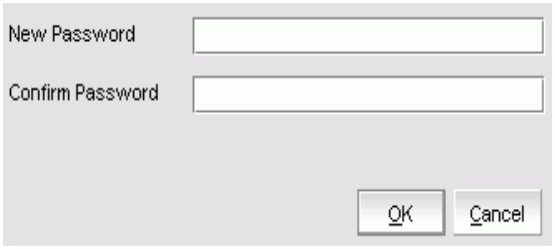

6. Enter the new password in the **New Password** and **Confirm Password** fields.
7. Click **OK**.

   The **Reset Password Summary** dialog box displays.

Frame viewer

NOTE
Frame viewer is only supported on devices running Fabric OS 7.1.0 or later.

Frame viewer enables you to view a list of devices with discarded frames due to c3 timeout, destination unreachable, and not routable. You can also view a summary of discarded frames for each device and clear the discarded frame log on the device.

Beginning with Fabric OS 8.1.0 or later, discarded frame log displays the source and destination entity IDs for VM applicable devices.

Viewing discarded frames from a device

1. Select a Fabric OS device running 7.1.0 or later and select **Monitor > Discarded Frames**.

   The **Discarded Frames** dialog box displays.

   FIGURE 202 Discarded Frames dialog box

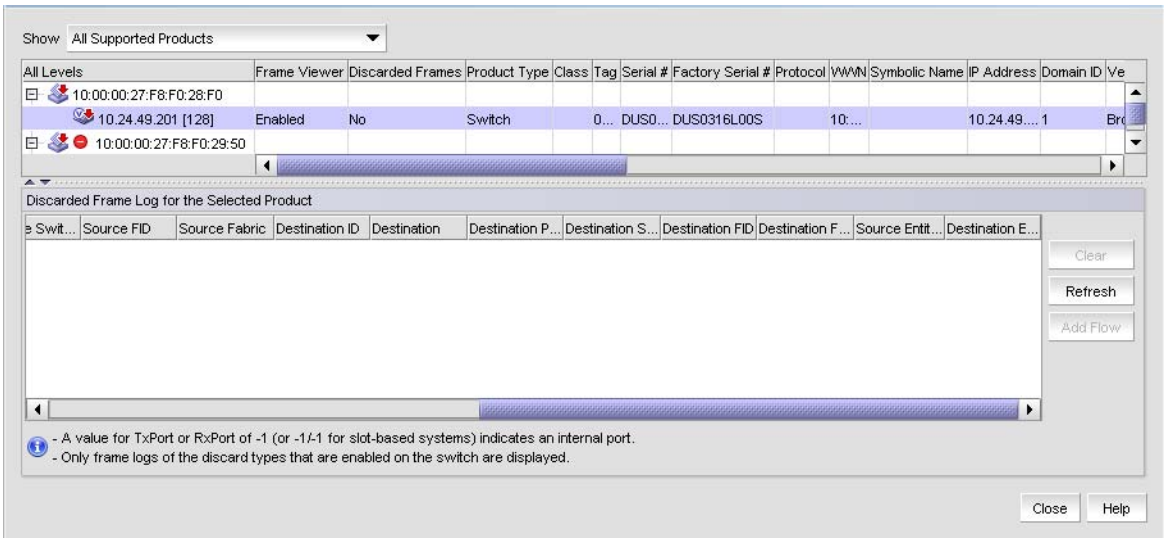

2. Select one of the following options from the **Show** list:

   - Select **Only Supported Products with Dropped Frames** in the Log.

     The top table displays Fabric OS devices running 7.1.0 or later that support frame viewer and have dropped frames.

   - Select **All Supported Products** to view all devices.

     The top table displays all Fabric OS devices running 7.1.0 or later that support frame viewer.

   The top table contains the same data as the Product List (refer to "Product List customization" on page 311) in addition to the following data:

   - **Frame Viewer** — Status of the feature. Valid values include enabled or disabled.
   - **Discarded Frames** — Whether the device frame log contains discarded frame records. Valid values include yes or no.

3. Select a device in the top table to view detailed data about the discarded frames on that device.

- **Discarded Frame History for the Selected Product** table — Summary of the discarded frames for the selected device.
  - **Count** – Number of discarded frames logged in the frame log with the same timestamp, Tx Port, Rx Port, SID, DID, SFID, and DFID. The maximum number of duplicate frames stored for any 1 second timestamp is 20.
  - **Date / Time** – Timestamp when the frames were discarded.
  - **Type** – Type of discard. Valid values include timeout, destination unreachable, and unroutable.
  - **Tx Port** – Egress port where the frames were bound to exit the chassis. When a port display as -1 (or -1/-1 on slot-based systems), this indicates an internal (backplane) port.
  - **Rx Port** – Ingress port where the frames entered the chassis. When a port displays as -1 (or -1/-1 on slot-based systems), this indicates an internal (backplane) port.
  - **Source ID** – Source ID in hex PID format.
  - **Source** – Source name. If the device port is an HBA managed by the Management application, the host name displays.
  - **Source Port** – Source port name.
  - **Source Switch-Port** – Source *Switch_Name – Port_Name*.
  - **Source FID** – Source fabric ID.
  - **Source Fabric** – Source fabric name.
  - **Source Entity ID** – Source entity ID.
  - **Destination ID** – Destination ID in hex PID format.
  - **Destination Entity ID** – Destination entity ID.
  - **Destination** – Destination name. If the device port is an HBA managed by the Management application, the host name displays.
  - **Destination Port** – Destination port name.
  - **Destination Switch-Port** – Destination *Switch_Name – Port_Name*.
  - **Destination FID** – Destination fabric ID.
  - **Destination Fabric** – Destination fabric name.

  The following label displays beneath the **Discarded Frame History for the Selected Product** table: A value for TxPort or RxPort of -1 (-1/-1 for slot-based systems) indicates an internal port.

- **Clear** button — Select a device in the upper table and click to clear the discarded frames from the frame log (refer to "Clearing the discarded frame log" on page 442). All discarded frame records from the frame log on the switch are cleared. The **Discarded Frames** column value in the upper table updates "No".
- **Refresh** button — Click to fetch new data from the frame log on the switch (refer to "Refreshing the discarded frame log" on page 443). Frame log records are not stored in the database.
- **Add Flow** button — Select a discarded frame in the **Discarded Frame History for the Selected Product** table and click to add a flow definition (refer to "Monitoring flows" on page 1021).

  NOTE
  Flow Vision is supported on devices running Fabric OS 7.2 and later.

  NOTE
  Only frame logs of the discard types that are enabled on the switch are displayed.

4. Click **Close**.

Viewing discarded frames from a port

1. Select a port on a Fabric OS device running 7.1.0 or later and select **Monitor > Discarded Frames**.

   The **Discarded Frames** dialog box displays.

2. Review the data for the discarded frames from the selected port.

- **Discarded Frame History for the Selected Product** table — Summary of the discarded frames for the selected port.
  - **Count** – Number of discarded frames logged in the frame log with the same timestamp, Tx Port, Rx Port, SID, DID, SFID, and DFID. The maximum number of duplicate frames stored for any 1 second timestamp is 20.
  - **Date / Time** – Timestamp when the frames were discarded.
  - **Type** – Type of discard. Valid values include **timeout**, **du** (destination unreachable), and **unroute** (not routable).
  - **Tx Port** – Egress port where the frames were bound to exit the chassis. When a port display as -1 (or -1/-1 on slot-based systems), this indicates an internal (backplane) port. For the **du** and **unroute** types, the column displays "--" because there is no Tx port value due to the discard type.
  - **Rx Port** – Ingress port where the frames entered the chassis. When a port displays as -1 (or -1/-1 on slot-based systems), this indicates an internal (backplane) port.
  - **Source ID** – Source ID in hex PID format.
  - **Source** – Source name. If the device port is an HBA managed by the Management application, the host name displays.
  - **Source Port** – Source port name.
  - **Source Switch-Port** – Source *Switch_Name – Port_Name*.
  - **Source FID** – Source fabric ID.
  - **Source Fabric** – Source fabric name.
  - **Source Entity ID** – Source entity ID.
  - **Destination ID** – Destination ID in hex PID format.
  - **Destination Entity ID** – Destination entity ID.
  - **Destination** – Destination name. If the device port is an HBA managed by the Management application, the host name displays.
  - **Destination Port** – Destination port name.
  - **Destination Switch-Port** – Destination *Switch_Name – Port_Name*.
  - **Destination FID** – Destination fabric ID.
  - **Destination Fabric** – Destination fabric name.

  The following label displays beneath the **Discarded Frame History for the Selected Product** table: A value for TxPort or RxPort of -1 (-1/-1 for slot-based systems) indicates an internal port.

- **Clear** button — Click to clear the discarded frames from the frame log ("Clearing the discarded frame log" on page 442). All discarded frame records from the frame log on the switch are cleared. The **Discarded Frames** column value in the upper table updates "No".
- **Refresh** button — Click to fetch new data from the frame log on the switch ("Refreshing the discarded frame log" on page 443). Frame log records are not stored in the database.
- **Add Flow** button — Select a device in the upper table and click to add a flow definition (refer to "Monitoring flows" on page 1021).

  NOTE
  Flow Vision is supported on platforms running Fabric OS 7.2 and later.

3. Click **Close**.

Clearing the discarded frame log

1. Open the **Discarded Frames** dialog box (refer to "Viewing discarded frames from a device" on page 440 or "Viewing discarded frames from a port" on page 441).
2. Select one of the following options:
   - If you are in switch view, select a device in the upper table and click **Clear** to clear the discarded frames from the frame log.
   - If you are in port view, click **Clear** to clear the discarded frames from the frame log.
3. Click **Close**.

Refreshing the discarded frame log

1. Open the **Discarded Frames** dialog box (refer to "Viewing discarded frames from a device" on page 440 or "Viewing discarded frames from a port" on page 441).
2. Select one of the following options:
   - If you are in switch view, select a device in the upper table and click **Refresh** to fetch new data from the switch.
   - If you are in port view, click **Refresh** to fetch new data from the switch.
3. Click **Close**.

Ports

You can enable and disable ports, as well as view port details, properties, type, status, and connectivity.

Viewing port connectivity

The connected switch and switch port information displays for all ports.

To view port connectivity, choose one of the following steps:

- Right-click a Fabric and select **Port Connectivity**.
- Right-click a product icon and select **Port Connectivity**.
- Select a product icon and select **Monitor > Port Connectivity**.

The **Port Connectivity View** dialog box displays (Figure 203).

FIGURE 203 Port Connectivity View dialog box

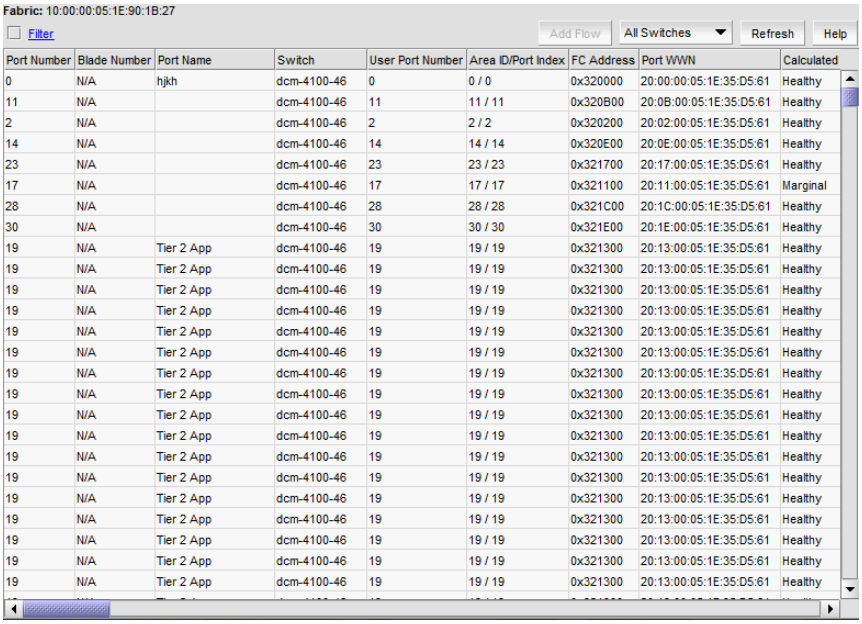

Port Number	Blade Number	Port Name	Switch	User Port Number	Area ID/Port Index	FC Address	Port WWN	Calculated
0	N/A	hjkh	dcm-4100-46	0	0 / 0	0x320000	20:00:00:05:1E:35:D5:61	Healthy
11	N/A		dcm-4100-46	11	11 / 11	0x320B00	20:0B:00:05:1E:35:D5:61	Healthy
2	N/A		dcm-4100-46	2	2 / 2	0x320200	20:02:00:05:1E:35:D5:61	Healthy
14	N/A		dcm-4100-46	14	14 / 14	0x320E00	20:0E:00:05:1E:35:D5:61	Healthy
23	N/A		dcm-4100-46	23	23 / 23	0x321700	20:17:00:05:1E:35:D5:61	Healthy
17	N/A		dcm-4100-46	17	17 / 17	0x321100	20:11:00:05:1E:35:D5:61	Marginal
28	N/A		dcm-4100-46	28	28 / 28	0x321C00	20:1C:00:05:1E:35:D5:61	Healthy
30	N/A		dcm-4100-46	30	30 / 30	0x321E00	20:1E:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy
19	N/A	Tier 2 App	dcm-4100-46	19	19 / 19	0x321300	20:13:00:05:1E:35:D5:61	Healthy

The following details the information located (in default order) on the **Port Connectivity View** dialog box.

- **Fabric / Switch Name** — If launched from a fabric, displays the fabric name. If launched from a switch, displays the fabric name and the switch name.

- **Filter** check box / link — Select to filter results (refer to "Filtering port connectivity" on page 446) in the **Port Connectivity View** dialog box.
- un-named list — Select to view port connectivity for all switches in the fabric or a specific switch. Default selection is All Switches.
- **Refresh** button — Click to refresh the dialog box.
- **Add Flow** button — Select a port and click to add a flow definition (refer to "Monitoring flows" on page 1021).

  NOTE
  Flow Vision is supported on platforms running Fabric OS 7.2 and later.

- Port connectivity table — Displays the ports connected to the selected fabric or device. Loop devices are displayed in multiple rows, one row for each related device port. If no switch or device is connected to the port, then the related fields are empty.
  - **Port Number** — The port's number. To enable or disable a port, refer to "Enabling a port" on page 446 or "Disabling a port" on page 446.
  - **Blade Number** — The number of the blade.
  - **Port Name** — The port's name.
  - **Switch** — The switch name.
  - **User Port Number** — The port number of the user's device.
  - **Area ID /Port Index** — The area ID and the port index of the port.
  - **FC Address** — The Fibre Channel address. Each FC port has both an address identifier and a world wide name (WWN).
  - **Port WWN** — The world wide name of the port.
  - **Calculated Status** — The operational status. There are four possible operation status values: Healthy, Down, Marginal, and Unmonitored.
  - **Status** — The port's status; for example, Enabled, Faulty, Healthy, Unknown, and so on.
  - **Switch Port Type** — The port type; for example, E-Port, F-Port, U-port, and so on.
  - **Speed** — The current port speed, in gigabits per second.
  - **Port Module** — The port's module.
  - **Prohibited** — Whether the allow/prohibit matrix is activated.
  - **Blocked** — Whether the selected port is blocked.
  - **Buffer Limited** — Whether buffers are limited.
  - **Actual Distance** — The actual distance for -end port connectivity.
  - **Buffers Needed/Allocated** — The ratio of buffers needed relative to the number of buffers allocated.
  - **Long Distance** — Whether the connection is considered to be normal or longer distance.
  - **Switch Domain Id —** The switch domain ID.
  - **Device Port/Switch WWN** — The device port and switch world wide name.
  - **Device Port/Switch Name** — The device port and switch name.
  - **Device Port/Switch State** — The device port and switch state; for example, Online.
  - **Device Port/Switch Manufacturer** — The device port and manufacturer of the switch.
  - **Serial #** — The port's serial number.
  - **Device Port / Switch Type Number** — The device port and switch type number.
  - **Switch/Device Model** — The model name and number of the device.
  - **Device Port/Switch Manufacturing Plant** — The device port and switch manufacturing plant.
  - **Device FC Address** — The port FC address of the connected Host or target device.
  - **Device Port Type**— The device port type; for example, U_Port (universal port), FL_Port (Fabric loop port), and so on.
  - **Device Node WWN** — The world wide name of the device node.
  - **Device Symbolic Name** — The symbolic name of the device node.
  - **Physical/Virtual/NPIV** — Whether the port is a physical port, a virtual port, or an NPIV_port.

- **Product Type** — The device type; for example, target or initiator.
- **FC4 Type** — The active FC4 type; for example, SCSI, FCP, and so on.
- **COS** — The class of service (CoS) value, which ranges between zero (low priority) and seven (high priority).
- **Port IP Address** — The port's IP address.
- **Hard Address** — The hard address of the device.
- **Tag** — The tag number of the port.
- **Flag** — Whether a flag is on or off.
- **Parameter** — Device parameters.
- **Unit Type** — The switch unit type.
- **Capability** — The device capability of the connected device port. The value is mapped depending on whether it is a name server (NS) or a FICON device.
- **Vendor** — The hardware vendor's name.
- **Host Name** — The name of the Host.
- **Switch IP** — The switch's IP address.
- **Switch Version** — The switch's version number.
- **Switch Role** — The role of the switch; for example, subordinate.
- **Switch FCS Role** — Whether the Fabric Configuration Server (FCS), which is the primary point of control that manages all the switches within a fabric, is enabled.
- **Switch Status** — The operational status. There are four possible operation status values:
  - Healthy — Operation is normal.
  - Down — The port is down or the route to the remote destination is disabled.
  - Marginal — Operational status is marginal.
  - Unknown — Operational status is unknown.
- **Switch Port Count** — The number of ports on the switch.
- **Switch Secure Mode** — Whether switch secure mode is enabled.
- **Switch FMS mode** — Whether the File Management Solution (FMS) mode is enabled.
- **Switch IDID** — Whether the switch's insistent domain ID (IDID) is enabled. If it is enabled, the IDID is the same ID that is requested during switch reboots, power cycles, CP failovers, firmware downloads, and fabric reconfiguration.
- **Switch Supplier Serial Number** — The serial number of the switch supplier.
- **Switch Has Certificate** — Whether the switch has a certificate (true or false).
- **Routing Policy** — The routing policy configured on the switch.
- **Switch Dynamic Load Sharing** — Whether switch dynamic load sharing is enabled.
- **Switch In Order Delivery** — Whether switch in-order delivery is enabled.
- **Connected Port WWN** — The world wide name of the connected port.
- **Connected Port Name** — The name of the connected port.
- **Connected User Port Number** — The port number of the connected user port.
- **Connected Port Area ID Port Index** — The area ID and the port index of the connected port.
- **Connected Port Speed** — The speed of the connected port.
- **Connected Blade Number** — The number of the connected blade.
- **Connected Port Number** — The number of the connected port.
- **Connected Port Status** — The connection status ; for example, online or offline.
- **Connected Port State** — The connected port's state; for example, online or offline.

Refreshing the port connectivity view

To obtain configuration changes that occurred since the **Port Connectivity View** dialog box opened, click **Refresh**.

Enabling a port

To enable a port from the port connectivity view, right-click the port you want to enable from the **Port Connectivity View** dialog box and select **Disable/Enable Port > Enable**.

Disabling a port

To disable a port from the port connectivity view, right-click the port you want to disable from the **Port Connectivity View** dialog box and select **Disable/Enable Port > Disable**.

Filtering port connectivity

To filter results from the port connectivity view, complete the following steps.

1. Click the **Filter** link from the **Port Connectivity View** dialog box

   The **Filter** dialog box displays (Figure 204).

   FIGURE 204 Filter dialog box

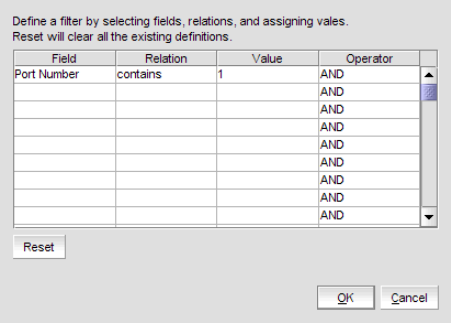

2. Click a blank cell in the **Field** column to select the property from which to filter the results.
3. Click a blank cell in the **Relation** column to select an action operation.

   The following actions are available:

   - ==
   - !=
   - <
   - >
   - <=
   - >=
   - contains
   - matches
4. Define a filter by entering a value that corresponds to the selected property in the **Value** column.

5. Repeat steps 2 through 4 as needed to define more filters.
6. Click **OK**.

   The **Port Connectivity View** dialog box displays. If filtering is already enabled, only those ports that meet the filter requirements display. To enable the filter, select the **Filter** check box.

Resetting the filter

Reset immediately clears all existing definitions. You cannot cancel the reset.

To reset the **Filter** dialog box, complete the following steps.

1. Click the **Filter** link from the **Port Connectivity View** dialog box.

   The **Filter** dialog box displays.

2. Click **Reset**.

   All existing definitions are cleared automatically. You cannot cancel the reset.

Enabling the filter

To enable the filter, select the **Filter** check box.

Disabling the filter

To disable the filter, clear the **Filter** check box.

Viewing port details

To view port details, complete the following steps.

1. Right-click the port for which you want to view more detailed information on the **Port Connectivity View** dialog box and select **Show Details**.

   The **Port Details** dialog box displays(Figure 203).

   FIGURE 205 Port Details dialog box

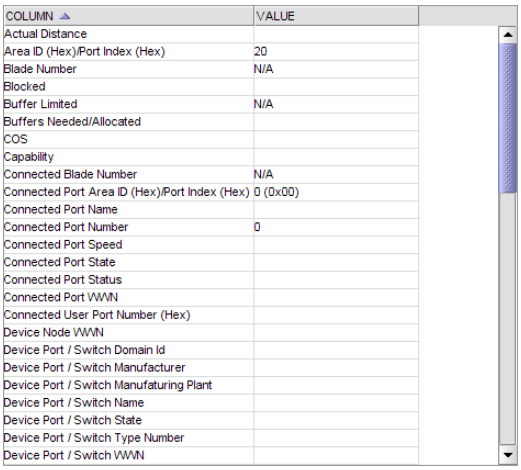

2. Review the port information.

   For the list of fields on the **Port Details** dialog box, refer to "Viewing port properties" on page 1360.

3. Sort the results by clicking on the column header.
4. Rearrange the columns by dragging and dropping the column header.
5. Click the close (X) button to close this dialog box.

Viewing ports

To view ports on the Connectivity Map, right-click a product icon and select **Show Ports**.

NOTE
Show Ports is not applicable when the map display layout is set to **Free Form** (default).

NOTE
This feature is only available for connected products. On bridges and CNT products, only utilized Fibre Channel ports display; IP ports do not display.

Port types

On the Connectivity Map, right-click a switch icon and select **Show Ports**. The port types display showing which ports are connected to which products.

NOTE
Show Ports is not applicable when the map display layout is set to **Free Form** (default).

NOTE
This feature is only available for connected products. On bridges and CNT products, only utilized Fibre Channel ports display. IP ports do not display.

TABLE 44 Port types

Port Type	Description
D	A port in diagnostic mode.
E	An expansion port connecting two Fibre Channel switches.
EX	On a Fibre Channel Router, a connection between a fibre channel router and a fibre channel switch
F	On a Fibre Channel switch, a port that supports an N_Port.
FL	An N_port or F_port that supports arbitrated loop functions associated with arbitrated loop topology.
VE	A virtual E_port configured for an FCIP Tunnel.
VEX	A virtual EX_port configured in an FCIP Tunnel.

Showing connected ports

You can jump from a port to its connected port.

1. Right-click the product whose port connection you want to determine and select **Show Ports**.

   The product's ports display.

2. Right-click a port and select **Connected Port**.

   The focus jumps to the connected port and the connection is highlighted.

Viewing port connection properties

You can view the information about products and ports on both sides of the connection.

1. Right-click the connection between two end devices on the Connectivity Map and select **Properties**.

   OR

   Double-click the connection between two devices on the Connectivity Map.

   The **Connection Properties** dialog box displays.

   NOTE
   If one of the devices is in an unknown state, the Product 1 and Product 2 information displays; however, the **Connections** table information does not display.

2. Review the following information:

   - Product properties for both devices.
   - Connection properties.
   - Selected connection port properties.

   Depending on the device type at either end of the connection, some of the following fields (Table 45) may not be available for all products.

TABLE 45 Port connection properties

Field	Description
Product Properties table	The product information for the two connected switches.
Domain ID	The domain ID of the selected switch and product in xxs(yy) format, where *xx* is the normalized value and *yy* is the actual value.
Fabric Name	The world wide name of the fabric.
IP Address	The IP address of the switch.
Name	The name of the switch.
WWN	The world wide name of the switch.
Connections table	One row for each circuit.
Status	Whether the connection is Active or Missing.
1-Port #	The port number of the first switch.
1-Port Type	The port type of the first switch.
1-WWPN	The world wide port number of the first switch.

TABLE 45 Port connection properties (Continued)

Field	Description
1-MAC Address	The media access control (MAC) address of the first switch.
1-IP Address	The IP address of the first switch.
1-Speed (Gbps)	The speed of the first switch.
1-Trunk	Whether there is a trunk on the first switch.
1-Tunnel ID	The tunnel ID of the first switch.
1-Circuit ID	The circuit ID of the first switch.
2-Port #	The port number of the second switch.
2-Port Type	The port type of the second switch.
2-WWPN	The world wide port number of the second switch.
2-MAC Address	The MAC address of the second switch.
2-IP Address	The IP address of the second switch.
2-Trunk	Whether there is a trunk on the second switch.
2-Speed (Gbps)	The speed of the second switch.
2-Tunnel ID	The tunnel ID of the second switch.
2-Circuit ID	The circuit ID of the second switch.
dB Loss (dB)	The power loss (dB) value between the source and destination ports. Only available when historical performance data collection is enabled. For Fabric OS devices, this field requires firmware version 7.0 or later. Does not display in Professional edition or if SNMP communication fails during discovery or if either switch is not reachable through ISL or IFL.
Selected Connection Properties table	The connected device port information.
Area ID (hex)/Port Index (hex)	The area identifier, in hexadecimal, of the switch-to-product connection.
Blocked	The configuration of the switch (blocked or unblocked).
Buffers Allocated	The number of buffers allocated.
Buffers Desired	The number of buffers required but not allocated.
Circuits	The circuit number of the connected switch.
Compression	Whether compression is enabled or disabled.
Connected Switch	The name of the connected switch.
Cost	The cost of the ISL link.
Distance Actual (km)	The actual distance (in km) for -end port connectivity.
Distance Estimated (km)	The estimated distance (in km) for -end port connectivity.
ED TOV	The Error Detect timeout value, in milliseconds, of the connected switch. This variable is used to flag a potential error condition when an unexpected response is not received.
Encryption	Whether encryption is enabled or disabled.
Fabric	The fabric name.
FC Address	The Fibre Channel (FC) address of the switch.
FC Port #	The FC port number of the switch.
FCIP Capable	Whether the switch is FCIP capable or not.

TABLE 45 Port connection properties (Continued)

Field	Description
Flag (FICON related)	Whether a FICON-related flag is on or off.
Forward Error Correction (FEC)	Whether FEC is enabled or disabled.
GE Port #	The GE port number of the switch.
InBand Management State	Whether inband management is enabled or disabled.
iSCSI Capable	Whether the switch is iSCSI capable or not.
L2 Mode	Whether the switch is in L2 mode or not.
LAG ID	The LAG identifier.
Locked Port Type	The port type of the locked product.
Long Distance Setting	Whether the connection is considered to be normal or longer distance.
MAC Address	The MAC address of the switch.
Manufacturer	The name of the manufacturer.
Manufacturer Plant	The name of the manufacturing plant.
Name	The name of the switch.
NPIV Enabled	Whether the NPIV port is enabled.
Parameter	The parameter of the switch.
Physical/Logical	Whether the port is a physical port or a logical port.
PID Format	The port ID format of the switch.
Port #	The port number.
Port Address	The address of the port.
Port Module	The port's module.
Port NPIV	The number of NPIV ports.
Port Type	The type of port.
Port State	Whether the port is online or offline.
Port Status	Whether the port is enabled or disabled.
Prohibited	Whether the port is prohibited.
Protocol	The network protocol, for example, Fibre Channel.
RA TOV	The resource allocation time out value, in milliseconds, of the connected switch. This variable works with the E D TOV variable to determine switch actions when presented with an error condition.
Sequence #	The sequence number of the switch.
Serial #	The serial number of the switch.
Slot #	The slot number of the switch.
Speed (Gb/s)	The speed in gigabytes per second.
State	The operational status of the port.
Status	The operational status of the switch
Switch	The switch name.
Tag	The tag number of the switch.
Trunking Enabled	Whether trunking is enabled on the switch.

TABLE 45 Port connection properties (Continued)

Field	Description
Tunnel Count	The number of tunnels on the switch.
Tunnel ID	The tunnel ID number of the switch.
User Port #	The user port number of the switch.
VLAN ID	The VLAN identifier.
VPWWN State	Whether the VPWWN state is enabled or disabled.
VPWWN Type	The VPWWN type: Auto or User.
Auto VPWWN	The automatically generated VPWWN.
User VPWWN	The user-defined VPWWN.

3. Click **Close** to close the dialog box.

Determining inactive iSCSI devices

For router-discovered iSCSI devices, you can view all of the inactive iSCSI devices in one list. To do this, use the **Ports Only** view and then sort the devices by FC Address. The devices that have an FC address of all zeros are inactive.

1. Select **View All**, **Levels**, and then **Ports Only** from the main window.
2. Use the scroll bar to view the columns to the right and locate the **FC Address** column in the **Ports Only** list.
3. Click the column label to sort the column in ascending order, if needed.

   iSCSI ports that have an FC Address of all zeros are inactive. All others are active.

Determining port status

You can determine whether a port is online or offline by looking at the Connectivity Map or the Product List.

To determine a port's status on the Connectivity Map, right-click on the product whose ports you want to view and select **Show Ports**.

To determine a port's status through the Product List, scroll down the Product List to the product whose ports you want to see and click the plus icon (+) to expand.

The following table lists the port status icons that display:

	Port added
	Port removed, missing, or segmented

Viewing port optics

Enables you to view port optics for FC, TE, GE, and XGE ports.

To view port optics, complete the following steps.

1. Right-click the switch for which you want to view port optic information on the Connectivity Map and select **Port Optics (SFP)**.

   The **Port Optics (SFP)** dialog box displays (Figure 206).

FIGURE 206 Port Optics dialog box

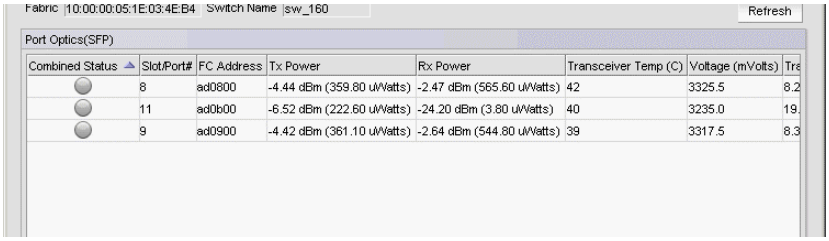

2. Review the port optics information.

- **Combined Status** — Displays the current status of the port.

  NOTE
  Requires a 16 Gbps capable port running Fabric OS 7.0 or later.

  NOTE
  For devices running Fabric OS 7.1 or earlier, the device must have a Fabric Watch license and threshold monitoring configured for the port. For more information, refer to the ***Fabric Watch Administrator's Guide***.

  NOTE
  For devices running Fabric OS 7.2 or later, the device must have a Fabric Vision license, MAPS must be enabled, and threshold monitoring configured for the port. For more information, refer to the "Monitoring and Alerting Policy Suite" on page 1201.

  If the port is online and port monitoring is active, displays the current status of the port based on these five parameters: **Transceiver Temp (C)**, **Rx Power**, **Tx Power**, **Transceiver Current (mAmps)**, and **Voltage (mVolts)**.

  If the port is offline, displays the current status of the port based on these two parameters: **Transceiver Temp (C)** and **Voltage (mVolts)**.

  Status icons:

  - Warning icon — One of the five parameters exceeds the threshold of that parameter. The corresponding parameter field displays with a yellow background.
  - No icon — No parameters exceed the threshold of that parameter.
  - Unknown icon — The port is not a 16 Gbps capable port or the device is running Fabric OS 6.4.X or earlier.
  - Error icon — Unable to retrieve status of the supported port.

- **Slot/Port #** — The slot and port number of the selected fabric. The port number includes the type of port (FC, TE, GE, or XGE).
- **FC Address** — The Fibre Channel address of the port.
- **TX Power** — The power transmitted to the SFP in dBm and uWatts.

NOTE
The uWatts display requires devices with Fabric OS 6.1.0 and later. Devices running Fabric OS 6.0.0 and earlier only display dBm.

- **RX Power** — The power received from the port in dBm and uWatts.

  NOTE
  The uWatts display requires devices with Fabric OS 6.1.0 and later. Devices running Fabric OS 6.0.0 and earlier only display dBm.

- **Transceiver Temp (C)** — The temperature of the SFP transceiver.
- **Voltage (mVolts)** — The voltage across the port in mVolts.
- **Transceiver Current (mAmps)** — The laser bias current value in mAmps.
- **Powered on Years (Hours)** — The powered on time in years and hours for 16 Gbps capable ports. Empty for unsupported ports.

  NOTE
  Requires a 16 Gbps capable port running Fabric OS 7.0 or later.

- **FC Speed (GB/s)** (Fabric OS 7.0 or later) — The FC port speed; for example, 4 Gbps.
- **FC Speed (MB/s)** (Fabric OS 6.4 or earlier) — The FC port speed; for example, 400 Mbps.
- **Distance** — The length of the fiber optic cable.
- **Vendor** — The vendor of the SFP.
- **Vendor OUI** — The vendor's organizational unique identifier (OUI).
- **Vendor PN** — The part number of the SFP.
- **Vendor Rev** — The revision number of the SFP.
- **Serial #** — The serial number of the SFP.
- **Data Code** — The data code.
- **Media Form Factor** — The type of media for the transceiver; for example, single mode.
- **Connector** — The type of port connector.
- **Wave Length** — The wave length.
- **Encoding** — Displays how the fiber optic cable is encoded.

3. To view port properties, select a row and click **Properties**.
4. Sort the results by clicking on the column header.
5. Rearrange the columns by dragging and dropping the column header.
6. Click **Close** to close the **Port Optics (SFP)** dialog box.

Refreshing port optics

To refresh port optics, click **Refresh**.

The Management application retrieves updated port optic information.

Administrative Domain-enabled fabric support

The Management application provides limited support for AD-enabled fabrics.

An ***Administrative Domain*** (Admin Domain or AD) is a logical grouping of fabric elements that defines which switches, ports, and devices you can view and modify. An Admin Domain is a filtered administrative view of the fabric.

> NOTE
> If you do not implement Admin Domains, the feature has no impact on users and you can ignore this section.

For more information about Admin Domains, refer to the ***Fabric OS Administrator's Guide***.

AD-enabled fabric discovery

The Management application enables you to discover AD-enabled fabrics using SAN fabric d*i*scovery. To discover AD-enabled fabrics, you must be a ***physical fabric administrator***. A physical fabric administrator is a user with admin permissions and access to all Admin Domains (AD0 through AD255). Only a physical fabric administrator can perform AD-enabled fabric discovery and management.

Discovery collects asset information using the AD255 (physical fabric) context. However, the Management application does not collect AD membership information.

Instructions for discovering AD-enabled fabrics are detailed in "Discovery" on page 33.

Management application behavior for AD-enabled fabrics

Note the following considerations and interactions that apply for AD-enabled fabrics.

- Does not display provisioned AD's in AD-enabled environments in the Product List. Provisioned AD's are available through Web Tools.
- Does not filter by AD membership in the Topology Map (for example, the Topology Map always displays the physical fabric connectivity and membership).
- Does not support zone Management (including LSAN management).
- Performs firmware management in a physical fabric context.
- Performs configuration upload and download in physical fabric context (per Virtual Fabrics capability is available in Virtual Fabrics environments).
- Performs basic user actions (for example, enabling or disabling a port or switch) in a physical fabric context.
- Supports fault and event management. Note that since AD's are not visualized in the Topology Map and Product List, the Master Log provides an unfiltered view of events for the entire AD-enabled fabric.
- Web Tools launch (with Single sign on support where applicable) defaults to Default AD (AD0). For AD life cycle management, you must switch the context to physical fabric (AD 255) using Web Tools.
- Does not support features dependent on the AD context and membership (for example, Troubleshooting and Diagnostics) for AD-enabled fabrics.
- If you try to enable Virtual Fabrics on an AD-enabled switch, that operation fails with the following message: "Failed to enable Virtual Fabric feature for Chassis (Remove All ADs before attempting to enable VF)."
- If you discover Fabric OS 8.0.1 switch with AD-enabled, the Management Application will not display any warning messages as it displays in the CLI.
- Performs performance management (including Advance Performance Monitoring and Top Talkers) data collection and reports in a physical fabric context.

- If AD is enabled any switch in a fabric, you cannot clear counters (performance management) on any switch in that fabric.

Management application support for AD-enabled fabrics

Table 46 details feature support for AD-enabled fabrics in the Management application.

TABLE 46 Feature support for AD-enabled fabrics

Feature	AD context				
	AD0	AD255	Not supported	All AD	User interface impact
Allow/Prohibit Matrix			X		Filters AD-enabled fabric from the Fabrics list.
Cascaded FICON/FICON Merge			X		Filters AD-enabled fabric from the Fabrics list.
Configuration Management		X			None.
Configuration Management > CEE FCoE Swap Blades			X		Filters AD-enabled fabric from the product tree.
Encryption			X		Filters AD-enabled fabric from the dialog box.
Fabric Binding			X		Filters AD-enabled fabrics from the Fabrics table. Displays all switches (including switches in an AD-enabled fabric) in the Available Switches table.
Fabric discovery (except zoning)		X			None.
Fault Management				Displays all events from the switch in the Master Log regardless of AD membership.	None.
FCIP Tunnels Configuration			X		Filters switches from an AD-enabled fabric from the dialog box.
Firmware Management		X			None.
High Integrity Fabric (HIF)			X		Filters AD-enabled fabric from the Fabrics list.
Logical Switches			X		None.
Names configuration		X			None.
Performance Management		X			None.
Performance Management > Configure Thresholds End-to-End Monitors Clear Counters			X		Filters AD-enabled fabric from the Fabrics list.
Port Auto Disable			X		Filters AD-enabled fabric from the dialog box.
Port Connectivity			X		Disables menu for a switch in an AD-enabled fabric.

TABLE 46 Feature support for AD-enabled fabrics (Continued)

Feature	AD context				
	ADO	AD255	Not supported	All AD	User interface impact
Port Fencing			X		Filters AD-enabled fabrics from the product tree.
Port Optics		X			None.
Product Administration (Switch Enable/Disable, Port Enable/Disable)		X			None.
Routing Configuration			X		Filters switches from an AD-enabled fabric from the dialog box.
SMI Agent			X		None.
SNMP Informs		X			None.
Syslog Registration		X			None.
Technical Support Save		X			None.
Technical Support Save > Auto Trace dump			X		Filters AD-enabled fabric from the Fabrics list.
Trap Registration		X			None.
Troubleshooting and Diagnostics			X		Filters AD-enabled fabrics from the Fabrics list.
Web Tools Launch	X				Launches Web Tools in ADO.
Zone DB collection	X				None.
Zoning dialog box			X		Filters AD-enabled fabric from the Fabrics list.

Port Auto Disable

NOTE
Port Auto Disable requires devices running Fabric OS 7.0 or later.

Port Auto Disable (PAD) allows you to enable and disable Port Auto Disable on individual FC_ports or on all ports on a selected device, as well as unblock currently blocked ports. Enabling port auto disable on a port or device configures ports to become blocked when any of the following five events occur:

- Loss of Sync
- Loss of Signal
- OLS (Offline Primitive Sequence)
- NOS (Not Operational Primitive Sequence)
- LIP (Loop Initialization Primitive Sequence)

For Fabric OS devices running 7.0 or later, you can configure ports to become blocked when a specific event is triggered (one or more of the events listed above).

You can also suspend or resume Port Auto Disable on a switch.

Viewing Port Auto Disable status

NOTE
Port Auto Disable requires devices running Fabric OS 7.0 or later.

To view the PAD status, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

   The **Port Auto Disable** dialog box displays.

   FIGURE 207 Port Auto Disable dialog box

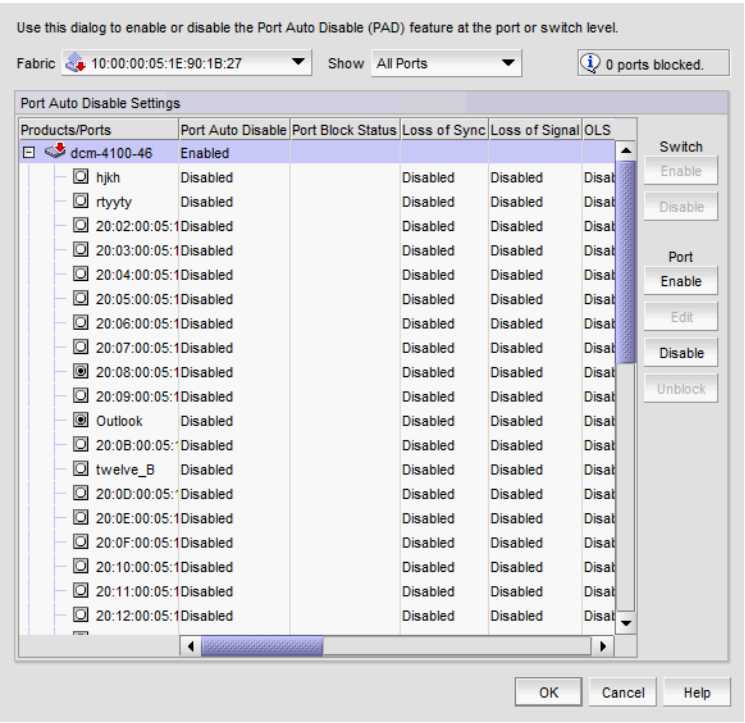

2. Select a fabric from the **Fabric** list.

   An information message displays the number of block ports for the fabric, if any.

3. Select one of the following from the **Show** list to determine what ports to display:

   - **All Ports** (default)
   - **Disabled PAD Ports**
   - **Enabled PAD Ports**
   - **Blocked Ports**

4. Review the port information:

   - **Products/Ports tree —** Displays devices and associated ports. Also, displays a Warning icon for blocked FC ports (displayed with the port icon). Click the plus sign (**+**) symbol to expand the view to display the ports.
   - **Port Auto Disable —** Displays whether Port Auto Disable is currently enabled or disabled on the device or port.
   - **Port Block Status —** Displays whether the port is currently blocked.

- **Loss of Sync —** Whether the Loss of Sync event is enabled or disabled.
- **Loss of Signal —** Whether the Loss of Signal event is enabled or disabled.
- **OLS —** Whether the Offline Primitive Sequence event is enabled or disabled.
- **NOS —** Whether the Not Operational Primitive Sequence event is enabled or disabled.
- **LIP —** Whether the Loop Initialization Primitive Sequence event is enabled or disabled.
- **Port Type —** Displays the port type.
- **Port # —** Displays the port number.
- **Port WWN —** Displays the port world wide name.
- **Port Name —** Displays the port name.
- **User Port # —** Displays the user port number.
- **PID —** Displays the port identifier.
- **Connected Port # —** Displays the connected port number.
- **Connected Port WWN —** Displays the connected port world wide name.
- **Connected Port Name —** Displays the connected port name.

5. Click **OK** on the **Port Auto Disable** dialog box.

Configuring Port Auto Disable event triggers

NOTE
To configure the specific events that trigger the Port Auto Disable, the device must be running Fabric OS 7.0 or later.

You can configure a port to become blocked when one or more of the following events occur on the configured port:

- Loss of Sync
- Loss of Signal
- OLS (Offline Primitive Sequence)
- NOS (Not Operational Primitive Sequence)
- LIP (Loop Initialization Primitive Sequence)

To configure the PAD event triggers, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

   The **Port Auto Disable** dialog box displays.

2. Select the fabric on which you want to configure the PAD event triggers from the **Fabric** list.
3. Select **All Ports** from the **Show** list to filter the port list:
4. Select one or more ports or devices on which you want to configure the PAD event triggers.
5. Click **Edit**.

   The **Edit Configuration** dialog box displays.

6. Select one or more of the following event triggers:

   - **Loss of Sync**

- **Loss of Signal**
- **OLS** (Offline Primitive Sequence)
- **NOS** (Not Operational Primitive Sequence)
- **LIP** (Loop Initialization Primitive Sequence)

7. Click **OK** on the **Edit Configuration** dialog box.
8. Click **OK** on the **Port Auto Disable** dialog box.

Enabling Port Auto Disable on individual ports

NOTE
Port Auto Disable requires devices running Fabric OS 7.0 or later.

To enable PAD on individual ports, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

   The **Port Auto Disable** dialog box displays.

2. Select the fabric on which you want to configure PAD from the **Fabric** list.
3. Choose one of the following options from the **Show** list to filter the port list:
   - **All Ports** (default) — Displays all ports in the fabric.
   - **Disabled PAD** — Displays only ports where PAD is disabled.
4. Select one or more ports on which you want to enable PAD.

   Press CTRL and click to select multiple ports.

5. To configure specific events to trigger PAD (device must be running Fabric OS 7.0 or later), refer to "Configuring Port Auto Disable event triggers" on page 459.
6. Click **OK** on the **Port Auto Disable** dialog box.

Enabling Port Auto Disable on all ports on a device

NOTE
Port Auto Disable requires devices running Fabric OS 7.0 or later.

To enable PAD on all ports on a device, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

   The **Port Auto Disable** dialog box displays.

2. Select the fabric on which you want to configure PAD from the **Fabric** list.
3. Select **All Ports** from the **Show** list.
4. Select the device on which you want to enable PAD on all ports.

   Press CTRL and click to select multiple devices.

5. To configure specific events to trigger PAD (device must be running Fabric OS 7.0 or later), refer to "Configuring Port Auto Disable event triggers" on page 459.
6. Click **Enable** (under **Port**).

   PAD is enabled on all ports on the selected device.
7. Click **OK** on the **Port Auto Disable** dialog box.

Disabling Port Auto Disable on individual ports

NOTE
Port Auto Disable requires devices running Fabric OS 7.0 or later.

To disable port auto disable on individual ports, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

   The **Port Auto Disable** dialog box displays.
2. Select the fabric on which you want to configure PAD from the **Fabric** list.
3. Choose one of the following options from the **Show** list to filter the port list:
   - **All Ports** (default) — Displays all ports in the fabric.
   - **Enabled PAD** — Displays only ports where PAD is enabled.
4. Select the ports on which you want to disable PAD.

   Press CTRL and click to select multiple ports.
5. Click **Disable** (under **Port**).

   PAD is disabled on the selected ports.
6. Click **OK** on the **Port Auto Disable** dialog box.

Disabling Port Auto Disable on all ports on a device

NOTE
Port Auto Disable requires devices running Fabric OS 7.0 or later.

To disable port auto disable on all ports on a device, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

   The **Port Auto Disable** dialog box displays.
2. Select the fabric on which you want to configure PAD from the **Fabric** list.
3. Select **All Ports** from the **Show** list.
4. Select the device on which you want to disable PAD on all ports.

   Press CTRL and click to select multiple devices.
5. Click **Disable** (under **Port**).

PAD is disabled on all ports of the selected device.

6. Click **OK** on the **Port Auto Disable** dialog box.

Stopping Port Auto Disable on a device

NOTE
Port Auto Disable requires devices running Fabric OS 7.2 or later.

You can disable PAD at the device level. This allows you stop PAD for the device regardless of the individual port setting.

To stop PAD on a device, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

   The **Port Auto Disable** dialog box displays.

2. Select the fabric on which you want to configure PAD from the **Fabric** list.
3. Select **All Ports** from the **Show** list, if necessary.
4. Select the device on which you want to stop PAD.
5. Click **Disable** (under **Switch**).

   PAD stops on all ports for the selected device.

6. Click **OK** on the **Port Auto Disable** dialog box.

Resuming Port Auto Disable on a device

NOTE
Port Auto Disable requires devices running Fabric OS 7.2 or later.

You can enable PAD at the device level. This allows you resume PAD for the device regardless of the individual port setting.

To resume PAD on a device, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

   The **Port Auto Disable** dialog box displays.

2. Select the fabric on which you want to configure PAD from the **Fabric** list.
3. Select **All Ports** from the **Show** list, if necessary.
4. Select the device on which you want to resume PAD.

   Press CTRL and click to select multiple devices.

5. Click **Enable** (under **Switch**).

   PAD resumes on the selected device.

6. Click **OK** on the **Port Auto Disable** dialog box.

Unblocking ports

NOTE
Port Auto Disable requires devices running Fabric OS 7.0 or later.

To unblock ports, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

   The **Port Auto Disable** dialog box displays.

2. Select the fabric on which you want to unblock ports from the **Fabric** list.
3. Select **Blocked Ports** from the **Show** list.
4. Select the device on which you want to unblock ports.
5. Click **Unblock** (under **Port**).
6. Click **OK** on the **Port Auto Disable** dialog box.

Host Port Mapping

Host port mapping overview

Host Bus Adapters (HBAs) and Hosts discovered through a fabric can be easily identified in the topology by their product icons. For a list of products and their icons, refer to "Icon legend" on page 303. Once identified in the topology, you can create Hosts and assign the HBAs to them and import an externally created Host port mapping file (.CSV) to the Management application.

NOTE
The Management application enables you to map HBAs from multiple fabrics (previous versions limited HBA mapping to one fabric).

NOTE
Converged Network Adapters (CNAs) and HBAs support Host port mapping.

NOTE
After Auto Enclosure, if you discover a host that has more or as many adapters as the Auto Enclosure, Host discovery succeeds.

The Management application also enables you to discover Hosts directly using Host discovery (for step-by-step instructions, refer to "Host discovery" on page 51). If you discover a Host directly, when you open the **Host Port Mapping** dialog box, the Management application automatically groups all HBAs under the discovered Host.

If you create a new Host and associate HBAs to it, and then try to discover a Host with the same HBAs using Host discovery, the HBAs discovered using Host discovery must match the HBAs associated to the Host exactly; otherwise, Host discovery will fail.

Creating a new Host

To create a new Host, complete the following steps.

1. Right-click an HBA icon in the Fabric topology and select **Host Port Mapping**.

   The **Host Port Mapping** dialog box displays.

FIGURE 208 Host Port Mapping dialog box

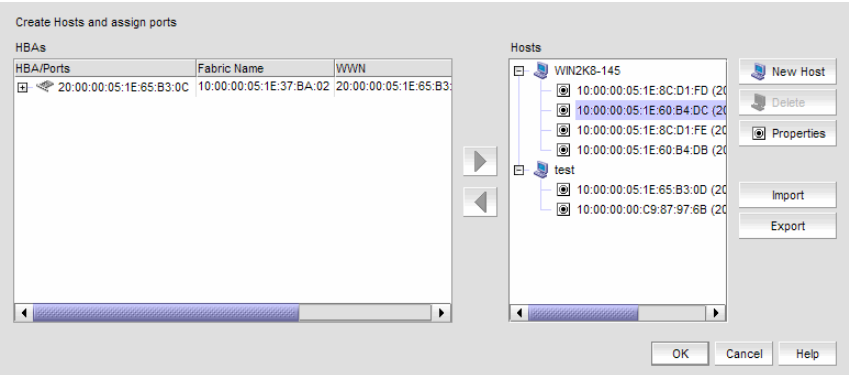

The **Host Port Mapping** dialog box includes the following details:

- **HBAs** table — All unassigned HBAs. Lists the following information for all available HBAs. You can sort the table by clicking once on any of the column titles.
  - **HBA** — The World Wide Name of the node.
  - **Fabric Name** — The fabric name.
  - **WWN** — The World Wide Name for the fabric.
  - **Connected Switch** — The label of the connected device.
  - **Connected Port #** — The port number of the connected device.
- **Hosts** list — All created Hosts.

2. Click **New Host**.

   A new Host displays in the **Hosts** list in edit mode.

3. Double-click the new Host name to make it editable, type a name for the new Host, and press **Enter**.

   The name of the new Host appears in the **Hosts** list in alphabetical order. To assign HBAs to this Host, refer to "Associating an HBA with a Host" on page 467.

4. Click **OK** to save your changes and close the **Host Port Mapping** dialog box.

Renaming an HBA Host

To rename a Host, complete the following steps.

1. Right-click an HBA icon in the Fabric topology and select **Host Port Mapping**.

   The **Host Port Mapping** dialog box displays.

2. Click the Host you want to rename in the **Hosts** table, wait a moment, and then click it again.

   The Host displays in edit mode.

3. Type a new name for the Host.

   The name of the Host appears in the **Hosts** table in alphabetical order with the new name. To assign HBAs to this Host, refer to "Associating an HBA with a Host" on page 467.

4. Click **OK** to save your changes and close the **Host Port Mapping** dialog box.

Deleting an HBA Host

To delete a Host, complete the following steps.

1. Right-click an HBA icon in the Fabric topology and select **Host Port Mapping**.

   The **Host Port Mapping** dialog box displays.

2. Select the Host you want to delete in the **Hosts** table.
3. Click **Delete**.

   The selected Host is deleted. Any HBAs associated with the Host are automatically moved from the **Host** table to the **HBAs** table.

4. Click **OK** to save your changes and close the **Host Port Mapping** dialog box.

Viewing Host properties

To view Host properties, complete the following steps.

1. Right-click an HBA icon in the Fabric topology and select **Host Port Mapping**.

   The **Host Port Mapping** dialog box displays.

2. Select the HBA Host port you want to view in the **Hosts** table.
3. Click **Properties**.

   The **Properties** dialog box for the selected port displays.

4. Click **OK** to close the **Properties** dialog box.
5. Click **OK** to close the **Host Port Mapping** dialog box.

Associating an HBA with a Host

ATTENTION
Discovered information overwrites your user settings.

To associate an HBA with a Host, complete the following steps.

1. Right-click an HBA icon in the Fabric topology and select **Host Port Mapping**.

   The **Host Port Mapping** dialog box displays.

2. Select the Host to which you want to assign HBAs in the **Hosts** table or click **New Host** to create a new Host.
3. Select the HBA from the **HBAs** table on the left and click the right arrow.

   NOTE
   If the HBA is part of more than one fabric, port nodes associated with the other fabrics will automatically be moved to the Host.

   The HBA displays in the **Hosts** table. The HBA is now associated with the selected Host.

4. Click **OK** to save your changes and close the **Host Port Mapping** dialog box.

   If the HBA is part of more than one fabric, a message displays: The selected *Host_Name*/*Host_WWN* is part of more than one fabric. The port nodes associated with the other fabrics will automatically be moved to the Host. Click **OK** to close the message.

   On the Connectivity Map, the HBA displays in the Host.

Importing HBA-to-Host mapping

The **Host Port Mapping** dialog box enables you to import externally created HBA ports-to-Host mapping information into the application. The imported file must be in CSV format. The first row must contain the headers (wwn, name) for the file.

Example

```
wwn,name
20:00:00:00:C9:69:D5:27, s1
20:00:00:05:1E:0A:35:0E, s2
```

To import Host port mapping, complete the following steps.

1. Right-click an HBA icon in the Fabric topology and select **Host Port Mapping**.

   The **Host Port Mapping** dialog box displays.

2. Click **Import**.

   The **Import** dialog box displays.

3. Browse to the file (CSV format only) you want to import.

4. Click **Open** on the **Import** dialog box.

   The file imports, reads, and applies all changes line-by-line and performs the following:

   - Checks for correct file structure and well-formed WWNs, and counts number of errors.

     If more than 5 errors occur, import fails and a 'maximum error count exceeded' message displays. Edit the Host port mapping file and try again.

   - Checks for duplicate HBAs.

     If duplicates exist, a message displays with the duplicate mappings detailed. Click **Yes** to continue. Click **No** to edit the Host port mapping file and try again.

   - Checks for existing mappings in the current map.

     If a mapping already exists, a message displays with the current mapping information. Click **Yes** to overwrite the current mapping. Click **Yes to All** to overwrite all mapping conflicts. Click **No** to leave the current mapping. Click **No to All** to leave all current mappings when conflict occurs. Click **Cancel** to cancel the import.

   When the import is complete a result summary displays with the information listed in Table 47.

TABLE 47 Import Results

Value	Definition
Total Valid Input Records	Number of lines identified in the CSV file without any errors (excluding the Header).
Unique HBA WWNs Recognized	Number of unique HBAs identified in the CSV file.

TABLE 47 Import Results

Value	Definition
Hosts Created or Identified	Number of Hosts identified in the CSV file already discovered, and which are either online or offline but not deleted.
Conflicting HBA Mappings	Number of occurrences where you were asked to decide whether to override previously discovered information. If you select Yes to All, or No to All, each occurrence where conflict resolution occurs automatically is counted as one conflict.
Overwritten HBA Mappings	Number of times a previously discovered mapping is overwritten during the import process.
Importing Errors	Number of errors encountered during the import.
Details table	Tabulates the error information with respect to the line number where it occurred. **Line #** column displays the line number where the erroneous information is located. **Contents** column displays the erroneous information.

5. Click **OK** to close the **Import Results** dialog box.
6. Click **OK** to close the **Host Port Mapping** dialog box.

Removing an HBA from a Host

To remove an HBA from a Host, complete the following steps.

1. Right-click an HBA icon in the Fabric topology and select **Host Port Mapping**.

   The **Host Port Mapping** dialog box displays.

2. Select the HBA from the **Hosts** table on the right and click the left arrow.

   The HBA you selected is removed from the **Hosts** table and the HBA is no longer associated with the Host.

   NOTE
   If the HBA is part of more than one fabric, port nodes associated with the other fabrics will automatically be moved to the Host.

3. Click **OK** to save your changes and close the **Host Port Mapping** dialog box.

   If the HBA is part of more than one fabric, a message displays: The selected *Host_Name*/*Host_WWN* is part of more than one fabric. The port nodes associated with the other fabrics will automatically be removed from the Host. Click **OK** to close the message.

   On the Connectivity Map, the HBA displays on its own.

Exporting Host port mapping

The **Host Port Mapping** dialog box enables you to export a Host port. The export file uses the CSV format. The first row contains the headers (HBA/Ports WWN, Host Name) and the switch to which the port is connected.

Example

```
HBA World Wide Name, Host Name
5005076717011E7D, Server1
50050767170A5AAF, Server1
```

To export a Host port, complete the following steps.

1. Open the **Host Port Mapping** dialog box by performing one of the following actions:
   - Select an HBA port icon in the Fabric topology , then select **Discover > Host Port Mapping**.
   - Right-click any HBA port icon in the Fabric topology and select **Host Port Mapping**.
   - Right-click any HBA port in the Device Tree on the SAN tab and select **Host Port Mapping**.

   The **Host Port Mapping** dialog box displays.
2. Select the Host port you want to export from the **HBA/Ports** list.

   To configure Host port mapping, refer to "Creating a new Host" on page 465 and "Associating an HBA with a Host" on page 467.
3. Click **Export**.

   The **Export** dialog box displays.
4. Browse to the location where you want to save the export file.

   Depending on your operating system, the default export location are as follows:
   - Desktop\My documents (Windows)
   - \root (Linux)
5. Enter a name for the files and click **Save**.
6. Click **OK** to close the **Host Port Mapping** dialog box.

Storage Port Mapping

Storage port mapping overview

The Management application enables you to see multiple ports on your storage devices in a SAN. It also displays the relationship between multiple ports and represents them as attached to a storage array (device) in the **Device Tree**, **Topology**, and **Fabric** views. Occasionally, there are cases where the Management application cannot see the relationship between ports attached to the same storage device. Therefore, the Management application allows you to manually associate the connections that the system is unable to make.

The Management application allows you to create and assign properties to a Storage Device during the mapping process using the **Storage Port Mapping** dialog box. Once a Storage Device has multiple ports assigned to it you cannot change the device type.

> NOTE
> When you open the **Storage Port Mapping** dialog box, Discovery is automatically turned off. When you close the **Storage Port Mapping** dialog box, Discovery automatically restarts.

During Discovery, if a previously mapped Storage Port is found to have a relationship with a port just discovered, the Management application automatically reassigns the Storage Port to the proper mapping. The two Ports are grouped together. This grouping is visually represented as a Storage Device. This Storage Device contains Node information from the discovered port and populates default information where available.

The Management application allows you to change the Device Type of a discovered device. Isolated Storage Ports are represented as Storage Devices. Using the Storage Port Mapping dialog you cannot change the device type to an HBA, JBOD, and so on. However, once a device has been identified as type Storage with ports assigned, you can no longer change its type.

Creating a storage array

To create a storage array, complete the following steps.

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays with the following information.

- **Storage Ports** table — Lists the following information for all available storage ports. You can sort the table by clicking once on any of the column titles.
  - **Fabric Name** — The fabric name.
  - **WWN** — The world wide name for the fabric.
  - **Connected Device** — The label of the connected device.
  - **Connected Port #** — The port number of the connected device.
- **Storage Array** list — Lists the following information for the Storage Array.
  - **Storage Array Name** — The name for the new Storage Array.
  - **Port Icon** — The icon for the port.
  - **Port Numbe**r — The number of the port.

2. Click **New Storage**.

   A new storage array displays in the **Storage Array** list in edit mode.

3. Rename the new storage array and press **Enter**.
4. Add storage ports to the new storage array.

   NOTE
   You must add at least one storage ports to the new storage array to save the new array in the system.

   For step-by-step instructions about adding ports to an array, refer to "Adding storage ports to a storage array" on page 472.

5. Click **OK** to save your work and close the **Storage Port Mapping** dialog box.

Adding storage ports to a storage array

To add storage ports to a storage array, complete the following steps.

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays.

2. Select a storage port from the **Storage Ports** table.

   To select more than one port, hold down the **CTRL** key while selecting multiple storage ports.

3. Select the storage array to which you want to assign the storage port in the **Storage Array** list.

   NOTE
   If the storage device is part of more than one fabric, port nodes associated with the other fabrics will automatically be moved to the storage array.

4. Click the right arrow.

   The storage port is added to the Storage Array.

5. Click **OK** to save your work and close the **Storage Port Mapping** dialog box.

   If the storage device is part of more than one fabric, a message displays: The selected *Storage_Name*/*Storage_WWN* is part of more than one fabric. The port nodes associated with the other fabrics will automatically be moved to the storage array. Click **OK** to close the message.

Unassigning a storage port from a storage array

To unassign a storage port from a storage array, complete the following steps.

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays.

2. Select the storage port you want to unassign from the **Storage Array** list.

   NOTE
   If the storage device is part of more than one fabric, port nodes associated with the other fabrics will automatically be removed from the storage array.

3. Click the left arrow button.

   The selected storage port is removed from the **Storage Array** list and added to the **Storage Ports** table.

4. Click **OK** to save your work and close the **Storage Port Mapping** dialog box.

   If the storage device is part of more than one fabric, a message displays: The selected *Storage_Name*/*Storage_WWN* is part of more than one fabric. The port nodes associated with the other fabrics will automatically be removed from the storage array. Click **OK** to close the message.

Reassigning mapped storage ports

To reassign a storage port, complete the following steps.

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays.

2. Select the storage port you want to unassign from the **Storage Array** list.
3. Click the left arrow button.

   The selected storage port is removed from the **Storage Array** list and added to the **Storage Ports** table.

4. Make sure the storage port you want to reassign is still selected.
5. Select the storage array to which you want to reassign the storage port in the **Storage Array** list.
6. Click the right arrow button.

   The storage port moves from the **Storage Ports** table to the selected storage array.

7. Click **OK** to save your work and close the **Storage Port Mapping** dialog box.

Editing storage array properties

To edit storage array properties, complete the following steps.

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays.

2. Select the storage array in the **Storage Array** list and click **Properties**.

The **Properties** dialog box appears.

3. Edit the property fields, as needed.

   Depending on which tab you select (Properties tab, Storage tab, Port tab), different fields will be available for editing. Editable fields have a green triangle in the lower right corner of the field.

4. Click **OK** on the **Properties** dialog box to save the storage array properties.
5. Click **OK** to save your work and close the **Storage Port Mapping** dialog box.

Deleting a storage array

To delete a storage array, complete the following steps.

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays.

2. Select a storage array in the **Storage Array** list.
3. Click **Delete**.

   The selected storage array and all storage ports assigned to the array are removed from **Storage Array** list. All Storage Ports assigned to the device are moved to the **Storage Ports** table.

4. Click **OK** to save your work and close the **Storage Port Mapping** dialog box.

Viewing storage port properties

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays.

2. Select a storage port from the **Storage Array** list.
3. Click **Properties**.

   The **Properties** dialog box displays.

4. Review the properties.
5. Click **OK** on the **Properties** dialog box.
6. Click **OK** on the **Storage Port Mapping** dialog box.

Viewing storage array properties

To view storage array properties, complete the following steps.

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays.

2. Select a storage array from the **Storage Array** list.
3. Click **Properties**.

The **Properties** dialog box displays.

4. Review the properties.
5. Click **OK** on the **Properties** dialog box.
6. Click **OK** on the **Storage Port Mapping** dialog box.

Importing storage port mapping

The **Storage Port Mapping** dialog box enables you to import externally created storage port mapping information into the application. The imported file must be in CSV format. The first row must contain the headers (wwn, name) for the file, which is ignored during the import.

Example

```
wwn,name
20:00:00:04:CF:BD:89:6E,name1
20:00:00:04:CF:BD:6F:32,name2
20:00:00:04:CF:BD:70:2F,name1
20:00:00:04:CF:BD:6F:52,name2
```

To import storage port mapping, complete the following steps.

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays.

2. Click **Import**.

   The **Import** dialog box displays.

3. Browse to the file (CSV format only) you want to import.
4. Click **Open** on the **Import** dialog box.

   The file imports, reads, and applies all changes line-by-line and performs the following:

   - Checks for correct file structure (first entry must be the storage node name (WWN) and second entry must be the storage array name), well formed WWNs, and counts number of errors

     If more than 5 errors occur, import automatically cancels. Edit the storage port mapping file and try again.

   - Checks for duplicate storage ports (the same storage port mapped to more than one storage array)

     If duplicates exist, a message displays with the duplicate mappings detailed. Click **Yes** to continue. Click **No** to edit the storage port mapping file and try again.

   - Checks if mapping exists in current map

     If mappings already exist, a message displays with the current mapping information. Click **Yes** to overwrite the current mapping. Click **Yes to All** to overwrite all mapping conflicts. Click **No** to leave the current mapping. Click **No to All** to leave all current mappings when conflict occurs. Click **Cancel** to cancel the import.

   When import is complete a result summary displays with the following information ("Import Results" on page 476).

TABLE 48 Import Results

Value	Definition
Total Valid Input Records	Number of lines identified in the CSV file without any errors (excluding the Header).
Unique storage port WWN's Recognized	Number of unique storage ports identified in the CSV file.
Storage Arrays Created or Identified	Number of storage ports identified in the CSV file already discovered and are either online or offline but not deleted.
Conflicting Port Mappings	Number of occurrences where you were asked to decide whether to override previously discovered information. If a you select Yes to All, or No to All, each occurrence where conflict resolution occurs automatically is counted as one conflict.
Overwritten Port Mappings	Number of times a previously discovered mapping is overwritten during the import process.
Importing Errors	Number of errors encountered during the import.
Details	Tabulates the error information with respect to the line number where it occurred.

5. Click **OK** to close the **Import Results** dialog box.
6. Click **OK** to close the **Storage Port Mapping** dialog box.

Exporting storage port mapping

The **Storage Port Mapping** dialog box enables you to export a storage port array. The export file uses the CSV format. The first row contains the headers (Storage Node Name (WWNN), Storage Array Name) for the file.

Example

```
Storage Node Name (WWNN), Storage Array Name
20000004CFBD7100,New Storage Array
20000004CFBD896E,New Storage Array
20000037E19CED,New Storage Array
```

To export a storage port array, complete the following steps.

1. Select a storage port icon in the topology view, then select **Discover > Storage Port Mapping**.

   The **Storage Port Mapping** dialog box displays.

2. Select the storage port array you want to export port from the **Storage Array** list.
3. Click **Export**.

   The **Export** dialog box displays.

4. Browse to the location where you want to save the export file.

   Depending on your operating system, the default export location are as follows:

   - Desktop\My documents (Windows)
   - \root (Linux)

5. Enter a name for the files and click **Save**.
6. Click **OK** to close the **Storage Port Mapping** dialog box.

Host Management

Host management

Extensive management operations are supported on the switches and fabrics of the SAN using the Management application. Adapters and hosts are visible as part of the fabrics managed by the Management application.

The Management application integrates with another manageability application called the Host Connectivity Manager (HCM) to provide complete management of the Host Bus Adapters (HBAs) and Converged Network Adapters (CNAs).

The Management application focuses on operations such as fault management, performance management, and configuration management for multiple adapters and adapter ports and security configuration using Fibre Channel Security Protocol (FC-SP) that is set up on the adapter port and the switch.

HCM supports management for individual adapters (4/8/16 Gbps HBAs), 10 Gbps CNAs, 10 Gbps or 16 Gbps Fabric Adapters, and other devices, such as the host, DCB ports, FCoE ports, and Ethernet ports.

The Management application, in conjunction with HCM, provides end-to-end management capability. For information about configuring, monitoring, and managing individual adapters using the HCM GUI or the Brocade Command Utility (BCU), refer to the *Adapters Administrator's Guide*.

Supported adapters

The following sections describe the supported adapter types:

- "Host Bus Adapters"
- "Converged Network Adapters"
- "Fabric Adapters"

Host Bus Adapters

Table 49 describes the Fibre Channel Host Bus Adapters (HBAs) that provide reliable, high-performance host connectivity for mission-critical SAN environments. The supported HBAs are listed in Table 49.

TABLE 49 Supported Fibre Channel HBA models

Model number	Description	Number of ports
Brocade	• Single or Dual-port stand-up HBA with a per-port maximum of 8 Gbps using an 8 Gbps SFP.[1] • Dual-port mezzanine HBA[2] with a per-port maximum of 8 Gbps. This HBA installs in server blades that install in supported blade system enclosures. • Single or Dual-port stand-up HBA with a per-port maximum of 4 Gbps using a 4 Gbps SFP.[3]	1 or 2 (depending on the model)
Emulex[4]	Single or Dual-channel stand-up HBA (depending on the model). Refer to the Emulex model data sheets for the maximum Gbps.	1 or 2 (depending on the model)
QLogic[4]	Single or Dual-channel stand-up HBA (depending on the model). Refer to the Qlogic model data sheets for the maximum Gbps.	1 or 2 (depending on the model)

[1] A 4 Gbps SFP installed in Brocade 815 or 825 HBAs allows 4, 2, or 1 Gbps speed only.

[2] Brocade 804 mezzanine cards connect to the embedded switch modules or embedded interconnect modules on the blade system chassis by way of an internal backplane and, therefore, no optical modules (SFP transceivers) are involved. With the exception of no SFP transceivers, the Brocade 804 mezzanine FC HBA card functions the same as the other Brocade HBAs.

[3] An 8 Gbps SFP installed in Brocade 425 or 415 HBAs allows 4 or 2 Gbps speed only.

[4] Third-party adapter model must support CIMOM (Common Information Model Object Manager) or WMI (Windows Management Instrumentation) based discovery.

Using Brocade or Emulex HBAs, you can connect your server (host system) to devices on the Fibre Channel SAN. The combined high performance and proven reliability of a single-ASIC design makes these HBAs ideal for connecting hosts to SAN fabrics based on Brocade Fabric operating systems.

Converged Network Adapters

Table 50 describes available Converged Network Adapters (CNAs) for PCIe x 8 host bus interfaces, hereafter referred to as CNAs. These adapters provide reliable, high-performance host connectivity for mission-critical SAN environments.

TABLE 50 Supported Fibre Channel CNA models

Model number	Port speed	Number of ports	Adapter type
1741M-k[1,2]	10 Gbps maximum	2	Expansion
1020	10 Gbps maximum	2	Stand-up
1010	10 Gbps maximum	1	Stand-up

TABLE 50 Supported Fibre Channel CNA models

Model number	Port speed	Number of ports	Adapter type
1007[1,2]	10 Gbps maximum	2	Expansion

[1]The Brocade 1741M-k and Brocade 1007 are two-port 10 GbE CNAs that mount on a blade server that installs in a system enclosure. The adapter uses FCoE to converge standard data and storage networking data onto a shared Ethernet link. Ethernet and Fibre Channel communication are routed through the DCB ports on the adapter to the blade system enclosure midplane and onto the installed switch modules installed in the enclosure.

[2]The Brocade 1741M-k and Brocade 1007 CNAs connect to the embedded switch modules or embedded interconnect modules on the blade system chassis by way of an internal backplane and, therefore, no optical modules (SFP transceivers) are involved. With the exception of no SFP transceivers, the Brocade 1741M-k and Brocade 1007 CNAs function the same as the other Brocade CNAs. For information on installing the Brocade CNAs on a blade server, refer to the *Adapters Installation and Reference Guide*.

Brocade CNAs combine the functions of a Host Bus Adapter (HBA) and Network Interface Card (NIC) on one PCIe x 8 card. The CNAs appear as NICs and Fibre Channel adapters to the host. These CNAs fully support FCoE protocols and allow Fibre Channel traffic to converge onto 10 Gbps Data Center Bridging (DCB) networks. FCoE and 10 Gbps DCB operations are simultaneous.

The combined high performance and proven reliability of a single-ASIC design makes these CNAs ideal for connecting host systems on Ethernet networks to SAN fabrics based on Brocade Fabric or M-Enterprise operating systems.

Fabric Adapters

Table 51 describes the available Fabric Adapter models.

TABLE 51 Supported Fabric Adapter models

Model number	Port speed	Number of ports
1860-1[1] 1860-2	16 Gbps FC HBA and 10 Gbps CNA or NIC	1 or 2
1867	16 Gbps FC mezzanine card	2

1. The Brocade 1860 provides dual mode support for the port. You can configure the port mode as a 16 Gbps Fibre Channel (FC) HBA and a 10 Gbps CNA using the Brocade Command Utility (BCU).

AnyIO™ technology

Although the Brocade 1860 Fabric Adapter can be shipped in a variety of small form-factor pluggable (SFP) transceiver configurations, you can change port function to the following modes using Brocade AnyIO™ technology, provided the correct SFP transceiver is installed for the port:

- HBA or Fibre Channel mode — This mode utilizes the Brocade Fibre Channel storage driver. An 8 or 16 Gbps Fibre Channel SFP transceiver can be installed for the port. The port provides Host Bus Adapter (HBA) functions on a single port so that you can connect your host system to devices on the Fibre Channel SAN. Ports with 8 Gbps SFP transceivers configured in HBA mode can operate at 2, 4, or 8 Gbps. Ports with 16 Gbps SFP transceivers configured in HBA mode can operate at 2, 4, 8, or 16 Gbps.

  Fabric Adapter ports set in HBA mode appear as “FC” ports when discovered in HCM. They appear as “FC HBA” to the operating system.

- Ethernet or NIC mode — This mode utilizes the Brocade network driver. A 10 GbE SFP+ transceiver must be installed for the port. This mode supports basic Ethernet, Data Center Bridging (DCB), and other protocols that operate over DCB to provide functions on a single port that are traditionally provided by an Ethernet Network Interface Card (NIC). Ports configured in this mode can operate at up to 10 Gbps. Fabric Adapters that ship from the factory with 10 GbE SFP transceivers installed or no SFP transceivers installed are configured for Ethernet mode by default.

Fabric Adapter ports set in NIC mode appear as Ethernet ports when discovered in HCM. These ports appear as “10 GbE NIC” to the operating system.

- CNA mode — This mode provides all functions of Ethernet or NIC mode, plus adds support for FCoE features by utilizing the Brocade FCoE storage driver. A 10 GbE SFP+ transceiver must be installed for the port. Ports configured in CNA mode connect to an FCoE switch. The port provides all traditional CNA functions for allowing Fibre Channel traffic to converge onto 10 Gbps DCB networks. The ports appear as Network Interface Cards (NICs) and Fibre Channel adapters to the host. FCoE and 10 GbE operations run simultaneously.

  Fabric Adapter ports set in CNA mode appear as FCoE ports when discovered in HCM. These ports appear as “10 GbE NIC” to the operating system.

HCM software

The Host Connectivity Manager (HCM) is a management software application for configuring, monitoring, and troubleshooting Brocade HBAs and CNAs in a SAN environment. For instructions about how to install the HCM software, refer to the *Adapters Installation and Reference Manual.*

You can manage the software on the host or remotely from another host. The communication between the management console and the agent is managed using JSON-RPC over HTTPS or CIM-XML over HTTPS.

NOTE
All HCM, utility, SMI-S Provider, boot software, and driver installation packages, as well as the Driver Update Disk (DUD), are described in the *Adapters Installation and Reference Manual*.

HCM features

Common HBA and CNA management software features include the following:

- Discovery using the agent software running on the servers attached to the SAN, which enables you to contact the devices in your SAN.
- Configuration management, which enables you to configure local and remote systems. With HCM, you can configure the following items:
  - Brocade 4 Gbps and 8 Gbps HBAs
  - HBA ports (including logical ports, base ports, remote ports, and virtual ports) associated with the local host
  - Brocade 10 Gbps single-port and 10 Gbps dual-port CNAs
  - Brocade 16 Gbps FC adapters
  - DCB ports (CNA only)
  - FCoE ports (CNA only)
  - Ethernet ports (CNA only)
- Diagnostics, which enables you to test the adapters and the devices to which they are connected:
  - Link status of each adapter and its attached devices
  - Loopback test, which is external to the adapter, to evaluate the ports (transmit and receive transceivers) and the error rate on the adapter
  - Read/write buffer test, which tests the link between the adapter and its devices
  - FC protocol tests, including echo, ping, and traceroute
  - Ethernet loopback test (CNA only)
  - Diagnostic Port (D-Port) test

- Monitoring, which provides statistics for the SAN components.
- Security, which enables you to specify a Challenge Handshake Authentication Protocol (CHAP) secret and configure authentication parameters.
- Event notifications, which provide asynchronous notification of various conditions and problems through a user-defined event filter.

Host adapter discovery

The Management application enables you to discover individual hosts, import a group of hosts from a CSV file, or import host names from discovered fabrics. The maximum number of host discovery requests that can be accepted is 1000. Host discovery requires HCM Agent 2.0 or later.

ESXi host adapter discovery requires the vendor-specific HBA CIM provider to be installed on the ESXi host.

The Management application supports CIMOM-based discovery for third-party adapters irrespective of the operating system and firmware version. For Windows, the third-party adapter discovery is based on Windows Management Instrumentation (WMI). A duplicate enclosure might be created when a third-party adapter is discovered through WMI and connected to a fabric. This duplicate enclosure is removed in the next discovery refresh cycle without any impact to the functionality.

NOTE
Pure Fabric discovery alone shows adapters behind Access Gateway and all adapter ports as virtual. When you discover an adapter and ports using host discovery, the adapter and all its ports are shown as physical.

Instructions for discovering hosts are detailed in "Discovery".

VM Manager

A vCenter server can be discovered by adding a VM Manager to the Management application. Refer to "Discovery" for information about discovering VM Managers.

Adding a VM Manager

1. Click **Add** on the **Discover VM Managers** dialog box.

   The **Add VM Manager** dialog box displays, as shown in Figure 209.

   FIGURE 209 Add VM Manager dialog box

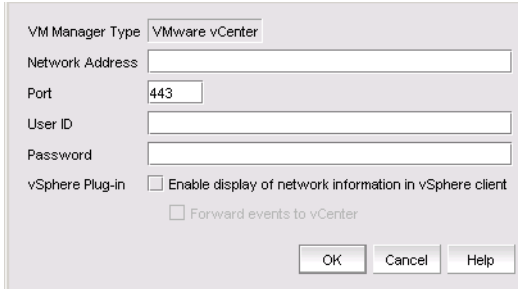

2. Enter the IP address or host name of the VM Manager (VMM) into the **Network Address** field. The maximum number of supported characters is 256.

3. Enter the VMM server port number into the **Port** field. The valid port number range is from 0 through 65536. The default port number is 443.
4. Enter the user ID into the **User ID** field to identify the user of the VMM. The maximum number of supported characters is 64.
5. Enter the password into the **Password** field. The maximum number of supported characters is 64.
6. Enable or disable the vSphere client plug-in registration. If you enable this plug-in, events are forwarded from the Management application to the vCenter server.
7. Click **OK**.

   The VMM discovery process begins. When complete, the vCenter server and all ESX and ESXi hosts managed by that vCenter display in the Host product tree.

Editing a VM Manager

The fields in the **Edit VM Manager** dialog box are identical to the fields in the **Add VM Manager** dialog box except for the **Network Address** field, which you cannot edit.

1. Click **Edit** on the **Discover VM Managers** dialog box.

   The **Edit VM Manager** dialog box displays.
2. Enter the VMM server port number into the **Port** field. The valid port number range is from 0 through 65536.
3. Enter the user ID into the **User ID** field to identify the user of the VMM. The maximum number of supported characters is 64.
4. Enter the password into the **Password** field. The maximum number of supported characters is 64.
5. Enable or disable the vSphere client plug-in registration. If you enable this plug-in, events are forwarded from the Management application to the vCenter server.
6. Click **OK**.

   The VMM discovery process begins. When complete, the vCenter server and all ESX and ESXi hosts managed by that vCenter display in the Host product tree.

Deleting a VM Manager

You cannot delete an ESX host. Hosts can only be excluded or included. If you select a host from the **Discovered VM Managers** list in the **Discover VM Managers** dialog box and click **Delete**, the host displays in the **Previously Discovered Addresses** list.

Adding an application name to a VM

NOTE
The vCenter must be discovered to add an application name to the VM.

To add one or more application names (running on the VM) to the VM Instance UUID and VM Name, complete the following steps.

1. Right-click the vCenter Host and select **Properties**.
2. Click the Virtual Machines tab, if necessary.
3. Enter the name of an application in the **Application Name** field.

If you are adding more than one application, enter the application names separated by commas. You can enter up to 255 characters.

4. Click **OK**.

HCM and Management application support on ESXi systems

Through the Brocade Adapters ESXi Management feature, ESXi systems support HCM and the Management application when CIM Provider is installed on these systems.

For installation and other information on CIM Provider, refer to the following publications:

- *CIM Provider for Brocade Adapters Developer's Guide*
- *CIM Provider for Brocade Adapters Installation Guide*

ESXi CIM listener ports

The Management application server uses two CIM indication listener ports to listen for CIM indications.

NOTE
s Management Application does not support CIM indications for Emulex Adapters.

- HCM Proxy Service CIM Indication Listener Port — This port is used to listen for CIM indications from ESXi hosts managed through HCM instances launched by the Management application. You can learn the value of these ports through the **Port Status** dialog box.
- Fault Management CIM Indication Listener Port — This port is used to listen for CIM indications from ESXi hosts managed through the Management application's host adapter discovery.

The two ports described above are part of the range of ports reserved for use by the Management application server, configurable during installation from the Server Configuration wizard. Refer to the *Installation and Migration Guide* for server configuration instructions.

Adding host adapter credentials for ESXi

CIM-based discovery is available for ESXi versions 4.1 and later. The CIM server transport does not support operating systems other than ESXi.

NOTE
CIM server credentials are optional. If you do not provide credentials, basic authentication on the CIM server is disabled and the Management application attempts discovery without authentication.

The Protocol, Port, User ID, and Password fields on the **Add Host Adapters** dialog box are persisted when changing from HCM agent to CIM Server (ESXi only).

1. Select **Discover > Host Adapters**.

   The **Host Adapters** dialog box displays.

2. Click **Add**.

   The **Add Host Adapters** dialog box, shown in Figure 210, displays.

FIGURE 210 Add Host Adapters dialog box

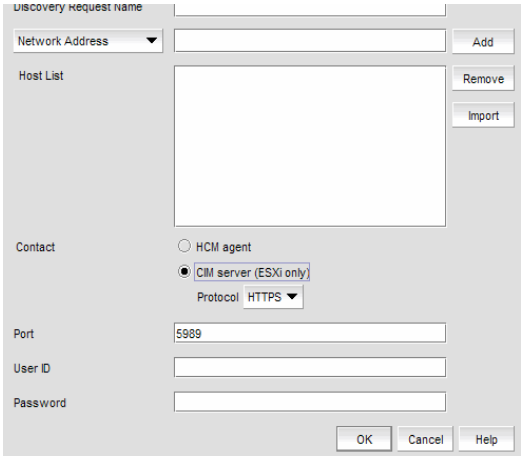

3. Select **CIM server (ESXi only)** as the **Contact** option.
4. (*Optional*) Select **HTTP** or **HTTPS** from the **Protocol** list. HTTPS is the default.
5. Click **OK**.

Connectivity map

The Connectivity Map, which displays in the upper right area of the main window, is a grouped map that shows physical and logical connectivity of Fabric OS components, including discovered and monitored devices and connections. These components display as icons in the Connectivity Map. For a list of icons that display in the Connectivity Map, refer to the following sections:

- "Host product icons" on page 304
- "Host group icons" on page 305
- "SAN port icons" on page 305

The Management application displays all discovered fabrics in the Connectivity Map by default. To display a discovered host in the Connectivity Map, you must select the host in the Product List. You can only view only one host and physical and logical connections at a time.

View management

You can customize the topology by creating views at the managed host level in addition to the fabric level views. If you discover or import a fabric with more than approximately 2,000 devices, the devices display on the Product List, but not on the Connectivity Map. Instead, the topology area shows a message stating that the topology cannot be displayed. To resolve this issue, create a new view to filter the number of devices being discovered.

Instructions for managing customized views of the topology are detailed in "View Management"

Host port mapping

HBAs and hosts discovered through one or more fabrics can be identified easily in the topology by their product icons. For a list of products and their icons, refer to "Host product icons" on page 304. Once identified in the topology, you can create hosts and assign the HBAs to them and import an externally created host port mapping file (.CSV) to the Management application.

> NOTE
> The Management application now enables you to map HBAs from multiple fabrics (previous versions limited HBA mapping to one fabric).

The Management application also enables you to discover hosts directly using host discovery (for step-by-step instructions, refer to "Host discovery" on page 51). If you discover a host directly, when you open the **Host Port Mapping** dialog box, the Management application automatically groups all HBAs under the host.

If you create a new host and associate HBAs to it, and then you try to discover a host with the same HBAs using Host discovery, the HBAs discovered using host discovery must match the HBAs associated to the host exactly; otherwise, host discovery will fail.

Instructions for mapping a host to HBAs are detailed in "Host Port Mapping".

Adapter software

The **Adapter Software** dialog box allows you to perform the following tasks:

- Select and import a driver file or delete existing drivers from the driver repository
- Update the driver to the hosts

> NOTE
> For Linux and Solaris systems, you cannot upgrade to driver file version 3.0.3.0. You must upgrade to version 3.0.3.1 or later.

The ability to update drivers to the hosts is available for hosts that are discovered through the Host Connectivity Manager (HCM) agent with driver version 2.3.0.0 or later. Driver updates cannot be performed for ESXi hosts, which are discovered using the CIM Server. Use the VMware vSphere Update Manager to update the drivers on ESXi hosts.

To update the drivers to selected hosts, complete the following steps.

1. Select **Host > Adapter Software** from the **Configure** menu.

   The **Adapter Software** dialog box, **Driver** tab, shown in Figure 211, displays.

FIGURE 211 Adapter Software dialog box, Driver tab

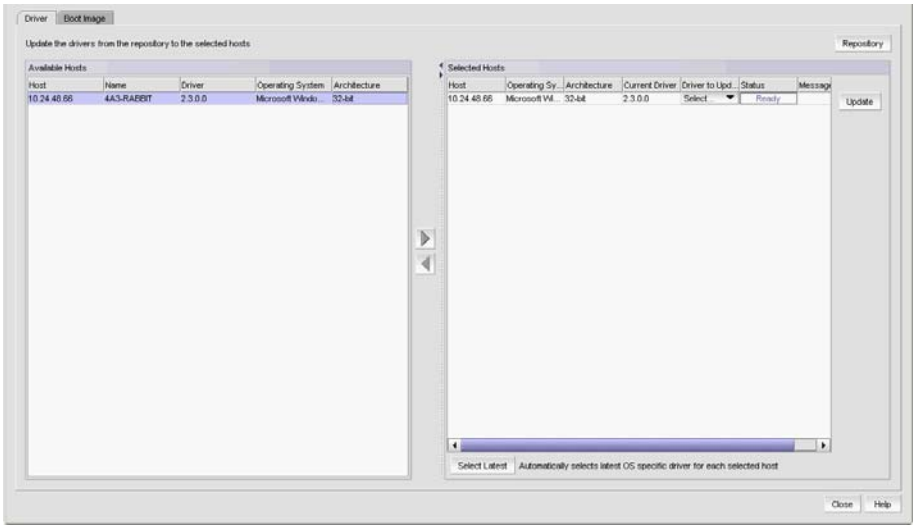

2. Select one or more hosts from the **Available Hosts** list and click the right arrow button to move the selected hosts to the **Selected Hosts** list.

   The **Available Host** list displays the following information for hosts that are discovered through the HCM agent with driver version 2.3.0.0 or later:

   - Hosts — The IP address of the host.
   - Name — The name of the host. The first three digits indicate the host's operating system; for example, WIN or LIN.
   - Operating System — The host operating system; for example, Microsoft Windows or Red Hat Linux.
   - Driver Version — The host's current driver version.
   - Architecture — The host's architecture; for example, 32-bit or 64-bit.

3. Select one or more hosts from the **Selected Hosts** list. You can select multiple hosts, but if the selected host count is greater than 20, a batch of 20 hosts is initiated for the driver update first and the remaining hosts are queued.

   The **Selected Hosts** list displays the following information for hosts that have been selected for the driver update:

   - Host — The IP address of the host.
   - Operating System — The host operating system; for example, Microsoft Windows or Red Hat Linux.
   - Driver to Update — Select the driver to update from the list.
   - Status — The ready status of the selected host.
   - Architecture — The host's architecture; for example, 32-bit or 64-bit.
   - Current Driver Version — The host's current driver version.
   - Message — Additional information pertaining to the selected host.

4. Select the host's corresponding driver to update from the **Driver to Update** list. Once the driver has been selected for each host, click **Update** .

   Alternatively, you can select one or more hosts from the **Selected Hosts** list and click **Select Latest** to automatically select the latest operating system-specific driver for each selected host. If you want to import a driver from another location, follow the instructions in "Driver repository" on page 487.

Driver repository

You can access the **Driver Repository** dialog box from the **Adapter Software** dialog box. Initially, the repository is empty. You must import files into the repository. Imported driver files are then displayed in the **Available Driver Files** list in the **Driver Repository** dialog box.

Importing a driver into the repository

To import drivers into the Management application, perform the following tasks.

1. From the **Adapter Software** dialog box, click the **Repository** button.

   The **Driver Repository** dialog box, shown in Figure 212, displays.

   FIGURE 212 Driver Repository dialog box

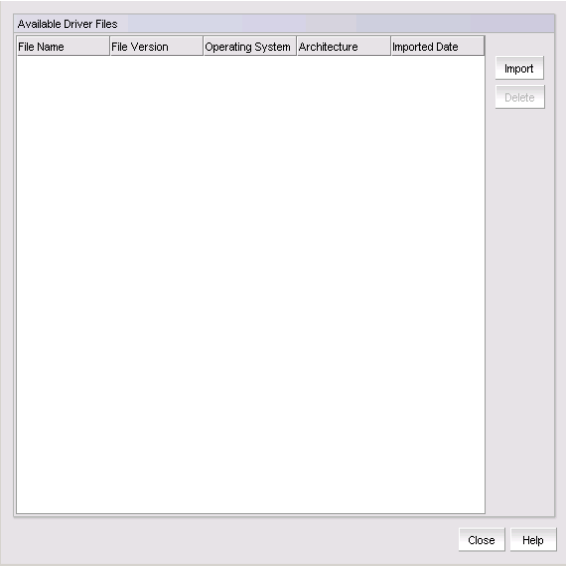

2. Click **Import** on the **Driver Repository** dialog box.

   The **Import Driver Repository** dialog box displays.

3. Locate the driver file using one of the following methods:
   - Search for the file you want from the **Look In** list.
   - Enter the name of the image file you want to import in the **File Name** field.
4. Click **Open**.

   After the import completes, you see a message that the driver imported successfully.

5. Click **OK**.

Deleting a driver file from the repository

1. Select one or more driver files from the **Available Driver Files** list on the **Driver Repository** dialog box.
2. Click **Delete**.

   The driver file is removed from the **Driver Repository** dialog box.

NOTE
Windows drivers (.exe files) cannot be imported into the server repository when the Management application server is running on Linux or Solaris platforms.

Boot image repository

The boot code image stored in the adapter's flash memory contains the instructions that enable the server to locate the boot disk in SAN. The boot code image contains the basic input/output system (BIOS), extensible firmware interface (EFI), and open firmware which enable the adapters to be compatible with any system platform.

Importing a boot image into the repository

Boot images are required for adapters that are shipped without a boot image or when it is necessary to overwrite images on adapters that contain older or corrupted boot image versions.

1. From the Management application menu bar, select **Configure > Host > Adapter Software**.
2. Click the **Boot Image** tab.

   The **Boot Image Management** dialog box, shown in Figure 213, displays.

   FIGURE 213 Boot Image Management dialog box

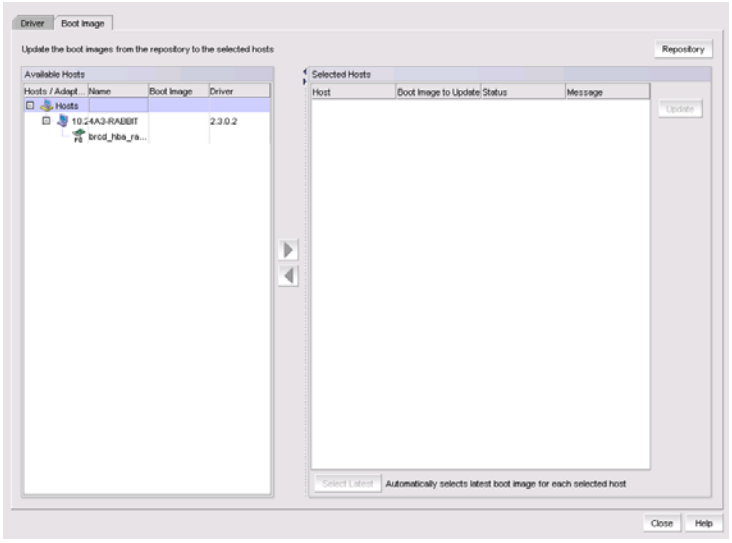

3. From the **Boot Image Management** dialog box, click the **Repository** button.

   The **Boot Image Repository** dialog box, shown in Figure 214, displays.

FIGURE 214 Boot Image Repository dialog box

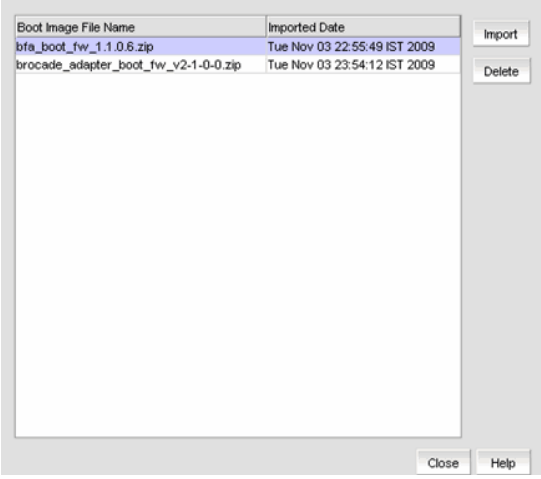

4. Click **Import** on the **Boot Image Repository** dialog box.
5. The **Import Boot Image** dialog box displays.
6. Locate the boot image file using one of the following methods:
   - Search for the file you want from the **Look In** list. Boot image files version 2.0.0.0 and 2.1.0.0 are .zip files and other boot image files are .tar files.
   - Enter the name of the image file you want to import in the **File Name** field.
7. Click **Open**.

   After the import completes, you see a message that the boot image imported successfully.

   NOTE
   The boot image file is imported to *Install_Server_Home*/data/*adapter_software*/adapter_boot_images.
8. Click **OK**.

Downloading a boot image to a selected host

To download boot images to a selected host, perform the following tasks.

1. Select one or more hosts from the **Available Hosts** list on the **Boot Image Management** dialog box, and click the right arrow button to move the selected hosts to the **Selected Hosts** list.

   You can select up to 50 hosts. The first 20 hosts execute the download concurrently. If you select more than 20 hosts, they will be queued and will start when the previous download completes.

   NOTE
   The boot image version must be 2.0.0.0 or later.
2. Click **Select Latest** to automatically select the latest boot image for the selected hosts.
3. From the **Boot Image Management** dialog box, click the **Update** button to download a boot image to one or more selected hosts.

   One of the following download status messages displays in the **Status** column of the **Selected Hosts** list:
   - Ready

- Queued
- In progress
- Failed — If the download failed, the failure reason displays in the **Message** column of the **Selected Hosts** list; for example, failed to connect to HCM agent, a checksum error occurred, or the file is invalid.
- Finished

4. Alternatively, you can click the **Select Latest** button to automatically select the latest boot image for the selected hosts.

Deleting a boot image from the repository

1. Select one or more boot images from the **Boot Image File Name** list on the **Boot Image Repository** dialog box.
2. Click **Delete**.

   The boot image is removed from the boot image repository.

Backing up boot image files

You can back up the boot image files from the repository using the **Options** dialog box. Refer to "Backup support" on page 504 for instructions.

Bulk port configuration

Use the **Adapter Host Port Configuration** dialog box to create and assign port-level configurations to either a single or multiple adapter ports at a time. You can save up to 50 port-level configurations.

The Management application supports the following default port configurations, which you can select and assign to one port or multiple ports. You cannot edit the default configurations, but you can delete them.

- Default Port — The port property. The default value is Enabled.
- Default FDFS — The Frame Data Field Size property. The default value is 2048.
- Default QoS — The Quality of Service property. The default value is Enabled.
- Default TRL — The Target Rate Limiting property. The default value is Enabled.

NOTE
You cannot configure Bulk port configuration through Management Application for Emulex Adapters.

Configuring host adapter ports

To create, edit, duplicate, or delete port configurations, complete the following steps.

Select **Host > Adapter Ports** from the **Configure** menu.

The **Configure Host Adapter Ports** dialog box, shown in Figure 215, displays.

FIGURE 215 Configure Host Adapter Ports dialog box

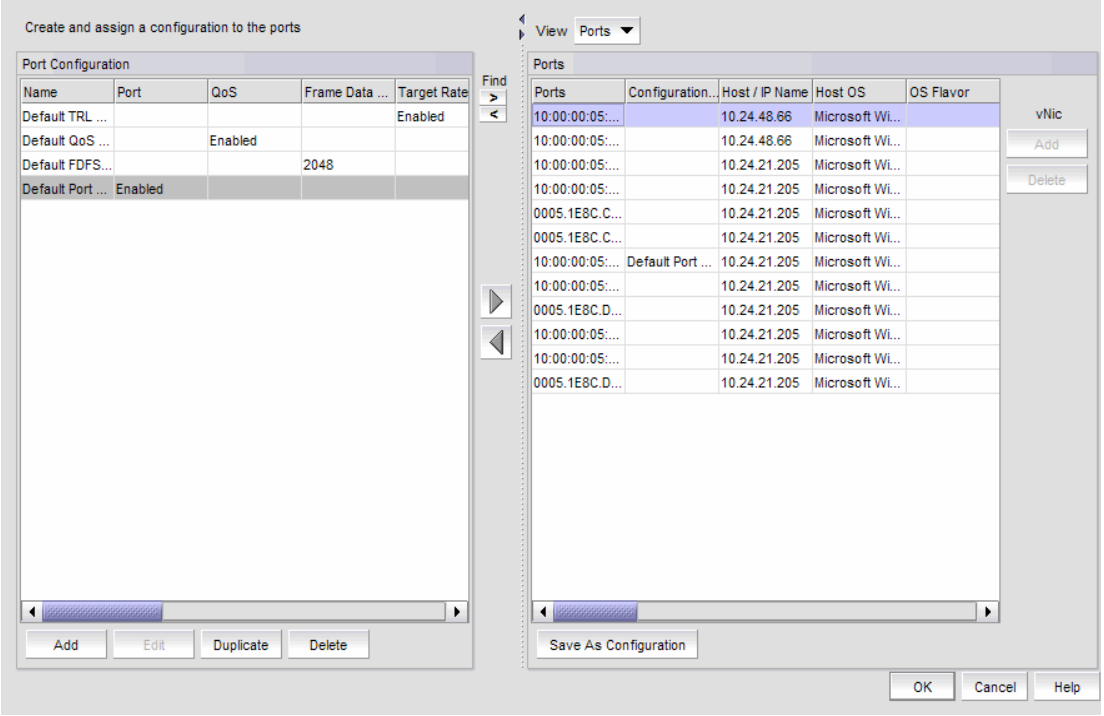

Adding a port configuration

The **Add Port Configuration** dialog box allows you to create a maximum of 50 customized port configurations which you can then select and assign to ports.

1. Click **Add** on the **Configure Host Adapter Ports** dialog box.

   The **Add Port Configuration** dialog box, shown in Figure 216, displays.

FIGURE 216 Add Port Configuration dialog box

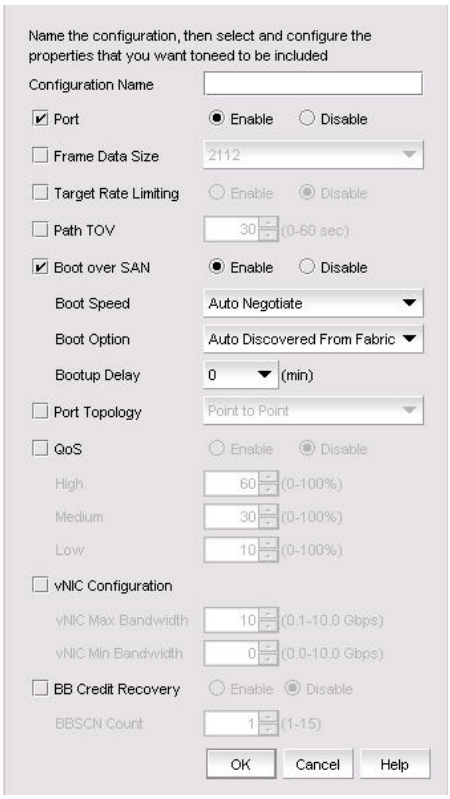

2. Enter a name for the port configuration in the **Configuration Name** field. A maximum of 128 alphanumeric characters is supported.

3. Configure at least one of the following port properties:

- **Port** — Enable or disable the port. Enable is the default.
- **Frame Data Size** — Select the frame data size, in bytes, of the port. Options include Auto, 512, 1024, 2112, and 2048; the default value is 2112. Select auto to set the frame data field size automatically. Buffer credits determine the maximum amount of frame data. If the number of buffer credits is not large enough to handle the link distance and speed, performance can be severely limited.
- **Target Rate Limiting** — Enable the Target Rate Limiting feature to minimize congestion at the adapter port. Limiting the data rate to slower targets ensures that there is no buffer-to-buffer credit back-pressure between the switch due to a slow-draining target.

  NOTE
  NOTE: Target Rate Limiting and **QoS** cannot be enabled at the same time.

- **Path TOV** — Enter a path timeout value (TOV) to either force an immediate failover (by setting the TOV to 0) or to specify a delay in seconds (1 through 60 seconds). The default value is 30 seconds.
- **Boot over SAN** — The Boot over SAN feature allows you to target remote boot devices (LUNs on SAN storage arrays) from which to boot the host system. Configure the following boot parameters:

  Boot Speed — Set the port speed. Possible values are Auto Negotiate (to auto-negotiate the speed) and 1, 2, 4, 8, and 16 Gbps and unknown speeds.

  Boot Option — From the list, select one of the following:

  - **Auto Discovered From Fabric** — Enables Boot over SAN using boot LUN information stored in the fabric. This is the default setting.

- **First Visible LUN** — Enables Boot over SAN from the first discovered LUN in the SAN.

Bootup Delay — Enter a bootup delay value. Valid values are 0, 1, 2, 5, and 10 minutes and the default value is 0 minutes. The Bootup Delay feature allows you to configure the delay to device discovery, offsetting the disk spinup delay time when servers and storage devices are powered on simultaneously.

- **Port Topology** — Specify the topology type. The supported topology mode is point-to-point (p2p) or loop. You can set the toplogy to loop only if QoS and Target Rate Limiting are disabled.
- **QoS** — Enable the Quality of Service (QoS) feature to assign traffic priority (high, medium, or low) for a given source and destination traffic flow. By default, all flows are marked as medium.

NOTE
NOTE: QoS and Target Rate Limiting cannot be enabled at the same time.

QoS Percentage — The QoS priority flow value extends QoS support by allowing the user to configure custom bandwidth values for High, Medium, and Low QoS priorities. The QoS % value represents the bandwidth in percentage for each of the priorities (high, medium, and low) and the three values must equal 100 percent.

The default priority flow settings of the switch are 60 (high), 30 (medium), and 10 (low). If QoS is disabled and enabled again without providing the high, medium, and low bandwidth values, the default values are applied.

- **vNIC Configuration** — Enables you to configure a single physical CNA Ethernet port into multiple virtual Network Interface Cards (vNICs).
  - Enter the maximum allowable output bandwidth in increments of 100 Mbps in the vNIC Max Bandwidth (Mbps) box. The maximum bandwidth is 10 Gbps and this is the default.
  - Enter the minimum allowable output bandwidth in the Min Bandwidth (Mbps) box. The minimum bandwidth is 0 Mbps. A zero value of minimum bandwidth (the default) implies that no bandwidth is guaranteed for that vNIC.
- **BB Credit Recovery** — Enables you to enable or disable buffer-to-buffer (BB) credits, which are a flow control mechanism that represent the availability of resources at the receiving port. Supported state change notification (BB_SCN) values are from 1 through 15 and the default is 1.

4. Click **OK**.

   The **Adapter Port Configuration Status** dialog box displays.

5. Click **Start**.

   The adapter port configuration is applied to the ports.

6. Click **Close** after the configuration is complete (indicated by "Completed" in the **Progress** list).

Editing a port configuration

The **Edit Port Configuration** dialog box allows you to modify port configuration parameters that were configured using the **Add Port Configuration** dialog box.

1. Click **Edit** on the **Configure Host Adapter Ports** dialog box.

   The **Edit Port Configuration** dialog box displays.

2. Modify the parameters that are described in "Adding a port configuration" on page 491.
3. Click **OK** to save the changes.

Duplicating a port configuration

1. Click **Duplicate** on the **Configure Host Adapter Ports** dialog box.

   The **Duplicate Port Configuration** dialog box displays. The default name of the configuration file is **source_name** copy1.

2. Change the name of the configuration and click **OK** to save the changes.

Deleting a port configuration

1. Select a configuration from the **Port Configuration** list in the **Configure Host Adapter Ports** dialog box.
2. Click the **Delete** button.

   The port configuration is removed from the list.

Adapter port WWN virtualization

Adapter port world wide name (WWN) virtualization enables the adapter port to use a switch-assigned WWN rather than the physical port WWN for communication, allowing you to preprovision the server with the following configuration tasks:

- Create the zones with the Fabric Assigned WWN (FAWWN) before the servers and devices are connected to the switches, before they are exposed to the SAN network.
- Create LUN mapping and LUN masking without the devices present in the network.
- Preconfigure boot LUN zoning. You can configure Solaris ports or Linux ports on the switch, enabling the server to boot automatically with the predefined boot LUNs.

  NOTE
  Fabric Assigned WWN (FAWWN) is not supported for base switches or FICON-enabled switches.

Configuring FAWWNs on switch ports

The **Configure Fabric Assigned WWNs** dialog box, shown in Figure 217, enables you to perform the following tasks:

- Enable and disable the Fabric Assigned WWN feature status on a switch or Access Gateway port.
- Set the type value to *auto* or *user-defined*. When the **User** button is clicked, the WWN is cleared from the table and editing is enabled.
- Delete the Fabric Assigned WWN from the **Fabric Assigned WWN - Configuration** list.

FIGURE 217 Configure Fabric Assigned WWNs dialog box

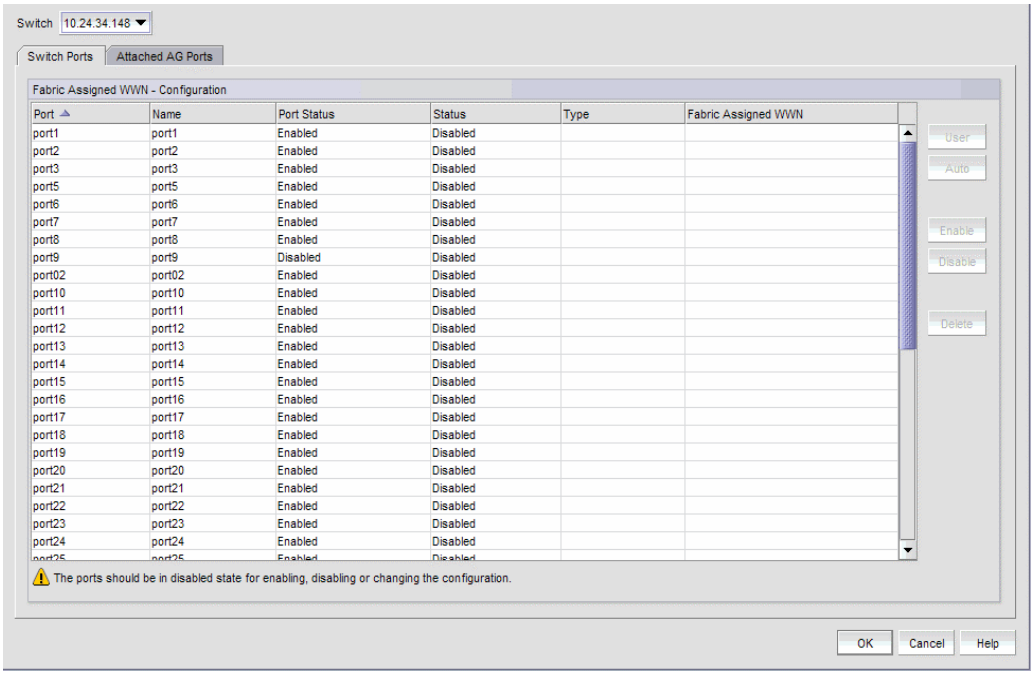

Enabling the FAWWN feature on a switch or AG ports

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select **Fabric Assigned WWN**.

   The **Configure Fabric Assigned WWNs** dialog box displays.

2. Select a switch port from the **Fabric Assigned WWN - Configuration** list.
3. Click the **Enable** button.

   The selected switch's port status is enabled.

4. Click **OK**.

   The **Fabric Assigned WWN Confirmation and Status** dialog box displays.

5. Click **Start** to save the changes to the switch.
6. Click **Close** on the **Fabric Assigned WWN Configuration and Status** dialog box.

Disabling the FAWWN feature on a switch or AG ports

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select **Fabric Assigned WWN**.

   The **Configure Fabric Assigned WWNs** dialog box displays.

2. Select a switch port from the **Fabric Assigned WWN - Configuration** list.

3. Click the **Disable** button.

   The selected switch's FAWWN feature status is disabled.

4. Click **OK**.

Auto-assigning a FAWWN to a switch or AG port

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select **Fabric Assigned WWN**.

   The **Configure Fabric Assigned WWNs** dialog box displays.

2. Select a switch port or AG port from the **Fabric Assigned WWN - Configuration** list.
3. Click the **User** button.

   The system sets the type to User and the Fabric Assigned WWN parameters are now editable.

4. Enter a valid WWN on the selected switch.
5. Click **OK**.

Manually assigning a FAWWN to a switch or AG port

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select **Fabric Assigned WWN**.

   The **Configure Fabric Assigned WWNs** dialog box displays.

2. Select a switch port or AG port from the **Fabric Assigned WWN - Configuration** list.
3. Click the **Auto** button.

   If the switch port does not have an Auto FAWWN map type and the FAWWN feature is not yet enabled on the port, a To Be Generated message displays.

4. Click **OK**.

Modifying a FAWWN on a switch or AG port

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select **Fabric Assigned WWN**.

   The **Configure Fabric Assigned WWNs** dialog box displays.

2. Select a switch port or AG port from the **Fabric Assigned WWN - Configuration** list.
3. Click the **User** button.

   The Fabric Assigned WWNs parameters are now editable.

Deleting a FAWWN from a switch or AG port

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select **Fabric Assigned WWN**.

   The **Configure Fabric Assigned WWNs** dialog box displays.

2. Select a switch port or AG port from the **Fabric Assigned WWN - Configuration** list.
3. Click the **Delete** button.

   The Fabric Assigned WWN row is deleted from the **Fabric Assigned WWN - Configuration** list for the selected switch port or AG port.

FAWWNs on attached AG ports

The **Configure Fabric Assigned Assigned WWNs** dialog box, shown in Figure 218, enables you to configure the Fabric Assigned WWN feature on a selected attached Access Gateway (AG) port.

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select Fabric Assigned WWN.

   The **Configure Fabric Assigned WWNs** dialog box displays.

2. Click the **Attached AG Ports** tab.

   The **Configure Fabric Assigned WWNs** dialog box — **Attached AG Ports** tab displays.

FIGURE 218 Configure Fabric Assigned WWNs dialog box--Attached AG Ports tab

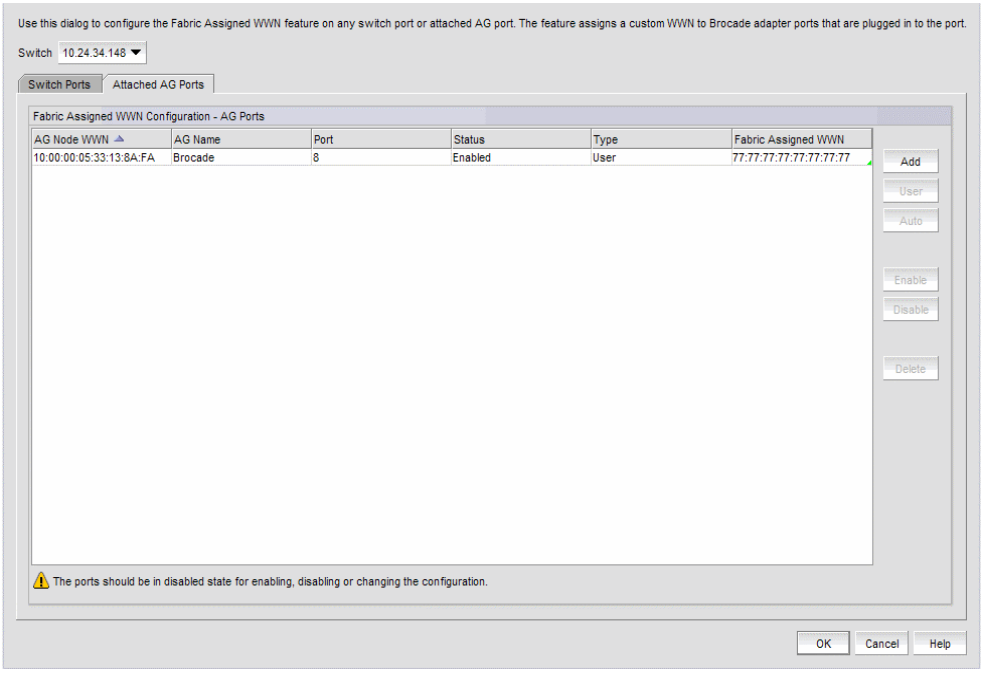

Adding AG port FAWWNs

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select **Fabric Assigned WWN**.

   The **Configure Fabric Assigned WWNs** dialog box displays.

2. Click the **Attached AG Ports** tab.
3. Select a row in the **Fabric Assigned WWN Configuration - AG Ports** list.
4. Click **Add**.

   The **Add AG Fabric Assigned WWN Configuration** dialog box displays.

5. Enter a valid world wide name (WWN), with or without colons, for the Access Gateway node. Optionally, you can select an existing AG Node WWN from the list. The **AG Node WWN** box includes all discovered AG Node WWNs that are connected to the selected switch.
6. Enter a port or a port range using numbers or a hyphen (-). For example, you can enter a range as 1-6 or you can separate values with a comma; for example: 1, 2, 5, 7-10, 20.
7. Click the **Enable** button to enable the FAWWN.
8. Set the FAWWN type to one of the following map types:
   - Auto — If the switch port does not have an Auto FAWWN map type and the FAWWN feature is not yet enabled on the port, a <To Be Generated> message displays.
   - User defined — If this option is selected, you must enter a valid world wide name, with or without colons. The User defined text box cannot be empty.
9. Click **OK** to add the rows for this configuration to the **Fabric Assigned WWN Configuration - AG Ports** list.

Deleting AG port FAWWNs

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select **Fabric Assigned WWN**.

   The **Configure Fabric Assigned WWNs** dialog box displays.

2. Click the **Attached AG Ports** tab.
3. Select an online AG FAWWN row and click the **Delete** button.

   The AG FAWWN row is cleared from the **Fabric Assigned WWN Configuration - AG Ports** list.

Moving an AG port FAWWN across switches

The AG port FAWWN can be online or offline when moved across switches.

1. Select **Configure > Fabric Assigned WWN**.

   or

   Right-click the switch and select **Fabric Assigned WWN**.

The **Configure Fabric Assigned WWNs** dialog box displays.

2. Click the **Attached AG Ports** tab.
3. Right-click the WWN row you want to move, select the **Copy Row** option, and paste the contents into a text editor.
4. Select an online AG FAWWN row and click the **Delete** button.
5. Select a switch from the **Switch** list and click **Add** to launch the **Add AG Fabric Assigned WWN Configuration** dialog box.
6. Using the information you copied to the text editor, configure the AG port FAWWN information to be moved to the selected switch.
7. Click **OK**.

   The specified AG FAWWN row is added to the new switch.

Role-based access control

The Management application enables you to create resource groups and assign users to the selected role within that group. This enables you to assign users to a role within the resource group.

The Management application provides one preconfigured resource group (All Fabrics). When you create a resource group, all available roles are automatically assigned to the resource group. Once the resource group is available, you can assign a user to a role within the resource group.

Host adapter management privileges

You can launch the Host Connectivity Manager (HCM) if you have read and write permissions to the Host Adapter Management privilege. Other HBA-related operations are controlled by the following privileges:

- The HBA technical support launch point is controlled by the Technical Support Data Collection privilege.
- The Fibre Channel Security Protocol (FC-SP) launch point is controlled by the Security privilege. Read-write (RW) and read-only (RO) permissions are required.
- The HBA performance monitoring launch point is controlled by the Performance privilege.

Host adapter administrator privileges

The Host Adapter Administrator role has the following privileges:

- Add and delete properties
- Discovery setup
- Host management
- Performance
- Properties edit
- Security
- Servers
- View management
- Port Mapping
- Virtual Network Management

Instructions for managing resource groups and users using roles and privileges are detailed in "User accounts,""Roles," and "Areas of responsibility," in "User Account Management".

Host performance management

Real-time performance enables you to collect data from managed HBA and CNA ports. You can use real-time performance to configure the following options:

- Select the polling rate from 20 seconds up to 1 minute.
- Select up to 32 ports total from a maximum of 10 devices for graphing performance.
- Choose to display the same Y-axis range for both the Tx MBps and Rx MBps measure types for easier comparison of graphs.

NOTE
In the **Port Picker** dialog box, the Brocade 1860 Fabric Adapter in AnyIO mode displays in both categories (HBA port measures and CNA port measures). The ports are properly filtered to display only the CNA or HBA port based on the selection.

Table 52 lists the counters that are supported for the FC ports and for the HBA and CNA ports.

TABLE 52 Counters

FC port measures	HBA port measures	CNA port measures
Tx % utilization	Tx % utilization	Tx % utilization
Rx % utilization	Rx % utilization	Rx % utilization
Tx MBps	Tx MBps	Tx MBps
Rx MBps	Rx MBps	Rx MBps
CRC errors	CRC errors	
Signal losses	Signal losses	
Sync losses	Sync losses	
Link failures	Link failures	
Sequence errors	Primitive sequence protocol errors	
Invalid transmissions		
Rx link resets		
Tx link resets		
	NOS count	
	Error frames	
	Dropped frames	
	Undersized frames	
	Oversized frames	
	Bad EOF frames	
	Invalid ordered sets	
	Non-frame coding error	
		Received paused frames
		Transmitted paused frames
		Received FCoE pause frames

TABLE 52 Counters (Continued)

FC port measures	HBA port measures	CNA port measures
		Transmitted FCoE pause frames
		Received FCS error frames
		Transmitted FCS error frames
		Received alignment error frames
		Received length error frames
		Received code error frames

Instructions for generating real-time performance data are detailed in "Generating a real-time performance graph" on page 966.

Host security authentication

Fibre Channel Security Protocol (FC-SP) is a mechanism used to secure communication between two switches or between a switch and a device such as an HBA port.

You can use either the Management application or the HCM GUI to display the authentication settings and status. *When you enable FC-SP authentication using the Management application, you can also set the authentication settings on the attached* 8 Gbps 8-FC port.

NOTE
FC-SP is only available for Brocade HBAs that are managed using the HCM agent and CIM Server. FC-SP is not available for virtual ports or unmanaged HBA ports. The user must have the Security privilege to use this feature. FC-SP is not supported for hosts connected to Access Gateway mode-enabled devices.

Configuring security authentication using the Management application

Access the **Fibre Channel Security Protocol Configuration** dialog box by selecting an adapter port from the device tree. Select the appropriate device based on how you want to configure security authentication.

1. Select **Configure > Element Manager > HCM**.

   The Host Connectivity Manager (HCM) launches.

2. From HCM, select **Configure > Authentication**.

   The **Fibre Channel Security Protocol Configuration** dialog box, shown in Figure 219, displays.

FIGURE 219 Fibre Channel Security Protocol Configuration dialog box

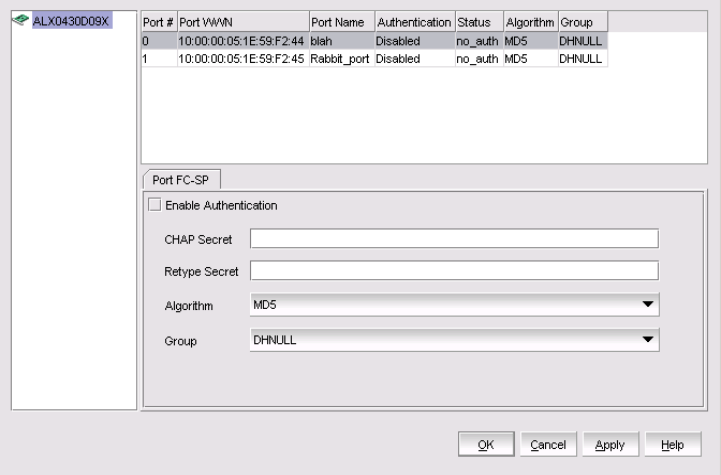

3. Configure the following parameters on the **Fibre Channel Security Protocol Configuration** dialog box:

   a. Select the **Enable Authentication** check box to enable the authentication policy.

      If authentication is enabled, the port attempts to negotiate with the switch. If the switch does not participate in the authentication process, the port skips the authentication process.

   b. In the **Algorithm** list, select one of the following options:

      - **MD5** - A hashing algorithm that verifies a message's integrity using Message Digest version 5. MD5 produces a 128-bit digest and is the required authentication mechanism for LDAP v3 servers.
      - **SHA1** - A secure hashing algorithm that computes a 160-bit message digest for a data file that is provided as input.
      - **MD5SHA1** - Similar to the MD5 hashing algorithm, but used for DH-CHAP authentication.
      - **SHA1MD5** - Similar to the SHA1 hashing algorithm, but used for DH-CHAP authentication.

   c. Enter a secret in the **CHAP Secret** field. Enter the secret again in the **Retype Secret** field.

      The length of the secret must be from 8 through 41 characters in length. The **Secret** field cannot be blank.

   d. From the **Group** list, select **DHNULL** as the DH-group type value.

4. Click **OK** to save the changes and close the dialog box.

   FC-SP settings are also applied to the attached switch.

supportSave on adapters

Host management features support capturing support information for managed Brocade adapters, which are discovered in the Management application. You can trigger supportSave for multiple adapters at the same time.
supportSave cannot be used to collect support information for ESXi hosts managed by a CIM Server. Refer to the *Adapters Administrator's Guide* for information about supportSave on ESXi hosts.

NOTE
You cannot schedule host supportSave information.

NOTE
You cannot schedule host supportSave information through Management Application for Emulex Adapters.

Instructions for scheduling and capturing technical support files are detailed in "Technical Support".

Host fault management

Fault management enables you to monitor your SAN using the following methods:

- Monitors logs for specified conditions and sends a notification or runs a script when the specified condition is met.
- Creates event-based policies, which contain an event trigger and action.
- Configures e-mail event notifications.
- Receives and forwards Syslog messages from Fabric OS switches and Brocade HBAs, managed using the Host Connectivity Manager (HCM).
- Through the Brocade Adapters ESXi Management feature, ESXi systems support the HCM and the Management application when the CIM provider is installed on these systems.

  NOTE
  The host name of the ESXi host being discovered through CIM discovery in the Management application should be configured such that it resolves to the same IP address used for discovering that ESXi host in the Management application.

Adapter events

You can configure triggers and actions for the following event types:

- Product Audit Event — Occurs when a target product is audited.
- Product Status Event — Occurs when a device or connection changes to up or down.
- Product Threshold Alert Event — Notifies you when a threshold alert has been reached.

Filtering event notifications

The Management application provides notification of many different types of SAN events. If a user wants to receive notification of certain events, you can filter the events specifically for that user.

NOTE
The e-mail filter in the Management application is overridden by the firmware e-mail filter. When the firmware determines that certain events do not receive e-mail notification, an e-mail notification is not sent for those events even when the event type is added to the **Selected Events** table in the **Define Filter** dialog box. Refer to "Setting up advanced event filtering" on page 1135 for more information.

To configure an e-mail event, use the instructions in "Configuring e-mail notification" on page 1132.

Syslog forwarding

NOTE
Syslog messages are only available on Fabric OS devices and HBAs (managed using the HCM Agent). CIM events are only logged in the master log and the forwarding of CIM events is not supported.

Syslog forwarding is the process by which you can configure the Management application to send Syslog messages to other computers. Switches only send the Syslog information through port 514; therefore, if port 514 is being used by another application, you must configure the Management application to listen on a different port. Then you must configure another Syslog server to listen for Syslog messages and forward the messages to the Syslog listening port of the Management application. Brocade HBAs only send the Syslog information through port 514; therefore, if port 514 is being used by another application, the Management application cannot send Syslog messages to another computer.

Syslog messages are persisted in the database. You can view the Syslog messages from the Management application. However, the Management application does not convert the Syslog messages into event objects except for the audit Syslog messages.

For more information about Syslog forwarding, refer to "Syslog forwarding" on page 1151.

Backup support

The Management application helps you to protect your data by backing it up automatically. The data can then be restored, as necessary.

Configuring backup to a hard drive

NOTE
Configuring backup to a hard drive requires a hard drive. The drive should not be the same physical drive on which your operating system or the Management application is installed.

To configure the backup function to a hard drive, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Server Backup** in the **Category** list.

   The currently defined directory displays in the **Output Directory** field.

3. Select the **Enable Backup** check box, if necessary.

4. Choose one or more of the following options:

   - Select the **Include Adapter Boot Image** check box to back up boot image files from the boot image repository.
   - Select the **Include FTP Root directory** check box.

     If you select the FTP Root directory, the FTP Root sub-directories, Technical Support, and Trace Dump are selected automatically and you cannot clear the sub-directory selections. If you do not select the FTP Root directory, the sub-directories can be selected individually.

5. Enter the time (using a 24-hour clock) you want the backup process to begin in the **Next Backup Start Time Hours** and **Minutes** fields.

6. Select an interval from the **Backup Interval** list to set how often backup occurs.

7. Browse to the hard drive and directory to which you want to back up your data.

8. Click **Apply** or **OK**.

   The application verifies that the backup device exists and that the server can write to it.

   If the device does not exist or is not writable, an error message displays that states you have entered an invalid device. Click **OK** to go back to the **Options** dialog box and fix the error.

   Backup occurs, if needed, at the interval you specified.

Enabling backup

Backup is enabled by default. However, if it has been disabled, complete the following steps to enable the function.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Server Backup** in the **Category** list.
3. Select the **Enable Backup** check box.
4. Click **Apply** or **OK**.

Disabling backup

Backup is enabled by default. If you want to stop the backup process, you must disable backup. To disable the backup function, complete the following steps.

1. Select **Server > Options**.

   The **Options** dialog box displays.

2. Select **Server Backup** in the **Category** list.
3. Clear the **Enable Backup** check box.
4. Click **Apply** or **OK**.

Fibre Channel over Ethernet

FCoE overview

Fibre Channel over Ethernet (FCoE) leverages Ethernet enhancements, called ***Data Center Bridging*** (DCB), to transport encapsulated Fibre Channel frames over Ethernet. Ethernet is the physical layer over which the encapsulated Fibre Channel frames are transported.

One of the barriers to using Ethernet as the basis for a converged network has been the limited bandwidth that Ethernet has historically provided. However, with 10 Gbps Ethernet, the available bandwidth offers the potential to consolidate all the traffic types over the same link.

Unlike Fibre Channel, Ethernet is not a peer-to-peer protocol. The mechanism used to discover new ports, MAC address assignments, and Fibre Channel logins and logouts is called the FCoE Initialization Protocol (FIP).

DCBX protocol

Data Center Bridging Exchange (DCBX) protocol allows enhanced Ethernet devices to convey and configure their DCB capabilities and ensures a consistent configuration across the network. DCBX protocol is used between DCB devices, such as a converged network adapter (CNA) and an FCoE switch, to exchange configuration with directly connected peers.

NOTE
When DCBX protocol is used, any other Link Layer Discovery Protocol (LLDP) implementation must be disabled on the host systems.

Enhanced Ethernet features

Data Center Bridging (DCB) is a set of IEEE 802 standard Ethernet enhancements that enable Fibre Channel convergence with Ethernet. The two basic requirements in a lossless Ethernet environment are Enhanced Transmission Selection (ETS) and priority-based flow control. These capabilities allow the Fibre Channel frames to run directly over 10 Gbps Ethernet segments without adversely affecting performance.

Enhanced Transmission Selection

Enhanced Transmission Selection (ETS) allows lower priority traffic classes to use available bandwidth that is not being used by higher priority traffic classes and maximizes the use of available bandwidth.

ETS allows configuration of bandwidth per priority group.

Priority group ID (PG ID) usage is defined as follows:

- PG ID 0, 7 are used when the priority group is limited for its bandwidth use.
- PG ID 8, 14 are reserved.
- PG ID 15.0 through 15.7 are used for priorities that are not limited for their bandwidth use.

The configured priority group percentage refers to the maximum percentage of available link bandwidth after PG ID 15.0 to 15.7 is serviced, assuming all priority groups are fully subscribed. If one of the priority groups does not consume its allocated bandwidth, then any unused portion is available for use by other priority groups.

Priority-based flow control

Priority-based flow control (PFC) allows the network to selectively pause different classes of traffic and create lossless lanes for Fibre Channel, while retaining packet drop congestion management for IP traffic. A high-level pause example follows:

- During periods of heavy congestion, the receive buffers reach high threshold and generate a pause.
- The pause tells transmission (Tx) queues to stop transmitting.
- After the receive (Rx) buffers reach low threshold, a zero pause is generated.
- The zero pause signals the Tx queues to resume transmitting.

Ethernet jumbo frames

The basic assumption underlying FCoE is that TCP/IP is not required in a local data center network and the necessary functions can be provided with Enhanced Ethernet. The purpose of an “enhanced” Ethernet is to provide reliable, lossless transport for the encapsulated Fibre Channel traffic. Enhanced Ethernet provides support for jumbo Ethernet frames and in-order frame delivery.

The Fabric OS FCoE 10 Gbps converged network adapter supports jumbo packets of up to 9 KB, compared to the original 1,518-byte maximum transmission unit (MTU) for Ethernet. The frame size increase allows the same amount of data to be transferred with less effort.

FCoE protocols supported

The Fabric OS FCoE converged network adapter supports two layers of protocols: Ethernet link layer and FCoE layer.

Ethernet link layer protocols supported

The following protocols support the Ethernet link layer:

- 802.1q (VLAN)
- 802.1Qaz (Enhanced Transmission Selection)
- 802.1Qbb (priority-based flow control)
- 802.3ad (link aggregation)
- 802.3ae (10 Gb Ethernet)
- 802.1p (priority encoding)
- IEEE 1149.1 (JTAG) for manufacturing debug and diagnostics
- IPv4 specification (RFC 793/768)
- IPv6 specification (RFC 2460)
- TCP/UDP specification (RFC 793/768)
- ARP specification (RFC 826)
- RSS with support for IPV4TCP, IPV4, IPV6TCP, IPV6 hash types
- HDS (Header-data split)

FCoE protocols

The following protocols support Fibre Channel over Ethernet:

- FIP (FC-BB5-compliant):
  - Support for FIP Discovery protocol for dynamic FCF discovery and FCoE link management
  - Support for FPMA and SPMA type FIP fabric login
- Support for Initiator mode only (FCP-3-compliant in Initiator mode)
- SCSI protection information support
- IP-over-FC
- NPIV support

FCoE licensing

The FCoE license enables Fibre Channel over Ethernet (FCoE) functionality on the following supported DCB switches:

- Network OS 10 GbE 24-port 8 GbE 8 FC port switch
- Network OS VDX 6710, 6720, and 6730 switches
- Network OS VDX 6740 and 6740T switches

- Network OS VDX 8770-series switches
- Network OS VDX 2730 10 GbE connection blade for the Fujitsu PRIMERGY BX900 and BX400 Blade Servers

Without the FCoE license, the DCB switches are pure Layer 2 Ethernet switches and do not allow FCoE bridging capabilities.

Saving running configurations

The **Save Running to Startup** dialog box lists discovered DCB switches with Fabric OS version 7.0 firmware or later. You can select available switches and move them to the **Selected Switches** list. Upon startup, the DCB switch configuration is copied to the selected switches.

NOTE
The **Save Running to Startup** dialog box launches if there is at least one DCB switch discovered. If no DCB switches exist, a warning message displays.

Copying switch configurations to selected switches

1. To access the **Save Running to Startup** dialog box, select **Configure > Configuration > Save Running to Startup**.

   The **Save Running to Startup** dialog box displays, as shown in Figure 220.

   FIGURE 220 Save Running to Startup dialog box

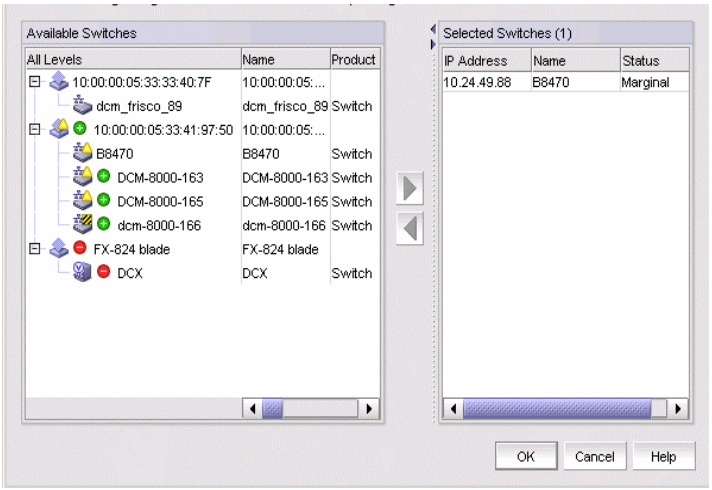

2. Highlight a discovered DCB switch from the **Available Switches** list, and click the right arrow button to move the switch to the **Selected Switches** list.
3. Highlight the selected switch and click **OK** to start the configuration.

   The running configuration is saved to the selected switch, effective on the next system startup. If you restore the DCB switch using the **Restore Switch Configuration** dialog box, you are prompted to select one of two restoration methods:

   - As the running configuration and reboot

     ATTENTION
     Rebooting a switch connected to a fabric will stop all traffic to and from the switch. All ports on the switch will become inactive until the switch comes back online.

- As the startup configuration (no reboot)

For instructions on how to restore a saved switch configuration, refer to the section "Adaptive backup" in the "Device Configuration" chapter.

DCB configuration management

Depending on the platform, the DCB switch has one of the configurations shown in Table 53.

TABLE 53 DCB configurations

Device type	Configuration possibilities
IBM blade server	• 14 internal 10-Gbps ports for IBM BladeCenter H (BCH) chassis type • 12 internal 10-Gbps ports IBM BladeCenter HT (BCHT) chassis type • 8 external 10-Gbps DCB ports • 8 8-Gbps FC ports
Dell embedded switch module	• 16 10-Gbps internal ports • 8 10-Gbps external ports • 4 8-Gbps FC ports
Fabric OS DCB switch	• 8 16-Gbps FC ports • 24 10-Gbps Ethernet ports
Fabric OS FCOE10-24 blade	24 10-Gbps Ethernet ports

You must configure DCB interfaces and ports differently than you configure Fibre Channel ports to effectively use the converged network features.

For example, priority-based flow control (PFC) and Enhanced Transmission Selection (ETS) are the two QoS policy enhancements you must configure to create a lossless Ethernet. You then use DCBX protocol on DCB-enabled devices to exchange configuration information.

The DCB ports of FOS DCB devices are categorized into two types:

- External ports - The eight external ports are the same as the original 10 Gbps Ethernet DCB ports. The default name in the device tree is ExT <slot>/<port>.
- Internal ports - The default name for the 12 or 14 internal ports is InT <slot>/<port>.
  802.1x, LAG configuration, and Spanning Tree Protocol (STP) are not supported on internal ports.

Switch policies

You can configure and enable a number of DCB policies on a switch, port, or link aggregation group (LAG).

The following switch policy configurations apply to all ports in a LAG:

- DCB map and Traffic Class map
- Link Layer Discovery Protocol (LLDP)

The switch policies are described in the following sections.

DCB map and Traffic Class map

With DCB, Fibre Channel uses a buffer management system based on buffer-to-buffer credits, with corresponding confirmation by the R-RDY frame. The flow control standard used for DCB is based on “pause” frames. Coupled with an appropriate input buffer, lossless transport of frames is possible.

Priority-based flow control (PFC) deals with the prioritization of frames. This standard IEEE 802.1Q allows application-specific bandwidth reservations in DCB. When you create a DCB map, you specify the precedence (priority) and then you map the priority groups with the Class of Service (CoS) and apply bandwidth percentages.

Refer to “QoS configuration” on page 524 for instructions on how to create DCB maps and Traffic Class maps.

LLDP profiles

Data Center Bridging Exchange (DCBX) protocol enables Enhanced Ethernet devices to discover whether a peer device supports particular features, such as Priority-based Flow Control (PFC) or Class of Service (CoS). In a Data Center Bridging (DCB) environment, LLDP is enhanced with DCBX protocol to further share or change the configured DCB enhancements.

Refer to “LLDP-DCBX configuration” on page 535 for instructions on how to configure LLDP for FCoE.

802.1x policy

802.1x is a standard authentication protocol that defines a client-server-based access control and authentication protocol. 802.1x restricts unknown or unauthorized clients from connecting to a LAN through publicly accessible ports.

Refer to “802.1x authentication” on page 539 for information on setting 802.1x parameters.

DCB configuration

To launch the **DCB Configuration** dialog box, select **Configure > DCB** from the menu bar.

The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

> NOTE
> For FOS DCB devices, the **Protocol Down Reason** column, shown in Figure 221, displays the values only for the external ports of embedded platforms but not for the internal ports.

FIGURE 221DCB Configuration dialog box

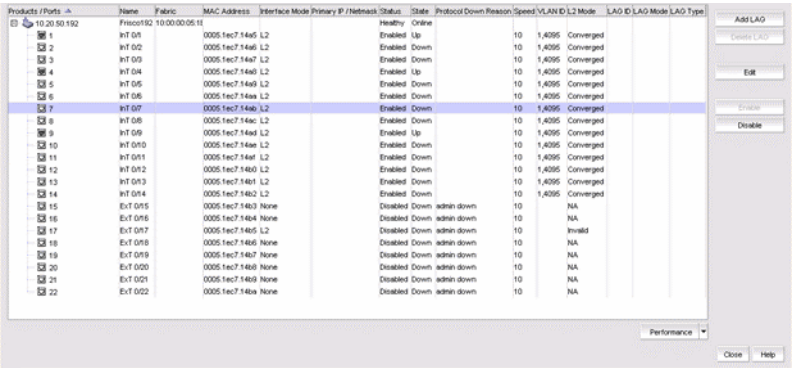

Minimum DCB configuration for FCoE traffic

You must complete the following procedures to create the basic configuration of DCB for FCoE traffic.

NOTE
This section is applicable for Fabric OS versions 6.3.0, 6.3.1, 6.3.2, 6.4.1, and 6.4.2. This section is not applicable for Fabric OS versions 6.3.1_dcb, 6.3.1_cee, 6.4.1_fcoe, and 7.0.*x*.

Creating a DCB map to carry the LAN and SAN traffic

To create a DCB map to carry the LAN and SAN traffic, complete the following steps.

NOTE
This procedure is applicable for Fabric OS versions earlier than Fabric OS 7.0. For Fabric OS versions 7.0 and later, you can only edit the default DCB map.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays.

2. Select the switch to edit from the **Products/Ports** list and click **Edit**.

   The **Edit Switch** dialog box displays.

3. Click the **QoS** tab.

   The **Edit Switch dialog box - QoS tab** displays, as shown in Figure 222.

   FIGURE 222 Edit Switch dialog box - QoS tab

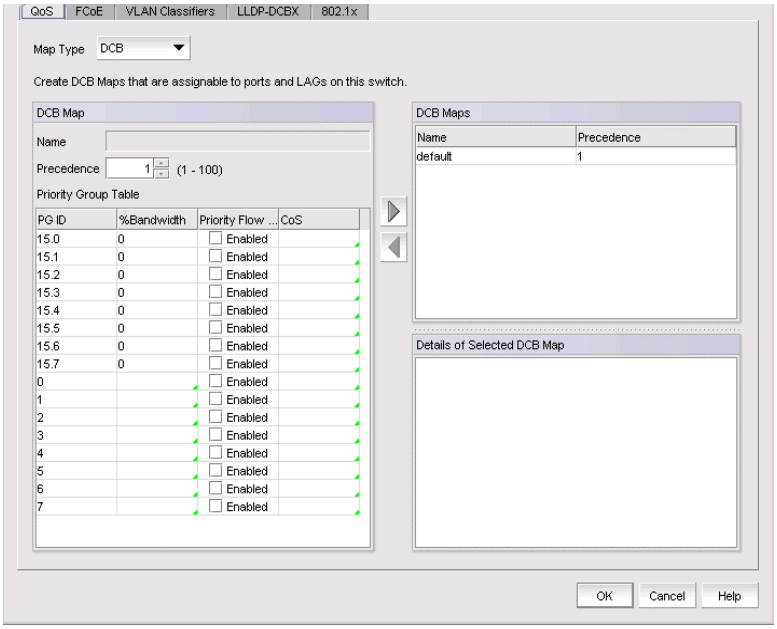

4. Select **DCB** from the **Map Type** list.
5. Configure the following DCB Map parameters in the **DCB Map** area:
   - **Name** - Enter a name to identify the DCB map.

- **Precedence** - Enter a value from 1 through 100. This number determines the map's priority.
- **Priority Flow Control** check box - Check to enable priority-based flow control on individual priority groups.
- **CoS** - Click the CoS cell to launch the **Edit CoS** dialog box, where you can select and assign one or more priorities (PG ID 15.0 through 15.7).

  All of the eight CoS values (0-7) must be used in a DCB map. Duplicate CoS values in two or more priority groups are not allowed.

  NOTE
  You can only edit CoS fields that are displayed with a green tick mark.

  % Bandwidth (optional) - While in the **Edit CoS** dialog box, enter a bandwidth value for PG IDs 15.0 through 15.7. You must map each CoS to at least one of the PG IDs.

  Note the following points:

  - You cannot define a bandwidth percentage for strict priorities (PG ID 15.0-15.7). The total bandwidth percentage for PG ID 15.0 through 15.7 must equal 0.
  - If you set a CoS value to one or more of the PG IDs 0-7, you must also enter a non-zero bandwidth percentage. The total bandwidth percentage must equal 100.
  - For PG IDs 0-7 that do not have an assigned CoS value or PFC enabled, the bandwidth percentage must be 0.

6. Click the right arrow button to add the map to the **DCB Maps** list.

   If a DCB map exists with the same name, a validation dialog box launches and you are asked if you want to overwrite the map.

7. Click **OK**.
8. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Configuring LLDP

To configure LLDP, complete the following steps.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays.

2. Select the switch to edit from the **Product/Ports** list and click **Edit**.

   The **Edit Switch** dialog box displays.

3. Click the **LLDP-DCBX** tab.

   The **Edit Switch dialog box - LLDP-DCBX** tab displays, as shown in Figure 223.

FIGURE 223 Edit Switch dialog box - LLDP-DCBX tab

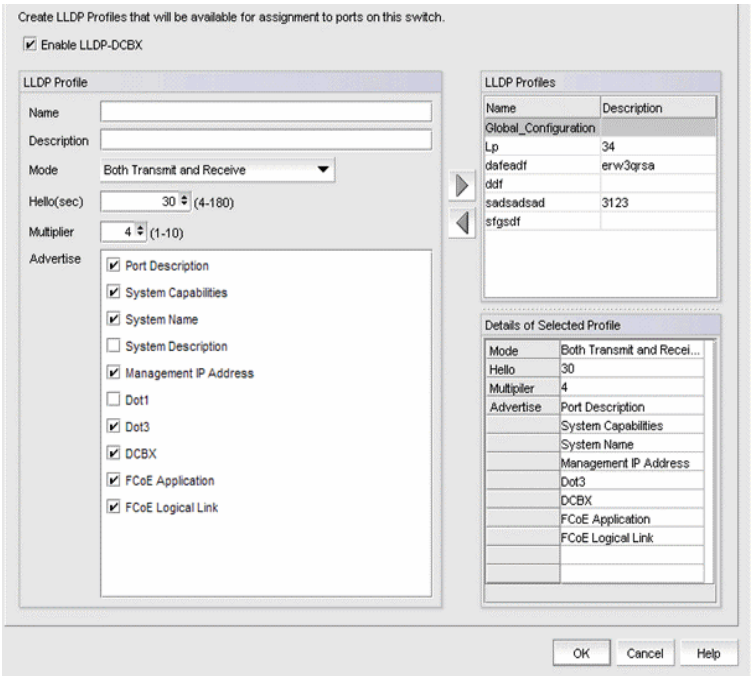

4. Select the **Global Configuration** LLDP profile in the **LLDP Profiles** list.
5. Click the left arrow button to edit.
6. Select the **FCoE Application** and **FCoE Logical Link** check boxes in the **Advertise** list to advertise them on the network.
7. Click **OK after changing the attributes of the current deployment**.

   The **Deployment Status** dialog box displays.
8. Click **Start** on the **Deployment Status** dialog box to save the changes to the switch.
9. Click **Close** to close the **Deployment Status** dialog box.

Configuring the DCB interface with the DCB map and global LLDP profile

To configure the DCB interface, complete the following steps.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays.
2. Select the Te port connected to the CNA from the **Product/Ports** list and click **Edit**..

   The **Edit Port** dialog box displays, as shown in Figure 226.
3. Select the **Port** tab, if necessary, and select the **Enable** check box.
4. Select **L2** from the **Interface Mode** list.
5. Select **Converged** (for a Brocade CNA) or **Access** (for a QLogic CNA) from the **L2 Mode** list.
6. Click the **QoS** tab and select the **Assign a map** check box.
7. Select **DCB** from the **Map Type** list.

8. Select the DCB map you created in "Creating a DCB map to carry the LAN and SAN traffic" on page 513 from the **Available DCB Maps** list.
9. Click the **LLDP-DCBX** tab and select the **Enable LLDP-DCBX on Te** *Port Number* check box.
10. Select **Assign the Global Configuration**.
11. Click **OK**.

    The **Deploy to Ports** dialog box displays.
12. Click **OK** after changing the attributes of the current deployment.

    The **Deployment Status** dialog box displays.
13. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected ports.
14. Click **Close** to close the **Deployment Status** dialog box.

Creating the FCoE VLAN to carry FCoE traffic

NOTE
You can complete this procedure using the Management application on embedded platforms such as the Fabric OS converged 10 GbE switch module for the IBM BladeCenter or the Dell M8428-k switch. You must use Web Tools to complete this procedure for the Fabric OS Fabric OS switch or the FCOE10-24 port blade.

To create the FCoE VLAN, complete the following steps. This procedure is applicable for Fabric OS versions earlier than Fabric OS 7.0.

1. Select the Fabric OS FCoE switch in the device tree.
2. Select **Configure > Element Manager > Admin**.

   The Web Tools application displays. You can also launch Web Tools by clicking the **Element Manager** button on the **DCB Configuration** dialog box.
3. Click the **DCB** tab.
4. Click the **VLAN** tab.
5. Click **Add**.

   The **VLAN Configuration** dialog box displays.
6. Enter the VLAN identifier in the **VLAN ID** field.
7. Click **OK** on the **VLAN Configuration** dialog box.
8. Select the VLAN you created and click **Edit** to convert the VLAN to FCoE VLAN.
9. Select the **FCoE** check box.
10. Select the DCB interface to carry the FCoE traffic from the **Selection List** and click **Add** to add it to the **Selected List**.
11. Click **OK** on the **VLAN Configuration** dialog box to save your changes.
12. Close the Web Tools application.

Creating and activating VLAN classifiers on the DCB interface

NOTE
You can complete this procedure using the Management application for Fabric OS versions 7.0 and later. For Fabric OS versions earlier than Fabric OS 7.0, you must use the CLI.

To create and activate the VLAN classifiers on the DCB interface, complete the following steps.

1. Log in to the switch and enter global configuration mode.

```
switch:<userid>>cmsh
switch#configure terminal
```

2. Create and apply VLAN classifiers to the DCB interface to classify Ethernet frames on an untagged interface to VLAN.

```
switch(config)#vlan classifier rule 1 proto fip encap ethv2
switch(config)#vlan classifier rule 2 proto fcoe encap ethv2
switch(config)#vlan classifier group 1 add rule 1
switch(config)#vlan classifier group 1 add rule 2
```

3. Apply the VLAN classifier group to the DCB interface.

```
switch(conf-if-te-0/7)#vlan classifier activate group 1 vlan 1002
```

4. Save the **running-config** file to the startup-config file.

```
switch#copy running-config startup-config
```

Adding a LAG

Link aggregation, based on the IEEE 802.3ad protocol, is a mechanism to bundle several physical ports together to form a single logical channel or trunk. The collection of ports is called a link aggregation group (LAG).

NOTE
An internal port cannot be part of a LAG. You can create LAGs with external ports only.

- The **Add LAG** button on the **DCB Configuration** dialog box is enabled when a single DCB switch or ports of a single DCB switch are selected.
- The **Add LAG** button is disabled when multiple switches are selected, ports from different switches are selected, or LAGs are selected.
- The **Edit LAG** button is enabled when a single LAG, port, or switch is selected.

NOTE
When LLDP-DCBX is disabled on the switch, a yellow banner displays on the **DCB Configuration** dialog box, indicating that LLDP-DCBX is not only disabled on the switch, but is also disabled for all ports and LAGs on the switch.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select the DCB switch or one or more DCB ports from the **Products/Ports** list to add to a link aggregation group (LAG).
3. Click **Add LAG or Edit LAG**.

   The **Add LAG** or **Edit LAG** dialog box displays, as shown in Figure 224.

FIGURE 224 Add LAG dialog box

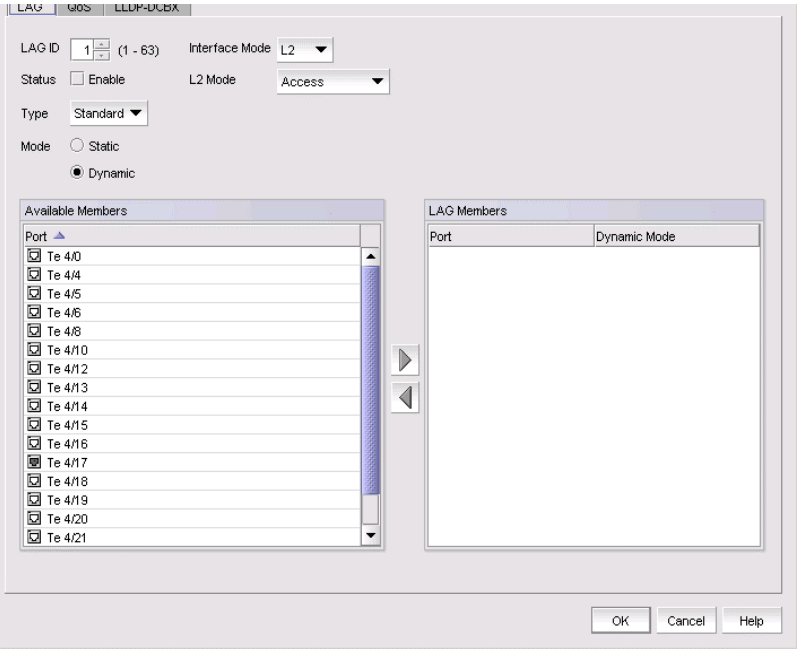

4. Configure the following LAG parameters:

   NOTE
   Ports with 802.1x authentication or ports that are enabled in L2 mode or L3 mode are not supported in a LAG.

- **LAG ID** - Enter the LAG identifier, using a value from 1 through 63. Duplicate LAG IDs are not allowed.
- **Status** - Click the **Enable** check box to enable the LAG. You must enable the LAG to use the DCB functionality.
- **Interface Mode - Select None or L2**. Ports that are in L2 mode cannot be added to a LAG. The L3 interface mode option is displayed in the **Edit LAG** dialog box only.
- **L2 Mode** - Select **Access** or **Trunk**:
  - Access mode allows only one VLAN and allows only untagged frames.
  - Trunk mode allows more than one VLAN association and allows tagged frames.
- **IP/Netmask** - The netmask is used to divide an IP address into subnets. It specifies which portion of the IP address represents the network and which portion represents the host, and can only be configured if the interface mode is L3.
  - **Primary** - The primary IP address assigned to a 10 Gbps DCB/FC switch module.
  - **Secondary** - The secondary IP address is optional. Secondary IP addresses are helpful when the interface port is part of multiple subnets.

5. Select at least one available DCB port from the **Available Members** list and click the right arrow button to move it to the **LAG Members** list.

   The DCB ports are now part of the link aggregation group.

6. Continue to configure the following LAG parameters. These parameters are always enabled.

- **Type** - Sets the limit on the size of the LAG. The type values include Standard, where the LAG is limited to 16 ports, and Brocade LAG, where the LAG is limited to 4 ports. The default is Standard.

  NOTE
  You cannot create Fabric OS-type LAGs from different anvil chips. If you do, an error message displays. Only the first port is considered as part of the LAG.

- **Mode -** Sets all ports added to the LAG members list in either Static or Dynamic mode. The default is Dynamic, Active, but LAG members can be Active or Passive if the LAG member is Dynamic.

7. When you have finished configuring the policies, click **OK**.

   The **Deploy to LAGs** dialog box displays.

8. Click **OK** after changing the attributes of the current deployment.

   The **Deployment Status** dialog box launches.

9. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected LAG or LAGs.
10. Click **Close** to close the **Deployment Status** dialog box.

Editing a DCB switch

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select the DCB switch from the **Products/Ports** list.
3. Click **Edit**.

   The **Edit Switch** dialog box displays (Figure 225).

   FIGURE 225 Edit Switch dialog box

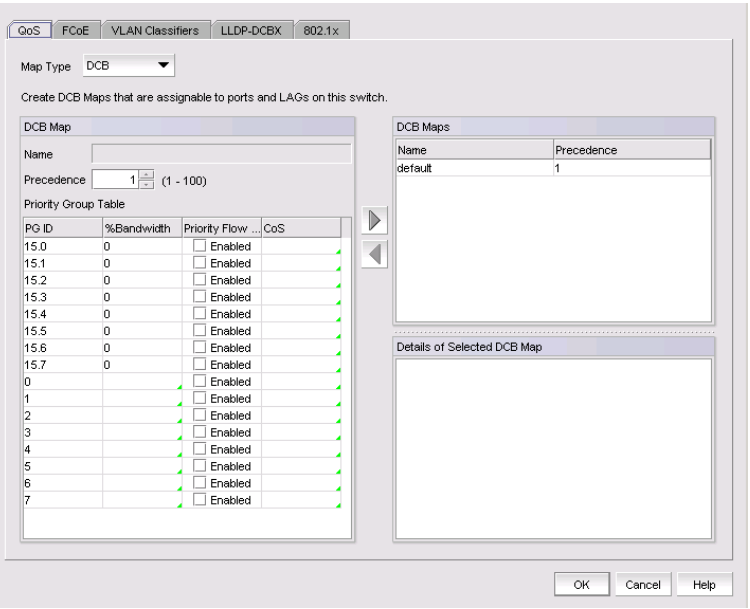

4. Configure the policies for the **Edit Switch** dialog box tabs, which are described in the following sections:
   - "QoS configuration" on page 524
   - "FCoE provisioning" on page 530
   - "VLAN classifier configuration" on page 532
   - "LLDP-DCBX configuration" on page 535

- "802.1x authentication" on page 539

5. When you have finished configuring the policies, apply the settings to the switch.

   NOTE
   Clicking **Cancel** when there are pending changes launches a pop-up dialog box.

6. Click **OK**.

   The **Deploy to Products** dialog box displays.

7. Click **OK** after changing the attributes of the current deployment.

   The **Deployment Status** dialog box launches.

8. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected devices.
9. Click **Close** to close the **Deployment Status** dialog box.

Editing a DCB port

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a DCB port from the **Products/Ports** list.
3. Click **Edit**.

   The **Edit Port** dialog box displays, as shown in Figure 226.

   FIGURE 226 Edit Port dialog box

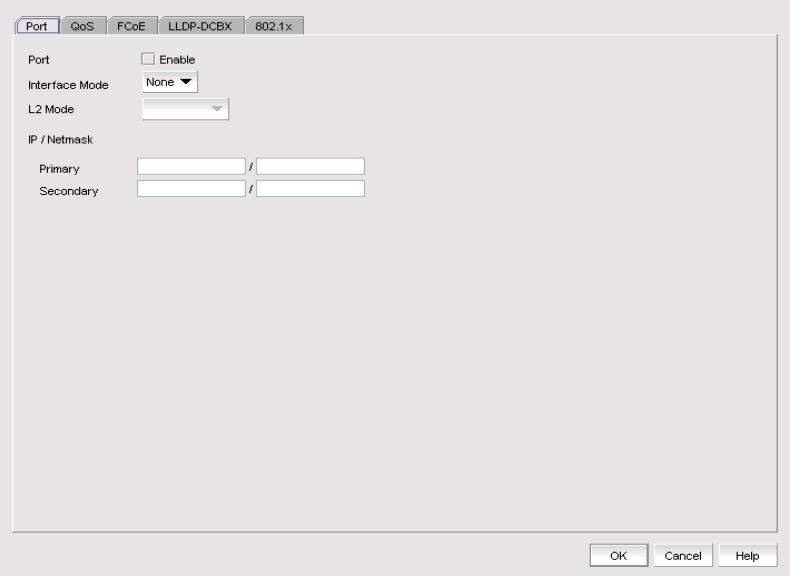

4. Modify the following DCB port parameters as required:

- **Interface Mode** - Select **None** or **L2**. For external ports, the **L3** interface mode displays in addition to **None** or **L2**. If you select **L3** as the interface mode, the **IP/Netmask** field is enabled and you can then assign the primary and secondary IP addresses.
  - **L2 mode** is enabled if you select L2 as the interface mode. If a DCB port is enabled on the 10 Gbps DCB/FC switch module, the L2 mode is disabled.
  - L3 mode appears only for the external ports of embedded platforms.

    NOTE
    You can change the interface mode from **L2** to **None** only if the port is assigned to the default VLAN 1.

- **IP/Netmask** - The netmask is used to divide an IP address into subnets. It specifies which portion of the IP address represents the network and which portion represents the host, and can only be configured if the interface mode is L3.
  - **Primary** - The primary IP address assigned to a 10 Gbps DCB/FC switch module.
  - **Secondary** - The secondary IP address is optional. Secondary IP addresses are helpful when the interface port is part of multiple subnets.

5. When you have finished configuring the policies, apply the settings to the DCB port.

   NOTE
   Clicking **Cancel** when there are pending changes launches a pop-up dialog box.

6. Click **OK when you have finished modifying the DCB port parameters**.

   The **Deploy to Ports dialog box** displays.

7. Click **OK** after changing the attributes of the current deployment.

   The **Deployment Status** dialog box launches.

8. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected port or ports.
9. Click **Close** to close the **Deployment Status** dialog box.

Editing a LAG

Use the following procedure to change members and policies in a link aggregation group (LAG).

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select the link aggregation group (LAG) from the **Products/Ports** list.
3. Click **Edit**.

   The **Edit LAG** dialog box displays, as shown in Figure 227.

FIGURE 227 Edit LAG dialog box

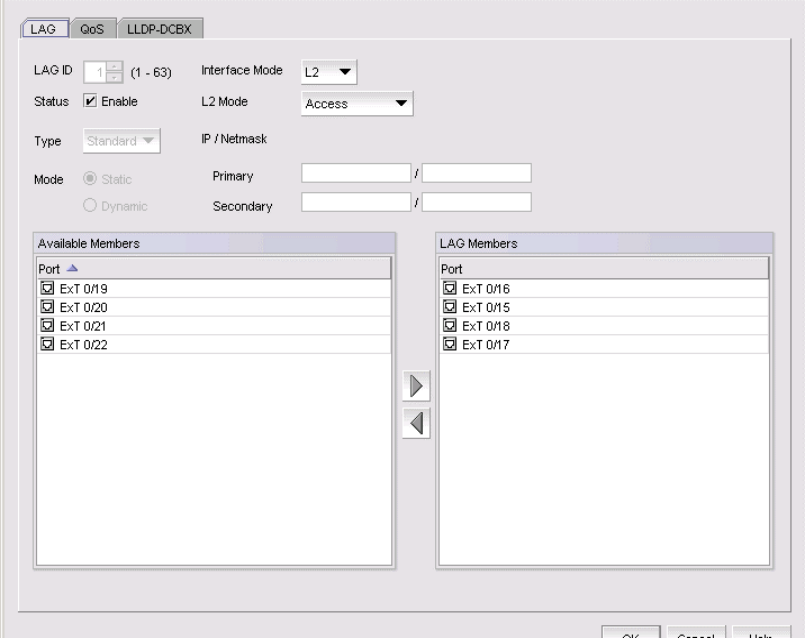

4. Configure the following LAG parameters, as required:

   NOTE
   Ports with 802.1x authentication or ports that are enabled in L2 mode or L3 mode are not supported in a LAG.

- **LAG ID** - The LAG identifier, which is not an editable field.
- **Status** - Click the **Enable** check box to enable the LAG. You must enable the LAG to use the DCB functionality.
- **Interface Mode** - Select **None** or **L2**. For external ports, the **L3** interface mode displays, in addition to **None** or **L2**. If you select **L3** as the interface mode, the **IP/Netmask** field is enabled and you can then assign the primary and secondary IP addresses.
  - A port must be in non-L2 mode if you are adding the port as a member of a LAG.
  - You cannot change the interface mode from **L2** to **None** if the LAG is assigned to a VLAN.
- **L2 Mode** - Select **Access** or **Trunk.**
  - Access mode allows only one VLAN and allows only untagged frames.
  - Trunk mode allows more than one VLAN association and allows tagged frames.
- **IP/Netmask** - The netmask is used to divide an IP address into subnets. It specifies which portion of the IP address represents the network and which portion represents the host, and can only be configured if the interface mode is L3. Primary and secondary IP address fields are applicable only to the external ports and the interface mode must be L3 to enable these fields.
  - **Primary** - Enter the primary IP address assigned to an L3 port.
  - **Secondary** - Enter the secondary IP address (optional). Multiple (secondary) IP addresses help when the interface and port are part of multiple subnets.

5. Continue to configure the following LAG parameters. These parameters are disabled until you add a DCB port to the **LAG Members** list.

- **Mode -** The ports that are LAG members are in either Static or Dynamic mode. You cannot change the mode on existing members of a LAG.

  If the mode is set as **Dynamic**, you can change the dynamic mode type (to Active or Passive) only for newly-added ports, not for existing port members of a LAG.
- **Type** - The type value options are **Standard**, where the LAG is limited to 16 ports, and **Brocade**, where the LAG is limited to four ports. The default is **Standard**. The type is set when you add a LAG; you cannot edit the type using the **Edit LAG** dialog box.

6. Click **OK**.

   The **Deploy to LAGs** dialog box displays.

7. Click **OK** after changing the attributes of the current deployment.

   The **Deployment Status** dialog box displays.

8. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected LAG or LAGs.

   NOTE
   If the primary or secondary IP address already exists on another interface, an error message displays in the **Status** area.

9. Click **Close** to close the **Deployment Status** dialog box.

Enabling a DCB port or LAG

If you select multiple switches or multiple ports and LAGs from two or more switches, both the **Enable** button and the **Disable button** are disabled.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select one or more DCB ports or LAGs (which can span multiple switches) that you want to enable.

   NOTE
   All selected LAGs must be in the same state (enabled or disabled); otherwise, both the **Enable** and **Disable** buttons are disabled.

3. Click **Enable**.

   The **Confirmation and Status** dialog box launches with the selected ports or LAGs.

4. Click **Start on the Confirmation and Status dialog box to save the changes to the selected ports or LAGs**.

   The selected DCB ports or LAGs are enabled in the **DCB Configuration** dialog box.

5. Click **Close** to close the **Confirmation and Status** dialog box.

Deleting a LAG

You can only delete a link aggregation group (LAG) that is selected from a single switch. If you select multiple switches or multiple ports from two or more switches, the **Delete** button is disabled.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select one or more LAGs (that can span multiple switches) that you want to delete from the **Products/Ports** list.
3. Click **Delete**.

   The **Confirmation and Status** dialog box launches with the selected LAGs.
4. Click **Start on the Confirmation and Status dialog box to save the changes to the DCB switches**.

   The selected LAGs are deleted in the **DCB Configuration** dialog box.
5. Click **Close** to close the **Confirmation and Status** dialog box.

QoS configuration

QoS configuration involves configuring packet classification, mapping the priority and traffic class, controlling congestion, and scheduling. The configuration of these QoS entities consists of DCB Map and Traffic Class Map configuration.

In a Data Center Bridging (DCB) configuration, Enhanced Transmission Selection (ETS) and priority-based flow control (PFC) are configured by utilizing a priority table, a priority group table, and a priority traffic table. The Traffic Class map is the mapping of user priority to traffic class.

Priority-based flow control

Priority-based flow control (PFC) is an enhancement to the existing pause mechanism in Ethernet. PFC creates eight separate virtual links on the physical link and allows any of these links to be paused and restarted independently, enabling the network to create a no-drop Class of Service (CoS) for an individual virtual link.

Table 54 shows examples of how priority grouping might be allocated in a 15-priority group scenario.

TABLE 54 Priority grouping allocated in a 15-priority group example

Priority group ID	Bandwidth (%)	Priority flow control
0	55	on
1	25	on
2	0	off
3	0	off
4	5	off
5	0	off
6	15	on
7	0	off
15.0-15.7	Strict priority No bandwidth % configuration allowed	on

Creating a DCB map

The procedure in this section applies only for Fabric OS versions earlier than Fabric OS 7.0.

When you create a DCB map, each of the Class of Service (CoS) options (0-7) must be mapped to at least one of the Priority Group IDs (0-7) and the total bandwidth percentage must equal 100. All QoS, DCB map, and Traffic Class map configurations apply to all ports in a LAG.

There can be, at the most, 16 entries in the Priority Group table. Eight of the entries are Strict Priority entries with a Priority Group ID (15.0-15.7) and eight are user-definable entries with a Priority Group ID of 0-7. Refer to Table 54 for an example of priority group configuration.

NOTE
The 10 Gbps DCB/FC switch module can have only one DCB map.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch, and click **Edit**.
3. Click the **QoS** tab on the **Edit Switch** dialog box.

   The **QoS** dialog box displays, as shown in Figure 228.

FIGURE 228 QoS, Create DCB Map dialog box

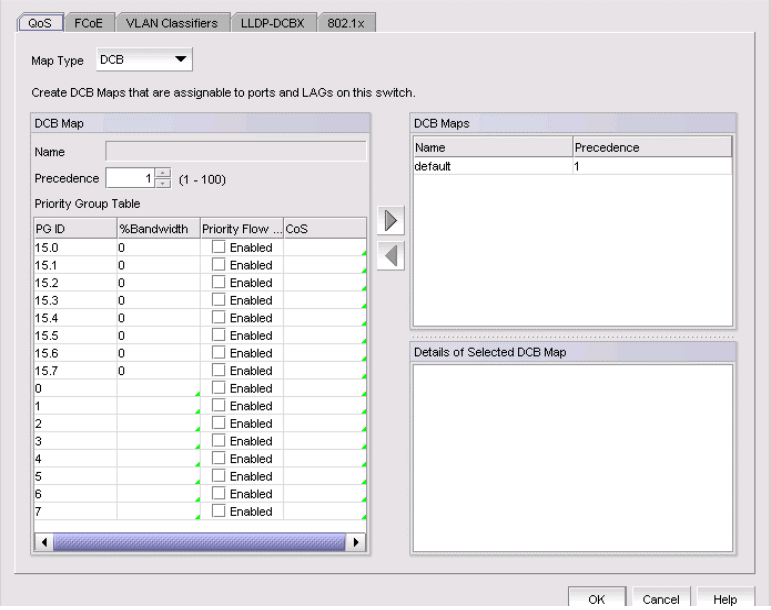

4. Select **DCB** from the **Map Type** list.
5. Configure the following DCB map parameters in the **DCB Map** area:

   - **Name** - Enter a name to identify the DCB map.

     NOTE
     Only one DCB map (the default) is supported on Fabric OS version 6.3.1_dcb and version 7.0.0 and later.

   - **Precedence** - Enter a value from 1 through 100. This number determines the map's priority.

- **Priority Flow Control** check box - Check to enable priority-based flow control on individual priority groups.
- **CoS** - Click the CoS cell to launch the **Edit CoS** dialog box, where you can select and assign one or more priorities (PG ID 15.0 through 15.7).

  All of the eight CoS values (0-7) must be used in a DCB map, separated with a comma and a space. Duplicate CoS values in two or more priority groups are not allowed.

  NOTE
  You can only edit CoS fields that are displayed with a green tick mark.

  % Bandwidth *(optional)* - While in the **Edit CoS** dialog box, enter a bandwidth value for priority group (PG) IDs 15.0 through 15.7. You must map each CoS to at least one of the PG IDs.

  Note the following points:

  - You cannot define a bandwidth percentage for strict priorities (PG ID 15.0-15.7). The total bandwidth percentage for PG ID 15.0 through 15.7 must equal 0.
  - If you set a CoS value to one or more of the PG IDs 0-7, you must also enter a non-zero bandwidth percentage. The total bandwidth percentage must equal 100.
  - For PG IDs 0-7 that do not have an assigned CoS value or PFC enabled, the bandwidth percentage must be 0.

6. Click the right arrow button to add the map to the **DCB Maps** list.

   If a DCB map exists with the same name, a validation dialog box launches and you are asked if you want to overwrite the map.

7. Click **OK**.
8. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Editing a DCB map

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch, and click **Edit**.
3. Click the **QoS** tab on the **Edit Switch** dialog box.

   The **QoS** dialog box displays.

4. Select a DCB map from the **DCB Maps** list and click the left arrow button to load its values in the left pane. The fields are now editable.
5. Keep the same DCB map name and modify the following values, as required. Refer to Table 54 for an example of priority group configuration.

   - **Name** - Enter a name to identify the DCB map.
   - **Precedence** - Enter a value from 1 through 100. This number determines the map's priority.
   - **% Bandwidth** - Enter a bandwidth value for priority group IDs 0-7. The total of all priority groups must equal 100 percent.
   - **Priority Flow Control** check box - Check to enable priority flow control on individual priority groups.
   - **CoS** - Click the CoS cell to launch the **Edit CoS** dialog box, where you can select and assign one or more priorities (PG ID 15.0 through 15.7).

6. Click the right arrow button to re-add the map to the **DCB Maps** list.

If the DCB map already exists, an overwrite message displays.

7. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Deleting a DCB map

You cannot delete the DCB map of a 10 Gbps DCB/FC switch module. To delete the DCB map of an 8 Gbps DCB switch, complete the following steps.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch, and click **Edit**.
3. Click the **QoS** tab on the **Edit Switch** dialog box.

   The **QoS** dialog box displays.

4. Select one or more DCB maps.
5. Click the left arrow button.

   The selected DCB map row is removed from the list.

6. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

   NOTE
   With Fabric OS version 7.0 and later, there is only one DCB map (default), that you cannot delete.

7. Click **OK** after changing the attributes of the current deployment.

   The **Deployment Status** dialog box displays.

8. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected devices.
9. Click **Close** to close the **Deployment Status** dialog box.

Assigning a DCB map to a port or link aggregation group

The **Edit Port** dialog box - **QoS** tab allows you to assign DCB maps to ports and LAGs on a selected switch.

NOTE
QoS maps are created using the **Edit Switch** dialog box, accessible from the **DCB Configuration** dialog box.

A port can have either a DCB map or a Traffic Class map assigned to it, but it cannot have both.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a port or LAG, and click **Edit**.
3. Click the **QoS** tab on the **Edit Port or Edit LAG** dialog box.

   The **QoS** dialog box displays.

4. Click the **Assign a map to <*device_name*>** check box to assign the selected port to a DCB map.

If you do not select this check box, all QoS edit features are disabled.

5. Select **DCB Map** in the **Map Type** list.
6. Select a DCB map in the **Available DCB Maps** list.

   If no DCB maps were created on the switch, the **Available DCB Maps** list is empty. Otherwise, the following DCB map details display:

   - **PG - ID** — Lists the priority group ID (15.0 through 15.7 and 0 through 7).
   - **% Bandwidth** — Lists the bandwidth value for priority group IDs 0-7. The total of all priority groups must equal 100 percent.
   - **Priority Flow** checkbox — Check to enable priority-based flow control on individual priority groups.
   - **CoS** — Lists the Class of Service (CoS) value that corresponds to the priority group ID rows. The CoS value must be mapped to at least one of the priority group IDs (0-7).
7. When you have finished the configuration, click **OK** to launch the **Deploy to Ports/LAGs** dialog box.

Creating a Traffic Class map

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.
2. Select a switch, and click **Edit**.
3. Click the **QoS** tab on the **Edit Switch** dialog box.

   The **QoS** dialog box displays.
4. Select **Traffic Class** from the **Map Type** list.
5. Name the Traffic Class map.
6. Click the Traffic Class cell in a CoS row and directly enter a value from 0-7. You can leave the cell empty to indicate zero (0).
7. Click the right arrow button to add the map to the **Traffic Class Maps** list.

   If the name of the Traffic Class map already exists, an overwrite warning message displays. Click **Yes** to overwrite the existing Traffic Class map.
8. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Editing a Traffic Class map

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.
2. Select a switch, and click **Edit**.
3. Click the **QoS** tab on the **Edit Switch** dialog box.

   The **QoS** dialog box displays.
4. Select a Traffic Class map from the **Traffic Class Maps** list and click the left arrow button to load its values in the left pane. The fields are now editable.

If the name of the Traffic Class map already exists, an overwrite warning message displays. Click **Yes** to overwrite the existing Traffic Class map.

5. Keep the same Traffic Class map name and modify the values, as required.
6. Click the right arrow button to re-add the map to the **Traffic Class Maps** list.
7. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Deleting a Traffic Class map

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch, and click **Edit**.
3. Click the **QoS** tab on the **Edit Switch** dialog box.

   The **QoS** dialog box displays.

4. Select a Traffic Class map that you want to delete from the **Traffic Class Maps** list.
5. Click the left arrow button.

   The selected Traffic Class map row is removed from the list.

6. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.
7. Click **OK** after changing the attributes of the current deployment.

   The **Deployment Status** dialog box displays.

8. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected devices.

Assigning a Traffic Class map to a port or link aggregation group

You can assign a Traffic Class map to a port or ports under the LAG; however, a port does not require a Traffic Class map be assigned to it. A port can have either a DCB map or a Traffic Class map assigned to it, but it cannot have both.

NOTE
You cannot configure QoS or LLDP-DCBX on a LAG.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a port or LAG, and click **Edit**.
3. Click the **QoS** tab on the **Edit Port or Edit LAG** dialog box.

   The **QoS** dialog box displays.

4. Click the **Assign a map** check box.
5. Select **Traffic Class** in the **Map Type** list.
6. Select a Traffic Class map in the **Traffic Class Map** list.

7. When you have finished the configuration, click **OK** to launch the **Deploy to Ports/LAGs** dialog box. Refer to "Switch, port, and LAG deployment" on page 541 for more information.

FCoE provisioning

The Management application supports FCoE provisioning only on Fabric OS version 6.3.1_dcb.

The command line interface (CLI) supports FCoE provisioning for the following versions of Fabric OS:

- Fabric OS 6.3.1_cee
- Fabric OS 6.3.1_del
- Fabric OS 6.4.1_fcoe
- Fabric OS 7.0.x

Refer to the *Fabric OS Command Reference* for CLI procedures.

FCoE provisioning simplifies the number of steps required to configure a DCB port to carry the FCoE traffic. The FCoE map contains the default DCB map and the VLAN ID. You can change the default VLAN ID using the **FCoE** tab of the **Edit Switch** dialog box, shown in Figure 225.

NOTE
For FOS DCB switches, the default DCB map associated with the default FCoE map can be edited on the switch from the **Edit Switch** dialog box - **QoS** tab.

Changing the VLAN ID on the default FCoE map

You can change the VLAN ID on the default FCoE map only when no ports or LAGs are participating as members of the switch. You must first manually remove the FCoE map option for each of the port members before you change the VLAN ID on the switch.

NOTE
You can complete this procedure using the Management application on embedded platforms such as the Fabric OS converged 10 GbE switch module for the IBM BladeCenter or the Dell M8428-k switch. You cannot perform this task on the Fabric OS Fabric OS switch or the FCOE10-24 port blade.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch and click **Edit**.
3. Click the **FCoE tab** on the **Edit Switch** dialog box.

   The **Edit Switch** dialog box, **FCoE** tab displays the following FCoE map parameters:

   NOTE
   The **FCoE** tab does not display for the Fabric OS Fabric OS switch or the FCOE10-24 port blade.

   - **Name** — The name of the FCoE map that will be available for assignment to ports on this switch. This is a read-only field.
   - **VLAN ID** — Enter an FCoE VLAN identifier to associate with the FCoE map. The values range from 2 through 3583, and 1002 is the default.
   - **DCB Map** — The DCB map that is associated with the FCoE map. This is a read-only field.
4. Accept the default VLAN ID of 1002, or change the value. The valid VLAN ID range is from 2 through 3583.

5. Click the right arrow button to move the FCoE map parameters into the **FCoE Maps** list.
6. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.
7. Click **OK** after changing the attributes of the current deployment.

   The **Deployment Status** dialog box displays.
8. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected devices.

Enabling or disabling the FCoE map on the port

You must first manually disable an FCoE map-enabled port if you want to edit the VLAN ID of the FCoE map. Refer to "Changing the VLAN ID on the default FCoE map" on page 530 for information on editing the VLAN ID using the **Edit Switch** dialog box, **FCoE** tab.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.
2. Select a port and click **Edit**.
3. Click the **FCoE tab** on the **Edit Port** dialog box.

   The **Edit Port** dialog box, **FCoE** tab displays the following parameters:

   - **FCoE Map** field — Displays the name of the FCoE map (read-only).
   - **VLAN ID** list — The FCoE VLAN identifier associated with the FCoE map. The values range from 2 through 3583, and 1002 is the default.
   - **DCB Map —** Displays the name of the DCB map (read-only).
   - Details of selected DCB Map list:
     - **PG - ID** — Lists the priority group ID (15.0 through 15.7 and 0 through 7)
     - **% Bandwidth** — Lists the bandwidth value for priority group IDs 0-7. The total of all priority groups must equal 100 percent.
     - **Priority Flow** check box — Check to enable priority-based flow control on individual priority groups.
     - **CoS** — Lists the Class of Service (CoS) value that corresponds to the priority group ID rows. The CoS value must be mapped to at least one of the priority group IDs (0-7).
4. If enabled, click the **Enable FCoE** check box to disable the port's membership on the FCoE map.
5. When you have finished the configuration, click **OK** to launch the **Deploy to Ports** dialog box.
6. Click **OK** after changing the attributes of the current deployment.

   The **Deployment Status** dialog box displays.
7. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected devices.

VLAN classifier configuration

The Management application supports VLAN classifier management only on Fabric OS 6.3.1_dcb and Fabric OS 7.0.0.

VLAN classifier rules are used to define specific rules for classifying untagged packets to selected VLANs based on protocol and MAC addresses. The classified frames are then tagged with a VLAN ID.

VLAN classifier rules can be categorized into the following areas:

- 802.1Q protocol-based classifier rules
- MAC address-based classifier rules

VLAN classifiers are created on a per-switch basis.

NOTE
The **VLAN Classifiers** tab on the **Edit Switch** dialog box displays only on switches with Fabric OS versions 7.0.0 and later.

Adding a VLAN classifier rule

The **Edit Switch** dialog box, **VLAN Classifiers** tab allows you to create rules and group them into VLAN classifiers, which can then be applied to access port and LAG VLAN members and converged port VLAN members.

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch and click **Edit**.
3. Click the **VLAN Classifiers tab** on the **Edit Switch** dialog box.

   The **Edit Switch** dialog box, **VLAN Classifiers** tab displays, as shown in Figure 229. The **Available Rules** list contains the following information:

   - **Rule ID** — The rule identifier. Valid rule ID values are from 1 through 256.
   - **Rule Type** — Valid rule types are **MAC** (MAC address-based rule) and Proto (802.1Q protocol-based rule).
   - **Encapsulation** — The encapsulation type (Ethv2, nosnaplic, or snaplic). The **Encapsulation** column only displays a value when Proto is the rule type.

FIGURE 229 Edit Switch dialog box, VLAN Classifiers tab

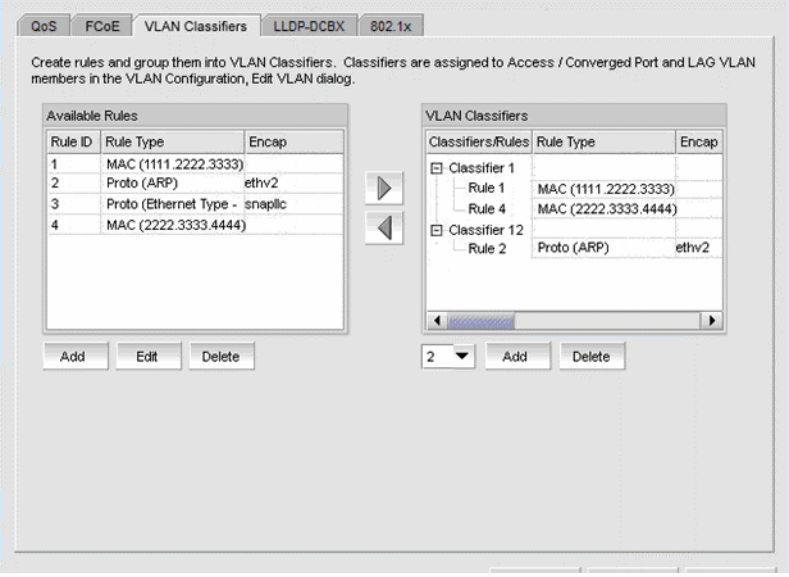

4. Click the **Add** button under the **Available Rules** list.

   The **Add Rules** dialog box displays, as shown in Figure 230.

FIGURE 230 Add Rules dialog box

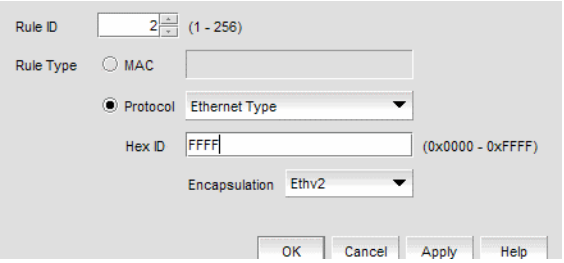

   The **Rule ID** field is pre-populated with the next available rule ID number.

5. Keep the rule ID number as it is, or change the number using a value from 1 through 256.
6. Select a rule type. Valid rule types are **MAC** (MAC address-based rule) and Proto (802.1Q protocol-based rule).
7. If **Ethernet Type** is selected as the protocol rule type, enter any valid four-digit hexadecimal value within the allowed range of 0x0000 through 0xFFFF. For the other Proto options, the hex ID value is hard-coded as follows:
   - ARP — 0x0808
   - IP — 0x8881
   - IPv6 — 0x86DD
8. Select an encapsulation type from the list. Options include Ethv2, nosnapllc, and snapllc. The **Encapsulation** list only accepts a value when **Protocol** is selected as the rule type.
9. Click **OK** to add the rule to the **Available Rules** list on the **VLAN Classifiers** tab of the **Edit Switch** dialog box and close the **Add Rules** dialog box.

NOTE
Clicking **Apply** also adds the rule to the **Available Rules** list on the **VLAN Classifiers** tab of the **Edit Switch** dialog box, and in addition, the **Add Rules** dialog box remains open and clears all entries for you to define the next rule.

10. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Editing a VLAN classifier rule

1. From the **VLAN Classifiers** tab of the **Edit Switch** dialog box, select a row in the **Available Rules** list and click **Edit**.

   The **Edit Rules** dialog box displays with the fields pre-populated with the rule details. The **Rule ID** field is disabled.

2. Select a rule type. Valid rule types are **MAC** (MAC address-based rule) and Proto (802.1q protocol-based rule).
3. If Ethernet is selected as the protocol-based rule type, enter any valid four-digit hexadecimal value within the allowed range of 0x0000 through 0xFFFF. For the other Proto options, the hex ID value is hard-coded as follows:
   - ARP — 0x0808
   - IP — 0x8881
   - IPv6 — 0x86DD
4. Select an encapsulation type from the list. Options include Ethv2, nosnapllc, and snapllc. The **Encapsulation** list only accepts a value when Protocol is selected as the rule type.
5. Click **OK** to add the edited rule to the **Available Rules** list on the **VLAN Classifiers** tab of the **Edit Switch** dialog box and close the **Edit Rules** dialog box.
6. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Deleting a VLAN classifier rule

1. From the **VLAN Classifiers** tab of the **Edit Switch** dialog box, select a row in the **Available Rules** list and click **Delete**.

   A message displays if the rules are participating in VLAN classifier groups that are currently associated with VLAN port or LAG members.

2. Click **Yes** to remove the selected rule row from the list.
3. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Creating a VLAN classifier group

You can assign existing rules to a selected VLAN classifier and form a VLAN classifier group. If no rules are available, you can add rules to a selected switch using the **Add Rules** dialog box.

1. Select **Configure > DCB** from the menu bar.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch and click **Edit**.
3. Click the **VLAN Classifiers tab** on the **Edit Switch** dialog box.

   The **Edit Switch** dialog box, **VLAN Classifiers** tab displays.

4. Select a classifier ID from the **VLAN Classifier** list. Values range from 1 through 16.

5. Click the **Add** button under the **VLAN Classifie**r list.

   The classifier with the selected ID is displayed in the **VLAN Classifier** list.

6. Select the classifier from the **VLAN Classifier** list and then select the rules you want to add under this classifier from the **VLAN Classifier Rules** list.

   - If no rules are available, the following error message displays: "No rules are available on this switch. Choose **Add** under the **Available Rules** list to add rules to this switch."
   - If no classifier group IDs are available, the list is disabled.

7. Click the right arrow button.

   The selected rules are assigned to the selected VLAN classifier ID in the **VLAN Classifier** list.

8. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Deleting a VLAN classifier group

1. Click the **VLAN Classifiers tab** on the **Edit Switch** dialog box.

   The **Edit Switch** dialog box, **VLAN Classifiers** tab displays.

2. Select a classifier from the **VLAN Classifiers** list.
3. Click **Delete**.

   The VLAN classifier group is deleted.

4. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

LLDP-DCBX configuration

Link Layer Discovery Protocol (LLDP) provides a solution for the configuration issues caused by increasing numbers and types of network devices in a LAN environment, because, with LLDP, you can statically monitor and configure each device on a network.

Data Center Bridging Exchange (DCBX) protocol enables Enhanced Ethernet devices to discover whether a peer device supports particular features, such as Priority-based Flow Control (PFC) or Class of Service (CoS). In a Data Center Bridging (DCB) environment, LLDP is enhanced with DCBX protocol to further share or change the configured DCB enhancements. You must enable the DCBX protocol and configure certain parameters in order to effectively utilize the benefits of a converged network.

Using the **LLDP-DCBX** dialog box, you can create and manage LLDP profiles and assign an LLDP profile to a port or link aggregation group (LAG).

Configuring LLDP for FCoE

To configure LLDP for FCoE, complete the following steps.

> NOTE
> When a TE port is selected to assign to an LLDP profile, a yellow banner displays with the following error message: "LLDP-DCBX is disabled on this switch. The configuration becomes functional when LLDP-DCBX is enabled on the switch."

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays.

2. Select the switch to edit in the **DCB Ports and LAGs** list and click **Edit**.

   The **Edit Switch** dialog box displays, as shown in Figure 231.

3. Click the **LLDP-DCBX** tab.

FIGURE 231 Edit Switch dialog box - LLDP-DCBX tab

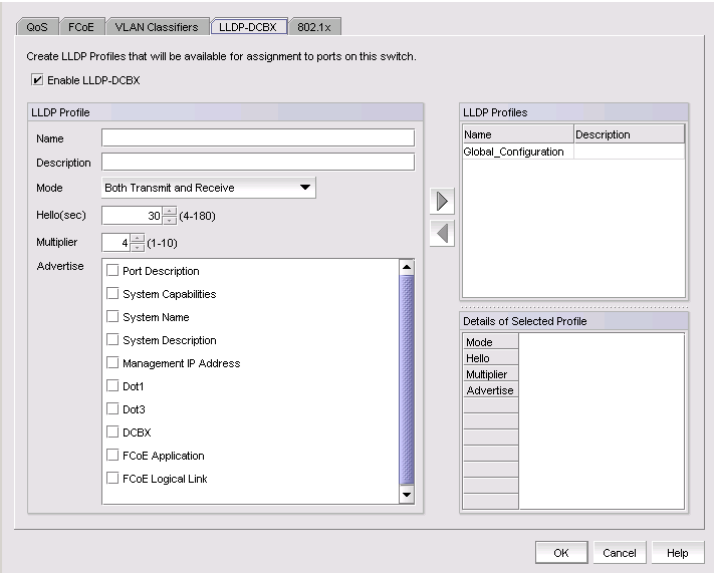

Adding an LLDP profile

NOTE
When a TE port is selected to assign to an LLDP profile, a yellow banner displays with the following error message: "LLDP-DCBX is disabled on this switch. The configuration becomes functional when LLDP-DCBX is enabled on the switch."

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch, and click **Edit**.
3. Click the **LLDP-DCBX** tab on the **Edit Switch** dialog box.

   The **LLDP-DCBX** dialog box displays.

4. Click the **Enable LLDP-DCBX** checkbox.
5. Configure the LLDP Profile parameters:

   - **Enter** a name for the LLDP profile.

     If the name of the LLDP profile already exists on the switch, an overwrite warning displays.

   - **Enter a** meaningful description of the LLDP profile.
   - Select a mode from the list: Both Tx (transmitted) or Rx (received), Tx only, or Rx only.
   - Enter a hello interval time (in seconds) for the bridge in the **Hello (secs)** field. The value range is from 4 through 180 and the default value is 30.

- Enter a multiplier (in seconds). The value range is from 1 through 10 and the default is 4.
- Check the profile parameters that you want to display as part of the LLDP profile from the **Advertise** list:
  - Port description - The user-configured port description.
  - System name - The user-configured name of the local system.
  - System capabilities - The system capabilities running on the system.
  - System description - The system description containing information about the software running on the system.
  - Management IP address - The management IP address of the local system.
  - Dot *x*
  - DCBX - The DCBX profiles.
  - FCoE application - The FCoE application feature.
  - FCoE logical link - The logical link level for the SAN network.

6. Click the right arrow button to move the newly created profile into the **LLDP Profiles** list.
7. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Editing an LLDP profile

1. Select **Configure > DCB**.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch, and click **Edit**.
3. Click the **LLDP-DCBX** tab on the **Edit Switch** dialog box.

   The **LLDP-DCBX Profile** dialog box displays.

4. Select an LLDP profile in the **LLDP Profile** list.

   NOTE
   You can edit the <Global Configuration> profile. You cannot, however, delete or duplicate global configurations.

5. Click the left arrow to load the LLDP profile's values in the left pane.
6. Modify the values, as described in "Adding an LLDP profile" on page 536. You are not allowed to modify the LLDP profile's name.
7. Click the right arrow to update the LLDP profile parameters.
8. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Deleting an LLDP profile

1. Select **Configure > DCB** from the menu bar.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch, and click **Edit**.
3. Click the **LLDP-DCBX** tab on the **Edit Switch** dialog box.
4. Select an existing LLDP profile from the **LLDP Profiles** list in the upper right pane.

NOTE
You cannot delete <Global Configurations>. You can, however, edit global configurations. For more information, refer to "Product configuration templates" on page 1456.

5. Click the left arrow button.

   The selected LLDP profile is removed from the list.

6. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.
7. Click **OK**.

   The **Deployment Status** dialog box launches.

8. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected devices.
9. Click **Close** to close the **Deployment Status** dialog box.

Assigning an LLDP profile to a port or ports in a LAG

You create LLDP profiles using the **Edit Switch** dialog box, which you access from the **DCB Configuration** dialog box. Global configuration parameters, which is the default selection, are displayed in the Assigned Profile table.

NOTE
A yellow banner displayed on the **LLDP-DCBX** dialog box indicates that LLDP-DCBX is disabled on the switch. The configuration options become functional when LLDP-DCBX is enabled on the switch.

1. Select **Configure > DCB** from the menu bar.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a port or link aggregation group (LAG), and click **Edit**.
3. Click the **LLDP-DCBX** tab on the **Edit Port or Edit LAG** dialog box.

   The **Assign an LLDP profile** dialog box displays.

4. Click **Assign an LLDP profile to <port name>** button to enable the feature.

   NOTE
   Assign the Global Configuration is the default. The **Available Profiles** list is disabled if global configuration is selected. In addition, the **Assign an LLDP profile** button is disabled if no LLDP profiles exist on the switch.

5. Select an LLDP profile from the **Available Profiles** list.
6. When you have finished the configuration, click **OK** to launch the **Deploy to Ports/LAGs** dialog box. Refer to "Switch, port, and LAG deployment" on page 541 for more information.

802.1x authentication

802.1x is a standard authentication protocol that defines a client-server-based access control and authentication protocol. 802.1x restricts unknown or unauthorized clients from connecting to a LAN through publicly accessible ports.

NOTE
802.1x is not supported for internal ports.

A switch must be enabled for 802.1x authentication before you configure its parameters. See "Setting 802.1x parameters for a port" for more information.

Enabling 802.1x authentication

802.1x authentication is enabled or disabled globally on the switch using the **Edit Switch** dialog box.

1. Select **Configure > DCB** from the menu bar.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch and click **Edit**.
3. Click the 802.1x tab on the **Edit Switch** dialog box.
4. Click the **Enable 802.1x** check box to enable 802.1x authentication, and click **OK**.
5. Configure the 802.1x parameters, which are described in "Setting 802.1x parameters for a port" on page 539.
6. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Disabling 802.1x authentication

1. Select **Configure > DCB** from the menu bar.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch and click **Edit**.
3. Click the 802.1x tab on the **Edit Switch** dialog box.
4. Clear the **Enable 802.1x** check box to disable 802.1x authentication.
5. When you have finished the configuration, click **OK** to launch the **Deploy to Products** dialog box.

Setting 802.1x parameters for a port

The 802.1x parameters can be configured whether or not the feature is enabled on the switch. The default parameters are initially populated when 802.1x is enabled, but you can change the default values as required.

1. Select **Configure > DCB** from the menu bar.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a port and click **Edit**.
3. Click the 802.1x tab on the **Edit Port** dialog box.

   The **Enable 802.1x** dialog box displays, as shown in Figure 232.

4. Click the **Enable 802.1x** check box to enable 802.1x authentication.

   The 802.1x parameters are enabled for editing.

   FIGURE 232 802.1x dialog box

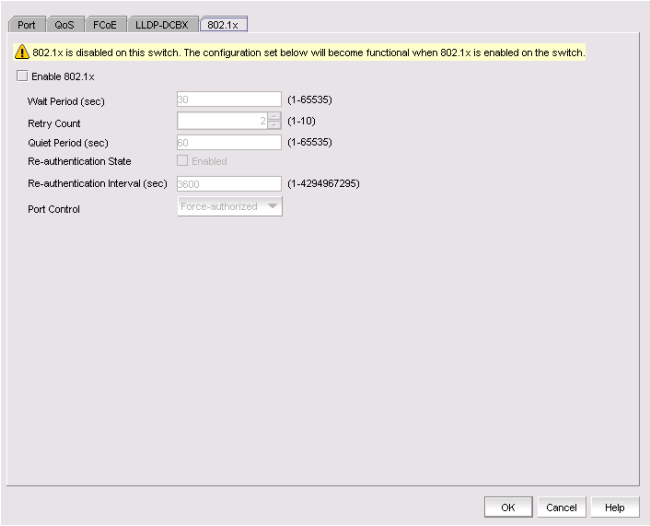

5. Configure the following 802.1x parameters:

   - **Wait Period** - The number of seconds the switch waits before sending an EAP request. The value range is 15 to 65535 seconds. The default value is 30.
   - **Retry Count** - The maximum number of times that the switch restarts the authentication process before setting the switch to an unauthorized state. The value range is 1 to 10. The default value is 2.
   - **Quiet Period** - The number of seconds that the switch remains in the quiet state after a failed authentication exchange with the client. The value range is 1 to 65535 seconds. The default value is 60.
   - **Re-authentication State** - Enable or disable the periodic re-authentication of the client. The default is Disable.
   - **Re-authentication Interval** - The number of seconds between re-authentication attempts. The value range is 1 to 4294967295. The default value is 3600 seconds. This feature is not dependent on the re-authentication state being enabled.
   - **Port Control** - Select an authorization mode from the list to configure the ports for authorization. Options include auto, force-authorized, or force-unauthorized and the default value is auto.

6. When you have finished the configuration, click **OK** to launch the **Deploy to Ports** dialog box. Refer to "Switch, port, and LAG deployment" on page 541 for more information.

Switch, port, and LAG deployment

The **Deploy to Products**, **Deploy to Ports**, and **Deploy to LAGs** dialog boxes provide the flexibility to commit DCB configurations either right away or at a scheduled time. These dialog boxes also allow you to commit the switch-level configuration changes to one or more target switches.

NOTE
Deployment from the Management application to a Network OS device is not supported.

Deploying DCB product, port, and LAG configurations

The switch, port, and LAG deployment dialog boxes provide common deployment options, save configuration options, and schedule options. Depending on which product, port, or LAG you select, the **Deploy to Products**, **Deploy to Ports**, or **Deploy to LAGs** dialog box displays upon deployment.

1. Select **Configure > DCB** from the menu bar.

   The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

2. Select a switch, port, or LAG, and click **Edit**.
3. Configure the switch, port, or LAG. When you have finished the configuration, click **OK** to launch the appropriate dialog box. Refer to Figure 290, Figure 291, and Figure 292.

4. Click one of the following deployment options:
   - Deploy now
   - Save and deploy now
   - Save deployment only
   - Schedule
5. Click one of the following save configuration options:
   - Save to running
   - Save to running and startup
   - Save to running and startup then reboot

   The name for the scheduled product deployment is pre-populated with a “DCB-MM-DD-YYYY-HR-MIN-SS” prefix. This is an editable field.
6. Provide a description for the product/port/LAG deployment.
7. If the **Schedule** option is selected, click the **Use** check box for one-time deployment. One-time deployment is the only option.

   The name of the origin product is a read-only field. The origin product receives the entire configuration, unless it is removed from the **Selected Targets** list.
8. Select one or more of the following configurations, to be deployed on the selected targets.

   NOTE
   These configurations can be pushed to target DCB switches, FOS version 6.3.1_cee or 6.3.1_del.

   For switches:

- QoS, DCB Map
- QoS, Traffic Class Map
- FCoE Map
- VLAN Classifiers and Rules
- LLDP Profiles
- 802.1x Configuration

  NOTE
  See "Source to target switch Fabric OS version compatibility for deployment" for restrictions.

For ports:

- Port attributes (interface mode, etc.)
- QoS, DCB Map / Traffic Class Map
- FCoE Map
- LLDP Profiles
- 802.1x Configuration

  NOTE
  On the **Deploy to Ports** dialog box, you can write port configurations to the switch by enabling the check box at the bottom of the dialog box.

For LAGs:

- LAG attributes (Interface Mode, etc.)
- QoS, DCB Map / Traffic Class Map
- LLDP Profiles

9. Click to move the available targets selected for configuration deployment to the **Selected Targets** list.
10. Click **OK**.

    The **Deployment Status** dialog box launches.
11. Click **Start** on the **Deployment Status** dialog box to save the changes to the selected devices.
12. Click **Close** to close the **Deployment Status** dialog box.

Source to target switch Fabric OS version compatibility for deployment

Table 55 lists the restrictions that exist when deploying source switches to target switches.

TABLE 55 Source to target switch Fabric OS version compatibility

Source Fabric OS version and device	Target Fabric OS version supported	Comments
Fabric OS DCB switch and FCOE10-24 DCB blade with Fabric OS version 6.4.2 or earlier.	Allows Fabric OS DCB switch and FCOE10-24 DCB blade with Fabric OS version 6.4.2 or earlier. Excludes Fabric OS Converged 10 Gbe switch module for IBM BladeCenter with Fabric OS 6.3.1_cee, Fabric OS 6.4.1_fcoe, and Fabric OS 6.3.1_dcb.	You cannot copy legacy configurations to Fabric OS version 7.0 switches, because these switches support FCoE maps and can have only one default DCB map. Legacy Fabric OS switches, however, can have more than one default map.
Fabric OS FCOE10-24 DCB blade with Fabric OS 6.4.1_fcoe	Allows FCOE10-24 DCB blade with Fabric OS 6.4.1_fcoe or Fabric OS 7.0.0. Allows Fabric OS Converged 10 Gbe switch module for IBM BladeCenter with Fabric OS 6.3.1_cee or Fabric OS 6.3.1_dcb. Excludes Fabric OS DCB switch and FCOE10-24 DCB blade with Fabric OS 6.4.2 or earlier.	Both the source and the target support only one default DCB map. You can copy QoS, LLDP, and 802.1x configurations from the source to the target.
Fabric OS DCB switch FCOE10-24 DCB blade with Fabric OS 7.0.	Allows Fabric OS DCB switch and FCOE10-24 DCB blade with Fabric OS 7.0.0. Excludes all others.	VLAN classifiers are supported, but the FCoE map is not supported on Fabric OS 7.0.0.
Fabric OS Converged 10 GbE switch module for IBM BladeCenter with Fabric OS 6.3.1_cee and 6.3.1_dcb	Allows Fabric OS Converged 10 Gbe switch module for IBM BladeCenter with Fabric OS 6.3.1_cee, Fabric OS 6.3.1_dcb. Allows Dell M8428-k switch with Fabric OS 6.3.1_dell, Fabric OS 6.3.1_dcb.	Both source and target switches must support the FCoE map and VLAN classifiers.
Dell M8428-k switch with Fabric OS 6.3.1_dell and 6.3.1_dcb	Allows Fabric OS Converged 10 Gbe switch module for IBM BladeCenter with Fabric OS 6.3.1_cee, Fabric OS 6.3.1_dcb. Allows Dell M8428-k switch with Fabric OS 6.3.1_dell, Fabric OS 6.3.1_dcb.	Both source and target switches must support the FCoE map and VLAN classifiers.

DCB performance

Performance monitoring provides details about the quantity of traffic and errors a specific port or device generates on the fabric over a specific time frame. You can also use Performance features to indicate the devices that create the most traffic and to identify the ports that are most congested.

The Performance menu items launch either SAN or IP performance dialog boxes based on which tab you select. Note the following points:

- The DCB configuration dialog box can be launched from either the SAN or IP tab.
- The appropriate IP Performance tab launches depending on whether you selected a port or a switch.

Real-time performance graph

You can monitor a device's performance through a performance graph that displays transmit and receive data. The graphs can be sorted by the column headers. You can create multiple real-time performance graph instances.

Generating a real-time performance graph from the SAN tab

To generate a real-time performance graph for a FOS device, complete the following steps.

1. Click the SAN tab.
2. Select a DCB port from the **DCB Configuration** dialog box, and select **Real Time Graph** from the Performance list.

   A message displays, prompting you to close the **DCB Configuration** dialog box.
3. Click **OK** to close the **DCB Configuration** dialog and open the Performance dialog box.

   The **Real Time Performance Graphs** dialog box displays.

FIGURE 233 Real Time Performance Graphs dialog box - SAN tab

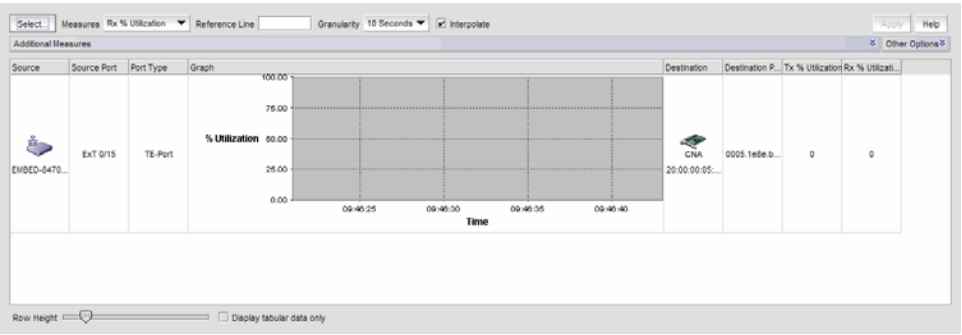

For complete information about Real Time Performance Graphs, refer to "SAN real-time performance data" on page 966.

Historical performance graph

The **Historical Performance Graph** dialog box enables you to customize how you want the historical performance information to display.

Generating a historical performance graph

You can generate a historical performance graph by selecting both Network OS and FOS DCB devices from the IP Tab or by selecting only Network OS DCB devices from the IP tab.

1. Select a DCB port from the **DCB Configuration** dialog box, and select **Historical Graph** from the Performance list.

   A message displays, prompting you to close the **DCB Configuration** dialog.

2. Click **OK** to close the **DCB Configuration** dialog and open the Performance dialog box.

   The **Historical Performance Graph** dialog box displays.

For complete information about Real Time Performance Graphs, refer to "SAN real-time performance data" on page 966.

Historical performance report

The **Historical Performance Report** dialog box enables you to customize how you want the historical performance information to display.

Generating a historical performance report

1. Select a DCB port from the **DCB Configuration** dialog box, and select **Historical Report** from the Performance list.

   A message displays, prompting you to close the **DCB Configuration** dialog box.

2. Click **OK** to close the **DCB Configuration** dialog and open the Performance dialog box.

   The **Historical Performance Report** dialog box displays, as shown in Figure 234.

FIGURE 234 Historical Performance Report dialog box

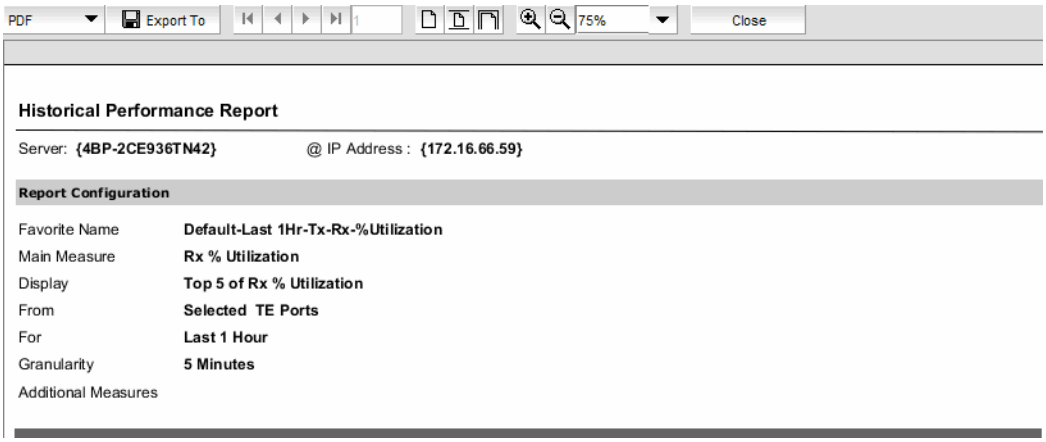

For complete information about Historical Performance Graphs, refer to "SAN historical performance data" on page 970.

FCoE login groups

The FCoE Configuration dialog box allows you to manage the FCoE login configuration parameters on the DCB switches in all discovered fabrics. FCoE login configuration is created and maintained as a fabric-wide configuration.

With the FCoE license, the **FCoE Configuration** dialog box displays virtual FCoE port information and enables you to manage the virtual port information. The topology displays directly connected converged network adapters (CNAs) and the **Properties** dialog box for the virtual FCoE port details.

Without the FCoE license, the virtual FCoE port displays in the device tree, but you cannot enable, disable, or view virtual FCoE port information.

NOTE
FCoE Login Group is not available for Network OS DCB devices.

1. Select **Configure > FCoE** from the menu bar.

   The **FCoE Configuration** dialog box displays all configured login groups and the following details associated with a selected device, shown in Figure 235.

   - FCoE login — Indicates whether the switch is FCoE enabled or disabled.
   - Group Status — Indicates whether the group is active or conflicted.
   - Member Status — Indicates whether the device associated with the group is active or conflicted.
   - Member WWN — Displays the world wide name (WWN) of the device associated with the group.
   - Type — Displays the model type.

   FIGURE 235 FCoE Configuration dialog box

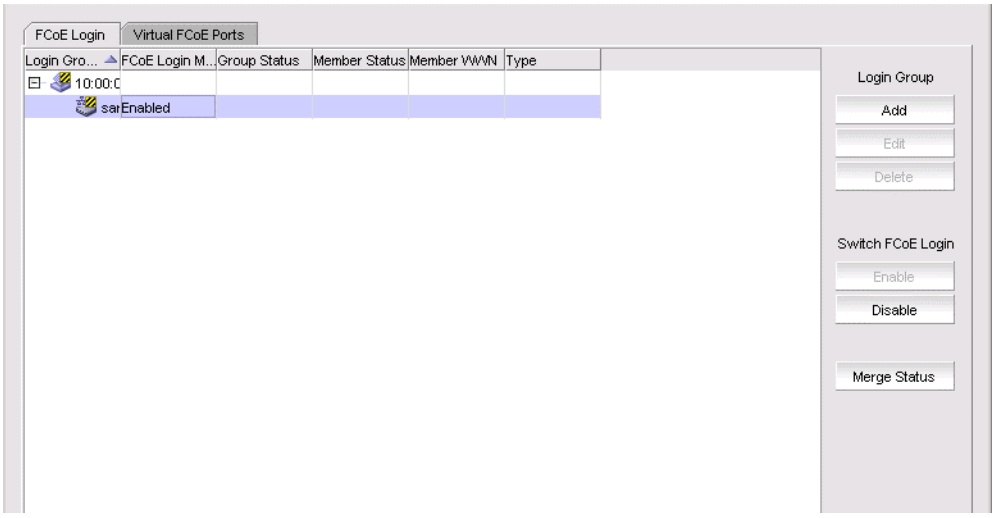

2. Perform one of the following tasks:

   Under Login Group:

   - Click **Add** to launch the Add Login Group dialog box, where you can select an existing switch or enter the WWN of a switch on which the FCoE login group will be created. See "Adding an FCoE login group" on page 547.
   - Click **Edit** to launch the Edit Login Group dialog box, where you can edit the login group parameters. See "Editing an FCoE login group" on page 548.
   - Click **Delete** to remove the login group from the list. See "Deleting one or more FCoE login groups" on page 549.

Adding an FCoE login group

Complete the following steps to add switches to a login group. You can manually add ports by entering the world wide name (WWN) or select available managed CNAs from all discovered hosts. Only directly-connected devices are supported.

1. Select **Configure > FCoE** from the menu bar.

   or

   Right-click the DCB device and select **FCoE**.

2. Click **Add**.

   The **Add Login Group** dialog box displays, as shown in Figure 236.

   FIGURE 236 Add Login Group dialog box

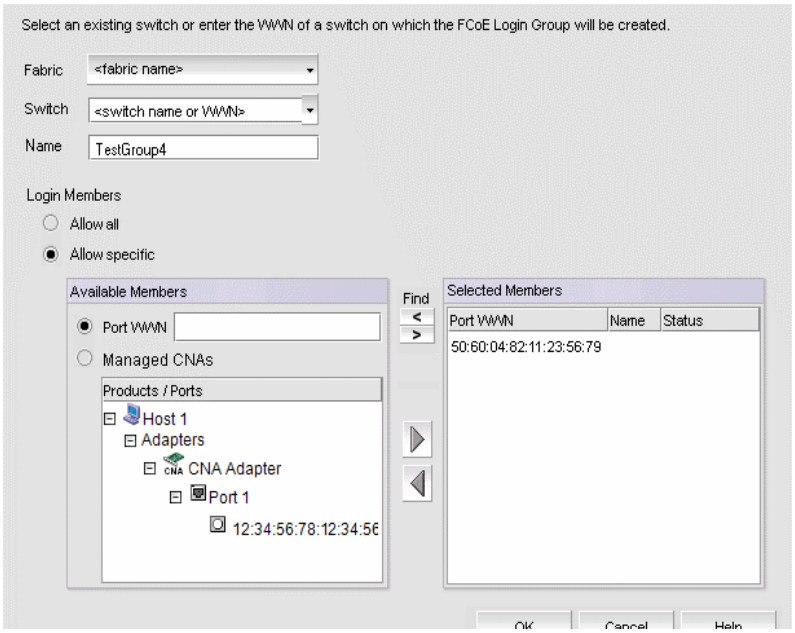

3. Select an existing switch from the **Switch** list, or enter the WWN of the switch that will be added to the FCoE login group.
4. Select one of the following Login Members options:
   - Allow all — Click to allow all login members into the Available Members list.
   - Allow specific — Click to allow specific login members into the Available Members list. If you select this option, you can add specific login members using the options in the **Available Members** area.
5. Select one of the following Available Member options:
   - Port WWN — Click to enter the world wide name (WWN) of the port to associate with the selected switch. The member port WWN text field allows a maximum of 16 digits.
   - Managed CNAs — Click to show a list of products and ports which can be selected as login group members.
6. Select available members from the **Products/Ports** list and click the right arrow button to move the available members to the **Selected Members** list.
7. Click **OK**.

   The **FCoE Login Group Confirmation and Status** dialog displays.

8. Review the changes carefully before you accept them.
9. Click **Start** to apply the changes, or click **Close** to abort the operation.

   On closing the FCoE Login Group Confirmation and Status dialog box, the **FCoE Configuration** Dialog refreshes the data and the latest information is displayed.

- "FCoE login groups"

Editing an FCoE login group

Complete the following steps to edit the name of a login group. You can manually add ports by entering the world wide name (WWN) or select available managed CNAs from all discovered hosts. Only directly-connected devices are supported.

1. Select **Configure > FCoE** from the menu bar.
2. Select a group from the Login Groups list and click **Edit**.

   The **Edit Login Group** dialog box displays, as shown in Figure 226.

   FIGURE 237 Edit Login Group dialog box

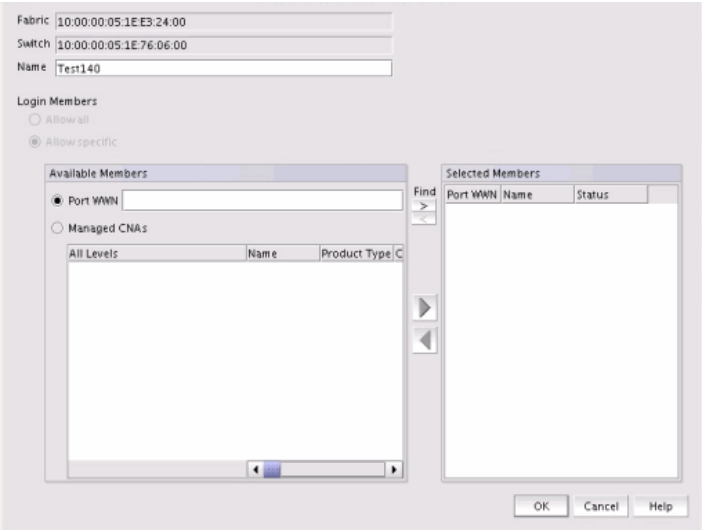

   NOTE
   The **Fabric** field and the **Switch** field are read-only fields.

3. Perform one of the following editing tasks:
   - Rename the login group by entering the new name into the **Name** field. The **Allow All** option must be selected to rename the login group.
   - Select one of the following options to add or remove login members into the **Available Members** list. The **Allow Specific** option must be selected to add or remove login members.
     - Port WWN — Click to enter the world wide name (WWN) of the port to associate with the selected switch. The member port WWN text field allows a maximum of 16 digits.
     - Managed CNAs — Click to show a list of products and ports which can be selected as login group members.
4. Select available members from the **Products/Ports** list and click the right arrow button to move the available members to the **Selected Members** list.

5. Click **OK**.

   The **FCoE Login Group Confirmation and Status** dialog displays.

6. Review the changes carefully before you accept them.
7. Click **Start** to apply the changes, or click **Close** to abort the operation.

   On closing the FCoE Login Group Confirmation and Status dialog box, the **FCoE Configuration** Dialog refreshes the data and the latest information is displayed.

Deleting one or more FCoE login groups

1. Select **Configure > FCoE** from the menu bar.

   or

   Right-click the DCB device and select **FCoE**.

   The **FCoE Configuration** dialog box displays.

2. Select a group from the Login Groups list and click **Delete**.

   The **FCoE Login Group Confirmation and Status** dialog displays.

3. Review the changes carefully before you accept them.
4. Click **Start** to apply the changes, or click **Close** to abort the operation.

   The login group is removed from the **Login Group** table.

Disabling the FCoE login management feature on a switch

1. Select **Configure > FCoE** from the menu bar.

   or

   Right-click the DCB device and select **FCoE**.

   The **FCoE Configuration** dialog box displays.

2. Select an FCoE-enabled switch from the Login Groups list and click **Disable**.

   The **FCoE Login Group Confirmation and Status** dialog displays.

3. Review the changes carefully before you accept them.
4. Click **Start** to apply the changes, or click **Close** to abort the operation.

   The FCoE login management feature is disabled and all login groups on the selected switch are deleted.

   The value in the FCoE Login Management State column for the selected switch is **Disabled** and no login groups appear under the switch after the FCoE Configuration dialog box refresh operation.

Enabling the FCoE login management feature on a switch

1. Select **Configure > FCoE** from the menu bar.

   or

   Right-click the DCB device and select **FCoE**.

   The **FCoE Configuration** dialog box displays.

2. Select an FCoE-disabled switch from the Login Groups list and click **Enable**.
3. The FCoE Login Group Configuration and Status dialog box displays.
4. Review the changes carefully before you accept them.
5. Click **Start** to apply the changes, or click **Close** to abort the operation.

   The FCoE login management feature is enabled on the selected switch.

   The value in the FCoE Login Management State column is **Enabled** after the **FCoE Configuration** dialog box refresh operation.

Virtual FCoE port configuration

The virtual FCoE port has the following configuration features:

- Displays the virtual FCoE ports on each of the DCB devices, which provides the Ethernet with bridging capability
- One-to-one mapping of FCoE ports with 10 Gbps Ethernet ports
- Option to enable or disable the virtual FCoE ports
- Option to view the end devices connected to a virtual FCoE port

Viewing virtual FCoE ports

Configuration of virtual FCoE ports requires installation of the FCoE license on the switch.

> NOTE
> For Network OS switches running the Network OS version 3.0 and later, the Management application retrieves all dynamically and statically bonded virtual FCoE ports in the virtual FCoE port pool and displays them. If there are no bonded virtual FCoE ports on any cluster member, then the cluster is not displayed.

The physical port and LAG details are displayed in the **Switch Port** column in the following circumstances:

- There is a dynamic binding between the virtual FCoE port and the physical port or LAG.
- There is a static binding between the virtual FCoE port and the physical port or lag and there are end devices connected to it.

To view the virtual FCoE ports, complete the following steps:

1. Select **Configure > FCoE** from the menu bar.

   or

   Right-click the DCB device and select **FCoE**.

   The **FCoE Configuration** dialog box displays.

2. Select the **Virtual FCoE Ports** tab.

   The **Virtual FCoE Ports** tab displays, as shown in Figure 238.

   FIGURE 238 Virtual FCoE Ports dialog box

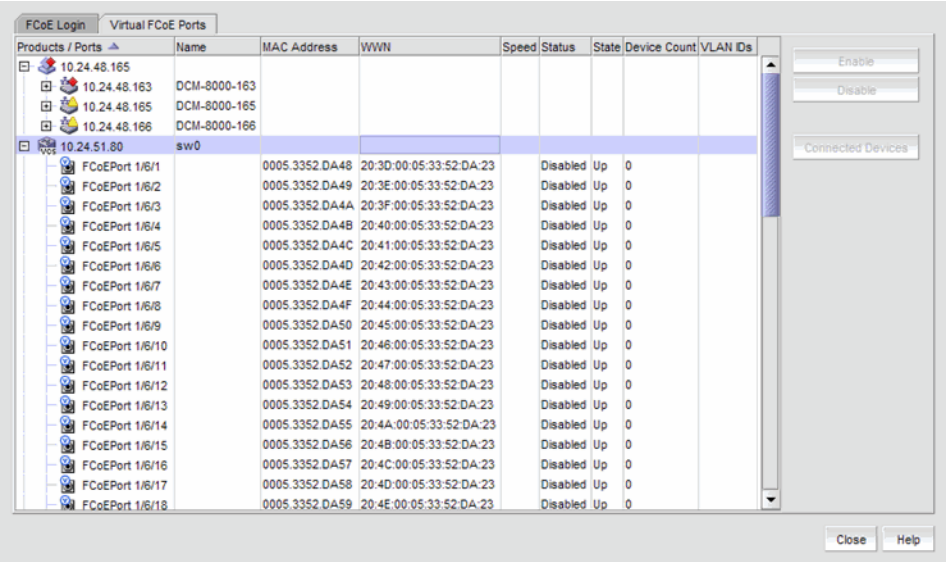

3. Select one or more virtual ports from the **Ports** list.
4. Perform one of the following tasks:
   - Click **Enable** to enable a selected virtual FCoE port from the **Virtual FCoE Ports** tab.
   - Click **Disable** to disable a selected virtual FCoE port from the **Virtual FCoE Ports** tab.
   - Click **Connected Devices** to view a list of FCoE virtual ports and to what they are directly connected.
5. Click **Close** to close the dialog box.

Clearing a stale entry

A stale entry is a device that logged in and logged off but, because a port went down after an FLOGI was received, the device failed to receive the message. The entry in the **FCoE Connected Devices** table becomes stale and you must clear it manually.

NOTE
Clearing a stale entry is not supported for Network OS devices.

1. Select a virtual FCoE port from the **FCoE Configuration** dialog box and click **Connected Devices**.

   The **Connected Devices** dialog box displays.

2. Select one or more rows from the **Connected Devices** table and click **Disconnect**.

   The **DCB Confirmation and Status** dialog displays.

   The selected connected device should be cleared from the switch cache and from the table. Note, however, that the connected devices might still be active and this operation could potentially stop traffic between the connected devices and the switch.

3. Review the changes carefully before you accept them.

4. Click **Start** to apply the changes, or click **Close** to abort the operation.

   On closing the DCB Confirmation and Status dialog box, the **FCoE Configuration** Dialog refreshes the data and the latest information about the FCoE ports are displayed.

Configuration and Operations Monitoring Policy Automation Services Suite

COMPASS overview

NOTE
Configuration and Operations Monitoring Policy Automation Services Suite (COMPASS) is supported in Professional Plus and Enterprise editions only. It is not supported in Professional edition.

COMPASS management requires the following privileges:

- Configuration File Management Read/Write privilege required for link and synchronization operations. Configuration File Management Read only privilege required to view the **COMPASS** dialog box.
- Fabric Configuration Read/Write privilege required for configuration block and template operations.

You can use COMPASS to monitor configuration drifts (changes) between a configured setting and the switch configuration to make it easier to provision new switches in a fabric.

COMPASS uses templates (user-defined) to compare to the configuration settings on a switch to configuration settings in a template to determine whether a configuration drift has occurred. Templates are made up of one or more configuration blocks. Configuration blocks are made up of one or more configuration settings. This enables you to monitor subsets of a switch configuration and receive notification when any of the defined configuration settings drifts occur.

Configuration blocks

You must have the Fabric Configuration read and write privilege to configure, edit, or delete configuration blocks. You can create configuration block from one or more pre-defined configuration settings. The Management application provides the following configuration settings:

- FTP Server — This setting details the following data:
  - **Host Name or IP** — The host name or IP address for the server.
  - **User Name** — The user name for the server.
  - **Password** — The password for the server.
  - **Remote Directory** — The directory path to the server.
  - **Type** — The server type (FTP, SCP, or SFTP).
- Syslog Destinations — This setting details the following data:
  - **New IP** — The IP address for the syslog destination.

- SNMPv3 Trap Destinations — This setting details the following data:
  - **Enable Informs** check box — Select to enable Informs.
  - **User Index** — Select 1 through 6 for the user index of the SNMPv3 trap destination.
  - **New IP** — The IP address for the SNMPv3 trap destination.
  - **Port #** — The port number for the SNMPv3 trap destination.
  - **Trap/Inform** — Select whether this is a Trap or Inform.
  - **Trap Level** — Select the trap level for the SNMPv3 trap destination.
- ACL Settings — This setting details the following data:
  - **New IP** — The IP address for the ACL host.
  - **Access** — Whether the ACL is set to Read Only or Read/Write.
- NTP Time Servers — This setting details the following data:
  - **New IP** — The IP address for the server.
- NTP Time Zone — This setting details the following data:
  - **Hour Offset** — The hour offset from GMT (-12 through 14).
  - **Minute Offset** — The minute offset from GMT (0 through 60).
- RADIUS — This setting details the following data:
  - **Server** — The IP address for the RADIUS server.
  - **Port** — The port number used by the RADIUS server.
  - **Timeout (s)** — The timeout timer value (in seconds) that specifies the amount of time to wait between retries when the server is busy.
  - **Encryption Level** — The encryption level (NONE or AES256) for the RADIUS server.
  - **Secret String** — The shared secret for the RADIUS server.
  - **Authentication** — The authentication type (CHAP, PAP, or PEAP-MSCHAPV2).
- AD/LDAP — This setting details the following data:
  - **Server** — The IP address for the AD/LDAP server.
  - **Port** — The port number used by the AD/LDAP server.
  - **Timeout (s)** — The timeout timer value (in seconds) that specifies the amount of time to wait between retries when the server is busy.
  - **Domain** — The domain of the AD/LDAP server.
- TACACS+ — This setting details the following data:
  - **Server** — The IP address for the TACACS+ server.
  - **Port** — The port number used by the TACACS+ server.
  - **Timeout (s)** — The timeout timer value (in seconds) that specifies the amount of time to wait between retries when the server is busy.
  - **Secret String** — The shared secret for the TACACS+ server.
  - **Authentication** — The authentication type (CHAP or PAP).
- MAPS Policy — This setting details the following data:
  - **Select Switch** button — Click to select a switch with MAPS policies configured.
  - **Select Policy** list — Select a MAPS policy from the list.MAPS Policy — This setting details the following data:
- Users — This setting details the users account details.
- User Credentials — This setting details the switch user account credentials.
  - **Select Switch** button — Click to select a switch to get the default and custom users configured.

NOTE
You must configure all three AAA servers (RADIUS, AD/LDAP, and TACACS+) together in a switch. If any one AAA server setting is configured, then existing AAA configurations on the switch are overwritten with the new one. You can retain the earlier AAA configuration on switch using the **Import from Switch** and **Edit** options in COMPASS.

Viewing configuration blocks

To view all defined configuration blocks, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

   FIGURE 239 Compass dialog box, Configuration tab

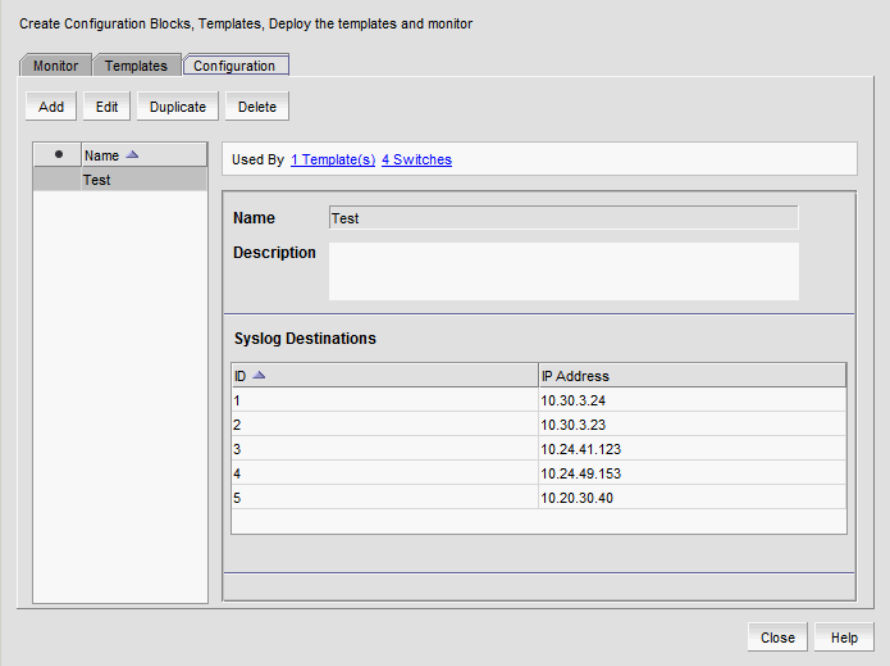

2. Click the **Configuration** tab.

   The left pane of the **Configuration** tab displays a list of existing configuration blocks. If the configuration block contains any settings that are incomplete, an incomplete icon ( ) displays before the configuration block name.

3. Select a configuration block in the left pane.

   The configuration block data displays in the right pane.

   - **Used by** — The number of templates and switches that use the selected configuration block. Click the Templates or Switches link to go to the appropriate tab with the associated templates or switches highlighted.
   - **Name** — The name of the configuration block.
   - **Description** — A description of the configuration block.
   - *defined_settings* — Any settings defined for the configuration block. For a list of configuration settings, refer to "Configuration blocks" on page 553.

4. To add a configuration block, refer to "Defining a configuration block" on page 556.

5. To edit a configuration block, refer to "Editing a configuration block" on page 567.
6. To duplicate a configuration block, refer to "Duplicating a configuration block" on page 568.
7. To delete a configuration block, refer to "Deleting a configuration block" on page 569.
8. Click **Close** on the **COMPASS** dialog box.

Defining a configuration block

A configuration block must contain at least one configuration settings. To define a configuration block, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Configuration** tab.
3. Click **Add**.

   FIGURE 240 Add Configuration

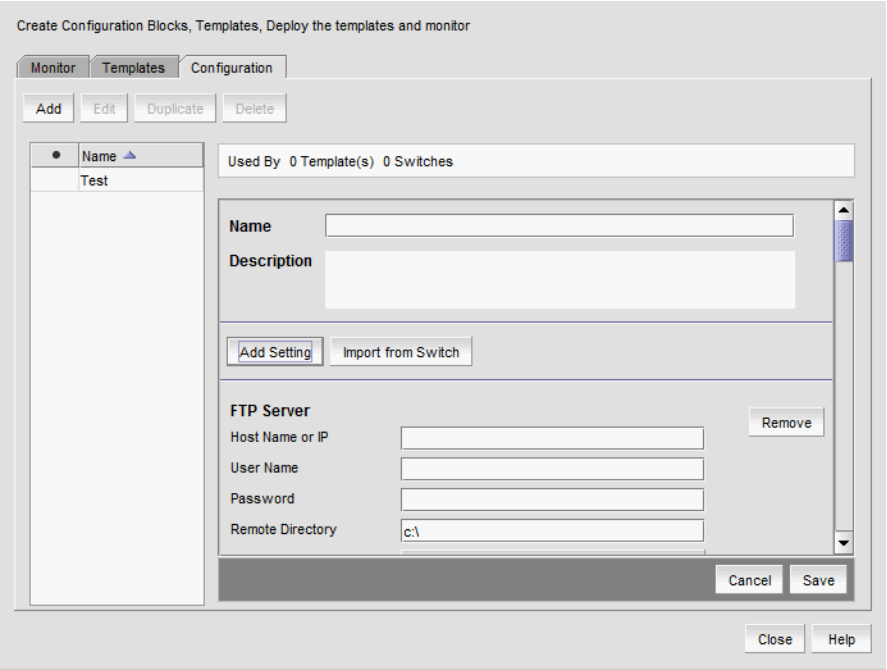

4. Enter a name (up to 128 characters) for the configuration block in the **Name** field.
5. Enter a description (up to 255 characters) for the configuration block in the **Description** field.
6. Click **Add Setting**.

   The **Add Settings** dialog box displays. For a list of configuration settings, refer to "Configuration blocks" on page 553.

7. Select the check box for each configuration setting you want to include in your configuration block.
8. Click **OK** on the **Add Settings** dialog box.

   To import configuration settings from a switch, refer to "Importing configuration settings" on page 557.

9. To configure the selected settings, choose from of the following options:
    - To configure an FTP server setting, refer to "Configuring FTP server settings" on page 558.
    - To configure a syslog destination setting, refer to "Configuring syslog destination settings" on page 558.
    - To configure a SNMPv3 inform or trap destination setting, refer to "Configuring SNMPv3 inform settings" on page 559.
    - To configure an access control list (ACL) settings setting, refer to "Configuring ACL settings" on page 561.
    - To configure an Network Time Protocol (NTP) time server setting, refer to "Configuring NTP time server settings" on page 561.
    - To configure the NTP time zone setting, refer to "Configuring NTP time zone settings" on page 562.
    - To configure the RADIUS server setting, refer to "Configuring RADIUS server settings" on page 562.
    - To configure the AD/LDAP server setting, refer to "Configuring AD/LDAP server settings" on page 563.
    - To configure the TACACS+ server setting, refer to "Configuring TACACS+ server settings" on page 563.
    - To configure the MAPS policy setting, refer to "Configuring MAPS policy settings" on page 564.
    - To configure the User settings, refer to "Configuring switch user account" on page 564
    - To configure the Users Credentials settings, refer to "Configuring switch user account credentials" on page 566
10. Repeat step 9 until you have configured all settings.
11. Remove a setting from the configuration block by click **Remove** in the associated setting area.
12. Click **Save** in the right pane of the **Configuration** tab.
13. Click **Close** on the **COMPASS** dialog box.

Importing configuration settings

To obtain configuration settings from the switch, complete the following steps.

1. Click **Import From Switch**.

    The **Import Configuration from switch** dialog box displays.

2. Select the switch from which you want to import the configuration settings in the **Available Switches** list and clicking the right arrow button.
3. Click **OK** on the **Import Configuration from switch** dialog box.
4. Click **Yes** on the confirmation message.

    The configuration settings defined on the switch display in the right pane. If you selected a configuration setting that was not defined on the switch, an incomplete icon ( ) displays next to the configuration setting name. To define configuration settings not imported from the switch, refer to "Editing a configuration block" on page 567.

Configuring FTP server settings

To obtain configuration settings from the switch, refer to "Importing configuration settings" on page 557.

To enter FTP server settings manually, complete the following steps.

FIGURE 241 FTP server settings

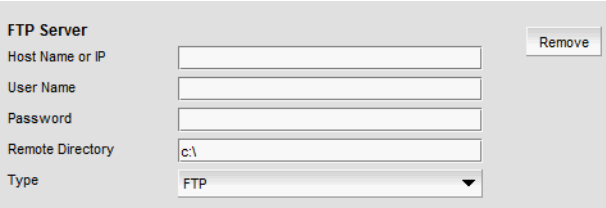

1. Enter the host name or IP address for the server in the **Host Name or IP** field.
2. Enter a user name in the **User Name** field.
3. Enter the password in the **Password** field.
4. Enter the directory path in the **Remote Directory** field.
5. Select the FTP server type (FTP, SCP, or SFTP) from the **Type** list.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring syslog destination settings

You can add up to 6 syslog destinations. To configure the syslog destination, complete the following steps.

1. Enter the IP address for the syslog destination in the **New IP** field.
2. Click **Add**.

   The syslog IP address displays in the **Syslog Destinations** table.

3. Repeat step 1 and step 2 for each syslog destination you want to configure.

   To delete a syslog destination, select the syslog IP address in the **Syslog Destinations** table and click **Delete**.

   To clear all syslog destinations, click **Delete All**.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring SNMPv3 inform settings

You can add up to 6 SNMPv3 informs. To configure the SNMPv3 inform, complete the following steps.

FIGURE 242SNMPv3 Inform / Trap Recipient settings

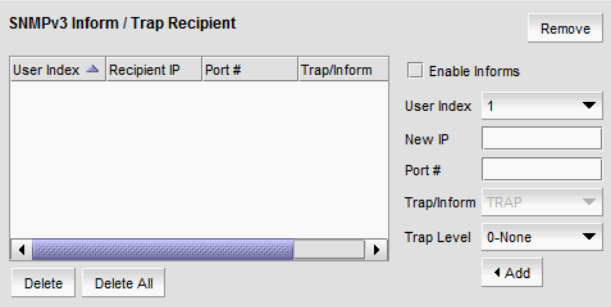

1. Select the **Enable Informs** check box.
2. Select a number (1 through 6) in the **User Index** list.
3. Enter the IP address for the SNMPv3 inform in the **New IP** field.
4. Enter the port number for the SNMPv3 inform port in the **Port #** field.
5. Select **INFORM** from the **Trap/Inform** list.
6. Select one of the severity levels from the **Trap Level** list.
   - 0-None
   - 1-Critical
   - 2-Error
   - 3-Warning
   - 4-Info
   - 5-Debug
7. Click **Add**.

   The SNMPv3 inform data displays in the **SNMPv3 Inform / Trap Recipient** table.
8. Repeat step 1 through step 7 for each SNMPv3 inform you want to configure.

   To delete a SNMPv3 inform, select the SNMPv3 inform in the **SNMPv3 Inform / Trap Recipient** table and click **Delete**.

   To clear all SNMPv3 inform, click **Delete All**.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring SNMPv3 trap destination settings

You can add up to 6 SNMPv3 inform or trap destinations. To configure the SNMPv3 inform or trap destination, complete the following steps.

FIGURE 243SNMPv3 Inform / Trap Destination settings

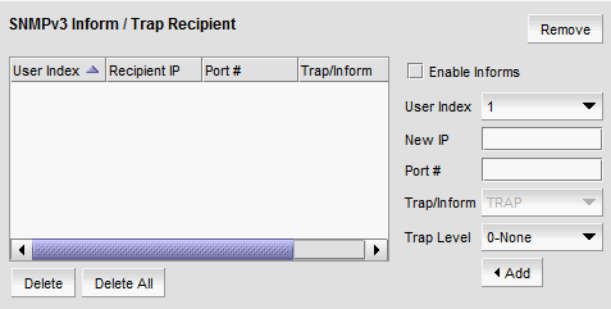

1. Make sure the **Enable Informs** check box is clear.
2. Select a number (1 through 6) in the **User Index** list.
3. Enter the IP address for the SNMPv3 trap destination in the **New IP** field.
4. Enter the port number for the SNMPv3 trap listening port in the **Port #** field.
5. Select one of the severity levels from the **Trap Level** list.
   - 0-None
   - 1-Critical
   - 2-Error
   - 3-Warning
   - 4-Info
   - 5-Debug
6. Click **Add**.

   The SNMPv3 trap destination data displays in the **SNMPv3 Trap Destinations** table.
7. Repeat step 1 through step 6 for each SNMPv3 trap destination you want to configure.

   To delete a SNMPv3 trap destination, select the SNMPv3 trap destination in the **SNMPv3 Trap Destinations** table and click **Delete**.

   To clear all SNMPv3 trap destinations, click **Delete All**.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring ACL settings

You can add up to 6 ACL settings. To configure the SNMPv3 trap destination, complete the following steps.

FIGURE 244ACL settings

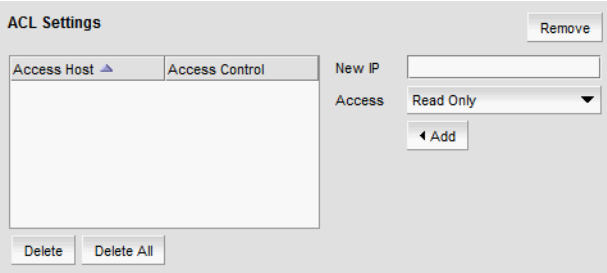

1. Enter the IP address for the ACL host in the **New IP** field.
2. Select **Read Only** or **Read/Write** from the **Access** list.
3. Click **Add**.

   The ACL settings data displays in the **ACL Settings** table.
4. Repeat step 1 through step 3 for each ACL setting you want to configure.

   To delete a ACL setting, select the ACL setting in the **ACL Settings** table and click **Delete**.

   To clear all ACL settings, click **Delete All**.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring NTP time server settings

To enter the NTP time server settings, complete the following steps.

1. Enter the IP address for the NTP time server in the **New IP** field.
2. Click **Add**.

   The NTP time server IP address displays in the **NTP Time Server** table.
3. Repeat step 1 and step 2 for each NTP time server you want to configure.

   To delete a NTP time server setting, select the NTP time server in the **NTP Time Server** table and click **Delete**.

   To clear all NTP time server settings, click **Delete All**.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring NTP time zone settings

You can configure the NTP time server to adjust the time offset, in hours and minutes, from Greenwich Mean Time (GMT). To set the NTP time zone setting, complete the following steps.

1. Select the hour offset from GMT (-12 through 14) in the **Hour Offset** list.
2. Select the minute offset from GMT (0 through 60) in the **Minute Offset** list.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring RADIUS server settings

To enter the RADIUS server settings, complete the following steps.

FIGURE 245RADIUS server settings

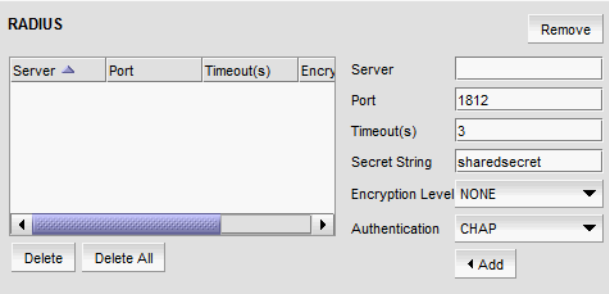

1. Enter the IP address for the RADIUS server in the **Server** field.
2. Enter the TCP port, if necessary, used by the RADIUS server in the **Port** field.

   Default is 1812.
3. Enter the timeout timer value (in seconds) that specifies the amount of time to wait between retries when the server is busy in the **Timeout(s)** field.

   Default is 3 seconds.
4. Enter the shared secret in the **Shared String** field.
5. Enter the encryption level in the **Encryption Level** field.
6. Select the authentication policy (CHAP, PAP, or PEAP-MSCHAPV2) from the **Authentication** field.

   Default is CHAP.
7. Click **Add**.

   The RADIUS server details display in the **RADIUS** table.
8. Repeat step 1 and step 7 for each RADIUS server you want to configure.

   To delete a RADIUS server setting, select the RADIUS server in the **RADIUS** table and click **Delete**.

   To clear all RADIUS server settings, click **Delete All**.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring AD/LDAP server settings

To enter the AD/LDAP server settings, complete the following steps.

FIGURE 246AD/LDAP server settings

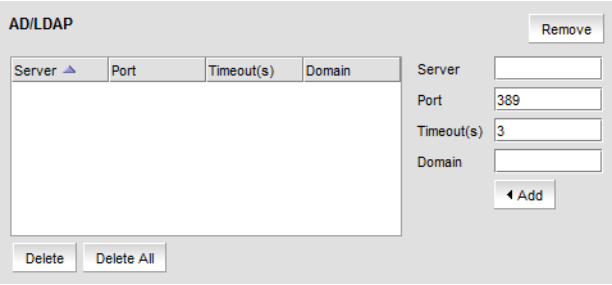

1. Enter the IP address for the AD/LDAP server in the **Server** field.
2. Enter the TCP port, if necessary, used by the AD/LDAP server in the **Port** field.

   Default is 1812.
3. Enter the timeout timer value (in seconds) that specifies the amount of time to wait between retries when the server is busy in the **Timeout (s)** field.

   Default is 3 seconds.
4. Enter the domain of the AD/LDAP server in the **Domain** field.
5. Click **Add**.

   The AD/LDAP server details display in the **AD/LDAP** table.
6. Repeat step 1 and step 5 for each AD/LDAP server you want to configure.

   To delete a AD/LDAP server setting, select the AD/LDAP server in the **AD/LDAP** table and click **Delete**.

   To clear all AD/LDAP server settings, click **Delete All**.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring TACACS+ server settings

To enter the TACACS+ server settings, complete the following steps.

FIGURE 247TACACS+ server settings

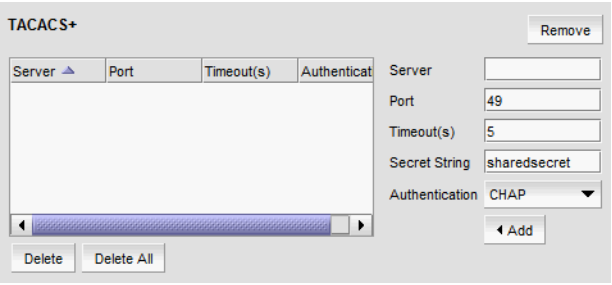

1. Enter the IP address for the TACACS+ server in the **Server** field.
2. Enter the TCP port, if necessary, used by the TACACS+ server in the **Port** field.

   Default is 1812.
3. Enter the timeout timer value (in seconds) that specifies the amount of time to wait between retries when the server is busy in the **Timeout(s)** field.

   Default is 5 seconds.
4. Enter the shared secret in the **Shared String** field.
5. Select the authentication policy (CHAP or PAP) from the **Authentication** field.

   Default is CHAP.
6. Click **Add**.

   The TACACS+ server details display in the **TACACS+** table.
7. Repeat step 1 and step 6 for each TACACS+ server you want to configure.

   To delete a TACACS+ server setting, select the TACACS+ server in the **TACACS+** table and click **Delete**.

   To clear all TACACS+ server settings, click **Delete All**.

   Return to step 9 of "Defining a configuration block" on page 556 to add additional configuration blocks and complete this procedure.

Configuring MAPS policy settings

To configure MAPS policy settings, complete the following steps.

1. Click **Select Switch** to select a switch with MAPS policies configured.
2. The **Select a switch for MAPS policies** dialog box displays.
3. Select a switch in the **Available Switches** list and click the right arrow.
4. Click **OK**.
5. Select a MAPS policy from the **Select Policy** list.

Configuring switch user account

You can configure an user account other than the default and custom users in the following ways.

- Import from Switch
- Manual

FIGURE 248Users Settings

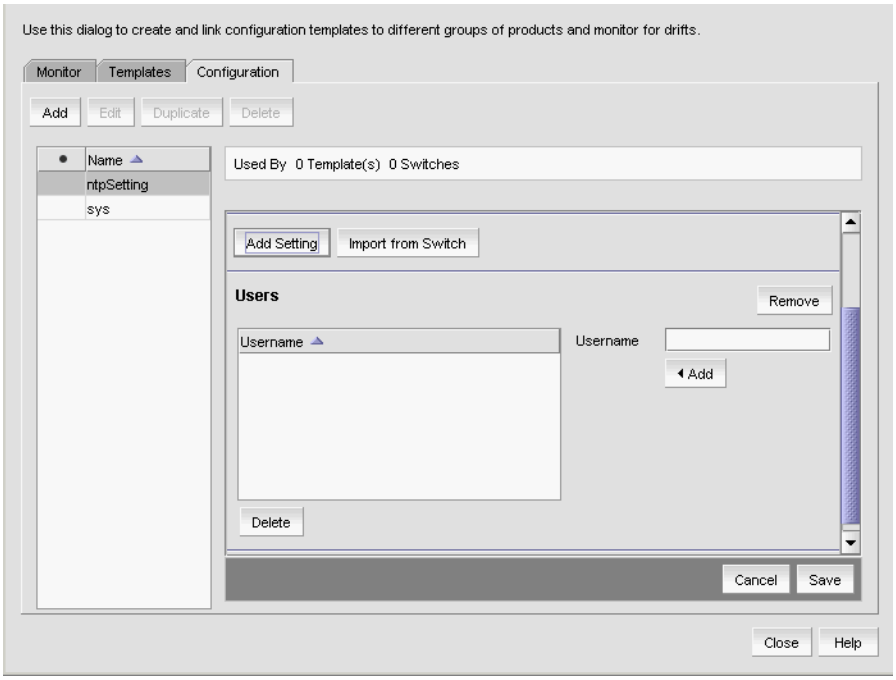

To track an user account by Import from Switch, complete the following steps.

1. Select **Configure** > **COMPASS.**

   The **COMPASS dialog box** displays.

2. Click the **Configuration** tab.
3. Click **Import From Switch**.

   The **Import Configuration from switch** dialog box displays.

4. Select the switch from which you want to import the configuration settings in the **Available Switches** list and click the right arrow button.

   The user account is copied or imported to the user table. All default and custom users are imported.

   NOTE
   You can select only one switch. If you select a switch with AAA configuration, a warning message displays.

To track an user account manually, complete the following steps.

1. Select **Configure** > **COMPASS.**

   The **COMPASS dialog box** displays.

2. Click the **Configuration** tab.
3. Under **Users**, enter the user account name in the **Username** field (case-sensitive) and click **Add**.

   The user account name is added to the **Username** table and the following validations are done.

   - Empty username
   - Duplicate username

4. Click **Save** in the right pane of the **Configuration** tab.
5. Click **Close** on the **COMPASS** dialog box.

   The user account is copied or imported to the user table. All default and custom users are imported.

   For a new user account, user configuration is created, mapped to a template, and linked to fabric or group for tracking. When you link the template to a fabric or group, you can calculate the drift. The drift status is updated in **Drift Status** column under the **Monitor** tab.

   For more information on drift, refer to "Viewing configuration drifts" on page 578.

   For more information on template, refer to "Templates" on page 569.

Configuring switch user account credentials

You can configure the switch user account credentials that are not compliant with pre-configured template settings. To configure an user account credential changes, complete the following steps.

FIGURE 249 User Credentials settings

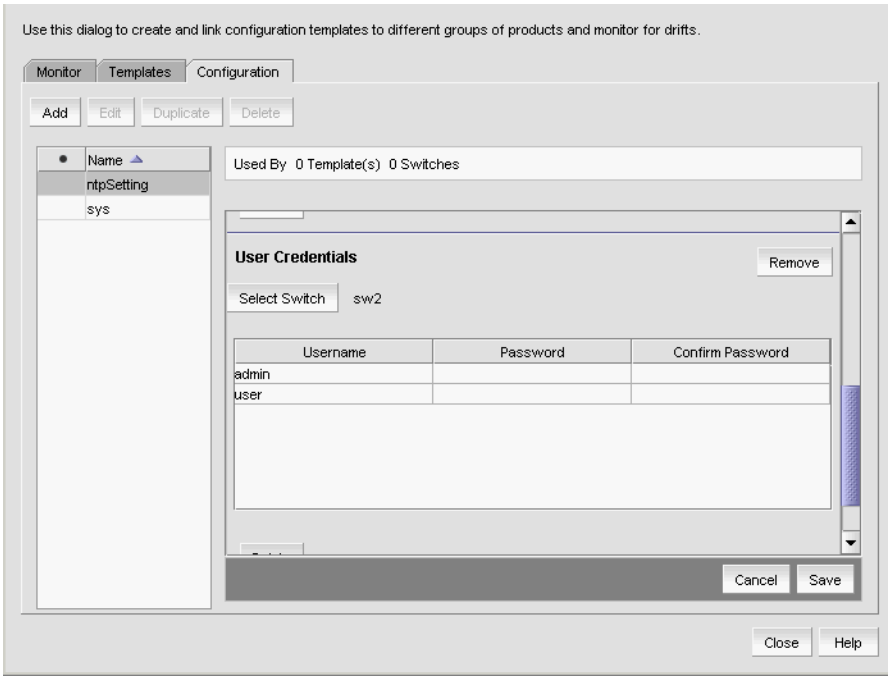

1. Select **Configure** > **COMPASS.**

   The **COMPASS dialog box** displays.
2. Click the **Configuration** tab.
3. Click **Import From Switch**.

   Or

   Click **Select Switch** under **User Credential**.

   The **Import Configuration from switch** dialog box displays.

4. Select the switch from which you want to import the configuration settings in the **Available Switches** list and click the right arrow button.

   The user is added to the User Credential table. All default and custom users are displayed except the root user. An admin user is not allowed to change the password of a root user.

5. Enter the password in the **Password** and **Confirm Password** columns.

   You can edit the Password and Confirm Password columns only when the switch configuration is imported using **Select Switch** option and not from **Import from Switch** option. When you use the **Import from Switch** option, the configuration is not complete. You can click the **Edit** button and enter the password to complete the configuration.

6. Click **Save** in the right pane of the **Configuration** tab.
7. Click **Close** on the **COMPASS** dialog box.

   You can view the configuration drift status in **Drift Status** column under the **Monitor** tab and the new password sync configuration details in the **Deployment Status** dialog box.

For more information on configuration drift, refer to "Viewing configuration drifts" on page 578.

For more information on configuration sync, refer to "Synchronizing a configuration" on page 577 and "Synchronizing all configurations" on page 578.

Editing a configuration block

To edit a configuration block, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Configuration** tab.
3. Select a configuration block in the **Name** list and click **Edit**.
4. Edit the name (up to 128 characters) for the configuration block in the **Name** field.
5. Edit the description (up to 255 characters) for the configuration block in the **Description** field.
6. Add additional settings by clicking **Add Setting**.

   The **Add Settings** dialog box displays. Any configuration settings previously selected display grayed out and cannot be cleared.

7. Select the check box for each additional configuration setting you want to include in your configuration block.
8. Click **OK** on the **Add Settings** dialog box.

   To import configuration settings from a switch, refer to "Importing configuration settings" on page 557.

9. To configure the selected settings, choose from of the following options:

   - To configure an FTP server setting, refer to "Configuring FTP server settings" on page 558.
   - To configure a syslog destination setting, refer to "Configuring syslog destination settings" on page 558.
   - To configure a SNMPv3 trap destination setting, refer to "Configuring SNMPv3 inform settings" on page 559.
   - To configure an access control list (ACL) settings setting, refer to "Configuring ACL settings" on page 561.
   - To configure an Network Time Protocol (NTP) time server setting, refer to "Configuring NTP time server settings" on page 561.

- To configure the NTP time zone setting, refer to "Configuring NTP time zone settings" on page 562.
- To configure the NTP time zone setting, refer to "Configuring NTP time zone settings" on page 562.
- To configure the RADIUS server setting, refer to "Configuring RADIUS server settings" on page 562.
- To configure the AD/LDAP server setting, refer to "Configuring AD/LDAP server settings" on page 563.
- To configure the TACACS+ server setting, refer to "Configuring TACACS+ server settings" on page 563.
- To configure the MAPS policy setting, refer to "Configuring MAPS policy settings" on page 564.
- To configure the User settings, refer to "Configuring switch user account" on page 564
- To configure the Users Credentials settings, refer to "Configuring switch user account credentials" on page 566

10. Repeat step 9 until you have configured all settings.
11. Remove a setting from the configuration block by click **Remove** in the associated setting area.
12. Click **Save** in the right pane of the **Configuration** tab.
13. Click **Close** on the **COMPASS** dialog box.

Duplicating a configuration block

To duplicate a configuration block, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Configuration** tab.
3. Select a configuration block in the **Name** list and click **Duplicate**.
4. Enter a name (up to 128 characters) for the configuration block in the **Name** field.
5. Enter a description (up to 255 characters) for the configuration block in the **Description** field.
6. Add additional settings by clicking **Add Setting**.

   The **Add Settings** dialog box displays. Any configuration settings previously selected display grayed out and cannot be cleared.

7. Select the check box for each additional configuration setting you want to include in your configuration block.
8. Click **OK** on the **Add Settings** dialog box.

   To import configuration settings from a switch, refer to "Importing configuration settings" on page 557.

9. To configure the selected settings, choose from of the following options:
   - To configure an FTP server setting, refer to "Configuring FTP server settings" on page 558.
   - To configure a syslog destination setting, refer to "Configuring syslog destination settings" on page 558.
   - To configure a SNMPv3 trap destination setting, refer to "Configuring SNMPv3 inform settings" on page 559.
   - To configure an access control list (ACL) settings setting, refer to "Configuring ACL settings" on page 561.
   - To configure an Network Time Protocol (NTP) time server setting, refer to "Configuring NTP time server settings" on page 561.
   - To configure the NTP time zone setting, refer to "Configuring NTP time zone settings" on page 562.
   - To configure the NTP time zone setting, refer to "Configuring NTP time zone settings" on page 562.

- To configure the RADIUS server setting, refer to "Configuring RADIUS server settings" on page 562.
- To configure the AD/LDAP server setting, refer to "Configuring AD/LDAP server settings" on page 563.
- To configure the TACACS+ server setting, refer to "Configuring TACACS+ server settings" on page 563.
- To configure the MAPS policy setting, refer to "Configuring MAPS policy settings" on page 564.
- To configure the User settings, refer to "Configuring switch user account" on page 564
- To configure the Users Credentials settings, refer to "Configuring switch user account credentials" on page 566

10. Repeat step 9 until you have configured all settings.
11. Remove a setting from the configuration block by click **Remove** in the associated setting area.
12. Click **Save** in the right pane of the **Configuration** tab.
13. Click **Close** on the **COMPASS** dialog box.

Deleting a configuration block

If you delete a configuration block associated with a template, the configuration block is removed from the template. To delete a configuration block, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.
2. Click the **Configuration** tab.
3. Select the configuration block you want to delete in the **Name** list and click **Delete**.
4. Click **OK** on the confirmation message.
5. Click **Close** on the **COMPASS** dialog box.

Templates

You must have the Fabric Configuration read and write privilege to configure, edit, or delete templates. Templates are made up of one or more configuration blocks. Configuration blocks are made up of one or more configuration settings. This enables you to monitor subsets of a switch configuration and receive notification when any of the defined configuration settings drifts occur.

You can create templates that only monitor specific configuration settings, such as management configuration settings or security configuration settings. For example:

- Management Template example — Create a configuration block ("Defining a configuration block" on page 556) using the following configuration settings:
  - FTP/SCP server settings
  - Syslog destination setting
  - SNMPv3 trap destination setting
  - ACL settings
  - NTP time server settings
  - NTP time zone settings

- Security Template example — Create a configuration block ("Defining a configuration block" on page 556) using the following configuration settings
  - AD/LDAP settings
  - RADIUS settings
  - TACACS+ settings

Viewing existing templates

If there is a configuration setting not defined in a template, an incomplete icon ( ) displays before the template name. To define configuration settings, refer to "Editing a configuration block" on page 567.

To view any configured template, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Template tab**.

   FIGURE 250 Compass dialog box, Templates tab

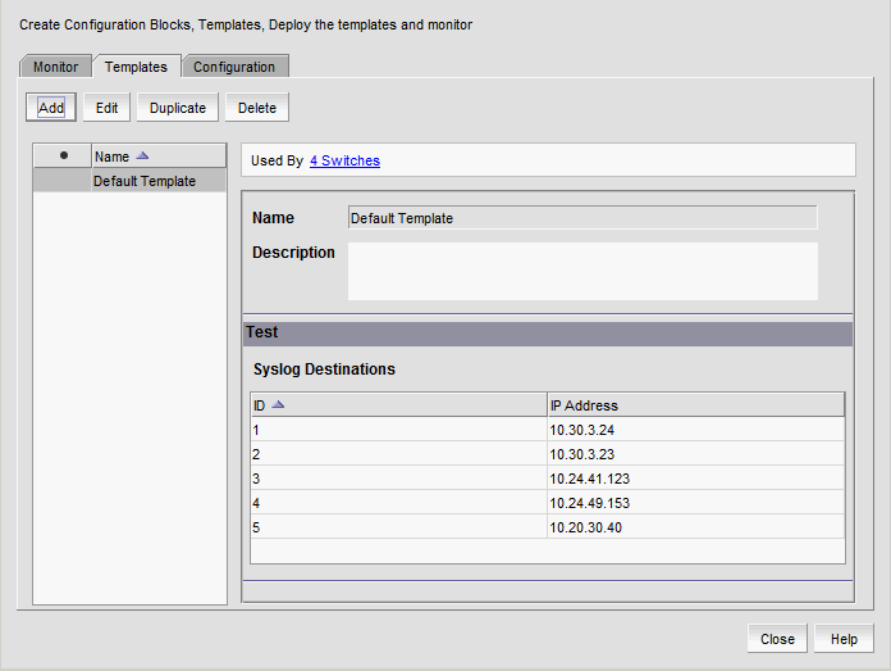

3. Select a template from the **Name** list.

   The template data displays in the right pane.

   - **Used by** — The number of switches that use the selected template. Click the Switches link to go to the **Monitor** tab with the associated switches highlighted.
   - **Name** — The name of the template.
   - **Description** — A description of the template.
   - *defined_configurations* — Any configuration blocks associated with the template. For a information about configuration blocks and settings, refer to "Configuration blocks" on page 553.

4. To add a template, refer to "Adding a template" on page 571.
5. To edit a template, refer to "Editing a template" on page 572.
6. To duplicate a template, refer to "Duplicating a template" on page 573.
7. To delete a template, refer to "Deleting a template" on page 574.
8. Click **Close** on the **COMPASS** dialog box.

Adding a template

To add a template, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Template** tab.
3. Click **Add**.

   The right pane displays the editable fields.

   FIGURE 251 Add templates

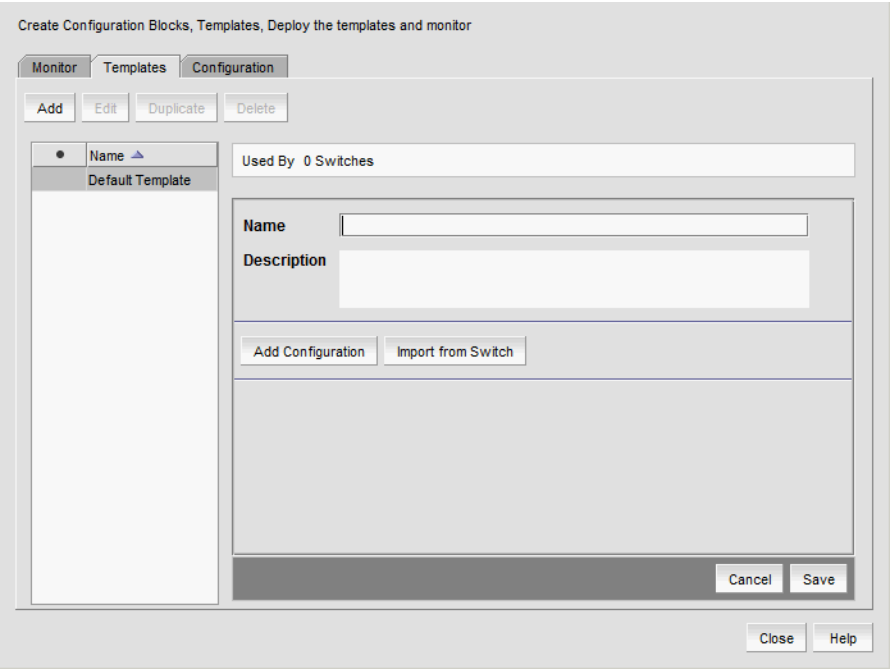

4. Enter a unique name (up to 128 characters) for the template in the **Name** field.
5. Enter a description (up to 255 characters) for the template in the **Description** field.
6. Add configurations from a switch by clicking **Import Settings From Switch**.

   The **Select Switch** dialog box displays with a list of switches that have a configuration file saved in the Management application.

7. Select a switch and click **OK**.

8. Click **Add Configuration** to add a configuration block to the template.

   The **Add Configuration** dialog box displays with the a list of available configuration blocks. To edit a configuration block, refer to "Editing a configuration block" on page 567.

   To import configuration settings from a switch, refer to "Importing configuration settings" on page 557.

9. Select one or more configuration blocks and click **OK**.

   The configuration blocks you selected display in the right pane.

10. Click **Save** in the right pane of the **Templates** tab.

    If any of the configuration blocks have incomplete configuration settings, a message displays. Click **OK** on the message. Complete the configuration setting (refer to "Removing a configuration block from a template" on page 572) or remove the configuration block from the template (refer to "Editing a template" on page 572).

11. Click **Close** on the **COMPASS** dialog box.

Removing a configuration block from a template

To remove a configuration block from a template, complete the following steps.

1. Select the template you want to edit in the **Name** list.
2. Click **Edit**.

   The right pane displays the current configuration settings for the template.

3. Click **Remove Configuration** next to the name of the configuration block you want to remove from the template.
4. Click **Save** in the right pane of the **Templates** tab.

Editing a template

To edit a template, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Template** tab.
3. Select the template you want to edit in the **Name** list.
4. Click **Edit**.

   The right pane displays the current configuration settings for the template.

5. Edit the description (up to 255 characters) for the template in the **Description** field.
6. Click **Import Settings From Switch** to add configuration settings from a switch.

   The **Select Switch** dialog box displays with a list of switches that have a configuration file saved in the Management application.

7. Select a switch and click **OK**.
8. Click **Add Configuration** to add a configuration block to the template.

The **Add Configuration** dialog box displays with the a list of available configuration blocks. To edit a configuration block, refer to "Editing a configuration block" on page 567.

To import configuration settings from a switch, refer to "Importing configuration settings" on page 557.

9. Select one or more configuration blocks and click **OK**.

   The configuration blocks you selected display in the right pane.

10. Click **Save** in the right pane of the **Templates** tab.

    If any of the configuration blocks have incomplete configuration settings, a message displays. Click **OK** on the message. Complete the configuration setting (refer to "Removing a configuration block from a template" on page 572) or remove the configuration block from the template (refer to "Editing a template" on page 572).

11. Click **Close** on the **COMPASS** dialog box.

Duplicating a template

To duplicate a template, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Template** tab.
3. Select the template you want to duplicate in the **Name** list.
4. Click **Duplicate**.

   The right pane displays the editable fields.

5. Enter a unique name (up to 128 characters) for the template in the **Name** field.
6. Edit the description (up to 255 characters) for the template in the **Description** field.
7. Click **Import Settings From Switch** to add a configuration settings from a switch.

   The **Select Switch** dialog box displays with a list of switches that have a configuration file saved in the Management application.

8. Select a switch and click **OK**.
9. Click **Add Configuration** to add a configuration block to the template.

   The **Add Configuration** dialog box displays with the a list of available configuration blocks. To edit a configuration block, refer to "Editing a configuration block" on page 567.

   To import configuration settings from a switch, refer to "Importing configuration settings" on page 557.

10. Select one or more configuration blocks and click **OK**.

    The configuration blocks you selected display in the right pane.

11. Click **Save** in the right pane of the **Templates** tab.

    If any of the configuration blocks have incomplete configuration settings, a message displays. Click **OK** on the message. Complete the configuration setting (refer to "Removing a configuration block from a template" on page 572) or remove the configuration block from the template (refer to "Editing a template" on page 572).

12. Click **Close** on the **COMPASS** dialog box.

Deleting a template

If you delete a template linked to a fabric or group, the template is unlinked from the fabric or group. To delete a template, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Template** tab.
3. Select the template you want to delete in the **Name** list.
4. Click **Delete**.
5. Click **OK** on the confirmation message.
6. Click **Close** on the **COMPASS** dialog box.

COMPASS monitoring

You can use the COMPASS monitoring feature to track whether a configuration has drifted from the configuration defined on a switch.

> NOTE
> Event triggered backup must be enabled to accurately track configuration drifts in COMPASS.

Viewing COMPASS monitors

To view all COMPASS monitors, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Monitor** tab.

FIGURE 252 Compass dialog box, Monitor tab

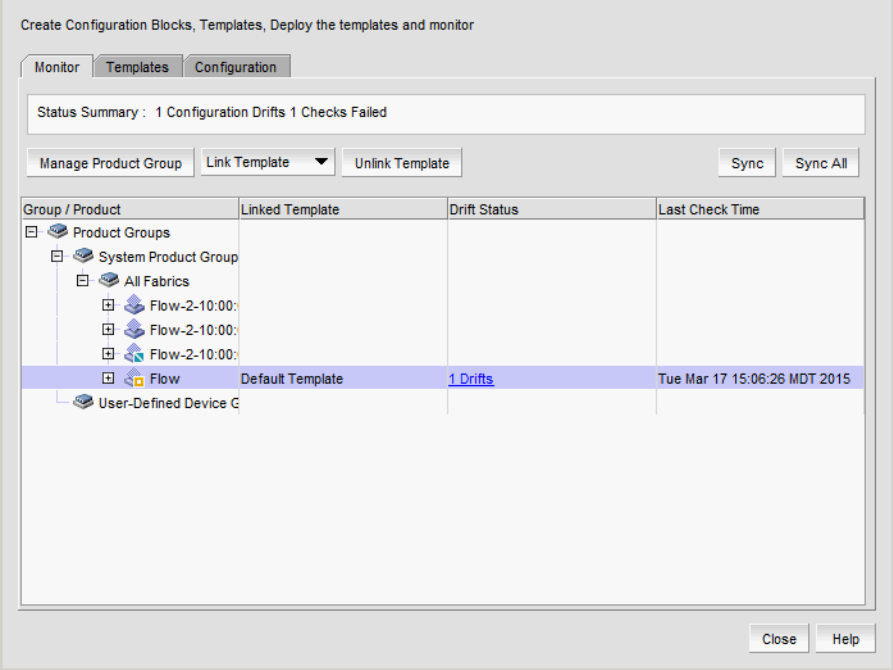

3. Review the configured COMPASS monitor details:

   A status summary displays at the top of the tab. The status summary data includes the following:

   - The number of configuration drifts.
   - The number of failed checks.

   The COMPASS monitor includes the following data:

   - **Group/Product** — List of all discovered and monitored fabrics and switches as well as any user-created switch groups.
   - **Linked Template** — The name of the linked template.
   - **Drift Status** — Whether the switch configuration is synchronized with the COMPASS template, the number of drifts (changes) between the COMPASS template and the switch configuration, or if the switch is unreachable or unmonitored (failed check). For fabrics and user-defined switch groups, the drift status is the aggregate status for all switches in the fabric or group. Click the drift status to view additional data (refer to "Viewing configuration drifts" on page 578).
   - **Last Check Time** — The date and time that the COMPASS monitor check ran last.

4. To create a static switch group, click **Manage Network Scope** (refer to "Creating a product group" on page 576).
5. To link a COMPASS template to a fabric or group, select a fabric or group, then select a template from the **Link Template** list (refer to "Linking a template" on page 577).
6. To remove a COMPASS template link from a fabric or group, select a fabric or group, then click **Unlink Template** (refer to "Unlinking a template" on page 577).
7. To synchronize a COMPASS template to all switches in a fabric or group, select a fabric or group and click **Sync** (refer to "Synchronizing a configuration" on page 577).
8. To synchronize all COMPASS templates (where the Management application detects a drift) with the associated switches, click **Sync All** (refer to "Synchronizing all configurations" on page 578).
9. Click **Close**.

Creating a product group

To create a static product group to which you can assign a COMPASS monitor, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Monitor** tab.
3. Click **Managed Product Group**.

   The **Manage Product Group** dialog box displays.

4. Enter a unique name (maximum 64 characters) for the product group in the **Name** field.
5. Add products to the group by selecting the product in the **Available Targets** list and clicking the right arrow button.

   The selected products move from the **Available Targets** list to the **Selected Targets** list.

6. Remove products from the group by selecting the product in the **Selected Targets** list and clicking the left arrow button.

   The selected products move from the **Selected Targets** list to the **Available Targets** list.

7. Click **Apply**.

   The new group displays in the **Product Groups** list.

8. Click **OK**.

   The new group displays in the **User-Defined Device Groups** folder of the **Group/Product** list.

9. Click **Close**.

Editing a product group

To edit a product group, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Monitor** tab.
3. Click **Managed Product Group**.

   The **Manage Product Group** dialog box displays.

4. Select the product group you want to edit in the **Product Groups** list.
5. Add products to the group by selecting the product in the **Available Targets** list and clicking the right arrow button.

   The selected products move from the **Available Targets** list to the **Selected Targets** list.

6. Remove products from the group by selecting the product in the **Selected Targets** list and clicking the left arrow button.

   The selected products move from the **Selected Targets** list to the **Available Targets** list.

7. Click **Apply**.

   The new group displays in the **Product Groups** list.

8. Click **OK**.

   The new group displays in the **User-Defined Device Groups** folder of the **Group/Product** list.

9. Click **Close**.

Linking a template

Only templates with a complete configuration block (all configuration settings are complete) display in the **Link Template** list. To link a template to a fabric or group, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Monitor** tab.
3. Select the fabric or group to which you want to link a template and select a template from the **Link Template** list.
4. Click **Apply**.
5. Click **Close**.

Unlinking a template

To remove a template link from a fabric or group, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Monitor** tab.
3. Select the fabric or group from which you want to remove the linked template and click **Unlink Template**.
4. Click **Apply**.
5. Click **Close**.

Synchronizing a configuration

When a switch drifts from the configuration specified in the COMPASS template, you can update the configuration on the switch using the COMPASS template settings.

To synchronize a template to all switches in a fabric or group, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Monitor** tab.
3. Select the fabric or group you want to synchronize and click **Sync**.

   The **Deploy to Products** dialog box displays.

4. Click **OK**.

   The **Deployment Status** dialog box displays.

5. Start the deployment by clicking **Start**.

   Click **Abort** to stop deployment on deployments that have not started.

   Click a switch in the **Deployment Status** list to display the status of the deployment in the **Status Details** area.

6. Click **Close**.

Synchronizing all configurations

To update all linked switches with the associated COMPASS template, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Monitor** tab.
3. Click **Sync All**.

   The **Deploy to Products** dialog box displays.

4. Click **OK**.

   The **Deployment Status** dialog box displays.

5. Start the deployment by clicking **Start**.

   Click **Abort** to stop deployment on deployments that have not started.

   Click a switch in the **Deployment Status** list to display the status of the deployment in the **Status Details** area.

6. Click **Close**.

Viewing configuration drifts

The configuration drift data is stored in the database for 30 days. The system purges old data (over 30 days) every night at 12:00 AM.

To view configuration drifts, complete the following steps.

1. Select **Configure > COMPASS**.

   The **COMPASS** dialog box displays.

2. Click the **Monitor** tab.
3. Click the drifts link in the **Drift Status** column.

The **Configuration Drifts** dialog box displays.

FIGURE 253 Configuration Drifts dialog box

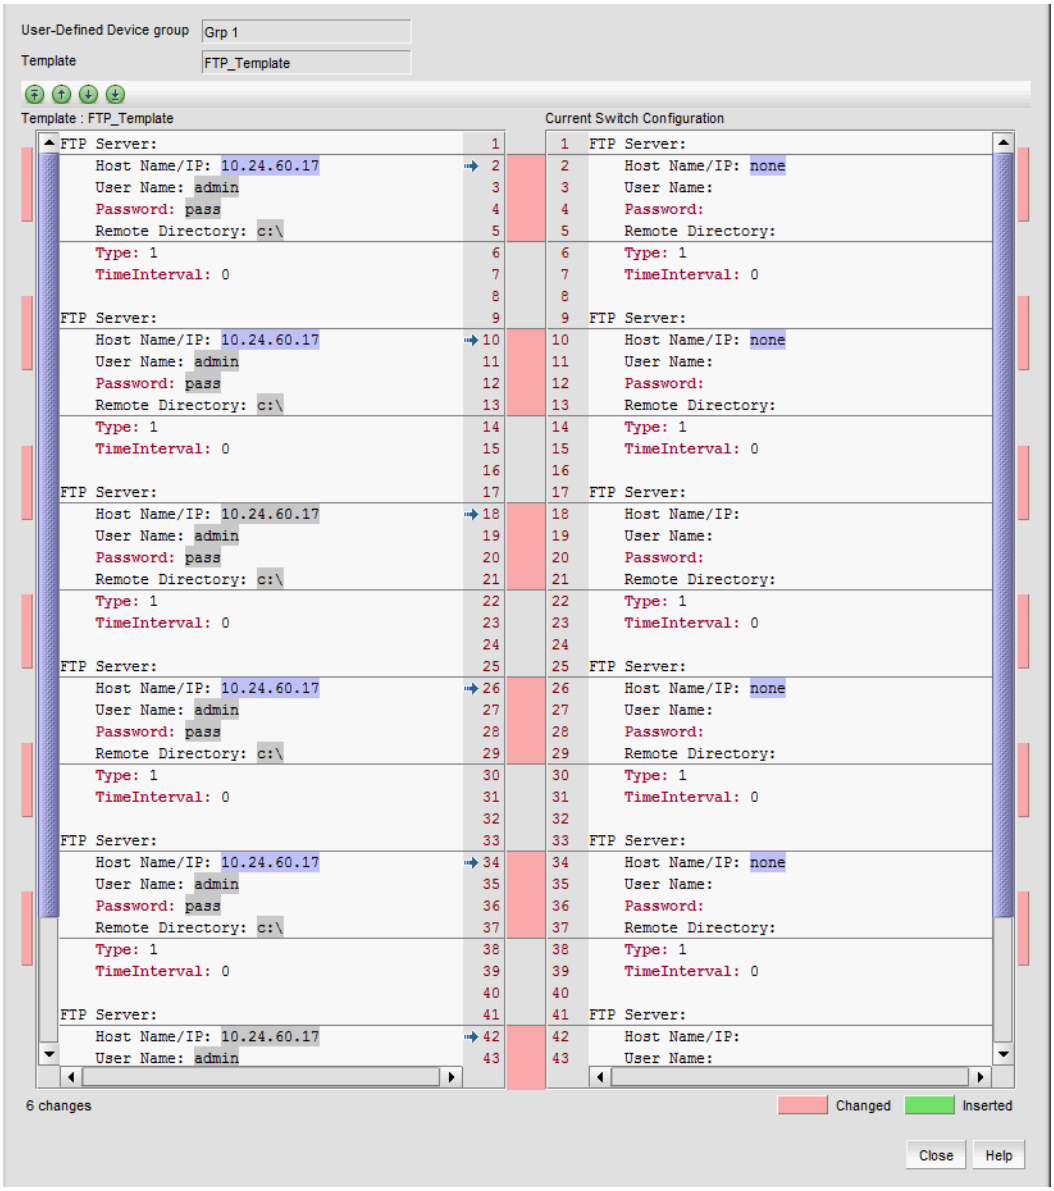

The configuration drift data includes the following:

- **Fabric/User-Defined Device group** — The name of the fabric or user-defined device group on which the drift occurred.
- **Template** — The name of the linked template.
- **Change Navigator** buttons — The **Change Navigator** buttons are enabled when there is at least one change between two compared files.
  - Go to first change button ( ) — Click to move to the first change.
  - Go to previous change button ( ) — Click to move to the previous change.
  - Go to next change button ( ) — Click to move to the next change.
  - Go to last change button ( ) — Click to move to the last change.
- **Template:** *template_name* text box — The configuration setting on the template.

- **Current Switch Configuration** — The configuration setting on the switch.
- **Change Navigator** legend — The **Change Navigator** legend displays when there is at least one change between two compared files.
  - Number of changes label — Indicates the number of changes. If there are no differences, displays “No change”.
  - Differences legend — Displays the color legend for differences:
    - Changed status displays in pink.
    - Inserted status displays in green.

4. Click **Close** on the **Configuration Drifts** dialog box.
5. Click **Close** on the **COMPASS** dialog box.

COMPASS dashboard widget

The Management application provides the following preconfigured COMPASS widget:

COMPASS Drifts widget

The **COMPASS Drifts** widget displays the number of fabric configuration drifts and failed checks for a specified network scope and time scope as a stacked bar graph.

FIGURE 254COMPASS Drifts widget

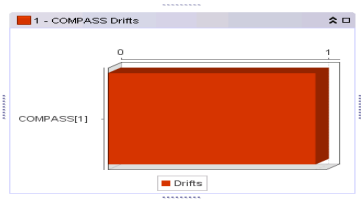

The **COMPASS Drifts** widget includes the following data:

- Widget title — The name of the widget.
- Widget summary — Displays the drift count beside the widget title.
- Bar chart — The number of fabric configuration drifts, failed checks, and violations using the color codes. The bar chart displays each group as a separate color on the chart. Tooltips showing the number of fabric configuration drifts and failed checks are shown when you pause on the bar.
- Color legend — Displays the color legend below the bar chart using the following color codes:
  - Red — Displays for drifts.
  - Blue — Displays for failed checks.

Accessing additional data from the COMPASS Drifts widget

The configuration drift data is stored in the database for 30 days. The system purges old data (over 30 days) every night at 12:00 AM.

To view configuration drifts, complete the following steps.

1. Double-click a bar in the **COMPASS Drifts** widget to navigate to the **COMPASS Drifts Detailed View** dialog box.

   The **COMPASS Drifts** dialog box displays.

2. Display data for a specific duration by selecting one of the following options from the Range list:

   - **30 Minutes** (default) — Displays data for the previous half hour beginning when the **Fabric Configuration Drifts** dialog box is displayed.
   - **1 Hour** — Displays data for the previous hour beginning when the **Fabric Configuration Drifts** dialog box is displayed.
   - **6 Hours** — Displays data for the previous 6 hours beginning when the **Fabric Configuration Drifts** dialog box is displayed.
   - **12 Hours** — Displays data for the previous 12 hours beginning when the **Fabric Configuration Drifts** dialog box is displayed.
   - **1 Day** — Displays data for the previous day beginning when the **Fabric Configuration Drifts** dialog box is displayed.
   - **3 Days** — Displays data for the previous 3 days beginning when the **Fabric Configuration Drifts** dialog box is displayed.
   - **1 Week** — Displays data for the previous week beginning when the **Fabric Configuration Drifts** dialog box is displayed.
   - **1 Month** — Displays data for the previous month beginning when the **Fabric Configuration Drifts** dialog box is displayed.

3. Review the configuration drift details.

   FIGURE 255 COMPASS Drifts dialog box

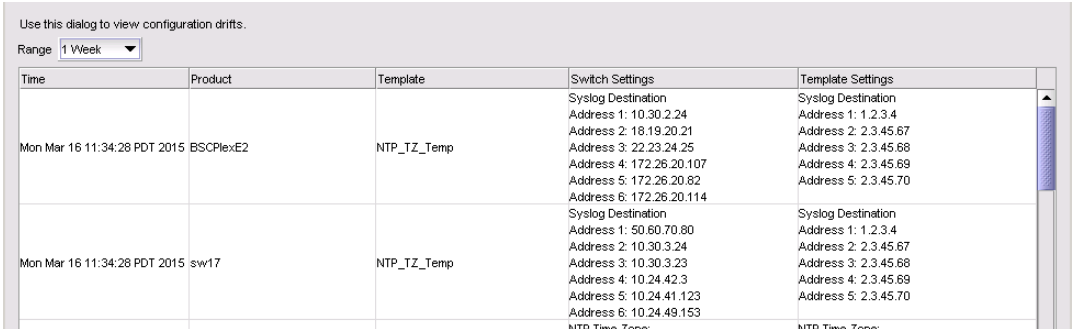

   The configuration drift data includes the following:

   - **Time** — The date and time that the drift occurred.
   - **Product** — The IP address of the product on which the drift occurred.
   - **Template** — The name of the linked template.
   - **Switch Setting** — The configuration setting on the switch.
   - **Template Setting** — The configuration setting on the template.

4. Click **Close** on the **COMPASS Drifts** dialog box.
5. Click **Close** on the **COMPASS** dialog box.

Security Management

Layer 2 access control list management

A Layer 2 access control list (ACL) enables you to filter traffic based on the information in the IP packet header using the MAC address and Ethernet type.

An ACL is a unique collection of permit and deny statements (rules) that apply to frames. You can use ACLs to permit or deny incoming frames from passing through an interface to which you assigned the ACLs. When the interface receives the frame, the device compares the fields in the frame against any ACLs assigned to the interface to verify that the frame has the required permissions to be forwarded. The device compares the frame, sequentially, against each rule in the assigned ACL. If the frame matches the permit rule, the traffic is forwarded; otherwise, the traffic is dropped.

You should configure the ACL on the device before you assign the ACL to an interface. You can create multiple ACLs and save them to the device configuration. However, the ACL does not filter traffic until you assign it to an interface. You can assign an ACL on a physical port, Virtual LAN (VLAN), or Link Aggregation Group (LAG).

For Fabric OS devices, you can create two types of ACLs:

- Standard ACL — Use to permit and deny traffic based on the source MAC address of incoming frames. You should use standard ACLs when you only need to filter traffic based on the source address.
- Extended ACL — Use to permit and deny traffic based on the source and destination MAC addresses and EtherType, of incoming frames.

Fabric OS Layer 2 ACL configuration

This section provides procedures for configuring a standard for extended Layer 2 ACL on a device, assigning the Layer 2 ACL to an interface, as well as clearing Layer 2 ACL assignments from a device.

Creating a standard Layer 2 ACL configuration (Fabric OS)

To create a standard Layer 2 ACL configuration, complete the following steps.

1. Select the device and select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* - **Layer 2 ACL Configuration** dialog box displays.

2. Select **New** from the **Add** list.

   The *Device_Name* - **Layer 2 ACL Configuration** dialog box displays.

FIGURE 256 *Device_Name* - Layer 2 ACL Configuration (Standard) dialog box

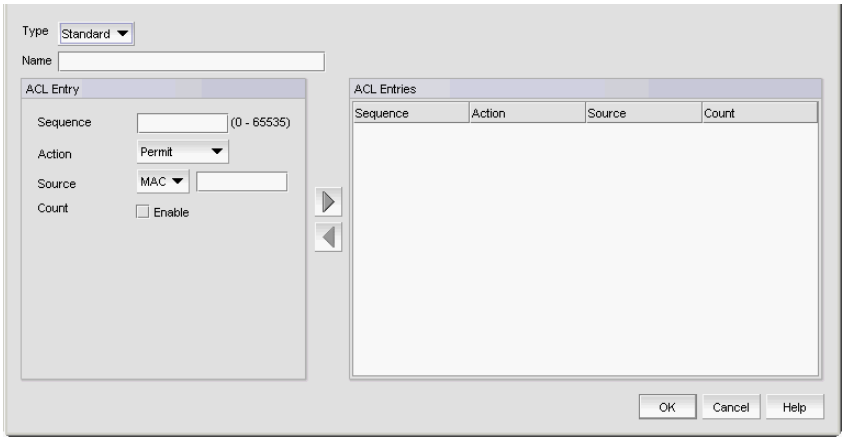

3. Select **Standard** from the **Type** list.
4. Enter a name for the ACL in the **Name** field.
5. Enter a sequence number for the ACL in the **Sequence** field.
6. Select **Permit** or **Deny** from the Action list.
7. In the **Source** list, select one of the following options:
   - **Any**
   - **MAC**

     Selecting MAC enables the **Source** field. Enter the source MAC address on which the configuration filters traffic in the **Source** field.
8. Select the **Count** check box to enable counting.

   Count specifies the number of times the ACL rule is applied.
9. Click the right arrow button.

   The new ACL entry displays in the **ACL Entries** list. To create additional ACL entries, repeat step 3 through step 9.
10. Click **OK** on the **Add - Layer 2 ACL Configuration** dialog box.

    The new ACL configuration displays in the **ACLs** list. To create additional ACLs, repeat step 2 through step 10.
11. Click **OK** on the *Device_Name* **- Layer 2 ACL Configuration** dialog box.

    The **Deploy to Products - Layer 2 ACL** dialog box displays. To save the configuration, refer to "Saving a security configuration deployment" on page 593

Editing a standard Layer 2 ACL configuration (Fabric OS)

To create a standard Layer 2 ACL configuration on a Fabric OS device, complete the following steps.

1. Select the device and select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* **- Layer 2 ACL Configuration** dialog box displays.

2. Select the ACL you want to edit in the **ACLs** list and click **Edit**.

   The *Configuration_Name* **Edit Standard Layer 2 ACL Configuration** dialog box displays.

3. To edit an existing ACL rule, complete the following steps.

   a. Select the rule you want to edit in the **ACL Entries** list and click the left arrow button.

   b. Complete step 5 through step 9 in "Creating a standard Layer 2 ACL configuration (Fabric OS)" on page 583.

      The updated ACL entry displays in the **ACL Entries** list. To edit additional ACL entries, repeat step 3.

4. To add a new ACL rule, complete step 4 through step 9 in "Creating a standard Layer 2 ACL configuration (Fabric OS)" on page 583.

   The new ACL entry displays in the **ACL Entries** list. To add additional ACL entries, repeat step 4.

5. To delete an existing ACL rule, select the rule you want to edit in the **ACL Entries** list and click the left arrow button.

6. Click **OK** on the **Edit - Layer 2 ACL Configuration** dialog box.

   The updated ACL configuration displays in the **ACLs** list. To edit additional ACLs, repeat step 2 through step 4.

7. Click **OK** on the *Device_Name* **- Layer 2 ACL Configuration** dialog box.

   The **Deploy to Products - Layer 2 ACL** dialog box displays. To save the configuration, refer to "Saving a security configuration deployment" on page 593

Duplicating a standard Layer 2 ACL configuration (Fabric OS)

To create a new Layer 2 ACL configuration from an existing Layer 2 ACL configuration a Fabric OS device, complete the following steps.

1. Select the device and select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* **- Layer 2 ACL Configuration** dialog box displays.

2. Select the ACL configuration from which you want to create a new Layer 2 ACL configuration in the **ACLs** list and click **Duplicate**.

   The **Duplicate - Layer 2 ACL Configuration** dialog box displays with the default name 'Copy of *Original_Name*'.

3. Enter a new name for the ACL in the **Name** field.

4. To edit an existing ACL rule, complete the following steps.

   a. Select the rule you want to edit in the **ACL Entries** list and click the left arrow button.

   b. Complete step 5 through step 9 in "Creating a standard Layer 2 ACL configuration (Fabric OS)" on page 583.

      The updated ACL entry displays in the **ACL Entries** list. To edit additional ACL entries, repeat step 4.

5. To add a new ACL rule, complete step 4 through step 9 in "Creating a standard Layer 2 ACL configuration (Fabric OS)" on page 583.

   The new ACL entry displays in the **ACL Entries** list. To add additional ACL entries, repeat step 5.

6. To delete an existing ACL rule, select the rule you want to edit in the **ACL Entries** list and click the left arrow button.

7. Click **OK** on the **Duplicate - Layer 2 ACL Configuration** dialog box.

   The new ACL configuration displays in the **ACLs** list. To copy additional ACLs, repeat step 2 through step 10.

8. Click **OK** on the *Device_Name* - **Layer 2 ACL Configuration** dialog box.

   The **Deploy to Products - Layer 2 ACL** dialog box displays. To save the configuration, refer to "Saving a security configuration deployment" on page 593

Creating an extended Layer 2 ACL configuration (Fabric OS)

To create an extended Layer 2 ACL configuration on a Fabric OS device, complete the following steps.

1. Select the device and select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* - **Layer 2 ACL Configuration** dialog box displays.

2. Select **New** from the **Add** list.

   The *Device_Name* - **Layer 2 ACL Configuration** dialog box displays.

3. Select **Extended** from the **Type** list.

   FIGURE 257 *Device_Name* - Layer 2 ACL Configuration (Extended) dialog box

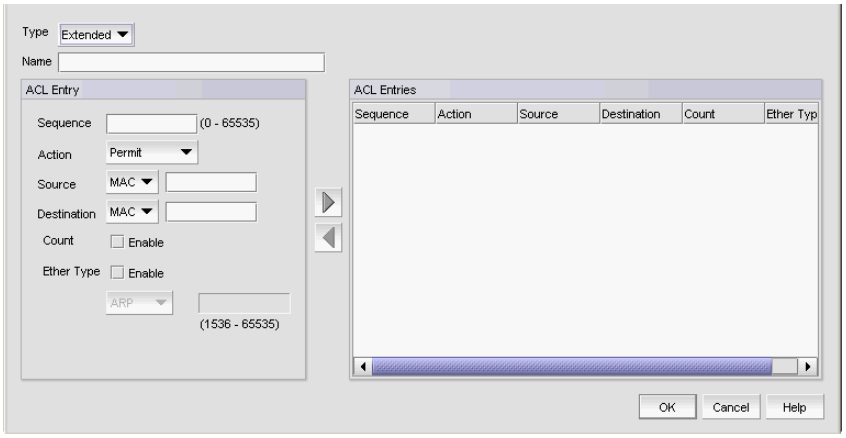

4. Enter a name for the ACL in the **Name** field.
5. Enter a sequence number for the ACL in the **Sequence** field.
6. Select **Permit** or **Deny** from the Action list.
7. In the **Source** list, select one of the following options:
   - **Any**
   - **Host**
   - **MAC**

     Selecting MAC or Host enables the **Source** field. Enter the source address on which the configuration filters traffic in the **Source** field.
8. In the **Destination Address** list, select one of the following options:
   - **Any**
   - **Host**

- **MAC**

  Selecting MAC or Host enables the **Destination** field. Enter the destination address on which the configuration filters traffic in the **Destination** field.

9. Select the **Count** check box to enable counting.

   Count specifies the number of packets filtered (allowed or denied) for the ACL rule.

10. Select the **Ether Type** check box to specify the Ethernet protocol.
11. In the **Ether Type** list, select one of the following to specify the Ethernet type being transferred in the Ethernet frame:
    - **ARP** — Address Resolution Protocol
    - **FCoE** — Fibre Channel over Ethernet
    - **IPV4**- — Internet Protocol, version 4
    - **Custom** — Enter a custom protocol. Valid values are 1536 through 65535.
12. Click the right arrow button.

    The new ACL entry displays in the **ACL Entries** list. To create additional ACL entries, repeat step 5 through step 12.

13. Click **OK** on the **Add - Layer 2 ACL Configuration** dialog box.

    The new ACL displays in the **ACL Entries** list. To create additional ACL entries, repeat step 2 through step 13.

14. Click **OK** on the *Device_Name* **- Layer 2 ACL Configuration** dialog box.

    The **Deploy to Products - Layer 2 ACL** dialog box displays. To save the configuration, refer to "Saving a security configuration deployment" on page 593

Editing an extended Layer 2 ACL configuration (Fabric OS)

To edit an extended Layer 2 ACL configuration on a Fabric OS device, complete the following steps.

1. Select the device and select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* **- Layer 2 ACL Configuration** dialog box displays.

2. Select the ACL you want to edit in the **ACLs** list and click **Edit**.

   The *Configuration_Name* **Edit Extended Layer 2 ACL Configuration** dialog box displays.

3. To edit an existing ACL rule, complete the following steps.
   a. Select the rule you want to edit in the **ACL Entries** list and click the left arrow button.
   b. Complete step 5 through step 12 in "Creating an extended Layer 2 ACL configuration (Fabric OS)" on page 586.

      The updated ACL entry displays in the **ACL Entries** list. To edit additional ACL entries, repeat step 3.

4. To add a new ACL rule, complete step 4 through step 12 in "Creating an extended Layer 2 ACL configuration (Fabric OS)" on page 586.

   The new ACL entry displays in the **ACL Entries** list. To add additional ACL entries, repeat step 4.

5. To delete an existing ACL rule, select the rule you want to edit in the **ACL Entries** list and click the left arrow button.
6. Click **OK** on the **Edit - Layer 2 ACL Configuration** dialog box.

   The updated ACL displays in the **ACL Entries** list. To edit additional ACLs, repeat step 2 through step 6.

7. Click **OK** on the *Device_Name* **- Layer 2 ACL Configuration** dialog box.

   The **Deploy to Products - Layer 2 ACL** dialog box displays. To save the configuration, refer to "Saving a security configuration deployment" on page 593

Duplicating an extended Layer 2 ACL configuration (Fabric OS)

To create a new extended Layer 2 ACL configuration from an existing extended Layer 2 ACL configuration, complete the following steps.

1. Select the device and select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* **- Layer 2 ACL Configuration** dialog box displays.

2. Select the extended Layer 2 ACL configuration from which you want to create a new extended Layer 2 ACL configuration in the **ACLs** list and click **Duplicate**.

   The **Duplicate - Layer 2 ACL Configuration** dialog box displays with the default name 'Copy of *Original_Name*'.

3. Enter a new name for the ACL in the **Name** field.

4. To edit an existing ACL rule, complete the following steps.

   a. Select the rule you want to edit in the **ACL Entries** list and click the left arrow button.

   b. Complete step 5 through step 12 in "Creating an extended Layer 2 ACL configuration (Fabric OS)" on page 586.

      The updated ACL entry displays in the **ACL Entries** list. To edit additional ACL entries, repeat step 4.

5. To add a new ACL rule, complete step 4 through step 12 in "Creating an extended Layer 2 ACL configuration (Fabric OS)" on page 586.

   The new ACL entry displays in the **ACL Entries** list. To add additional ACL entries, repeat step 5.

6. To delete an existing ACL rule, select the rule you want to edit in the **ACL Entries** list and click the left arrow button.

7. Click **OK** on the **Duplicate - Layer 2 ACL Configuration** dialog box.

   The new ACL displays in the **ACL Entries** list. To copy additional ACLs, repeat step 2 through step 7.

8. Click **OK** on the *Device_Name* **- Layer 2 ACL Configuration** dialog box.

   The **Deploy to Products - Layer 2 ACL** dialog box displays. To save the configuration, refer to "Saving a security configuration deployment" on page 593

Assigning a Layer 2 ACL configuration to an interface (Fabric OS)

To assign Layer 2 ACL configuration to a interface, complete the following steps.

1. Select **Configure > Security > Layer 2 ACL > Port**.

   The **Port Selection - Layer 2 ACL** dialog box displays.

2. Select a port or Link Aggregation Group (LAG) in the **Available Ports** list and click the right arrow button.

   LAGs display in the **Available Ports** list using the following convention: Po *LAG_Number.*

3. Click **OK**.

   The *Device_Name* **-** *Port_Number*/**LAG** *LAG_Number***- Layer 2 ACL Configuration** dialog box displays.

FIGURE 258 *Device_Name - Port_Number* - Layer 2 ACL Configuration dialog box

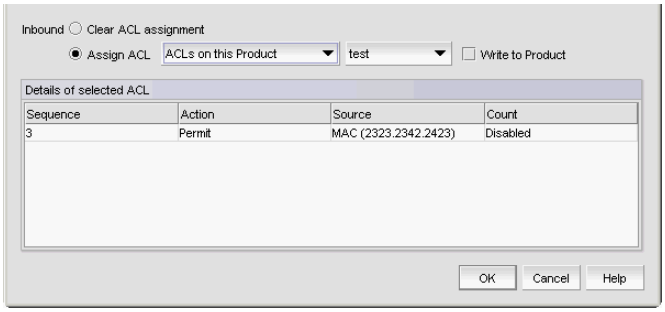

4. Select the **Assign ACL** option and choose one of the following options from the first **Assign ACL** list:
   - Select **ACLs on this Product** to assign ACLs deployed on the product to the port.

     The second list is populated with the ACLs deployed on the switch or associated with a save deployment object.
   - Select **ACLs bound to this port** to assign ACLs bound to the interface to the port.

     The second list is populated with the ACLs bound to the interface.
   - Select *Deployment_Name* (a user-configured deployment) to assign a user-configured deployment on the port.
5. Select the ACL you want to assign to the port from the second **Assign ACL** list.
6. Select the **Write to Product** check box to create the selected ACL on the device if it does not already exist.
7. Click **OK** on the *Device_Name - Port_Number* - **Layer 2 ACL Configuration** dialog box.

   The **Deploy to Ports - Layer 2 ACL** dialog box displays. To deploy the configuration, refer to "Security configuration deployment" on page 591.

Clearing Layer 2 ACL assignments (Fabric OS)

To clear Layer 2 ACL configuration from interfaces, complete the following steps.

1. Select **Configure > Security > Layer 2 ACL > Port**.

   The **Port Selection - Layer 2 ACL** dialog box displays.
2. Select a port or LAG in the **Available Ports** list and click the right arrow button.

   LAGs display in the **Available Ports** list using the following convention: Po *LAG_Number.*
3. Click **OK**.

   The *Device_Name - Port_Number*/LAG *LAG_Number* - **Layer 2 ACL Configuration** dialog box displays.
4. Select the **Clear ACL Assignment** option.
5. Click **OK** on the *Device_Name - Port_Number*/LAG *LAG_Number* - **Layer 2 ACL Configuration** dialog box.

The **Deploy to Ports - Layer 2 ACL** dialog box displays. To deploy the configuration, refer to "Security configuration deployment" on page 591.

Creating a Layer 2 ACL from a saved configuration

To create a Layer 2 ACL from a saved configuration, complete the following steps.

1. Select the device and select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* **- Layer 2 ACL Configuration** dialog box displays.

2. Select **From Saved Configurations** from the **Add** list.

   The **Layer 2 ACL Saved Configurations** dialog box displays.

3. Select one or more configurations to add to the new Layer 2 ACL configuration.
4. Click **OK** on the **Layer 2 ACL Saved Configurations** dialog box.

   The new ACL displays in the **ACLs** list.

5. Click **OK** on the *Device_Name* **- Layer 2 ACL Configuration** dialog box.

   The **Deploy to Products - Layer 2 ACL** dialog box displays. To save the configuration, refer to "Saving a security configuration deployment" on page 593

Deleting a Layer 2 ACL configuration from the application

To delete a Layer 2 ACL configuration from the application, complete the following steps.

1. Select the device and select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* **- Layer 2 ACL Configuration** dialog box displays.

2. Select the Layer 2 ACL you want to delete in the **ACLs** list and click **Delete**.

   This deletes the Layer 2 ACL configuration from the application.

3. Click **Yes** on the confirmation message.
4. Click **OK** on the *Device_Name* **- Layer 2 ACL Configuration** dialog box.

   NOTE
   The Layer 2 ACL configuration is not deleted from the switch until you deploy the configuration to the switch.

   The **Deploy to Products - Layer 2 ACL** dialog box displays. To save the configuration, refer to "Saving a security configuration deployment" on page 593

Deleting a Layer 2 ACL configuration from the switch

To delete a Layer 2 ACL configuration from the switch, complete the following steps.

1. Select the device and select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* **- Layer 2 ACL Configuration** dialog box displays.

2. Select the **Incremental** option as the configuration type.

3. Select **Delete** from the **Operation** list for the Layer 2 ACL configuration you want to delete.
4. Click **OK** on the *Device_Name* **- Layer 2 ACL Configuration** dialog box.

   NOTE
   The Layer 2 ACL configuration is not deleted from the switch until you deploy the configuration to the switch.

   The **Deploy to Products - Layer 2 ACL** dialog box displays. To save the configuration, refer to "Saving a security configuration deployment" on page 593.

Security configuration deployment

Figure 259 shows the standard interface used to deploy security configurations.

FIGURE 259 Deploy to Product/Ports dialog box

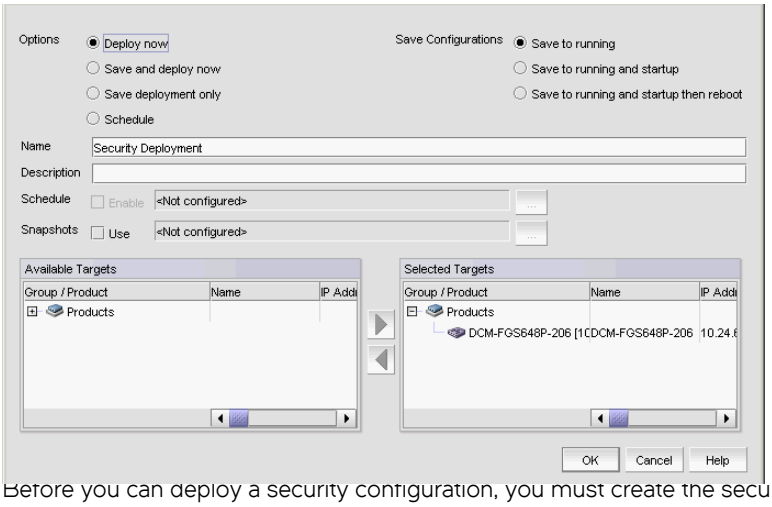

Before you can deploy a security configuration, you must create the security configuration. For step-by-step instructions, refer to the following sections:

Security Management enables you to configure, persist, and manage a security configuration as a "deployment configuration object". A deployment configuration object is comprised of the following parts:

- Security configuration (Layer 2 ACL)
- Target information
- Deployment option
- Persistence option
- Scheduling option
- Snapshot option

To create a deployment configuration object, you must save the deployment. Once you create a deployment configuration object, you can access the security configuration from the Deployment manager. For more information about the Deployment manager, refer to "Deployment Manager" on page 939.

Deploying a security configuration on demand

To deploy a security configuration immediately, complete the following steps.

FIGURE 260 Deploy to Product/Ports dialog box

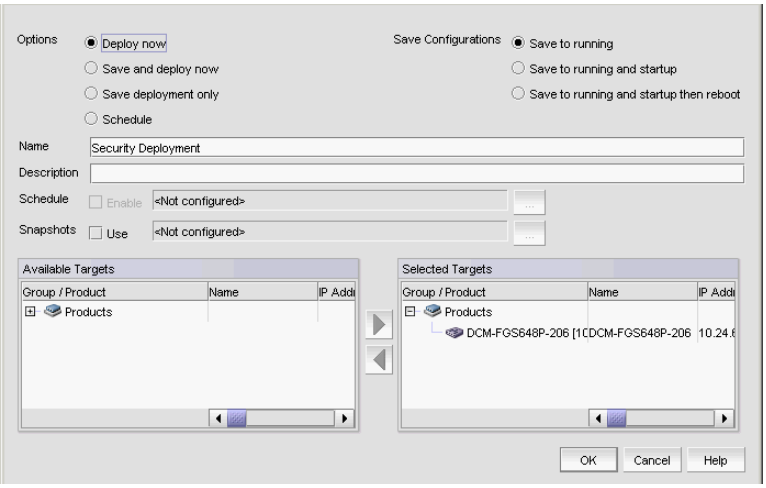

1. Choose one of the following options:
   - **Deploy now** — Select to deploy the configuration immediately on the product or port without saving the deployment definition.
   - **Save and deploy now** — Select to deploy the configuration immediately on the product or port and save the deployment definition for future deployment.
2. Select one of the following save configuration options:
   - **Save to running** — Select to update the running configuration; however, the deployment is not saved to the product's flash memory.
   - **Save to running and startup** — Select to update the running configuration as well as save the deployment configuration to the product's flash memory. Selecting this option is the equivalent to a write memory command on the product CLI.
   - **Save to running and startup then reboot** — Select to update the running configuration, save the deployment configuration to the product's flash memory, and reboot the product. Selecting this option is the equivalent to entering a write memory and a reload command on the product CLI.
3. Enter a name for the deployment in the **Name** field.
4. Enter a description for the deployment in the **Description** field.
5. Select one or more ports or products to which you want to deploy the configuration in the **Available Targets** list and click the right arrow button to move them to the **Selected Targets** list.
6. Click **OK** on the **Deploy to Products - Layer 2 ACL** dialog box.

Saving a security configuration deployment

To save a security configuration deployment, complete the following steps.

FIGURE 261 Deploy to Product/Ports dialog box

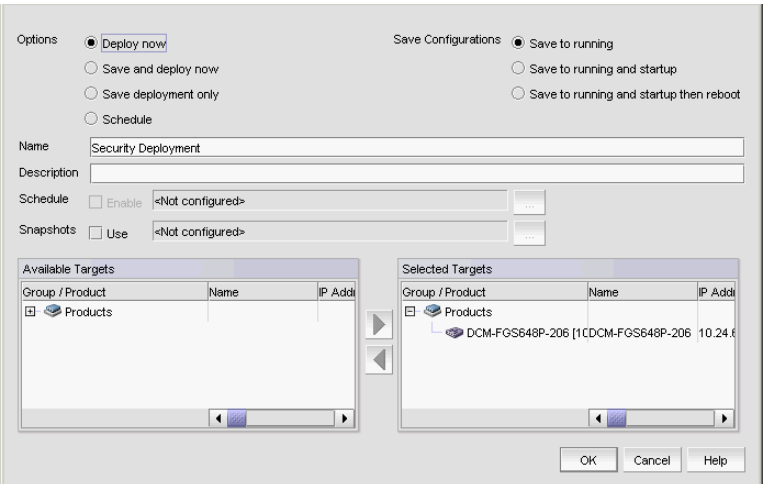

1. Select the **Save deployment only** option to save the deployment definition for future deployment.
2. Select one of the following save configuration options:
   - **Save to running** — Select to update the running configuration; however, the deployment is not saved to the product's flash memory.
   - **Save to running and startup** — Select to update the running configuration as well as save the deployment configuration to the product's flash memory. Selecting this option is the equivalent to a write memory command on the product CLI.
   - **Save to running and startup then reboot** — Select to update the running configuration, save the deployment configuration to the product's flash memory, and reboot the product. Selecting this option is the equivalent to entering a write memory and a reload command on the product CLI.
3. Enter a name for the deployment in the **Name** field.
4. Enter a description for the deployment in the **Description** field.
5. Select one or more ports or products to which you want to deploy the configuration in the **Available Targets** list and click the right arrow button to move them to the **Selected Targets** list.
6. Click **OK** on the **Deploy to Products - Layer 2 ACL** dialog box.

Scheduling a security configuration deployment

To schedule a security configuration deployment, complete the following steps.

FIGURE 262 Deploy to Product/Ports dialog box

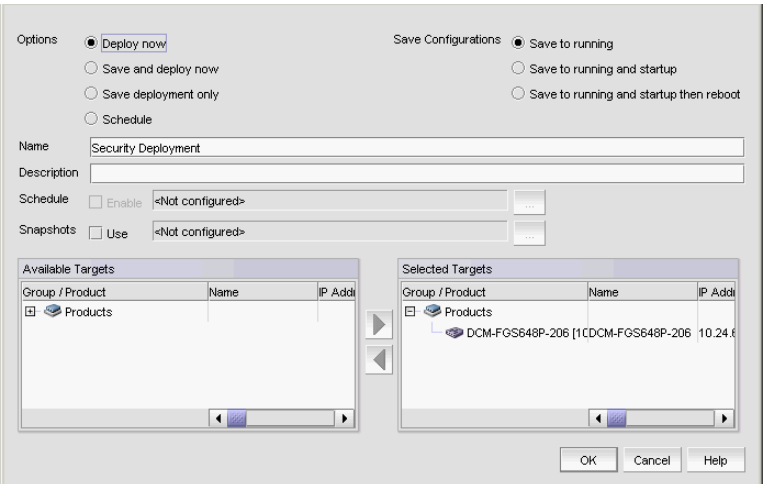

1. Select **Configure > Security > Layer 2 ACL > Product**.

   The *Device_Name* **- Layer 2 ACL Configuration** dialog box displays.

2. Choose one of the following options:

   - Select **New** from the **Add** list.

     The **Add - Layer 2 ACL Configuration** dialog box displays.

   - Select an ACL in the list and click **Edit**.

     The **Edit - Layer 2 ACL Configuration** dialog box displays.

3. Configure the Layer 2 ACL and click **OK** on the **Add/Edit - Layer 2 ACL Configuration** dialog box.
4. Click **OK** on the *Device_Name* **- Layer 2 ACL Configuration** dialog box.

   The **Deploy to Products - Layer 2 ACL** dialog box displays.

5. Select the **Schedule** option.
6. Select one of the following save configuration options:

   - **Save to running**
   - **Save to running and startup**
   - **Save to running and startup then reboot**

7. Enter a name for the deployment in the **Name** field.
8. Enter a description for the deployment in the **Description** field.
9. Click the **Schedule Enable** check box and click the ellipsis button to schedule deployment.

   The **Schedule Properties** dialog box displays.

10. Choose one of the following options to configure the frequency at which deployment runs for the schedule:
    - To configure deployment to run only once, refer to "Configuring a one-time deployment schedule" on page 595.
    - To configure hourly deployment, refer to "Configuring an hourly deployment schedule" on page 595.
    - To configure daily deployment, refer to "Configuring a daily deployment schedule" on page 595.
    - To configure weekly deployment, refer to "Configuring a weekly deployment schedule" on page 595.
    - To configure monthly deployment, refer to "Configuring a monthly deployment schedule" on page 596.
11. Click **OK** on the **Schedule Properties** dialog box.
12. Select one or more ports or products to which you want to deploy the configuration in the **Available Targets** list and click the right arrow button to move them to the **Selected Targets** list.
13. Click **OK** on the **Deploy to Products - Layer 2 ACL** dialog box.

Configuring a one-time deployment schedule

To configure a one-time schedule, complete the following steps.

1. Select **One Time** from the **Frequency** list.
2. Select the time of day you want deployment to run from the **Time (hh:mm)** lists.

   Where the hour value is from 1 through 12, the minute value is from 00 through 59, and the day or night value is AM or PM.
3. Click the **Date** list to select a date from the calendar.

   To configure security configuration schedule, refer to step 11 of "Scheduling a security configuration deployment" on page 594.

Configuring an hourly deployment schedule

To configure an hourly schedule, complete the following steps.

1. Select **Hourly** from the **Frequency** list.
2. Select the minute past the hour you want deployment to run from the **Minutes past the hour** list.

   Where the minute value is from 00 through 59.

   To configure security configuration schedule, refer to step 11 of "Scheduling a security configuration deployment" on page 594.

Configuring a daily deployment schedule

To configure a daily deployment schedule, complete the following steps.

1. Select **Daily** from the **Frequency** list.
2. Select the time of day you want deployment to run from the **Time (hh:mm)** lists.

   Where the hour value is from 1 through 12, the minute value is from 00 through 59, and the day or night value is AM or PM.

   To configure security configuration schedule, refer to step 11 of "Scheduling a security configuration deployment" on page 594.

Configuring a weekly deployment schedule

To configure a weekly schedule, complete the following steps.

1. Select **Weekly** from the **Frequency** list.
2. Select the time of day you want deployment to run from the **Time (hh:mm)** lists.

   Where the hour value is from 1 through 12, the minute value is from 00 through 59, and the day or night value is AM or PM.
3. Select the day you want deployment to run from the **Day of the Week** list.

   To configure security configuration schedule, refer to step 11 of "Scheduling a security configuration deployment" on page 594.

Configuring a monthly deployment schedule

To configure a monthly schedule, complete the following steps.

1. Select **Monthly** from the **Frequency** list.
2. Select the time of day you want deployment to run from the **Time (hh:mm)** lists.

   Where the hour value is from 1 through 12, the minute value is from 00 through 59, and the day or night value is AM or PM.
3. Select the day you want deployment to run from the **Day of the Month** list (1 through 31).

   To configure security configuration schedule, refer to step 11 of "Scheduling a security configuration deployment" on page 594.

FC-FC Routing Service Management

Devices that support Fibre Channel routing

The FC-FC Routing Service is supported only on the following devices:

- 40-port, 8 Gbps FC Switch
- 80-port, 8 Gbps FC Switch
- 48-port, 16 Gbps FC Switch
- 96-port, 16 Gbps FC switch
- 8 Gbps Extension Switch
- 16 Gbps Extension Switch
- 64 port, 32 Gbps FC Switch
- Any of the following blades on a Backbone chassis:
  - 4 Gbps Router, Extension Blade
  - FC 8 GB 16-port Blade
  - FC 8 GB 32-port Blade
  - FC 8 GB 32-port Enhanced Blade (16 Gbps or 32 Gbps Backbone Chassis only)
  - FC 8 GB 48-port Blade
  - FC 8 GB 48-port Enhanced Blade (16 Gbps or 32 Gbps Backbone Chassis only)
  - FC 8 GB 64-port Blade
  - 8 Gbps Extension Blade
  - 16 Gbps 32-port Blade
  - 16 Gbps 48-port Blade
  - 16 Gbps 64-port Blade
  - 32 Gbps 48-port blade
  - 32 Gbps, Router Extension blade

Fibre Channel routing overview

Fibre Channel (FC) routing provides connectivity to devices in different fabrics without merging the fabrics. Using Fibre Channel routing, you can share tape drives across multiple fabrics without the administrative overhead, such as change management and network management, and scalability issues that might result from merging the fabrics.

Fibre Channel routing allows you to create logical storage area networks (LSANs) that can span fabrics. These LSANs allow Fibre Channel zones to cross physical SAN boundaries without merging the fabrics and while maintaining the access controls of zones.

Refer to the *Fabric OS Administrator's Guide* for detailed information about Fibre Channel routing.

The following terminology is used in this chapter:

FC router	A switch running the FC-FC Routing Service.
Interfabric link (IFL)	The link between an E_Port and an EX_Port, or a VE_Port and a VEX_Port.
Edge fabric	A standard Fibre Channel fabric with targets and initiators connected through an FC router to another Fibre Channel fabric.
Backbone fabric	The fabric to which the FC router belongs. An FC router connects two or more edge fabrics; a *backbone fabric* connects FC routers. A backbone fabric consists of at least one FC router and possibly a number of Fabric OS-based Fibre Channel switches. Initiators and targets in the edge fabric can communicate with devices in the backbone fabric through the FC router.
LSAN	A logical SAN that spans fabrics. An LSAN is defined by zones in two or more edge or backbone fabrics that contain the same devices. LSANs enable Fibre Channel zones to cross physical SAN boundaries without merging the fabrics while maintaining the access controls of zones.
metaSAN	The collection of all SANs interconnected with FC routers.

Figure 263 on page 599 shows a metaSAN with a backbone fabric and three edge fabrics. The backbone consists of one 4 Gbps Router, Extension Switch connecting hosts in Edge fabrics 1 and 3 with storage in Edge fabric 2 and the backbone fabric. LSANs provide device sharing between the following pairs of fabrics:

- The backbone fabric and Edge fabric 1
- Edge fabric 1 and Edge fabric 2
- Edge fabric 2 and Edge fabric 3

FIGURE 263A metaSAN with edge-to-edge and backbone fabrics

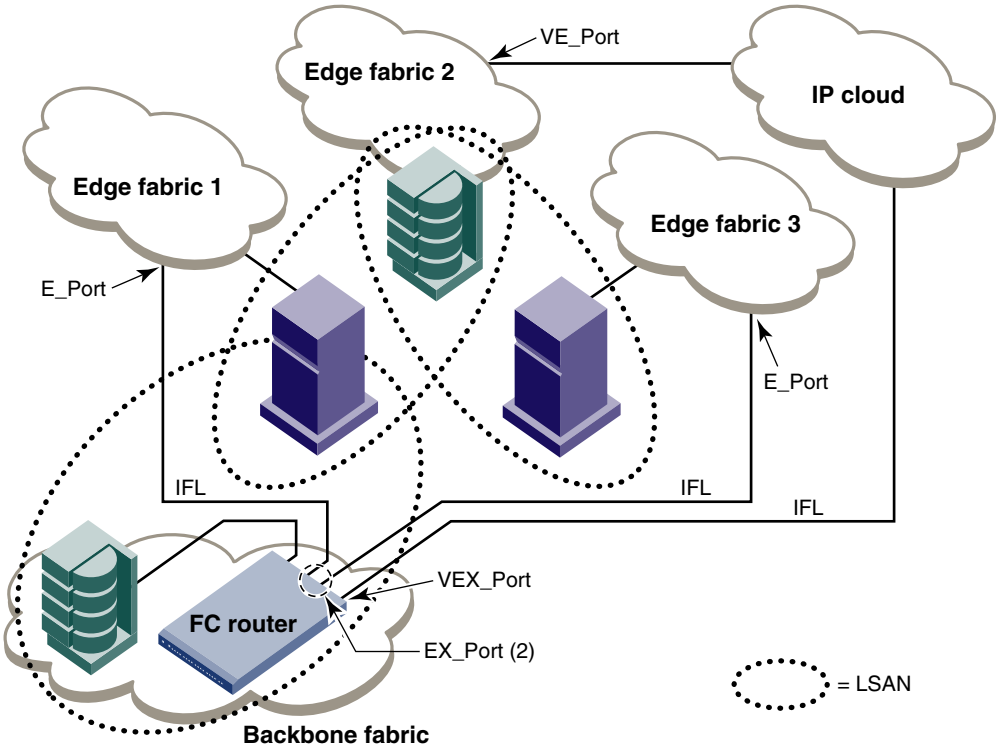

Guidelines for setting up Fibre Channel routing

The following are some general guidelines for setting up Fibre Channel routing:

- Ensure that the backbone fabric ID of the FC router is the same as that of other FC routers in the backbone fabric.
- On the FC router, ensure that the ports to be configured as EX_Ports are either disabled or not connected.
- When configuring EX_Ports, supply a fabric ID for the fabric to which the port will be connected. You can choose any unique fabric ID as long as it is consistent for all EX_Ports that connect to the same edge fabric.
- For Virtual Fabric (VF)-enabled fabrics, only the base switch can be configured as the FC router; for example, EX_Ports can be configured only on a base switch for a VF-enabled switch.

Connecting edge fabrics to a backbone fabric

The following procedure explains how to set up FC-FC routing on two edge fabrics connected through an FC router using E_Ports and EX_Ports.

> NOTE
> To configure an EX_Port, switches running Fabric OS 7.0.0 or earlier must have an FCR license. Switches running Fabric OS 7.0.1 or later configured in Brocade Native mode (IM0) or Brocade NOS mode (IM5) do not require an FCR license.

You must have an FCR license to display interfabric link (IFL). However, you do not need an IR license to display routing-enabled switches in the **Routing Configuration** and **Routing Domain Ids** dialog boxes.

For Enterprise Edition only: If you are connecting Fibre Channel SANs through an IP-based network, see "Configuring an FCIP tunnel" on page 856 for instructions on setting up an FCIP tunnel between a VE_Port and a VEX_Port.

ATTENTION
Be sure that you do not physically connect a port to the remote fabric before configuring it as an EX_Port; otherwise, the two fabrics merge and you lose the benefit of FC-FC routing.

1. Select the edge fabric you want to connect to an FC router from the Connectivity Map or Product List.
2. Right-click the edge fabric in the Connectivity Map or Product List and select **Router Configuration**.

   The **Router Configuration-Connect Edge Fabric** dialog box is displayed (Figure 264). The edge fabric you selected is also displayed in the title of the dialog box. Discovered extension switches capable of FC routing are displayed in the **Available Routers** list.

   NOTE
   If the configuration includes virtual fabrics, only the base switch displays in the **Available Routers** list.

FIGURE 264 Router Configuration-Connect Edge Fabric dialog box

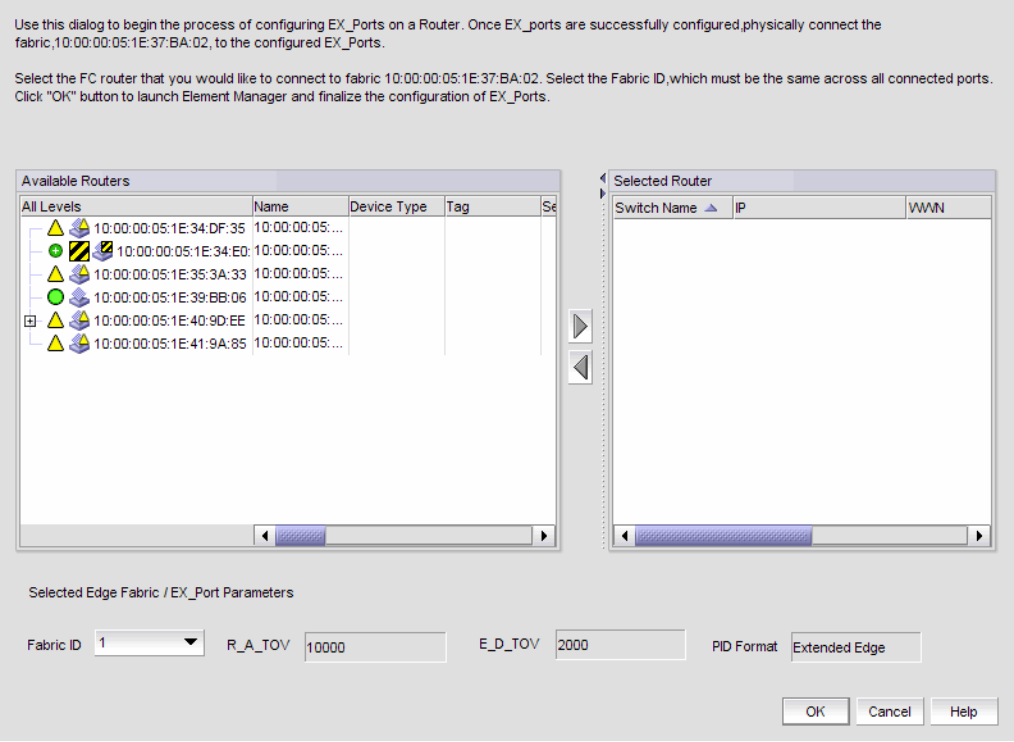

3. Select the FC router from the **Available Routers** list.
4. Click the right arrow button to move the FC router you selected to the **Selected Router** list.
5. Select a valid fabric ID from the **Fabric ID** list.

   You can choose any unique fabric ID as long as it is consistent for all EX_Ports that connect to the same edge fabric. If the edge fabric is already configured with the backbone fabric, the **Fabric ID** list is disabled and populated with the pre-selected value.

6. Click **OK** on the **Router Configuration-Connect Edge Fabric** dialog box.

   The Element Manager launches automatically and opens the **FC Router** dialog box and Port Configuration wizard. For more information, refer to the *Web Tools Administrator's Guide.*

7. Follow the instructions in the Port Configuration wizard to configure the EX_Port:

   a. Select the port to be configured as an EX_Port.

   b. Ensure the backbone fabric ID of the switch is the same as that of other FC routers in the backbone fabric. The backbone fabric ID is the fabric ID that was selected in the **Router Configuration-Connect Edge Fabric** dialog box.

   c. Complete the wizard to configure the EX_Port.

   d. Physically connect the EX_Port to the edge fabric, if it is not already connected.

8. Repeat step 1 through step 7 to connect a second edge fabric to the FC router, if your configuration involves two edge fabrics.

   A logical domain, or ***front domain***, is added in the edge fabric and is given a name in the format fcr_fd_***domainID***. For example, if the domain ID is 3, the name of the front domain is fcr_fd_3.

9. Configure LSAN zones in each fabric that will share devices.

   For specific instructions, refer to "Configuring LSAN zoning" on page 805.

Configuring routing domain IDs

Logical (phantom) domains are automatically created to enable routed fabrics. Two types of logical domains are created:

- A front domain is created in edge fabrics for every interfabric link (IFL).
- A translate (Xlate) domain is created in routed fabrics that share devices.

You can change the domain IDs of these logical domains.

1. In the Product List or Connectivity Map, right-click the fabric for which you want to configure logical domains, and select **Routing Domain IDs**.

   The **Configure Routing Domain IDs** dialog box is displayed (Figure 265).

FIGURE 265 Configure Routing Domain IDs dialog box

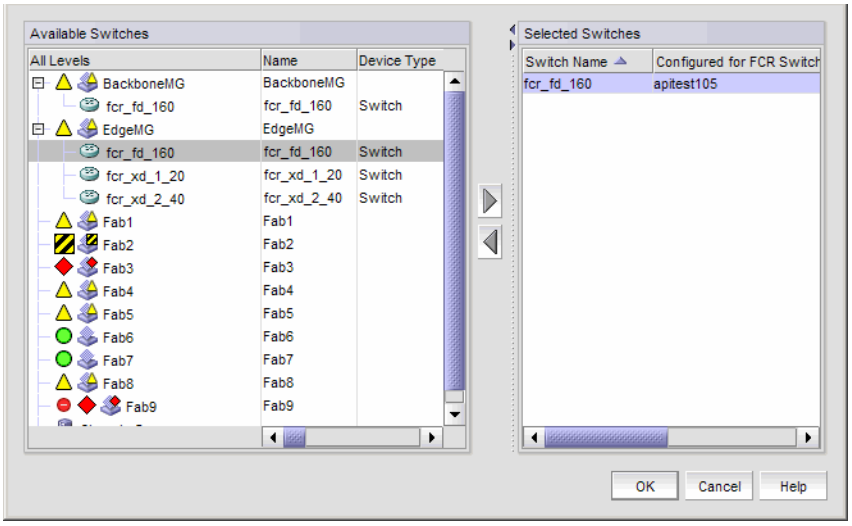

2. Right-click anywhere in the **Available Switches** list and select **Expand All** in the right-click menu.

   The switch group for the fabric expands to display the logical domains.

3. Select a logical domain, and click the right arrow button to move the switch to the **Selected Switches** list.
4. Select a domain ID number from the **Domain ID** column in the **Selected Switches** list. The **Domain ID** column lists unused domain IDs.

   You may need to scroll right or drag the dialog box open further to see the **Domain ID** column.

5. Click **OK**.

Virtual Fabrics

Virtual Fabrics overview

NOTE
Virtual Fabrics requires that you have at least one Virtual Fabrics-enabled physical chassis running Fabric OS 7.0 or later in your SAN.

Virtual Fabrics enables you to divide one physical chassis into multiple logical switches that can be managed by separate administrators. Logical switches consist of one or more ports that act as a single FC switch. You can interconnect logical switches to create a logical fabric.

The following lists the benefits of using the Management application to manage Virtual Fabrics:

- Enables you to view your entire SAN (both physical and virtual) at a glance.
- Enables you to manage a logical switch the same as a physical switch, so that fewer physical chassis are required for Management application deployment.
- Enables you to use a logical switch for discovery and eliminate the requirement for one physical chassis for each fabric.
- Enables you to manage multiple Virtual Fabrics-capable physical chassis from the same interface.
- Enables you to provide logical isolation of data, control, and management paths at the port level.

You can easily determine which devices in your SAN are logical switches. Logical switches are shown with a Virtual Fabrics icon ( ) next to the switch icon. For example, in Figure 266, Switch_2 is a logical switch.

FIGURE 266 Virtual Fabrics icon in Product List

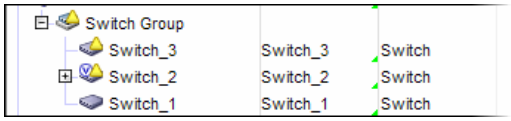

Before using the Management application to manage Virtual Fabrics, you should familiarize yourself with Virtual Fabrics concepts, as described in the *Fabric OS Administrator's Guide*.

Terminology for Virtual Fabrics

Table 56 lists definitions of Virtual Fabrics terms.

TABLE 56 Virtual Fabrics terms

Term	Definition
Physical chassis	The physical switch or chassis from which you create logical switches and fabrics.
Logical switch	A collection of ports that act as a single Fibre Channel (FC) switch. When Virtual Fabrics is enabled on the chassis, there is always at least one logical switch: the default logical switch. You must assign each logical switch (default or general) in the same chassis to a different logical fabric. The logical switch supports all E_Ports and F_Ports. Note that EX_Ports are only allowed on the base switch.

TABLE 56 Virtual Fabrics terms (Continued)

Term	Definition
Default logical switch	A logical switch that is created automatically when the Virtual Fabrics feature is enabled in a physical chassis. Initially, all ports in a chassis belong to the default logical switch. The default logical switch always exists as long as Virtual Fabrics is enabled. You cannot delete the default logical switch. The default logical switch supports all E_Ports and F_Ports.
Base switch	A special logical switch used to communicate among different logical switches. The legacy EX_Port is connected to the base logical switch. Inter-Switch Links (ISLs) connected to the base switch are used to communicate among different fabrics. The base switch supports E_Ports and EX_Ports.
Fabric ID (FID)	An identifier you assign to a logical switch (default or general) or a base switch to designate to which logical or base fabric it belongs.
Logical fabric	A fabric with at least one logical switch.
Base fabric	A fabric formed from base switches that have the same FID. The base fabric provides the physical connectivity across multiple segments of a fabric over which logical switches in the fabric can establish logical connectivity.
Extended ISL (XISL)	An ISL physically connected between two base switches that carries traffic for multiple logical fabrics. By default, logical switches are configured to not use XISLs. XISL use is not supported in the following cases: • Logical switches in an edge fabric connected to an FC router. • A logical switch in InteropMode 2 or InteropMode 3. • The logical switch has VE_Ports and is running Fabric OS 6.4.x or earlier. • The logical switch has lossless DLS and is running Fabric OS 7.0.x or earlier. For switches running Fabric OS 7.1.0 or later, XISL use is supported with lossless DLS. • FICON logical fabrics, for switches running Fabric OS 7.0.x or earlier. For switches running Fabric OS 7.1.0 or later, XISL use is supported when FMS mode is enabled.

Virtual Fabrics requirements

To configure Virtual Fabrics, you must have at least one Virtual Fabrics-enabled physical chassis running Fabric OS 7.0 or later in your SAN. Use one of the following options to discover a Virtual Fabrics-enabled physical chassis on the Management application topology:

- Discover a Virtual Fabrics-capable seed physical chassis running Fabric OS 7.0 or later. Virtual Fabrics is disabled by default. This physical chassis displays as a legacy switch. Once discovered, you must enable Virtual Fabrics.
- Discover a Virtual Fabrics-enabled seed physical chassis running Fabric OS 7.0 or later with Virtual Fabrics enabled, and at least one logical switch defined on the core switch. The physical chassis displays as a virtual switch.
- Upgrade a physical chassis already in your SAN to Fabric OS 7.0 or later. Virtual Fabrics is disabled by default. This switch displays as a legacy switch. Once upgraded, you must enable Virtual Fabrics.

For more information about enabling Virtual Fabrics on a physical chassis, refer to "Enabling Virtual Fabrics" on page 608.

Table 57 lists the Virtual Fabrics-capable physical chassis and the number of logical switches allowed for each of those physical chassis.

TABLE 57 Maximum number of logical switches per chassis

Physical chassis	Number of logical switches allowed
40-port, 8 Gbps FC Switch	3
80-port, 8 Gbps FC Switch	4
48-port, 16 Gbps FC Switch	4[1]
96-port, 16 Gbps FC Switch	4
64-port, 32 Gbps FC Switch	4
8 Gbps Extension Switch	4
16 Gbps Extension Switch	4

TABLE 57 Maximum number of logical switches per chassis (Continued)

8 Gbps Backbone Chassis	8
16 Gbps Backbone Chassis	8
32 Gbps Backbone Chassis	16

1. The maximum is 3 logical switches if you are using FC-FC routing.

NOTE
The 8 Gbps Extension Switch does not support base switches.

For the 8 Gbps Extension Switch, any port can be assigned to the logical switch or default logical switch. For the other switches, any port can be assigned to any logical switch (logical switch, default logical switch, or base switch).

Depending on the logical switch type, the backbone chassis have the port requirements shown in Table 58.

TABLE 58 Blade and port types supported on logical switches for backbone chassis

Logical switch type	Ports
Default logical switch	• Extension Blade — E_Ports, F_Ports, GE_Ports, and VE_Ports • FC 10-6 ISL Blade — E_Ports and F_Ports • FC 8 GB Port Blade — E_Ports and F_Ports • FC 16 GB Port Blade — E_Ports and F_Ports • 10 Gig FCoE port Blade — E_Ports and F_Ports • 8 Gbps Extension Blade - FC ports: E_Ports, F_Ports, and VE_Ports - GE ports: VE_Ports • 32 Gbps 48 Port blade — E_Ports and EX_Ports • 32 Gbps Extension Blade — GE_Ports and VE_Ports • 16 Gbps and 32 Gbps Backbone Chassis — ICL ports
Logical switch	• Extension Blade — GE_Ports and VE_Ports • FC 8 GB Port Blade — E_Ports and F_Ports • FC 16 GB Port Blade — E_Ports and F_Ports • 8 Gbps Extension Blade - FC ports: E_Ports, F_Ports, and VE_Ports - GE ports: VE_Ports • 32 Gbps 48 Port blade — E_Ports and EX_Ports • 32 Gbps Extension Blade — GE_Ports and VE_Ports • 16 Gbps and 32 Gbps Backbone Chassis — ICL ports
Base switch	• Extension Blade — GE_Ports and VEX_Ports • FC 8 GB Port Blade — E_Ports and EX_Ports • FC 16 GB Port Blade — E_Ports and F_Ports • 8 Gbps Extension Blade - FC ports: E_Ports, EX_Ports, VE_Ports, and VEX_Ports - GE ports: VE_Ports • FC 16 Gbps Port Blade — E_Ports and EX_Ports • 16 Gbps Extension Switch — EX_Ports • 32 Gbps 48 Port blade — E_Ports and EX_Ports • 32 Gbps Extension Blade — GE_Ports and VE_Ports • 16 Gbps and 32 Gbps Backbone Chassis — ICL Ports

NOTE
In the 8-slot Backbone Chassis, ports 48–63 of the FC 8 GB 64-port and FC 16 GB 64-port blades are not supported in the base switch, and ports 56–63 are not supported as E_Ports on the default logical switch. The 4-slot Backbone Chassis does not have these limitations.

FICON best practices for Virtual Fabrics

Use the following recommended best practices and considerations for configuring Virtual Fabrics in a FICON environment when following the procedures under "Configuring Virtual Fabrics" on page 607:

- When configuring the logical switch in the **New Logical Fabric Template or New Logical Switch** dialog box (**Fabric** tab), use the following parameters. Note that the **New Logical Fabric Template** dialog box creates a fabric template. You can always rename the fabric and change parameters after the new fabric is created.
  - **Logical Fabric ID (FID)** – Use any FID as long as all switches in a fabric have the same Fabric ID. The default Fabric ID for the default switch is 128, which leaves 1 through 127 for newly created fabrics.
  - **256 Area Limit** – "Disabled" is not supported for FICON. As a recommended best practice, use "Zero Based Area Assignment" as this will work for any configuration.
  - **R_A_TOV**, **E_D_TOV**, **WAN_TOV**, **Maximum Hops**, **BB Credit**, **Data Field Size** – Do not change these parameters unless otherwise directed by your switch service provider. Any change to these parameters is a rare case.
  - **Interoperability Mode** – With Fabric OS 7.0.0 and later, only "Brocade Native" mode is supported, so this parameter cannot be changed.
  - **Base Switch** or **Base Fabric for Transport (XISL)** - Do not select these check boxes as they are not supported for FICON.
  - **Allow XISL Use (XISL)** – Select this check box to enable "Disable LISL Ports" check box.
  - **Disable LISL Ports** – Select this check box to create LISL ports and continue to be in offline state.
  - **Sequence Level Switching**, **Per-Frame Routing Priority**, **Suppress Class F Traffic** – Do not select these check boxes.
  - **Disable Device Probing** – When selected, third-party software, except for CUP, is prohibited from managing the switch. This check box should be selected unless otherwise advised by your switch service provider.
  - **Long Distance Fabric** – This parameter sets E_Ports to LD mode (increases BB credits for long distance performance). Select this check box only when ISLs between the switch and a connected device exceed 10 Km. Dense wave division multiplexing (DWDM) equipment usually provides BB credits, so there is typically no reason for additional BB credits unless there are direct ISLs between switches or course wave division multiplexing (CWDM) is being used. Long Distance Fabric requires a license.
  - **FICON** – Select this check box to create a FICON logical switch and to change an existing logical switch to a FICON logical switch. FICON logical switch creation is independent of FMS mode.
- When configuring the logical fabric in the **New Logical Fabric Template or New Logical Switch** dialog box (**Switch** tab), use the following parameters:
  - **Preferred Domain ID** – Use a unique domain ID for all switches. Domain IDs are entered in either decimal or hexadecimal. If you enter the domain ID in decimal, ensure you use the correct hexadecimal equivalent. For example, if the first byte of the link address is 33, then the domain ID in decimal is 51. Also, use a domain ID that is the hexadecimal equivalent of the Switch ID in the input/output completion port (IOCP). For example, for Switch ID 1F, set Domain ID to 31 in decimal or 1F in hexadecimal.
  - **Insistent** – As a best practice, select this check box to not allow the domain ID to be changed when a duplicate domain ID exists. Although an insistent domain ID is only required when 2-byte link addressing is used on the host, setting **Insistent** for all environments is the recommended best practice. Setting this parameter does not cause any problems, but not selecting it can cause problems if 2-byte addressing is used in the future.
- When the **Logical Switch Change Conformation and Status** dialog box displays after configuring logical switches through the **Logical Switches** dialog box, be sure the following parameters are selected:
  - **Re-Enable ports after moving them (If the selected ports have more than 125 devices, the Re-enable Ports after moving them option is disabled.)**
  - **Unbind Port Addresses while moving them**
  - **QoS disable the ports while moving them**.

  If you do not select the **Unbind Port Addresses while moving them** check box, the port address is "remembered" by the switch from where it was moved and cannot be assigned to another port. This is rarely desired when configuring switches for FICON applications. Also, because it is not obvious that the address is in memory, not selecting this option can cause confusion when making future changes.

- Configure at least one logical switch and move all FICON ports to that logical switch, even if that means moving all ports in the chassis.
- Enabling or disabling Virtual Fabrics is disruptive as it requires you to reboot the switch. If the switch is in a production environment, make sure all channel connections to the switch have been configured offline first.
- As a best practice, do not change a production fabric unless there is a compelling reason to do so. For new installations, the recommended best practice is to enable Virtual Fabrics. Even if you do not plan to use the chassis for more than a single logical switch, you have the option of adding a logical switch in the future without an outage.
- Create at least one logical switch for FICON connections.
- Fibre Channel ports on the 8 Gbps Extension Blade can be placed in any logical switch. The default switch should only be used for FICON connections when FC ports on a 4 Gbps Router, Extension blade are required for FICON. FICON connections are not supported in the default switch for 48-port blades in a 4-slot or 8-slot Backbone Chassis.

Configuring Virtual Fabrics

The Management application allows you to discover, enable, create, and manage Virtual Fabrics-capable physical chassis from the same interface.

NOTE
For FICON, it is recommended to limit it to four operations at a time.

NOTE
Although Fabric OS v8.1.0 supports 16 Virtual Fabrics per Gen6 chassis, you can only configure maximum of 8 virtual fabrics in a single operation using Logical Switches dialog box.

This procedure describes the general steps you take to enable the Virtual Fabrics feature and configure logical fabrics. The logical fabrics in this example span multiple physical chassis, and the logical switches in each fabric communicate using an XISL in the base fabric.

1. Enable Virtual Fabrics in each physical chassis.

   Refer to "Enabling Virtual Fabrics" on page 608 for instructions.

2. Set up base switches in each physical chassis.

   a. Create base switches in each physical chassis and assign ports to them.

      Refer to "Creating a logical switch or base switch" on page 608 for instructions.

   b. Disable the base switches in each physical chassis.

      Right-click each base switch in the Connectivity Map or Product List and select **Enable/Disable > Disable**.

   c. Physically connect ports in the base switches to form XISLs.

   d. Enable all of the base switches. This forms the base fabric.

      Right-click each base switch in the Connectivity Map or Product List and select **Enable/Disable > Enable**.

3. Set up logical switches in each physical chassis.

   a. Create logical switches in each physical chassis and assign ports to them. Make sure the logical switches are configured to allow XISL use.

      Refer to "Creating a logical switch or base switch" on page 608 for instructions.

   b. Disable all of the logical switches in each physical chassis.

      Right-click each logical switch in the Connectivity Map or Product List and select **Enable/Disable > Disable**.

c. Physically connect devices and ISLs to the ports on the logical switches.

   You can connect ISLs from one logical switch to another logical switch in a different physical chassis only if the two logical switches have the same FID (and are thus in the same logical fabric). Traffic between these logical switches can travel over either this ISL or the XISL in the base fabric. The physical ISL path is favored over the XISL path because it has a lower cost.

d. Enable all logical switches in each chassis.

   Right-click each logical switch in the Connectivity Map or Product List and select **Enable/Disable > Enable**.

The logical fabric is formed.

Enabling Virtual Fabrics

For a list of platforms that are Virtual Fabrics-capable, refer to "Virtual Fabrics requirements" on page 604.

ATTENTION
If the physical chassis is participating in a fabric, the affected fabric will be disrupted.

1. Select the physical chassis in the topology and select **Configure > Virtual Fabric > Enable**.

   Alternatively, you can right-click the physical chassis and select **Enable Virtual Fabric**.

2. Read the warning message and click **OK**.

Disabling Virtual Fabrics

ATTENTION
Disabling Virtual Fabrics deletes all logical switches, returns port management to the physical chassis, and reboots the physical chassis. If these logical switches are participating in a fabric, all affected fabrics will be disrupted.

1. Select the physical chassis in the Chassis Group and select **Configure > Virtual Fabric > Disable**.

   Alternatively, you can right-click the physical chassis in the Chassis Group and select **Disable Virtual Fabric**.

2. Read the warning message and click **OK**.

Creating a logical switch or base switch

Before you can create a logical switch, you must enable Virtual Fabrics on at least one physical chassis in your fabric.

Optionally, you can define the logical switch to be a base switch. Each chassis can have only one base switch.

NOTE
The 8 Gbps Extension Switch does not support base switches.

1. Select **Configure > Virtual Fabric > Logical Switches**.

   The **Logical Switches** dialog box displays.

FIGURE 267 Logical Switches dialog box

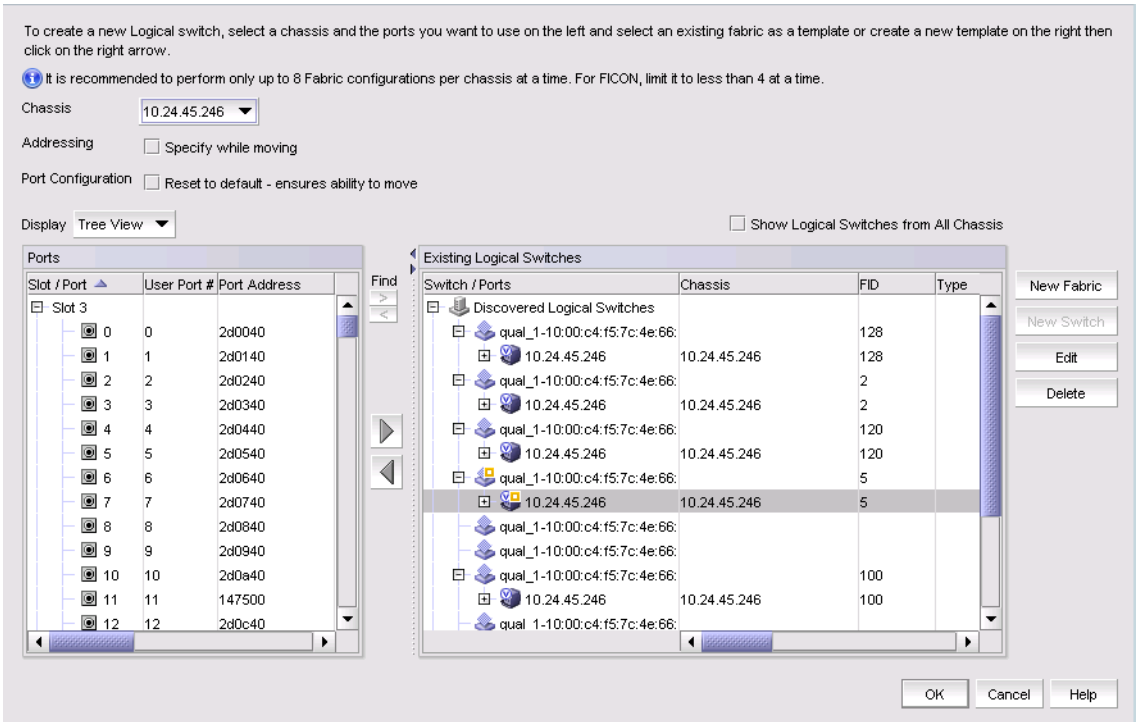

2. Select the physical chassis from which you want to create a logical switch in the **Chassis** list.

   You can display all logical switches from all chassis by selecting the **Show Logical Switches from all Chassis** check box.

3. Select one of the following in the **Existing Logical Switches** list:

   - A physical chassis in the Discovered Logical Switches node
   - A NewFabric logical switch template in the Discovered Logical Switches node
   - The Undiscovered Logical Switches node

   If you select a logical switch template, the fabric-wide settings for the logical switch are obtained from the settings in the template.

   If you select a physical chassis or the Undiscovered Logical Switches node, the fabric-wide settings for the logical switch are the default settings.

4. Click **New Switch**.

   The **New Logical Switch** dialog box displays.

FIGURE 268 New Logical Switch dialog box

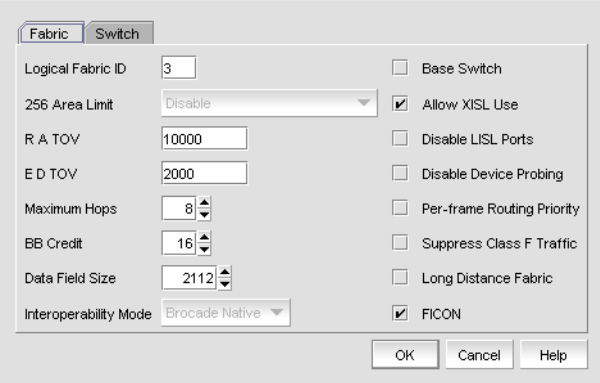

5. Click the **Fabric** tab and enter fabric-wide parameters.
   a. Enter a fabric identifier in the **Logical Fabric ID** field.

      This assigns the new logical switch to a logical fabric.

      If the logical fabric does not exist, this creates a new logical fabric as well as assigning the new logical switch.

   b. Enter new values for the fabric-wide parameters or leave the parameters unchanged to accept the current values.

      Click the **Help** button for detailed information on each parameter.

   c. (*Optional*) To configure the switch to use XISLs, select the **Base Fabric for Transport** check box.

      In the following cases, make sure the **Base Fabric for Transport** check box is cleared, because XISL use is not supported:

      - FICON logical fabrics, for switches running Fabric OS 7.0.1 or earlier
      - Logical switches in an edge fabric connected to an FC router
      - A logical switch in InteropMode 2 or InteropMode 3
      - The logical switch has VE_Ports and is running Fabric OS 6.4.x or earlier
      - The logical switch has lossless DLS and is running Fabric OS 7.0.x or earlier

      NOTE
      For switches running Fabric OS 7.0.0 or later, VE_Ports on the 8 Gbps Extension Blade are supported on logical switches that use XISLs.

   d. (*Optional*) To make the logical switch a base switch, clear the **Base Fabric for Transport** check box and select the **Base Switch** check box.

      The **Base Fabric for Transport** check box is not relevant for base switches because all base switches can use XISLs.

   e. (*Optional*) For Backbone Chassis only, select an option in the **256 Area Limit** list to use 256-area addressing mode (zero-based or port-based) or to disable this mode (default).

      The 256-area addressing mode can be used in FICON environments, which have strict requirements for 8-bit area FC addresses.

6. Click the **Switch** tab and enter switch parameters.
   a. Enter a name for the logical switch in the **Name** field.
   b. Select either **Decimal** or **Hex** and enter a preferred domain ID for the logical switch.

      In a FICON environment, select a domain ID that is not in use by the default or another logical switch in the same chassis.

c. (*Optional*) Select the **Insistent** check box to not allow the domain ID to be changed when a duplicate domain ID exists.

If you select this check box and a duplicate domain ID exists, the switch will segment from the fabric instead of changing the domain ID.

NOTE
If you disable the Insistent Domain ID (IDID) on the FMS-enabled switches, the following error message displays: The operation cannot be performed. Disable HIF mode and try again.

d. (Optional) Select **Disable LISL Ports** check box to disable the LISL port.

e. (Optional) Select **FICON** check box to change an existing logical switch to a FICON logical switch. To perform this operation HIF mode must be enabled and the switch address mode must be in zero-based addressing.

7. Click **OK** on the **New Logical Switch** dialog box.

The new logical switch displays in the **Existing Logical Switches** list (already highlighted). This logical switch has no ports.

The newly created logical switch has no ports. To assign ports to the logical switch, refer to "Assigning ports to a logical switch" on page 612.

If the newly created logical switch is not part of a discovered fabric, then you must undiscover and rediscover the switch.

- To undiscover the physical chassis, refer to "Deleting a fabric" on page 42 for instructions.
- To rediscover the physical chassis, refer to "Discovering fabrics" on page 35 for instructions.

  When entering the IP address, use the IP address of the physical fabric.

Finding the physical chassis for a logical switch

The Management application enables you to locate the physical chassis in the Product List from which the logical switch was created.

To find the physical chassis for a logical switch, right-click the logical switch in the Connectivity Map or Product List and select **Virtual Fabric > Chassis**.

The physical chassis is highlighted in the Product List.

Finding the logical switch from a physical chassis

The Management application enables you to locate the logical switch from the physical chassis.

1. Expand the Chassis Group node in the Product List.
2. Right-click the physical chassis within the Chassis Group.
3. Select **Virtual Fabric > Logical Switches >** *Logical_Switch_Name*.

The logical switch you selected is highlighted in the Product List and Connectivity Map.

Assigning ports to a logical switch

When you create a logical switch, it has no ports and you must explicitly assign ports to it.

When you assign a port to a logical switch, it is removed from the original logical switch and assigned to the new logical switch. All ports are initially assigned to the default logical switch.

A port can be assigned to only one logical switch.

1. Select **Configure > Virtual Fabric > Logical Switches**.

   The **Logical Switches** dialog box displays.

2. Select the physical chassis from which you want to assign ports in the **Chassis** list.
3. (*Optional*) Select the **Addressing** check box to specify the starting port address for the ports that will be moved.

   If this check box is cleared, the port addresses are set to "unassigned". The ports are assigned a system-generated port address when they are configured on the destination logical switch.

   This option is supported only for FC ports in zero-based addressing mode or 10-bit addressing mode.

4. (*Optional*) Select the **Port Configurations** check box to clear the port configurations prior to the move.

   Clearing the port configurations ensures that the port move is not blocked by configuration-related validation checks.

5. Select the ports you want to include in the logical switch from the **Ports** list.

   You can configure the **Ports** list by selecting **Table View** (list of all ports) or **Tree View** (list of ports grouped by slot) from the **Display** list.

6. Select the logical switch in the **Existing Logical Switches** list.

   To see all of the items in the **Existing Logical Switches** list, you can right-click anywhere in the list and select **Table > Expand All**.

7. Click the right arrow button to move the selected ports to the logical switch.

   If you selected the **Addressing** check box, enter the starting port address in the **Bind Port Address** dialog box.

   The ports display in the selected logical switch node in the **Existing Logical Switches** list.

8. Click **OK** on the **Logical Switches** dialog box.

   The **Logical Switch Change Confirmation and Status** dialog box displays with a list of all changes you made in the **Logical Switches** dialog box.

   The **Re-Enable ports after moving them** and **QoS disable the ports while moving them** check boxes are selected by default. If the selected ports have more than 125 devices, the **Re-enable Ports after moving them option** is disabled.

   NOTE
   Ports are disabled before moving from one logical switch to another.

9. (*Optional*) Select the **Unbind Port Addresses while moving them** check box.
10. Click **Start** to send these changes to the affected chassis.

    NOTE
    Most changes to logical switches will disrupt data traffic in the fabric.

    The status of each change is displayed in the **Status** column and **Status** area in the dialog box.

11. When the changes are complete, click **Close**.

Removing ports from a logical switch

1. Select **Configure > Virtual Fabric > Logical Switches**.

   The **Logical Switches** dialog box displays.

2. Select the physical chassis to which the ports belong in the **Chassis** list.
3. Select the ports you want to remove from the logical switches from the **Existing Logical Switches** list.

   To see all of the ports in the **Existing Logical Switches** list, you can right-click anywhere in the list and select **Table > Expand All**.

4. Click the left arrow button.

   A message displays indicating that the ports will be moved to the default logical switch.

5. Click **OK** on the warning message.

   The selected ports are removed from the logical switch and automatically reassigned to the default logical switch. The selected ports are highlighted in the **Ports** list.

6. (*Optional*) Perform the following steps to assign the ports to a logical switch other than the default logical switch.
   a. Select the destination logical switch in the **Existing Logical Switches** list.
   b. Click the right arrow button.

      The ports display in the selected logical switch node in the **Existing Logical Switches** list.

7. Click **OK** on the **Logical Switches** dialog box.

   The **Logical Switch Change Confirmation and Status** dialog box displays with a list of all changes you made in the **Logical Switches** dialog box.

   The **Re-Enable ports after moving them** and **QoS disable the ports while moving them** check boxes are selected by default. If the selected ports have more than 125 devices, the **Re-enable Ports after moving them option** is disabled.

   NOTE
   Ports are disabled before moving from one logical switch to another.

8. (*Optional*) Select the **Unbind Port Addresses while moving them** check box.
9. Click **Start** to send these changes to the affected chassis.

   NOTE
   Most changes to logical switches will disrupt data traffic in the fabric.

   The status of each change is displayed in the **Status** column and **Status** area in the dialog box.

10. When the changes are complete, click **Close**.

Deleting a logical switch

1. Select **Configure > Virtual Fabric > Logical Switches**.

   The **Logical Switches** dialog box displays.

2. Right-click anywhere in the **Existing Logical Switches** list and select **Table > Expand All**.
3. Select the logical switch you want to delete from the **Existing Logical Switches** list and click **Delete**.

   All ports in the deleted logical switch are reassigned to the default logical switch.

4. Read the confirmation message and click **Yes**.
5. Click **OK** on the **Logical Switches** dialog box.

   The **Logical Switch Change Confirmation and Status** dialog box displays with a list of all changes you made in the **Logical Switches** dialog box.

   The **Re-Enable ports after moving them** and **QoS disable the ports while moving them** check boxes are selected by default. If the selected ports have more than 125 devices, the **Re-enable Ports after moving them option** is disabled.

   NOTE
   Ports are disabled before moving from one logical switch to another.

6. (*Optional*) Select the **Unbind Port Addresses while moving them** check box.
7. Click **Start** to send these changes to the affected chassis.

   NOTE
   Most changes to logical switches will disrupt data traffic in the fabric.

   The status of each change is displayed in the **Status** column and **Status** area in the dialog box.

8. When the changes are complete, click **Close**.

Configuring fabric-wide parameters for a logical fabric

When you create a logical switch, you must assign it to a fabric and configure fabric-wide parameters. All the switches in a fabric must have the same fabric-wide settings.

Instead of configuring these settings separately on each logical switch, you can create a ***logical fabric template***, which defines the fabric-wide settings for a logical fabric. Then, when you create logical switches for that fabric, these fabric-wide settings are used automatically and you do not need to re-enter them.

Creating a logical fabric template does ***not*** create a logical fabric. A logical fabric is created only when you assign logical switches to a fabric ID (FID).

The logical fabric template exists only in the lifetime and scope of the **Logical Switches** dialog box. When you exit this dialog box, the logical fabric templates are deleted.

1. Select **Configure > Virtual Fabric > Logical Switches**.

   The **Logical Switches** dialog box displays.

2. Select the physical chassis from which you want to create a logical fabric in the **Chassis** list.

3. Click **New Fabric**.

   The **New Logical Fabric Template** dialog box displays.

4. Enter a new identifier in the **Logical Fabric ID** field to create a new logical fabric template.

   This identifier is how you distinguish among multiple logical fabric templates in the **Logical Switches** dialog box. If you create more than one logical fabric template, give them different fabric IDs.

5. Enter new values for the fabric parameters or leave unchanged to accept the default values.

   Click the **Help** button for detailed information on each parameter.

   NOTE
   If you set the long distance fabric, it must be set on all devices in the fabric.

6. Click the **Switch** tab.

7. Select the **Insistent Domain ID** check box to guarantee that a switch operates only with its preassigned domain ID. If a duplicate domain ID exists, the switch will segment from the fabric instead of changing the domain ID.

   Leave this check box blank to allow the domain ID to be changed if a duplicate address exists.

8. Click **OK** on the **New Logical Fabric Template** dialog box.

   The new logical fabric template displays under the **Discovered Logical Switches** node in the **Existing Logical Switches** list (already highlighted).

   All of the logical fabric templates have the same name, "NewFabric". You can differentiate among the templates by the FID number.

   You can now create logical switches using the fabric-wide settings in the logical fabric template. To assign logical switches, refer to "Creating a logical switch or base switch" on page 608.

   NOTE
   When you close the **Logical Switches** dialog box, the logical fabric templates are automatically deleted. Create the logical switches first, before closing the dialog box, to use the template.

Applying logical fabric settings to all associated logical switches

You can apply a selected logical switch configuration to all logical switches in the same fabric. This configures the fabric parameters for the selected logical switch to all logical switches in the fabric.

1. Select **Configure > Virtual Fabric > Logical Switches**.

   The **Logical Switches** dialog box displays.

2. Right-click anywhere in the **Existing Logical Switches** list and select **Table > Expand All**.

3. Right-click the logical switch for which you have configured logical fabric settings from the **Existing Logical Switches** list and select **Configure All**.

   The logical fabric configuration settings (**Fabric** tab) are applied to all logical switches in the same fabric (determined by FID).

4. Click **OK** on the **Logical Switches** dialog box.

   The **Logical Switch Change Confirmation and Status** dialog box displays with a list of all changes you made in the **Logical Switches** dialog box.

   The **Re-Enable ports after moving them** and **QoS disable the ports while moving them** check boxes are selected by default. If the selected ports have more than 125 devices, the **Re-enable Ports after moving them option** is disabled.

   NOTE
   Ports are disabled before moving from one logical switch to another.

5. (*Optional*) Select the **Unbind Port Addresses while moving them** check box.
6. Click **Start** to send these changes to the affected chassis.

   NOTE
   Most changes to logical switches will disrupt data traffic in the fabric.

   The status of each change is displayed in the **Status** column and **Status** area in the dialog box.

7. When the changes are complete, click **Close**.

Moving a logical switch to a different fabric

You can move a logical switch from one fabric to another by assigning a different fabric ID.

1. Select **Configure > Virtual Fabric > Logical Switches**.

   The **Logical Switches** dialog box displays.

2. Right-click anywhere in the **Existing Logical Switches** list and select **Table > Expand All**.
3. Select the logical switch you want to move to another logical fabric.
4. Click **Edit**.

   The **Edit Properties** dialog box displays.

5. Change the FID in the **Logical Fabric ID** field.
6. Click **OK** on the **Edit Properties** dialog box.

   The logical switch displays under the new logical fabric node in the **Existing Logical Switches** list.

7. Click **OK** on the **Logical Switches** dialog box.

   The **Logical Switch Change Confirmation and Status** dialog box displays with a list of all changes you made in the **Logical Switches** dialog box.

   The **Re-Enable ports after moving them** and **QoS disable the ports while moving them** check boxes are selected by default. If the selected ports have more than 125 devices, the **Re-enable Ports after moving them option** is disabled.

   NOTE
   Ports are disabled before moving from one logical switch to another.

8. (*Optional*) Select the **Unbind Port Addresses while moving them** check box.

9. Click **Start** to send these changes to the affected chassis.

   NOTE
   Most changes to logical switches will disrupt data traffic in the fabric.

   The status of each change is displayed in the **Status** column and **Status** area in the dialog box.

10. When the changes are complete, click **Close**.
11. If the newly created switch is not part of a discovered fabric, then you must discover the switch.
    a. Undiscover the physical chassis. Refer to "Deleting a fabric" on page 42 for instructions.
    b. Rediscover the physical chassis. Refer to "Discovering fabrics" on page 35 for instructions.

       When entering the IP address, use the IP address of the physical fabric.

Changing a logical switch to a base switch

The **Base Switch** column in the **Existing Logical Switches** list indicates whether a logical switch is a base switch.

1. Select **Configure > Virtual Fabric > Logical Switches**.

   The **Logical Switches** dialog box displays.

2. Right-click anywhere in the **Existing Logical Switches** list and select **Table > Expand All**.
3. Select the logical switch you want to change to a base switch.
4. Click **Edit**.

   The **Edit Properties** dialog box displays.

5. Clear the **Base Fabric for Transport** check box.

   This check box is applicable only to logical switches that are *not* base switches.

6. Select the **Base Switch** check box.
7. Click **OK** on the **Edit Properties** dialog box.

   The **Base Switch** column in the **Existing Logical Switches** list now displays **Yes** for the logical switch.

8. Click **OK** on the **Logical Switches** dialog box.

   The **Logical Switch Change Confirmation and Status** dialog box displays with a list of all changes you made in the **Logical Switches** dialog box.

   The **Re-Enable ports after moving them** and **QoS disable the ports while moving them** check boxes are selected by default. If the selected ports have more than 125 devices, the **Re-enable Ports after moving them option** is disabled.

   NOTE
   Ports are disabled before moving from one logical switch to another.

9. (*Optional*) Select the **Unbind Port Addresses while moving them** check box.

10. Click **Start** to send these changes to the affected chassis.

    NOTE
    Most changes to logical switches will disrupt data traffic in the fabric.

    The status of each change is displayed in the **Status** column and **Status** area in the dialog box.

11. When the changes are complete, click **Close**.

SAN Encryption Configuration

Encryption Center features

The **Encryption Center** dialog box is the single launching point for all encryption-related configuration in the Management application. (Refer to Figure 269.) It also provides a table that shows the general status of all encryption-related hardware and functions at a glance. To open the dialog box, select **Configure > Encryption**.

FIGURE 269 Encryption Center dialog box

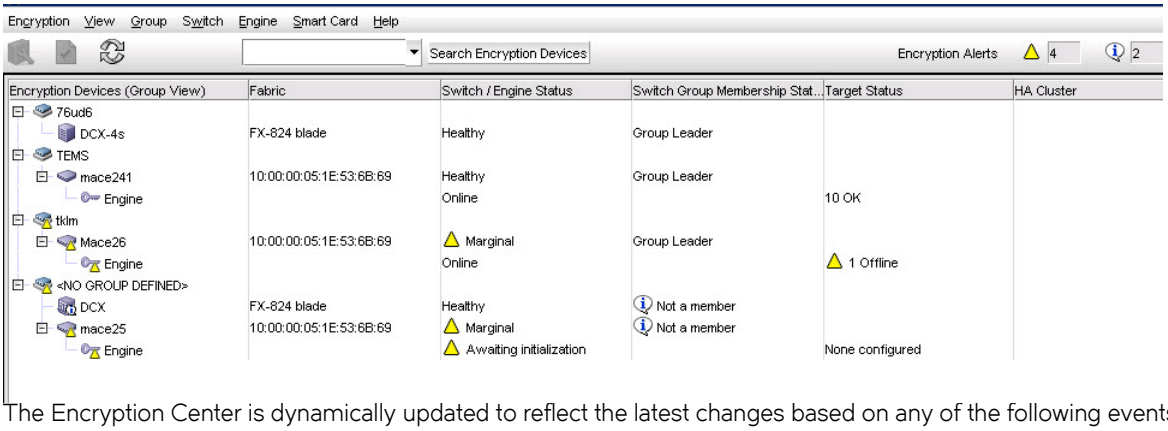

The Encryption Center is dynamically updated to reflect the latest changes based on any of the following events:

- Encryption group creation or deletion.
- A change in encryption group status or encryption engine status
- Addition or removal of an encryption group member or encryption engine

If you are using the Encryption Center for the first time, please read the following topics before you begin to perform encryption operations:

- "Encryption user privileges" on page 620 describes the Role-based Access Control privileges that are specific to encryption.
- "Smart card usage" on page 621 and the topics that follow describe the options available for the use of Smart Cards for user authentication, system access control, and storing backup copies of data encryption master keys.
- "Network connections" on page 631 describes the network connections that must be in place to enable encryption.
- "Blade processor links" on page 631 describes the steps for interconnecting encryption switches or blades in an encryption group through a dedicated LAN. This must be done before the encryption engines are enabled. Security parameters and certificates cannot be exchanged if these links are not configured and active.
- "Encryption node initialization and certificate generation" on page 632 lists the security parameters and certificates that are generated when an encryption node is initialized.
- "Supported encryption key manager appliances" on page 636 lists the supported key manager appliances, and lists topics that provide additional detail.

Encryption user privileges

In the Management application, resource groups are assigned privileges, roles, and fabrics. Privileges are not directly assigned to users; users get privileges because they belong to a role in a resource group. A user can only belong to one resource group at a time.

The Management application provides three pre-configured roles:

- Storage encryption configuration
- Storage encryption key operations

- Storage encryption security

Table 59 lists the associated roles and their read/write access to specific operations. The functions are enabled from the **Encryption Center** dialog box:

TABLE 59 Encryption privileges

Privilege	Read/Write
Storage Encryption Configuration	• Launch the Encryption center dialog box. • View switch, group, or engine properties. • View the Encryption Group Properties Security tab. • View encryption targets, hosts, and LUNs. • View LUN centric view • View all rekey sessions • Add/remove paths and edit LUN configuration on LUN centric view • Rebalance encryption engines. • Clear tape LUN statistics • Create a new encryption group or add a switch to an existing encryption group. • Edit group engine properties (except for the Security tab) • Add targets. • Select encryption targets and LUNs to be encrypted or edit LUN encryption settings. • Edit encryption target hosts configuration. • Show tape LUN statistics.
Storage Encryption Key Operations	• Launch the Encryption center dialog box. • View switch, group, or engine properties, • View the Encryption Group Properties Security tab. • View encryption targets, hosts, and LUNs. • View LUN centric view. • View all rekey sessions. • Initiate manual rekeying of all disk LUNs. • Initiate refresh DEK. • Enable and disable an encryption engine. • Decommission LUNs. • Zeroize an encryption engine. • Restore a master key. • Edit key vault credentials. • Show tape LUN statistics.
Storage Encryption Security	• Launch the Encryption center dialog box. • View switch, group, or engine properties. • View Encryption Group Properties Security tab. • View LUN centric view. • View all rekey sessions. • View encryption targets, hosts, and LUNs. • Create a master key. • Backup a master key. • Edit smart card. • View and modify settings on the Encryption Group Properties Security tab (quorum size, authentication cards list and system card requirement). • Establish link keys for LKM/SSKM key managers. • Show tape LUN statistics.

Smart card usage

Smart cards are credit card-sized cards that contain a CPU and persistent memory. Smart cards can be used as security devices. You must have ***Storage Encryption Security*** user privileges to activate, register, and configure smart cards.

Smart cards can be used to do the following:

- Control user access to the Management application security administrator roles

- Control activation of encryption engines
- Securely store backup copies of master keys

Smart card readers provide a plug-and-play interface that allows you to read and write to a smart card. The following smart card readers are supported:

- GemPlus GemPC USB

  http://www.gemalto.com/readers/index.html
- Indentive

  http://www.indentive-infrastructure.com

  NOTE
  Only the Brocade smart cards that are included with the encryption switches are supported.

Using authentication cards with a card reader

When authentication cards are used, one or more authentication cards must be read by a card reader attached to a Management application workstation to enable certain security-sensitive operations. These include the following:

- Performing master key generation, backup, and restore operations.
- Registering or deregistering and replacement of authentication cards.
- Enabling and disabling the use of system cards.
- Changing the quorum size for authentication cards.
- Establishing a trusted link with the NetApp LKM/SSKM key vault.
- Decommissioning a LUN.

When a quorum of authentication cards is registered for use, authentication must be provided before you are granted access.

Registering authentication cards from a card reader

To register an authentication card or a set of authentication cards from a card reader, have the cards physically available. Authentication cards can be registered during encryption group or member configuration when running the configuration wizard, or they can be registered using the following procedure.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption group from the **Encryption Center Devices** table, then select **Group > Security** from the menu task bar to display the **Encryption Group Properties** dialog box. The **Security** tab is selected. (Refer to Figure 270.)

FIGURE 270 Security tab - registering authentication cards

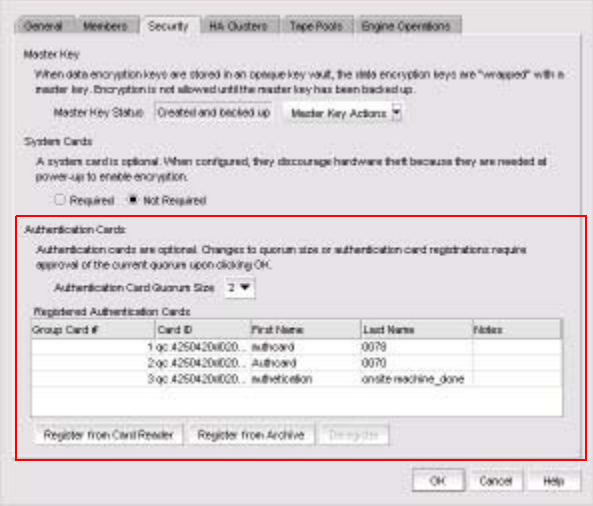

The **Authentication Cards** section contains the following information:

- **Group Card#:** A number assigned to the card as it is registered.
- **Card ID:** The serial number read from the smart card.
- **First Name:** The first name of the person assigned to the card.
- **Last Name:** The last name of the person assigned to the card.
- **Notes:** An optional entry of information.
- **Register from Card Reader** button: Launches the **Add Authentication Card** dialog box.
- **Register from Archive** button: Launches the **Add Authentication Card** dialog box.
- **Deregister** button: Deregisters a card selected from the **Registered Authentication Cards** table, which enables the cards to be removed from the switch and the database.

3. Locate the **Authentication Card Quorum Size** and select the **quorum size from the list**.

   The quorum size is the minimum number of cards necessary to enable the card holders to perform the security sensitive operations listed above. The maximum quorum size is five cards. The actual number of authentication cards registered is always more than the quorum size, so if you set the quorum size to five, for example, you will need to register at least six cards in the subsequent steps.

   NOTE
   Ignore the **System Cards** setting for now.

4. Click **Register from Card Reader** to register a new card.

   The **Add Authentication Card** dialog box displays. (Refer to Figure 271.)

FIGURE 271 Add Authentication Card dialog box

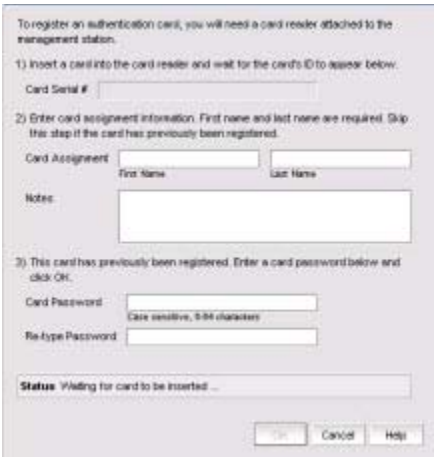

The **Add Authentication Card** dialog box contains the following information:

- **Card Serial#:** A serial number read from the smart card.
- **Card Assignment:** The first and last name of the person assigned to the card.
- **Notes:** An optional entry of information.
- **Card Password:** Create a password for the card holder to enter for user verification.
- **Re-type Password:** Re-enter the password in this field.
- **Status:** Indicates the status when a card is being registered.

5. Insert a smart card into the card reader. Wait for the card serial number to appear, enter card assignment information as directed, then click **OK**.
6. Wait for the confirmation dialog box indicating initialization is done, then click **OK**.

   The card is added to the **Registered Authentication Cards** table.
7. Repeat step 5 and step 6 until you have successfully registered all cards. Ensure that the number of cards registered equals at least the quorum size plus one.

Registering authentication cards from the database

Smart cards that are already in the Management program's database can be registered as authentication cards.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption group from the **Encryption Center Devices** table, then select **Group > Security** from the menu task bar to display the **Encryption Group Properties** dialog box. The **Security** tab is selected. (Refer to Figure 272.)

FIGURE 272 Encryption Group Properties dialog box - Security tab

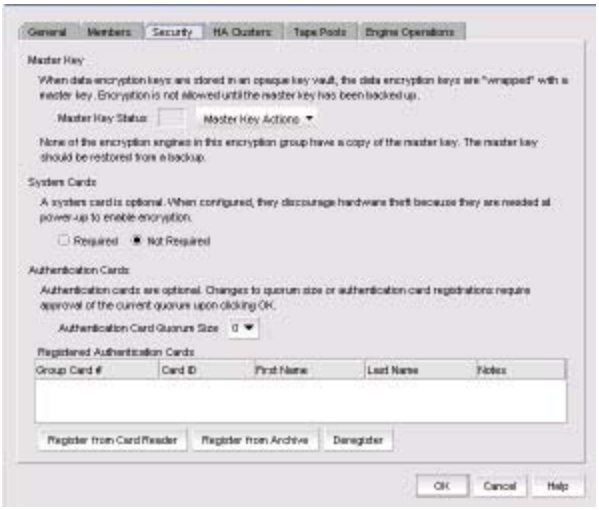

3. Click **Register from Archive**.

   The **Authentication Cards** dialog box displays. (Refer to Figure 273.) The table lists the smart cards that are in the database.

FIGURE 273 Authentication Cards dialog box - Registering smart cards from archive

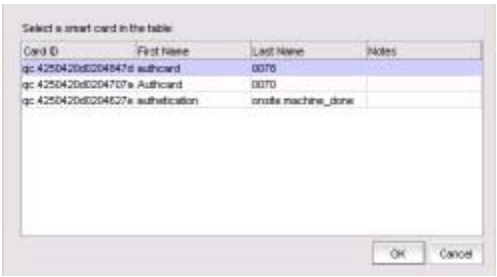

4. Select a card from the table, then click **OK**.
5. Wait for the confirmation dialog box indicating initialization is done, then click **OK**.

   The card is added to the **Registered Authentication Cards** table.

Deregistering an authentication card

Authentication cards can be removed from the database and the switch by deregistering them. Complete the following procedure to deregister an authentication card.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption group from the **Encryption Center Devices** table, then select **Group > Security** from the menu task bar to display the **Encryption Group Properties** dialog box. The **Security** tab is selected.
3. Select the desired authentication card in the **Registered Authentication Cards** table, then click **Deregister**.
4. Click **Yes** to confirm deregistration.

   The **registered authentication card is removed from the table.**

5. Click **OK**.

   The card is deregistered from the group.

Setting a quorum for authentication cards

To authenticate using a quorum of authentication cards, complete the following steps:

1. When using the **Authenticate** dialog box, gather the number of cards needed according to the instructions in the dialog box. The registered cards and the assigned owners are listed in the table near the bottom of the dialog box.

   The **Authenticate** dialog box contains the following information:

   - **Card ID:** Insert a smart card into an attached card reader, and wait for the card ID to appear in this field.
   - **Password:** The card holder must enter a password for the card.
   - **Authenticate button:** Authenticates the card after entering the password.
   - **Currently registered authentication cards table:** Lists the currently registered cards, showing the card ID and the name of the person assigned to the card.
   - **Status:** Displays the status of the card authentication operation.

2. Insert a card, then wait for the ID to appear in the **Card ID** field.
3. Enter the assigned password, then click **Authenticate**.
4. Wait for the confirmation dialog box, then click **OK**.
5. Repeat step 2 through step 4 for each card until at least the quorum plus one is reached, then click **OK**.

Using system cards

System cards are smart cards that can be used to control activation of encryption engines. You can choose whether the use of a system card is required or not. Encryption switches and blades have a card reader that enables the use of a system card. System cards discourage theft of encryption switches or blades by requiring the use of a system card at the switch or blade to enable the encryption engine after a power off.

When the switch or blade is powered off, the encryption engine will not work without first inserting a system card into its card reader. If someone removes a switch or blade with the intent of accessing the encryption engine, it will function as an ordinary FC switch or blade when it is powered up, but use of the encryption engine is denied.

To register a system card from a card reader, the smart card must be physically available. (Refer to Figure 274.)

FIGURE 274 System Cards dialog box

The **System Cards** dialog box can be accessed by selecting a switch from the **Encryption Center Devices** table, then selecting **Switch > System Cards** from the menu task bar. The **Register System Card** dialog box displays.

The dialog box contains the following information:

- **Group System Card:** Identifies if smart cards are used to control activation of encryption engines.
- **Registered System Cards table:** Lists all currently registered system card serial numbers and to whom the cards are assigned by first and last name. Also included are any free-form notes related to the cards.
- **Register from Card Reader button:** Launches the **Register from Card Reader** dialog box.
- **Deregister** button: Launches the **Deregister** dialog box.

Enabling or disabling the system card requirement

To use a system card to control activation of an encryption engine on a switch, you must enable the system card requirement. If a system card is required, it must be read by the card reader on the switch. You access the system card GUI from the **Security** tab.

Complete the following procedure to enable or disable the system card requirement.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices** table, then select **Group > Security** from the menu task bar.

   The **Properties** dialog box displays with the **Security** tab selected.
3. Under **System Cards**, select **Required** or **Not Required** as needed.
4. Click **OK**.

Registering system cards from a card reader

To register a system card from a card reader, a smart card must be physically available. System cards can be registered during encryption group creation or member configuration when running the configuration wizard, or they can be registered using the following procedure:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a switch from the **Encryption Center Devices** table that is not already in an encryption group, then select **Switch > System Cards** from the menu task bar.

   The **System Cards** dialog box displays. (Refer to Figure 274 on page 626.) The **Registered System Cards** table lists all currently registered system card serial numbers and to whom they are assigned. Also included are any notes related to the cards.
3. Click **Register from Card Reader**.
4. Insert a smart card into the card reader.
5. Wait for the card serial number to appear, then enter card assignment information as directed and click **OK**.
6. Wait for the confirmation dialog box indicating initialization is done, then click **OK**.

   The card is added to the **Registered System Cards** table.

   NOTE
   Store the card in a secure location, not in proximity to the switch or blade.

Deregistering system cards

System cards can be removed from the database by deregistering them. Use the following procedure to deregister a system card:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select the switch from the **Encryption Center Devices** table, then select **Switch > System Cards** from the menu task bar.

   The **System Cards** dialog box displays. (Refer to Figure 274 on page 626.)
3. Select the system card to deregister, then click **Deregister**.
4. A confirmation dialog box displays. Click **OK** to confirm deregistration.

   The card is removed from the **Registered System Cards** table.

Using smart cards

Smart cards can be used for user authentication, master key storage and backup, and as a system card for authorizing use of encryption operations. Card types identify if the smart card is a system card, authentication card, or recovery set.

The Smart Card Asset Tracking dialog box displays two tables: **Smart Cards** table and **Card Details** table.

- Selecting an authentication in the **Smart Cards** table, displays all group names for which the card is registered in the **Card Details** table.
- Selecting a system cards in the **Smart Cards** table displays all encryption engines for which the card is registered by switch name and, for encryption blades, slot number in the **Card Details** table.
- Selecting a recovery card in the **Smart Cards** table displays, the group name, the card creation date, and the position of the card in the set (for example, Card 1 of 3) in the **Card Details** table.

Tracking smart cards

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select **Smart Card > Smart Card Tracking** from the menu task bar to display the **Smart Card Asset Tracking** dialog box. (Refer to Figure 275.)

FIGURE 275 Smart Card Asset Tracking dialog box

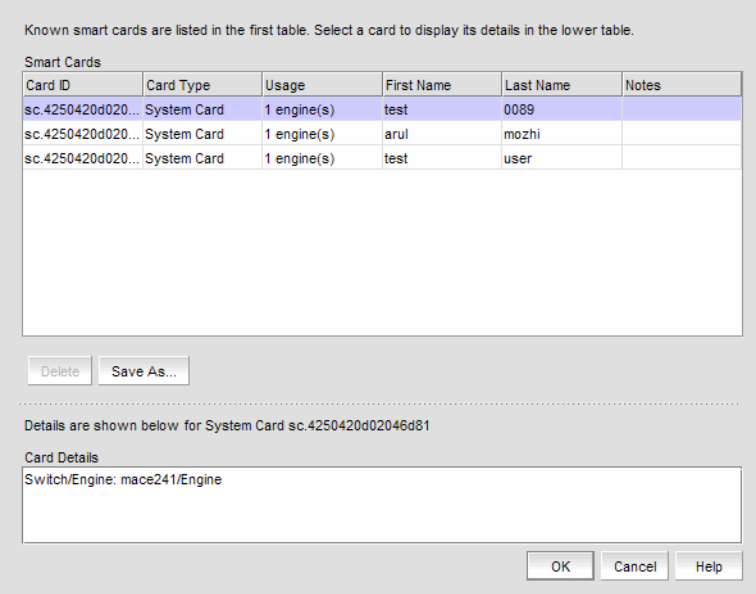

The **Smart Cards** table lists the known smart cards and the details for the smart cards. These details include the following:

- **Card ID**: Lists the smart card ID, prefixed with an ID that identifies how the card id used. For example, rc.123566b700017818, where rc stands for recovery card.
- **Card Type**: Options are: System card, Authentication card, and Recovery set.
- **Usage**: Usage content varies based on the card type.
  - For Authentication cards, the **Usage** column shows the number of groups for which the card is registered.
  - For System cards, the **Usage** column shows the number of encryption engines for which the card is registered.
  - For Recovery cards, the **Usage** column shows the group name and the creation date.
- **First Name**: The first name of the person (up to 64 characters) to whom the smart card is assigned. All characters are valid in the editable columns, including spaces. Editing these values in the Management application does not modify the information that is stored on the card.
- **Last Name**: The last name of the person (up to 64 characters) to whom the smart card is assigned. All characters are valid in the editable columns, including spaces. Editing these values in the Management application does not modify the information that is stored on the card.
- **Notes**: Miscellaneous notes (up to 256 characters) related to the smart card. Editing these values in the Management application does not modify the information that is stored on the card. Notes are optional.
- **Delete** button: Deletes a selected smart card from the Management application database.

  NOTE
  You can remove smart cards from the table to keep the **Smart Cards** table at a manageable size, but removing the card from the table does not invalidate it; the smart card can still be used.

- **Save As** button: Saves the entire list of smart cards to a file. The available formats are comma-separated values (.csv) and HTML (.html).
- **Card Details** table: Card details vary based on the card type.
  - For Authentication cards, the **Card Details** table shows all group names for which the card is registered.
  - For system cards, the **Card Details** table shows all encryption engines for which the card is registered by switch name and, for encryption blades, slot number.
  - For recovery cards, the **Card Details** table shows the group name, the card creation date, and the position of the card in the set (for example, Card 1 of 3).

3. Select a smart card from the table, then do one of the following:

- Click **Delete** to remove the smart card from the Management application database. Deleting smart cards from the Management application database keeps the **Smart Cards** table at a manageable size, but does not invalidate the smart card. The smart card can still be used. You must deregister a smart card to invalidate its use.

  NOTE
  The Delete operation applies only to recovery cards.

- Click **Save As** to save the entire list of smart cards to a file. The available formats are comma-separated values (.csv) and HTML files (.html).

Editing smart cards

Smart cards can be used for user authentication, master key storage and backup, and as a system card for authorizing use of encryption operations.

1. From the **Encryption Center dialog box**, select **Smart Card > Edit Smart Card** from the menu task bar to display the **Edit Smart Card** dialog box. (Refer to Figure 276.)

FIGURE 276 Edit Smart Card dialog box

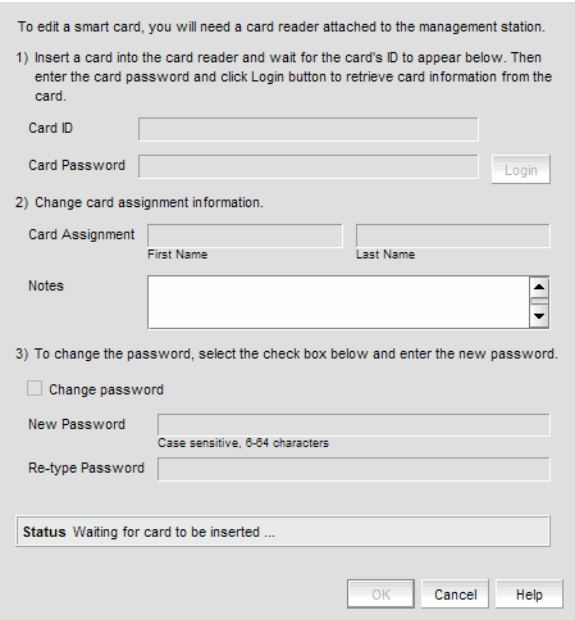

2. Insert the smart card into the card reader.

3. After the card's ID is displayed by the card reader in the **Card ID** field, enter the security administrator password used to allow editing of the smart card, then click **Login**.

   NOTE
   The **Card Password** field is activated after the card ID is read, and the **Login** button is activated after the password is entered in the **Card Password** field.

4. Edit the card as needed. Note the following:

   - **Card Assignment**: A maximum of 64 characters is permitted for the user first and last name to whom the card is assigned. All characters are valid in the editable columns, including spaces.
   - **Notes**: A maximum of 256 characters is permitted for any miscellaneous notes. Editing these values in the Management application does not modify the information that is stored on the card. Notes are optional.
   - The **Change Password** check box must be selected before you can enter the new password information. You must re-enter the new password for verification.

5. Click **OK**.

   NOTE
   You can view the status indicator at the bottom of the dialog box to determine card reader status.

Network connections

Before you use the encryption setup wizard for the first time, you must have the following required network connections:

- The management ports on all encryption switches and DCX Backbone Chassis CPs that have Encryption Blades installed must have a LAN connection to the SAN management program, and must be available for discovery.
- A supported key management appliance must be connected on the same LAN as the management port, which supports the encryption switches, DCX Backbone Chassis CPs, and the SAN Management program.
- In some cases, you might want to have an external host available on the LAN to facilitate certificate exchange between encryption nodes and the key management appliance. You may use the SAN management program host computer rather than an external host.
- All switches in the planned encryption group must be interconnected on a private LAN using the eth-0 and eth-1 ports located on the encryption switch or encryption blade. (We refer to these ports as RJ-45 gigabit Ethernet ports (labeled eth0 and eth1) for clustering and centralized management of multiple encryption switches through a group leader.)

Blade processor links

Each encryption switch or blade has two GbE ports labeled Ge0 and Ge1. The Ge0 and Ge1 ports are Ethernet ports that connect encryption switches and blades to other encryption switches and blades. Both ports of each encryption switch or blade must be connected to the same IP network and the same subnet. Static IP addresses should be assigned. Neither VLANs nor DHCP should be used. These two ports are bonded together as a single virtual network interface to provide link layer redundancy.

All encryption switches and blades in an encryption group must be interconnected by these links through a dedicated LAN before their encryption engines are enabled. Both ports of each encryption switch or blade must be connected to the same IP network and the same subnet. Static IP addresses should be assigned. VLANs should not be used, and DHCP should not be used. Security parameters and certificates cannot be exchanged if these links are not configured and active.

The **Blade Processor Link** dialog box can be launched from the following locations:

- Select an encryption group from the **Encryption Center Devices** table, then select **Group > HA Clusters from the menu task bar. The Properties dialog box displays with the HA Clusters tab selected.** Select a device from the **Non-HA Encryption Engines** table, then click **Configure Blade Processor Link**.
- Select a group, switch, or engine from the **Encryption Center Devices** table, then select **Group/Switch/Engine > Targets from the menu task bar. Select a container from the Encryption Targets table, click LUNs**, then click **Configure Blade Processor Link**.
- Select an engine from the **Encryption Center Devices** table, then select **Engine > Blade Processor Link**.

Configuring blade processor links

To configure blade processor links, complete the following steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select the encryption engine from the **Encryption Center Devices** table, then select **Engine > Blade Processor Link from the menu task bar** to display the **Blade Processor Link** dialog box. (Refer to Figure 277.)

   FIGURE 277 Blade Processor Link dialog box

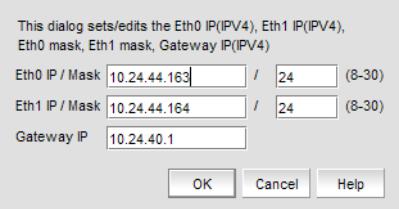

3. Enter the link IP address and mask, and the gateway IP address.
   - **Eth0 IP /Mask** identifies the Ge0 interface IP address and mask.
   - **Eth1 IP /Mask** identifies the Ge1 interface IP address and mask.
   - The **Gateway IP** address is optional.
4. Click **OK**.

Encryption node initialization and certificate generation

When an encryption node is initialized, the following security parameters and certificates are generated:

- FIPS crypto officer
- FIPS user
- Node CP certificate
- A signed Key Authentication Center (KAC) certificate
- A KAC Certificate Signing Request (CSR)

From the standpoint of external SAN management application operations, the FIPS crypto officer, FIPS user, and node CP certificates are transparent to users. The KAC certificates are required for operations with key managers. In most cases, KAC certificate signing requests must be sent to a Certificate Authority (CA) for signing to provide authentication before the certificate can be used. In all cases, signed KACs must be present on each switch.

Setting encryption node initialization

Encryption nodes are initialized by the **Configure Switch Encryption** wizard when you confirm a configuration. Encryption nodes may also be initialized from the **Encryption Center** dialog box.

1. Select a switch from the **Encryption Center Devices table**, then select **Switch > Init Node from the menu task bar**.
2. Select **Yes** after reading the warning message to initialize the node.

Key Management Interoperability Protocol

The Key Management Interoperability Protocol (KMIP) standardizes the communication between an Enterprise key management system and an encryption device. The same key vault servers can be used, only in a different mode. Currently, KMIP versions 1.0 and 1.1 are supported.

The initial deployment of the KMIP client is on the Encryption switch, where it will replace multiple third-party implementations/vendor APIs. The interfaces of the KMIP client are generic and are not tied to the key record formats used by the Encryption switch. Any encryption solution should be able to use the KMIP client to communicate to a key server by compiling it on Linux-based PPC or X 86 environments.

Currently, the Encryption switch supports the KMIP servers from SafeNet Key Secure 6.1 and TEKA 4.0. All nodes in an encryption group should be running Fabric OS 7.1.0 and later for the key vault type to be set to KMIP.

Although KMIP support is available from multiple key vaults, each key vault implementation is different in terms of High Availability (HA) clustering support, certificate exchange, and authentication. In the current Fabric OS implementation, each key vault uses a separate adapter at the Key Authentication Center (KAC), which is implemented to suit the key vault feature implementation.

> NOTE
> Currently, KMIP with SafeNet KeySecure 6.1 in native KMIP mode and Thales e-Security keyAuthority running version 4.0 with the Encryption switch in KMIP mode are supported.

A generic KMIP 1.0 or 1.1 server is supported. The following KMIP servers can be configured on the Encryption switch:

- SafeNet KeySecure. The KeySecure is a KMIP-compliant server. (SSKM is the trusted mode version of SafeNet which continues to use the LKM OpenKey Interfaces. These are mutually exclusive use scenarios and cannot be used interchangeably.) This configuration is allowed only for new installations. Refer to "Steps for connecting to a KMIP-compliant SafeNet KeySecure" on page 663.
- TEKA 4.0. The Thales keyAuthority is a KMIP-compliant server that can be configured with the Encryption switch; however, backward compatibility for keys created with Fabric OS versions earlier than v7.2.0 is not supported. This configuration is allowed only for new installations. For more information about configuring a KMIP-compliant keyAuthority, refer to Chapter 3 of the *Fabric OS Encryption Administrator's Guide Supporting Key Management Interoperability Protocol (KMIP) Key-Compliant Environments*.

Ensure that KMIP server is running on the key vault in order for the key vault to be configured as a KMIP type on the Encryption switch.

Configuration parameters

The encryption group object has three additional properties that can be configured when the key vault (KV) type is KMIP. These additional properties must be set by the user:

- High availability
- User credentials

- Certificate type

High availability

The KMIP Key Authentication Center (KAC) adapter provides configurable HA support. HA for the key vault should be set before you register the key vault. Three settings are supported; however, certain settings are determined by the compliant key vault type that is being used:

- **Transparent**: The client assumes the entire HA replication is implemented on the key vault. Key archival and retrieval is performed without any additional key hardening or integrity checks.
- **Opaque:** The primary and secondary key vaults are both registered on the Encryption switch. The client archives the key to a single (primary) key vault and lets the KV pair internally perform the replication. For disk operations, an additional key hardening and integrity check is done on the secondary key vault before the key is used for encryption.
- **None:** If no HA is selected, the primary and secondary key vaults are both registered on the Encryption switch. The client archives keys to both key vaults and ensures that the archival process succeeds before the key is used for encryption, including hardening and integrity checks.

By default, the HA mode is disabled and KAC login is not used. All parameters except log level are configurable on the group leader only. All parameters except for logging are distributed to all nodes in the encryption group. Log level, however, is configurable on a per-node basis.

User credentials

The Encryption switch has support for the optional credential structure used for username and password. Username authentication can be defined after TLS connectivity to a client device is requested. Three modes are available:

- **User Name:** Only a user name is required to identify the client device.
- **User Name and Password:** Both a user name and a password are required to identify the client device.
- **None:** No authentication is required.

Certificate type

The TLS certificates used between the Encryption switch and the key vault are either **Self Signed** or **CA Signed**.

Key vault type and vendor

The key vault type for any KMIP-compliant key vault is shown on the Encryption switch as “KMIP” in the **groupcfg** output. The key vault vendor or key manager name is displayed under “Server SDK Version”.

Sample **groupCfg** output for SafeNet KeySecure is provided:

SafeNet

```
switch:root> cryptocfg --show -groupcfg
Encryption Group Name:          CRYPTO_LSWAT
    Failback mode:              Auto
    Replication mode:           Disabled
    Heartbeat misses:           3
    Heartbeat timeout:          2
    Key Vault Type:             KMIP
    System Card:               Disabled

Primary Key Vault:
    IP address:                10.38.145.10
    Certificate ID:            LKM10_CA
```

```
    Certificate label:         SSKM_10
    State:                     Connected
    Type:                      KMIP

Secondary Key Vault:
    IP address:                10.38.145.17
    Certificate ID:            LKM10_CA
    Certificate label:         SSKM_17
    State:                     Connected
    Type:                      KMIP


Additional Primary Key Vault Information::
                Key Vault/CA Certificate Validity:      Yes
                Port for Key Vault Connection:          5696
                Time of Day on Key Server:              N/A
                Server SDK Version:                     SafeNet, Inc.


Additional Secondary Key Vault Information:
                Key Vault/CA Certificate Validity:      Yes
                Port for Key Vault Connection:          5696
                Time of Day on Key Server:              N/A
                Server SDK Version:                     SafeNet, Inc.

Encryption Node (Key Vault Client) Information:
                Node KAC Certificate Validity:          Yes
                Time of Day on the Switch:              2012-12-20 07:33:44
                Client SDK Version:                     N/A
                Client Username:                        brcduser
                Client Usergroup:                       brocade
                Connection Timeout:                     10 seconds
                Response Timeout:                       10 seconds
                Connection Idle Timeout:                N/A

Key Vault configuration and connectivity checks successful, ready for key operations.


Authentication Quorum Size:       0
Authentication Cards not configured


NODE LIST
Total Number of defined nodes:             2
Group Leader Node Name:                    10:00:00:05:1e:53:ae:4c
Encryption Group state:                    CLUSTER_STATE_CONVERGED
Crypto Device Config state:                In Sync
Encryption Group Config state:             In Sync

      Node Name                 IP address           Role
10:00:00:05:1e:b6:68:80        10.37.36.128        MemberNode
                    EE Slot:                                    1
                    SP state:                                   Online

10:00:00:05:1e:53:ae:4c        10.37.39.111        GroupLeader (current node)
                    EE Slot:                                    0
                    SP state:
```

Supported encryption key manager appliances

As stated under "Network connections" on page 631, a supported key management appliance must be connected on the same LAN as the management port of the encryption switches, or of the Backbone Chassis Control Processors (CPs) in the case of the encryption blade.

Secure communication between encryption nodes in an encryption group, and between encryption nodes and key manager appliances requires an exchange of certificates that are used for mutual authentication. Each supported key manager appliance has unique requirements for setting up a secure connection and exchanging certificates.

The following key manager appliances are supported:

- RSA Data Protection Manager (DPM). Refer to "Steps for connecting to a DPM appliance" on page 636.
- NetApp Lifetime Key Manager (LKM) and SafeNet KeySecure for key management (SSKM). Refer to "Steps for connecting to an LKM/SSKM appliance" on page 641.
- HP Secure Key Manager (SKM) and Enterprise Secure Key Manager (ESKM). Refer to "Steps for connecting to an ESKM/SKM appliance" on page 646.
- Thales e-Security keyAuthority (TEKA). Refer to "Steps for connecting to a TEKA appliance" on page 655.
- Tivoli Key Lifecycle Manager (TKLM). Refer to "Steps for connecting to a TKLM appliance" on page 660.
- Key Management Interoperability Protocol (KMIP). Refer to "Steps for connecting to a KMIP-compliant SafeNet KeySecure" on page 663.

Steps for connecting to a DPM appliance

All switches that you plan to include in an encryption group must have a secure connection to the RSA Data Protection Manager (DPM). The following is a suggested order of steps needed to create a secure connection to the DPM.

> NOTE
> The Encryption switch uses the manual enrollment of identities with client registration to connect with DPM 3.x servers. Client registration is done automatically when you upgrade to Fabric OS 7.1.0 from an earlier version and no additional user interaction is needed during the upgrade scenario.

Once completed, client registration occurs after key vault registration, when the Encryption switch attempts to connect to the DPM server for the first time.

1. Export the Key Authentication Center (KAC) CSR to a location accessible to a CA for signing. Refer to "Exporting the KAC certificate signing request (CSR)" on page 637.
2. Submit the KAC CSR for signing by a CA. Refer to "Submitting the CSR to a certificate authority" on page 637.
3. Set the KAC certificate registration expiry. Refer to "KAC certificate registration expiry" on page 638.
4. Import the signed certificate into the Fabric OS encryption node. Refer to "Importing the signed KAC certificate" on page 638.
5. Upload the signed KAC and CA certificates onto the DPM appliance and select the appropriate key classes. Refer to the following:
   - "Uploading the CA certificate onto the DPM appliance (and first-time configurations)" on page 639.
   - "Uploading the KAC certificate onto the DPM appliance (manual identity enrollment)" on page 640.
6. If dual DPM appliances are used for high availability, the DPM appliances must be clustered, and must operate in maximum availability mode, as described in the DPM appliance user documentation. Refer to "DPM key vault high availability deployment" on page 640.

Exporting the KAC certificate signing request (CSR)

1. Export the Key Authentication Center (KAC) CSR to a temporary location prior to submitting the KAC CSR to a CA for signing.
2. Synchronize the time on the switch and the key manager appliance. Time settings should be within one minute of each other. Differences in time can invalidate certificates and cause key vault operations to fail.
3. Select a switch from the **Encryption Center Devices** table, then select **Switch > Properties** from the menu task bar to display the **Properties** dialog box.

   NOTE
   You can also select a switch from the **Encryption Center Devices** table, then click the **Properties** icon.

4. Do one of the following:
   - If a CSR is present, click **Export**.
   - If a CSR is not present, select a switch from the **Encryption Center Devices** table, then select **Switch > Init Node** from the menu task bar**.** This generates switch security parameters and certificates, including the KAC CSR.
5. Save the file. The default location for the exported file is **in the Documents** folder.

   NOTE
   The CSR is exported in Privacy Enhanced Mail (.pem) format. This is the format required in exchanges with Certificate Authorities (CAs).

Submitting the CSR to a certificate authority

The CSR must be submitted to a Certificate Authority (CA) to be signed. The CA is a trusted third-party entity that signs the CSR. Several CAs are available and procedures vary, but the general steps are as follows:

1. Open an SSL/TLS connection to an X.509 server.
2. Submit the CSR for signing.
3. Request the signed certificate.

   Generally, a public key, the signed Key Authentication Center (KAC) certificate, and a signed CA certificate are returned.

4. Download and store the signed certificates.

The following example submits a CSR to the demoCA from RSA:

```
cd /opt/CA/demoCA
openssl x509 -req -sha1 -CAcreateserial -in certs/<Switch CSR Name> -days 365 -CA cacert.pem -CAkey
private/cakey.pem -out newcerts/<Switch Cert Name>
```

NOTE
You can change the number of days that a certificate will expire based on your site's security policies. For more information on changing the certificate expiry date, refer to "KAC certificate registration expiry" on page 638.

KAC certificate registration expiry

It is important to keep track as to when your signed Key Authentication Center (KAC) certificates will expire. Failure to work with valid certificates causes certain commands to not work as expected. If you are using the certificate expiry feature and the certificate expires, the key vault server will not respond as expected. For example, the Group Leader in an encryption group might show that the key vault is connected; however, a member node reports that the key vault is not responding.

To verify the certificate expiration date, use the following command:

openssl x509 –in newcerts/<Switch Cert Name> -dates –noout

```
Output:
            Not Before: Dec  4 18:03:14 2009 GMT
            Not After : Dec  4 18:03:14 2010 GMT
```

In the example above, the certificate validity is active until “Dec 4 18:03:14 2010 GMT.” After the KAC certificate has expired, the registration process must be redone.

NOTE
In the event that the signed KAC certificate must be re-registered, you will need to log in to the key vault web interface and upload the new signed KAC certificate for the corresponding Encryption switch Identity.

You can change the value of the certificate expiration date using the following command:

openssl x509 -req -sha1 -CAcreateserial -in certs/<Switch CSR Name> -days 365 -CA cacert.pem -CAkey private/cakey.pem -out newcerts/<Switch Cert Name>

In the example above, the certificate is valid for a period of one year (365 days). You can increase or decrease this value according to your own specific needs. The default is 3649 days, or 10 years.

Importing the signed KAC certificate

After a Key Authentication Center (KAC) CSR has been submitted and signed by a CA, the signed certificate must be imported into the switch.

1. Select a switch from the **Encryption Center Devices table**, then select **Switch > Import Certificate** from the menu task bar to display the **Import Signed Certificate** dialog box. (Refer to Figure 278.)

   FIGURE 278 Import Signed Certificate dialog box

   

2. Browse to the location where the signed certificate is stored, then click **OK**.

   The signed certificate is stored on the switch.

Uploading the CA certificate onto the DPM appliance (and first-time configurations)

After an encryption group is created, you need to install the signing authority certificate (CA certificate) onto the DPM appliance.

1. Open a web browser and connect to the DPM appliance setup page. You will need the URL and have the proper authority level, user name, and password.
2. Select the **Operations** tab.
3. Select **Certificate Upload**.
4. In the **SSLCAcertificateFile** field, enter the full local path of the CA certificate. Do not use the UNC naming convention format.
5. Select **Upload**, **Configure SSL**, and **Restart Webserver**.
6. After the web server restarts, enter the root password.
7. Open another web browser window, and start the RSA management user interface.

   You will need the URL, and have the proper authority level, user name, and password.

   NOTE
   The Identity Group name used in the next step might not exist in a freshly installed DPM. To establish an Identity Group name, click the **Identity Group** tab, and create a name. The name **Hardware Retail Group** is used as an example in the following steps.

8. Select the **Key Classes** tab. The key classes must be created only once, regardless of the number of nodes in your encryption group or the number of encryption groups that will be sharing this DPM.

   kcn.1998-01.com.brocade:DEK_AES_256_XTS

   kcn.1998-01.com.brocade:DEK_AES_256_CCM

   kcn.1998-01.com.brocade:DEK_AES_256_GCM

   kcn.1998-01.com.brocade:DEK_AES_256_ECB

   a. Click **Create**.
   b. Type the key name string into the **Name** field.
   c. Select **Hardware Retail Group** for **Identity Group**.
   d. Deselect **Activated Keys Have Duration**.
   e. Select **AES** for **Algorithm**.
   f. Select **256** for **Key Size**.
   g. Select the **Mode** for the respective key classes as follows:

      XTS for Key Class "kcn.1998-01.com.brocade:DEK_AES_256_XTS"

      CBC for Key Class "kcn.1998-01.com.brocade:DEK_AES_256_CCM"

      CBC for Key Class "kcn.1998-01.com.brocade:DEK_AES_256_GCM"

      ECB for Key Class "kcn.1998-01.com.brocade:DEK_AES_256_ECB"

   h. Click **Next**.
   i. Repeat step a through step h for each key class.
   j. Click **Finish**.

Uploading the KAC certificate onto the DPM appliance (manual identity enrollment)

NOTE
The Encryption switch will not use the Identity Auto Enrollment feature supported with DPM 3.x servers. You must complete the identity enrollment manually to configure the DPM 3.x server with the Encryption switch as described in this section.

You need to install the switch public key certificate (KAC certificate). For each encryption node, manually create an identity as follows:

1. Select the **Identities** tab.
2. Click **Create**.
3. Enter a label for the node in the **Name** field. This is a user-defined identifier.
4. Select the **Hardware Retail Group** in the **Identity Groups** field.
5. Select the **Operational User** role in the **Authorization** field.
6. Click **Browse** and select the imported certificate as the **Identity certificate**.
7. Click **Save**.

The CA certificate file referenced in the **SSLCAcertificateFile** field (step 4) must be imported and registered on the switch designated as an encryption Group Leader. You may want to note this location before proceeding to "Loading the CA certificate onto the encryption group leader" on page 640.

DPM key vault high availability deployment

When dual DPM appliances are used for high availability, the DPM appliances must be clustered and must operate in maximum availability mode, as described in the DPM appliance user documentation.

When dual DPM appliances are clustered, they are accessed using an IP load balancer. For a complete high availability deployment, the multiple IP load balancers are clustered, and the IP load balancer cluster exposes a virtual IP address called a floating IP address. The floating IP address must be registered on the encryption Group Leader.

Neither the secondary DPM appliance nor individual DPM appliance IP addresses should be registered.

Loading the CA certificate onto the encryption group leader

The certificate for the CA that signed the switch KAC CSRs must be loaded onto the encryption Group Leader. The Group Leader can then distribute the CA certificate to the encryption group members.

1. From the **Encryption Center**, select a group from the **Encryption Center Devices** table, then select **Group > Properties** from the menu task bar to display the **Encryption Group Properties** dialog box. The **General** tab is selected. (Refer to Figure 279.)

   If groups are not visible in the **Encryption Devices** table, select **View > Groups** from the menu bar.

FIGURE 279 Encryption Group Properties with Key Vault Certificate

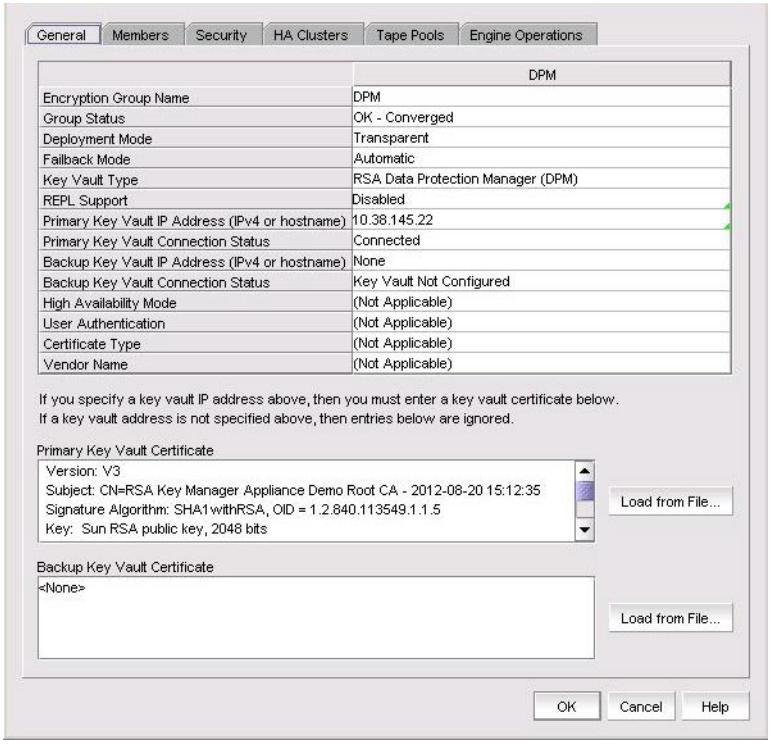

2. Select **Load from File** and browse to the location on your client PC that contains the downloaded CA certificate in .pem format.

Steps for connecting to an LKM/SSKM appliance

The NetApp Lifetime Key Manager (LKM) resides on an FIPS 140-2 Level 3-compliant network appliance. The encryption engine and LKM appliance communicate over a trusted link. A trusted link is a secure connection established between the Encryption switch or blade and the NetApp LKM/SSKM appliance, using a shared secret called a link key.

The following configuration steps are performed from the NetApp DataFort Management Console **(DMC)** and from the Management application:

- Install and launch the NetApp DataFort Management Console. Refer to "Launching the NetApp DataFort Management Console" on page 642.
- Establish the trusted link. Refer to "Establishing the trusted link" on page 642.
- Obtain and import the LKM/SSKM certificate. Refer to "Obtaining and importing the LKM/SSKM certificate" on page 643.
- Export and register encryption node certificates on LKM/SSKM. Refer to "Exporting and registering the switch KAC certificates on LKM/SSKM" on page 643.
- If required, create an LKM/SSKM cluster for high availability. Refer to "LKM/SSKM key vault high availability deployment" on page 644.
- Understanding Data Encryption Keys (DEKs). Refer to "Data Encryption Keys" on page 645.

Launching the NetApp DataFort Management Console

The **NetApp DataFort Management Console (DMC) must be installed on your PC or workstation to complete certain procedures described in this chapter. Refer to the appropriate DMC product documentation for DMC installation instructions. After you install the DMC, complete the following steps:**

1. Launch the DMC.
2. Click the **Appliance** tab on the top panel.
3. Add the NetApp LKM/SSKM appliance IP address or hostname.
4. Right-click the added IP address and log in to the NetApp LKM/SSKM key vault.

Establishing the trusted link

You must generate the trusted link establishment package (TEP) on all nodes to obtain a **trusted acceptance package (TAP) before you can** establish a trusted link between each node and the NetApp LKM/SSKM appliance.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an LKM/SSKM group from the **Encryption Center Devices table**, then select **Group > Link Keys** from the menu task bar.

   The switch name displays in the link status table under **Switch**, with a **Link Key Status** of **Link Key requested, waiting for LKM approval**.
3. Select the switch, then click **Establish**.

   This sends a Trust Establishment Package (TEP) message to the LKM/SSKM, which is needed to establish the trusted link between the switch and the LKM/SSKM appliance.
4. Launch the **NetApp DataFort Management Console (**DMC) and click the **View Unapproved Trustees** tab.

   The switch is listed as openkey_trustee_<ip address>, where the IP address is the switch IP address.
5. Select the switch, then click **Approve and Create TAP**.

   The **Approve TEP** dialog box displays. The TEP must be approved before a TAP can be created.
6. Provide a label in the dialog box, then click **Approve** to approve the TEP.

   A list of recovery cards and recovery officers is displayed. TEP approval is done by a quorum of recovery officers, using assigned recovery cards. Each recovery officer must individually insert one of the listed recovery cards into a card reader attached to the PC or workstation, then enter the password for that card and click **Start**. The procedure is repeated until a quorum of recovery officers has approved the TEP.
7. Save the TAP to a file (location does not matter).
8. Select the **Link Keys** tab from the **Encryption Group Properties** dialog box.
9. Select the switch in the link key status table, then click **Accept to retrieve the TAP from the LKM/SSKM appliance.**
10. Repeat the above steps for each of the remaining member nodes.

Obtaining and importing the LKM/SSKM certificate

Certificates must be exchanged between the LKM/SSKM appliance and the encryption switch to enable mutual authentication. You must obtain a certificate from the LKM/SSKM appliance and import it into the encryption Group Leader. The encryption Group Leader exports the certificate to other encryption group members.

To obtain and import an LKM/SSKM certificate, complete the following steps:

1. Open an SSH connection to the NetApp LKM/SSKM appliance and log in.

```
host$ssh admin@10.33.54.231
admin@10.33.54.231's password:

Copyright (c) 2001-2009 NetApp, Inc.
All rights reserved
+--------------------------------+
| NetApp Appliance Management CLI |
|      Authorized use only!      |
+--------------------------------+
Cannot read termcapdatabase;
using dumb terminal settings.
Checking system tamper status:
No physical intrusion detected.
```

2. Add the Group Leader to the LKM/SSKM key sharing group. Enter **lkmserver add --type third-party --key-sharing-group "/"** followed by the Group Leader IP address.

```
lkm-1>lkmserver add --type third-party --key-sharing-group \
    "/" 10.32.244.71
NOTICE: LKM Server third-party 10.32.244.71 added.
Cleartext connections not allowed.
```

3. On the NetApp LKM appliance terminal, enter **sys cert getcert-v2 to** display the LKM certificate content.

```
lkm-1> sys cert getcert-v2
-----BEGIN CERTIFICATE-----
[content removed]
-----END CERTIFICATE-----
```

4. Copy and paste the LKM/SSKM certificate content from the NetApp LKM/SSKM appliance terminal into an editor buffer. Save the file as **lkmcert.pem** on the SCP-capable host. Save the entire certificate, including the lines `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----`.
5. If you are using the Management application, the path to the file must be specified in the **Select Key Vault** dialog box when creating a Group Leader. If the proper path is entered, the file is imported.

Exporting and registering the switch KAC certificates on LKM/SSKM

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a switch from the **Encryption Center Devices table**, then select **Switch > Export Certificate** from the menu task bar.

   The **Export Switch Certificate** dialog box allows you to export a switch public key certificate signing request (CSR) to a location you specify. (Refer to Figure 280.) The procedures for submitting a CSR for signing are determined by the Certificate Authority (CA).

The CSR must be submitted to a Certificate Authority CA for signing, then imported into the switch and the key vault. The signed switch certificate may be imported directly by a key vault.

FIGURE 280 Export switch certificate dialog box

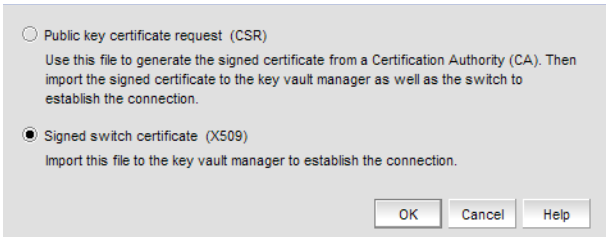

3. Select **Signed switch certificate (X.509)**, which allows you to export a signed switch certificate to a location of your choosing. The default location is **My Documents** on your client PC. In most cases, this certificate file should be in privacy email (.pem) format.
4. Click **OK**.

   You are prompted to save the CSR, which can be saved to your SAN Management Program client PC, or an external host of your choosing.
5. Register the signed KAC certificate that you exported from the member node with the NetApp LKM/SSKM appliance.

LKM/SSKM key vault high availability deployment

LKM/SSKM appliances can be clustered to provide high availability capabilities. You can deploy and register one LKM/SSKM with an encryption switch or blade and later deploy and register another LKM/SSKM at any time if LKM/SSKMs are clustered or linked together. Refer to LKM/SSKM documentation to link or cluster the LKM/SSKMs.

When LKM/SSKM appliances are clustered, both LKM/SSKMs in the cluster must be registered and configured with the link keys before starting any crypto operations. If two LKM/SSKM key vaults are configured, they must be clustered. If only a single LKM/SSKM key vault is configured, it may be clustered for backup purposes, but it is not directly used by the switch.

When dual LKM/SSKMs are used with the encryption switch or blade, the dual LKM/SSKMs must be clustered. There is no enforcement done at the encryption switch or blade to verify whether or not the dual LKM/SSKMs are clustered, but key creation operations will fail if you register non-clustered dual LKM/SSKMs with the encryption switch or blade.

Regardless of whether you deploy a single LKM/SSKM or clustered dual LKM/SSKMs, register only the primary key vault with the encryption switch or blade. You do not need to register a secondary key vault.

Data Encryption Keys

The following sections describe Data Encryption Key (DEK) behavior during DEK creation, retrieval, and updates as they relate to disk keys and tape pool keys, and tape LUN and DF-compatible tape pool support:

Disk keys and tape pool keys (Brocade native mode support)

Data Encryption Key (DEK) creation, retrieval, and update for disk and tape pool keys in Brocade native mode are as follows:

- **DEK creation**: The DEK is archived into the primary LKM/SSKM. Upon successful archival of the DEK onto the primary LKM/SSKM, the DEK is read from the secondary LKM/SSKM until it is either synchronized to the secondary LKM/SSKM, or a timeout of 10 seconds occurs (2 seconds with 5 retries).
  - If key archival of the DEK to the primary LKM/SSKM is successful, the DEK that is created can be used for encrypting disk LUNs or tape pools in Brocade native mode.
  - If key archival of the DEK to the primary LKM/SSKM fails, an error is logged and the operation is retried. If the failure occurs after archival of the DEK to the primary LKM/SSKM, but before synchronization to the secondary LKM/SSKM, a VAULT_OFFLINE error is logged and the operation is retried. Any DEK archived to the primary LKM/SSKM in this case is not used.
- **DEK retrieval**: The DEK is retrieved from the primary LKM/SSKM if the primary LKM/SSKM is online and reachable. If the registered primary LKM/SSKM is not online or not reachable, the DEK is retrieved from a clustered secondary LKM/SSKM.
- **DEK update**: DEK update behavior is the same as DEK creation.

Tape LUN and DF -compatible tape pool support

Data Encryption Key (DEK) creation, retrieval, and update for tape LUN and DF-compatible tape pool support are as follows:

- **DEK creation**: The DEK is created and archived to the primary LKM/SSKM only. Upon successful archival of the DEK to the primary LKM/SSKM, the DEK can be used for encryption of a Tape LUN or DF-Compatible tape pool. The DEK is synchronized to a secondary LKM/SSKM through LKM/SSKM clustering.

  If DEK archival onto the primary LKM/SSKM fails, DEK archival is retried to the clustered secondary LKM/SSKM. If DEK archival also fails to the secondary LKM/SSKM, an error is logged and the operation is retried.
- **DEK retrieval**: The DEK is retrieved from the primary LKM/SSKM if the primary LKM/SSKM is online and reachable. If the primary LKM/SSKM is not online or reachable, the DEK is retrieved from the clustered secondary LKM/SSKM.
- **DEK update**: DEK update behavior is the same as DEK creation.

LKM/SSKM key vault deregistration

Deregistration of either the primary or secondary LKM/SSKM key vault from an encryption switch or blade is allowed independently.

- **Deregistration of Primary LKM/SSKM:** You can deregister the Primary LKM/SSKM from an encryption switch or blade without deregistering the backup or secondary LKM/SSKM for maintenance or replacement purposes. However, when the primary LKM/SSKM is deregistered, key creation operations will fail until either the primary LKM/SSKM is reregistered, or the secondary LKM/SSKM is deregistered and reregistered as the primary LKM/SSKM.

  When the primary LKM/SSKM is replaced with a different LKM/SSKM, you must first synchronize the DEKs from the secondary LKM/SSKM before reregistering the primary LKM/SSKM.
- **Deregistration of Secondary LKM/SSKM:** You can deregister the secondary LKM/SSKM independently. Future key operations will use only the primary LKM/SSKM until the secondary LKM/SSKM is reregistered on the encryption switch or blade.

  When the secondary LKM/SSKM is replaced with a different LKM/SSKM, you must first synchronize the DEKs from the primary LKM/SSKM before reregistering the secondary LKM/SSKM.

Steps for connecting to an ESKM/SKM appliance

The ESKM/SKM management web console can be accessed from any web browser with Internet access to the ESKM/SKM appliance. The URL for the appliance is as follows:

```
https://<appliance hostname>:<appliance port number>
```

Where:

- `<appliance hostname>` is the hostname or IP address when installing the ESKM/SKM appliance.
- `<appliance port number>` is 9443 by default. If a different port number was specified when installing the ESKM/SKM appliance, use that port number.

The following configuration steps are performed from the ESKM/SKM management web console and from the Management application:

- Configure a Brocade group on the ESKM/SKM. Refer to "Configuring a Brocade group on ESKM/SKM" on page 646.
- Register the Brocade group user name and password on the encryption node. Refer to "Registering the ESKM/SKM Brocade group user name and password" on page 647.
- Set up a local CA on the ESKM/SKM. Refer to "Setting up the local Certificate Authority (CA) on ESKM/SKM" on page 648.
- Download the CA certificate. Refer to "Downloading the local CA certificate from ESKM/SKM" on page 649.
- Create and install an ESKM/SKM server certificate. Refer to "Creating and installing the ESKM/SKM server certificate" on page 649.
- Enable an SSL connection. Refer to "Enabling SSL on the Key Management System (KMS) Server" on page 651.
- Configure a cluster of ESKM/SKM appliances for high availability. Refer to the following sections:
  - "Creating an ESKM/SKM High Availability cluster" on page 651
  - "Copying the local CA certificate for a clustered ESKM/SKM appliance" on page 651
  - "Adding ESKM/SKM appliances to the cluster" on page 652
- Export and sign the encryption node certificate signing requests. Refer to "Signing the encryption node KAC certificates" on page 653.
- Import the signed certificates into the encryption node. Refer to "Importing a signed KAC certificate into a switch" on page 653.

Configuring a Brocade group on ESKM/SKM

A Brocade group is configured on ESKM/SKM for all keys created by encryption switches and blades. This needs to be done only once for each key vault.

1. Log in to the ESKM/SKM management web console using the admin password.
2. Select the **Security** tab.
3. Select **Local Users & Groups** under **Users and Groups**.
4. Select **Add** under **Local Users.**
5. Create a Brocade user name and password.
6. Select the **User Administration Permission** and **Change Password Permission** check boxes, then click **Save**.
7. Select **Add** under **Local Groups**.
8. Add a Brocade group under **Group**, then click **Save**.
9. Select the new Brocade group name, then select **Properties**.

Local **Group Properties** and a **User List** are displayed.

10. In the **User List** section, select or type the Brocade user name under **Username**, then click **Save**.

    The Brocade user name and password are now configured on ESKM/SKM.

Registering the ESKM/SKM Brocade group user name and password

The Brocade group user name and password you created when configuring a Brocade group on ESKM/SKM must also be registered on each encryption node.

NOTE
This operation can be performed only after the switch is added to the encryption group.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select the Group Leader switch from the **Encryption Center Devices** table, then select **Switch > Key Vault Credentials** from the menu task bar.

   The **Key Vault Credentials** dialog box displays. (Refer to Figure 281.)

FIGURE 281 Key Vault Credentials dialog box

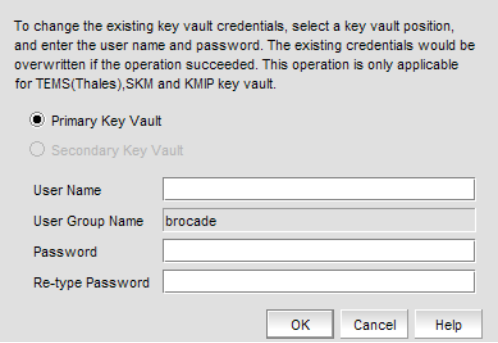

The dialog box contains the following information:

- **Primary Key Vault: P**reselected. ESKM/SKM key vaults are clustered, so only one set of credentials is needed.
- **Secondary Key Vault:** The selection is inactive.
- **User Name**: Enter a user name for the Group Leader.
- **User Group Name: Displays the selected User Group Name.**
- **Password:** Enter a password for the Group Leader.
- **Re-type Password:** Re-enter the password for verification.

3. Enter the Brocade user name and password, then re-enter the password for verification.
4. Repeat the procedure for each node.

General rules when creating user names and passwords

When creating user names and passwords for ESKM/SKM, the following rules apply:

- Initially, the user name and password are created when a Brocade user group is created on ESKM/SKM. The switch user name and password must match the user name and password specified for the Brocade group.

- The same user name and password must be configured on all nodes in an encryption group. This is not enforced or validated by the encryption group members, so use care when configuring the user name and password to ensure they are the same on each node.
- Different user names and passwords can never be used within the same encryption group, but each encryption group may have its own user name and password.
- If you change the user name and password, the keys created by the previous user become inaccessible. The Brocade group user name and password must also be changed to the same values on ESKM/SKM to make the keys accessible.
- When storage is moved from one encryption group to another, and the new encryption group uses a different user name and password, the Brocade group user name and password must also be changed to the same values on ESKM/SKM to make the keys accessible.

Setting up the local Certificate Authority (CA) on ESKM/SKM

To create and install a local CA, complete the following steps:

1. Log in to the ESKM/SKM management web console using the admin password.
2. Select the **Security** tab.
3. Under **Certificates & CAs**, click **Local CAs**. (Refer to Figure 282.)
4. Enter information required by the **Create Local Certificate Authorit**y section of the window to create your local CA.
   - Enter a **Certificate Authority Name** and **Common Name**. These may be the same value.
   - Enter your organizational information.
   - Enter the **Email Address** to receive messages for the Security Officer.
   - Enter the **Key Size**. HP recommends using 2048 for maximum security.
   - Select **Self-signed Root CA**.
   - Enter the **CA Certification Duration** and **Maximum User Certificate Duration**. These values determine when the certificate must be renewed and should be set in accordance with your company's security policies. The default value for both is 3650 days or 10 years.
5. Click **Create**.

   The new local CA displays under **Local Certificate Authority List**.

   NOTE
   Fabric OS 7.1.0 will use SHA256 signatures for the TLS certificates used to connect to the ESKM 3.0.

FIGURE 282 Creating an HP ESKM/SKM local CA

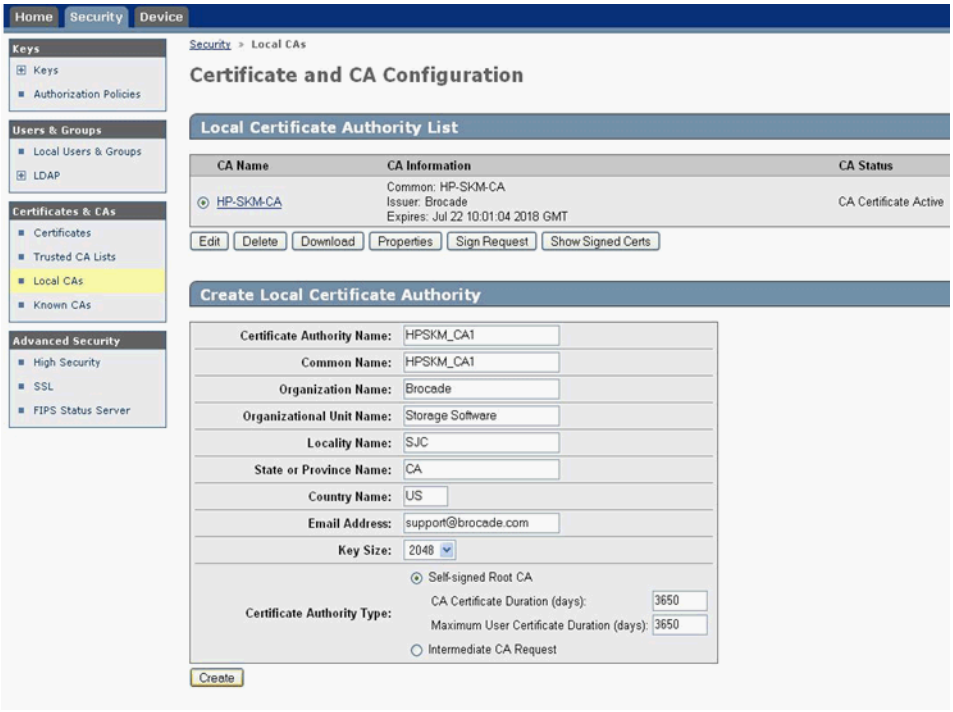

5. Under **Certificates & CAs**, select **Trusted CA Lists** to display the **Trusted Certificate Authority List Profiles**.
6. Click on **Default** under **Profile Name**.
7. In the **Trusted Certificate Authority List**, click **Edit**.
8. From the list of **Available CAs** in the right panel, select the CA you just created.

Repeat these steps any time another local CA is needed.

Downloading the local CA certificate from ESKM/SKM

The local CA certificate you created using the procedure for "Setting up the local Certificate Authority (CA) on ESKM/SKM" on page 648 must be saved to your local system. Later, this certificate must be imported onto the Brocade encryption Group Leader nodes.

1. From the **Security** tab, select **Local CAs** under **Certificates and CAs**.
2. Select the CA certificate you created and click **Download**, then save the certificate file on your local system.
3. Rename the downloaded file, changing the .cert extension to a .pem extension.

Creating and installing the ESKM/SKM server certificate

To create the ESKM/SKM server certificate, complete the following steps:

1. Click the **Security** tab.
2. Under **Certificates and CAs**, select **Certificates**.

3. Enter the required information under **Create Certificate Request**.
   - Enter a **Certificate Name** and **Common Name**. The same name may be used for both.
   - Enter your organizational information.
   - Enter the **E-mail Address** where you want messages to the Security Officer to go.
   - Enter the **Key Size**. HP recommends using the default value: 1024.
4. Click **Create Certificate Request**.

   Successful completion is indicated when the new entry for the server certificate displays on the **Certificate List** with a **Certificate Status** of **Request Pending.**
5. Select the newly created server certificate from the **Certificate List**.
6. Select **Properties**.

   The pending request displays under **Certificate Request Information**.
7. Copy the certificate data from -----BEGIN CERTIFICATE REQUEST----- to -----END CERTIFICATE REQUEST----- lines. Be careful to exclude extra carriage returns or spaces after the data.
8. Under **Certificates & CAs**, select **Local CAs**.

   The **Certificate and CA Configuration** page is displayed.
9. From the **CA Name** column, select the name of the local CA you just created in "Setting up the local Certificate Authority (CA) on ESKM/SKM" on page 648.
10. Click **Sign Request**.
11. Enter the required data in the **Sign Certificate Request** section of the window.
    - Select the CA name from the **Sign with Certificate Authority** drop-down list.
    - Select **Server** as the **Certificate Purpose**.
    - Enter the number of days before the certificate must be renewed based on your site's security policies. The default value is 3649 or 10 years.
12. Paste the copied certificate request data into the **Certificate Request** box.
13. Click **Sign Request**.

    The signed certificate request data displays under **Sign Certificate Request**.
14. Click **Download** to download the signed certificate to your local system.
15. Copy the signed certificate data, from -----BEGIN to END----- lines. Be careful to exclude extra carriage returns or spaces after the data.
16. From the **Security** tab select **Certificates** under **Certificates & CAs**.
17. Select the server certificate name you just created from the certificate list, and select **Properties**.

    The **Certificate Request Information** window displays.
18. Click **Install Certificate**.

    The **Certificate Installation** window displays.
19. Paste the signed certificate data you copied under **Certificate Response**, then click **Save**.

    The status of the server certificate should change from **Request Pending** to **Active**.

Enabling SSL on the Key Management System (KMS) Server

The KMS Server provides the interface to the client. Secure Sockets Layer (SSL) must be enabled on the KMS Server before this interface will operate. After SSL is enabled on the first appliance, it will be enabled automatically on the other cluster members.

To configure and enable SSL, complete the following steps:

1. Select the **Device** tab.
2. In the **Device Configuration** menu, click **KMS Server** to display the **Key Management Services Configuration** window.
3. In the **KMS Server Settings** section of the window, click **Edit**.
4. Configure the KMS Server Settings. Ensure that the port and connection timeout settings are 9000 and 3600, respectively. For **Server Certificate**, select the name of the certificate you created in "Creating and installing the ESKM/SKM server certificate" on page 649.
5. Click **Save**.

Creating an ESKM/SKM High Availability cluster

The HP ESKM/SKM key vault supports clustering of HP ESKM/SKM appliances for high availability. If two ESKM/SKM key vaults are configured, they must be clustered. If only a single ESKM/SKM appliance is configured, it may be clustered for backup purposes, but the backup appliance will not be directly used by the switch. The procedures in this section will establish a cluster configuration on one ESKM/SKM appliance and then transfer that configuration to the remaining appliances.

- Create the cluster on one ESKM/SKM appliance that is to be a member of the cluster.
- Copy the local CA certificate from the first ESKM/SKM appliance or an existing cluster member.
- Paste the local CA certificate into the management console for each of the ESKM/SKM appliances added to the cluster.

To create a cluster, complete the following steps on one of the HP ESKM/SKM appliances that is to be a member of the cluster:

1. From the ESKM/SKM management console, click the **Device** tab.
2. In the **Device Configuration** menu, click **Cluster**.

   The **Create Cluster** section displays.
3. Select and note the **Local IP** address. You will need this address when you add an appliance to the cluster.
4. For **Local Port**, use the default value of 9001 unless you are explicitly directed to use a different value for your site.
5. Type the cluster password in the **Create Cluster** section of the main window to create the new cluster, then click **Create**.
6. In the **Cluster Settings** section of the window, click **Download Cluster Key** and save the key to a convenient location, such as your computer's desktop. The cluster key is a text file and is only required temporarily. It may be deleted from your computer's desktop after all ESKM/SKM appliances have been added to the cluster.

Copying the local CA certificate for a clustered ESKM/SKM appliance

Before adding an ESKM/SKM appliance to a cluster, you must obtain the local CA certificate from the original ESKM/SKM or from an ESKM/SKM that is already in the cluster.

1. Select the **Security** tab.
2. Select **Local CAs** under **Certificates & CAs**.
3. Select the name of the local CA from the **Local Certificate Authority** list.

The **CA Certificate Information** is displayed.

4. Copy the certificate request, beginning with `---BEGIN CERTIFICATE REQUEST---` and ending with `---END CERTIFICATE REQUEST---`. Be careful not to include any extra characters.

Adding ESKM/SKM appliances to the cluster

If you are adding an appliance to an existing cluster, select the Cluster Settings section of the window, click **Download Cluster Key**, then save the key to a convenient location, such as your computer's desktop.

To add ESKM/SKM appliances to the cluster you are creating, you will need the original cluster member's local IP address and port number, and the location of the cluster key you downloaded, as specified in "Creating an ESKM/SKM High Availability cluster" on page 651.

Complete the following steps on each ESKM/SKM appliance you want to add to the cluster:

1. Open a new browser window, keeping the browser window from **Copying the Local CA certificate** open.
2. In the new browser window, log in to the management console of the ESKM/SKM appliance that is being added to the cluster, then click the **Security** tab.
3. In the **Certificates & CAs** menu, click **Known CAs**.
4. Enter the information required in the **Install CA Certificate** section near the bottom of the page.
    a. Enter the **Certificate Name** of the certificate being transferred from the first cluster member.
    b. Paste the copied certificate data into the **Certificate** box.
5. Click **Install**.
6. In the **Certificates & CA** menu, click **Trusted CA Lists**.
7. Click **Default Profile Name**, then click **Edit**.
8. Select the name of the CA from the list of **Available CAs** in the right panel, then click **Add**.
9. Click **Save**.
10. Select the **Device** tab.
11. In the **Device Configuration** menu, click **Cluster**.
12. Click **Join Cluster**. In the **Join Cluster** section of the window, leave **Local IP** and **Local Port** set to their default settings.
13. Enter the original cluster member's local IP address into **Cluster Member IP**.
14. Enter the original cluster member's local Port into **Cluster Member Port**.
15. Click **Browse**, then select the **Cluster Key File** you saved.
16. Enter the cluster password, then click **Join**.
17. After adding all members to the cluster, delete the cluster key file from the desktop.
18. Create and install an ESKM/SKM server certificate. Refer to "Creating and installing the ESKM/SKM server certificate" on page 649 for a description of this procedure.

Signing the encryption node KAC certificates

The KAC certificate signing request generated when the encryption node is initialized must be exported for each encryption node and signed by the Brocade local CA on ESKM/SKM. The signed certificate must then be imported back into the encryption node.

1. **Select Configure > Encryption from the menu task bar to display the** The **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. **Select a switch from the Encryption Center Devices table, then select Switch > Export Certificate, from the menu task bar.**

   The **Export Switch Certificate** dialog box displays.
3. Select **Public Key Certificate Request (CSR)**, then click **OK**.

   You are prompted to save the CSR, which can be saved to your SAN Management Program client PC, or an external host of your choosing.

   Alternatively, you may select a switch, then select **Switch > Properties**. Click the **Export** button beside the **Public Key Certificate Request**, or copy the CSR for pasting into the **Certificate Request Copy** area on the ESKM/SKM **Sign Certificate Request** page.
4. Launch the ESKM/SKM administration console in a web browser and log in.
5. Select the **Security** tab.
6. Select **Local CAs** under **Certificates & CAs**.

   The **Certificate and CA Configuration** page displays.
7. Under **Local Certificate Authority List**, select the Brocade CA name.
8. Select **Sign Request**.

   The **Sign Certificate Request** page displays.
9. Select **Sign with Certificate Authority** using the Brocade CA name and maximum of 3649 days.
10. Select **Client** as **Certificate Purpose**.
11. Allow Certificate **Duration** to default to 3649.
12. Paste the file contents that you copied in step 3 in the **Certificate Request Copy** area.
13. Select **Sign Request**.
14. Download the signed certificate to your local system as signed_kac_eskm_cert.pem or signed_kac_skm_cert.pem, depending on your key vault type.

    This file is ready to be imported to the encryption switch or blade.

Importing a signed KAC certificate into a switch

After a KAC CSR has been submitted and signed by a CA, the signed certificate must be imported into the switch.

NOTE
This operation can be performed only after the switch is added to the encryption group.

1. **Select Configure > Encryption from the menu task bar to display** the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)

1. Select a switch from the **Encryption Center Devices** table, then select **Switch > Import Certificate** from the menu task bar.

   The **Import Signed Certificate** dialog box displays. (Refer to Figure 283.)

FIGURE 283 Import Signed Certificate dialog box

2. Browse to the location where the signed certificate is stored, then click **OK**.

   The signed certificate is stored on the switch.

ESKM/SKM key vault high availability deployment

The ESKM/SKM key vault has high availability clustering capability. ESKM/SKM appliances can be clustered together in a transparent manner to the end user. Encryption keys saved to one key vault are synchronously hardened to the cluster pairs. Refer to the HP ESKM/SKM Appliance user documentation for configuration requirements and procedures.

The configured primary and secondary HP ESKM/SKM appliances must be registered with the encryption switch or blade to begin key operations. You can register only a single ESKM/SKM if desired. In that case, the HA features are lost, but the archived keys are backed up to any other non-registered cluster members. Note that the primary and secondary appliances must be clustered.

Both ESKM/SKM appliances in the cluster can be registered using the following command.

cryptocfg --reg -keyvault <**cert label**> <**certfile**> <**hostname/ip address**> <**primary** | **secondary**>

Data Encryption Keys

The following sections describe Data Encryption Key (DEK) behavior during DEK creation, retrieval, and updates as they relate to disk keys and tape pool keys, and tape LUN and DF-compatible tape pool support:

Disk keys and tape pool keys support

Data Encryption Key (DEK) creation, retrieval, and update for disk and tape pool keys are as follows:

- **DEK creation**: The DEK is first archived using the session list available for the configured ESKMs/SKMsin the cluster. After the DEK is archived successfully, it gets synchronized with other ESKMs/SKMs in the cluster. If archival is succesful, the DEK is then read from both the primary and secondary ESKMs/SKMs in the cluster until the DEK is read successfully from both.If the set of operations is successful, the DEK created can be used for encrypting disk LUNs or tape pools in Brocade native mode. If key archival of the DEK to the ESKM/SKM cluster fails, an error is logged and the operation is retried. If the failure occurs during DEK retrieval after successful archival to one of the ESKMs/SKMs, or synchronization to any ESKMS/SKMs in the cluster times out, an error is logged and the operation is retried. Any DEK archived in this case is not used.
  - If key archival of the DEK to the ESKM/SKM cluster is successful, the DEK is read from either the primary or secondary ESKMs or SKMs in the cluster until the DEK is read successfully from both. If successful, then the DEK created can be used for encrypting disk LUNs or tape pools in Brocade native mode.
  - If key archival of the DEK to the ESKM/SKM cluster fails, an error is logged and the operation is retried. If the failure occurs after archival to one of the ESKMs or SKMs, but synchronization to all ESKMS or SKMs in the cluster times out, then an error is logged and the operation is retried. Any DEK archived in this case is not used.

- **DEK retrieval**: The DEK is retrieved from the ESKM/SKM cluster using the session list available from the configured ESKMs/SKMs in the cluster. If the DEK retrieval fails, it is retried.
- **DEK update**: DEK update behavior is the same as DEK creation.

Tape LUN support

Data Encryption Key (DEK) creation, retrieval, and update for tape LUNs are as follows:

- **DEK creation**: The DEK is created and archived to the ESKM/SKM cluster using the using the session list available for configured ESKMs/SKMs in the cluster. The DEK is synchronized with other ESKMs/SKMs in the cluster. Upon successful archival of the DEK to the ESKM/SKM cluster, the DEK can be used for encryption of the tape LUN. If archival of the DEK to the ESKM/SKM cluster fails, an error is logged and the operation is retried.
- **DEK retrieval**: The DEK is retrieved from the ESKM/SKM cluster using the session list available for configured SKM/ESKM in the cluster. If the DEK retrieval fails, it is retried.
- **DEK update**: DEK update behavior is the same as DEK creation.

ESKM/SKM key vault deregistration

Deregistration of either the primary or secondary ESKM/SKM key vault from an encryption switch or blade is allowed independently.

- **Deregistration of primary ESKM**: You can deregister the primary ESKM/SKM from an encryption switch or blade without deregistering the backup or secondary ESKM/SKM for maintenance or replacement purposes. Future key operations will use only the secondary ESKM/SKM until the primary ESKM/SKM is reregistered on the Brocade Encryption Switch or blade.

  When the primary ESKM/SKM is replaced with a different ESKM/SKM, you must first synchronize the DEKs from the secondary ESKM/SKM before reregistering the primary ESKM/SKM.
- **Deregistration of secondary ESKM**: You can deregister the secondary ESKM/SKM independently. Future key operations will use only the primary ESKM/SKM until the secondary ESKM/SKM is reregistered on the encryption switch or blade.

  When the secondary ESKM/SKM is replaced with a different ESKM/SKM, you must first synchronize the DEKs from primary ESKM/SKM before reregistering the secondary ESKM/SKM.

Steps for connecting to a TEKA appliance

TEKA provides a web user interface for management of clients, keys, admins, and configuration parameters. A Thales officer creates domains, groups, and managers (a type of administrator), assigns groups to domains, and assigns managers to manage groups. Managers are responsible for creating clients and passwords for the groups they manage.

The following configuration steps are performed from the TEKA web user interface and from the Management application:

1. Set up network connections to TEKA. Refer to "Setting up TEKA network connections" on page 656.
2. Create a TEKA client. Refer to "Creating a client on TEKA" on page 657.
3. Establish TEKA key vault credentials. Refer to "Establishing TEKA key vault credentials on the switch" on page 658.
4. Sign encryption node certificate signing requests. Refer to "Exporting the Fabric OS node self-signed KAC certificates" on page 660.
5. Import the signed requests onto the encryption nodes. Refer to "Converting the KAC certificate format" on page 661.

Setting up TEKA network connections

Communicating to TEKA is enabled over an SSL connection. Two IP addresses are needed. One IP address is used for the management interface, and a second IP address is used for communication with clients. These IP addresses are typically assigned during the initial setup of the TEKA appliance.

1. Log in to the Thales management program as admin and select the **Network** tab. (Refer to Figure 284.)

FIGURE 284 TEKA Network Settings

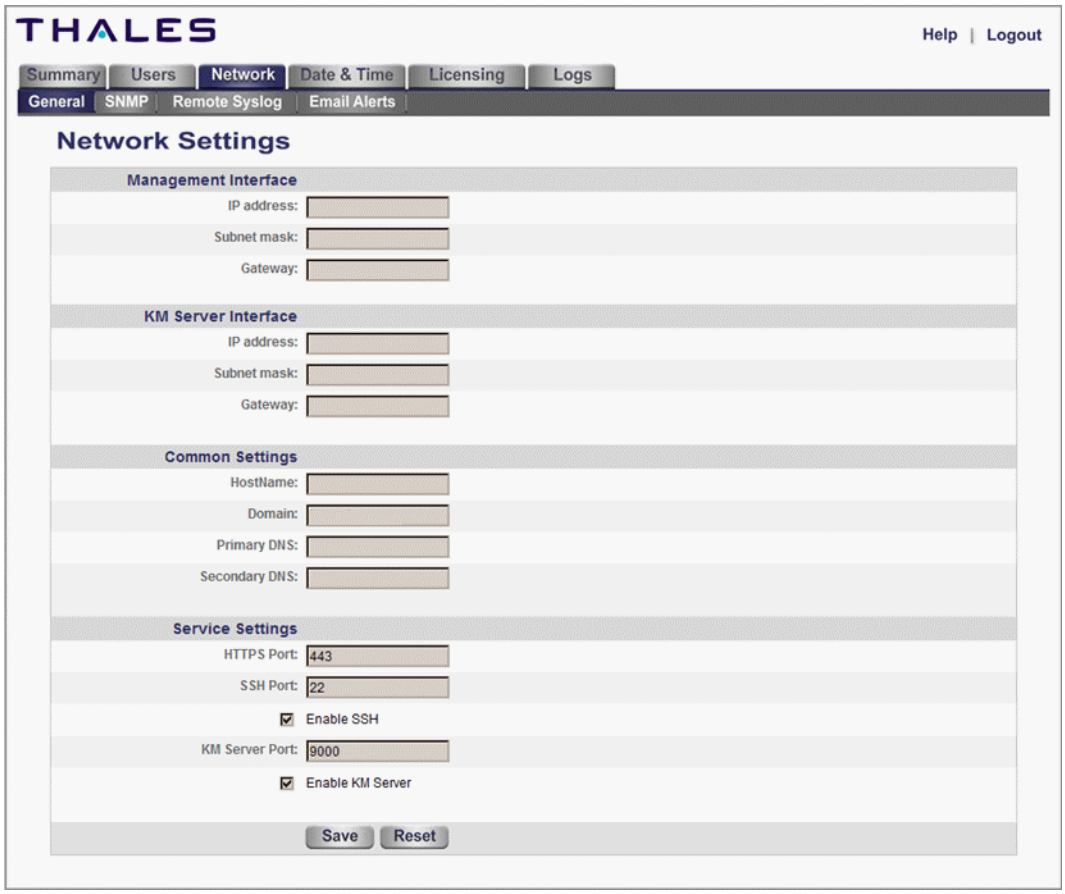

2. Enter the management IP address information under **Management Interface**.
3. Enter the client IP address information under **KM Server Interface**.
4. Enter a host name for the appliance, Internet or intranet domain, and, if used, the primary and secondary DNS IP address under **Common Settings**.
5. Set **Service Settings**.
   - HTTPS Port 433
   - SSH Port 22
   - Enable SSH
   - KM Server Port 9000
   - Enable KM Server

Creating a client on TEKA

This step assumes the group **brocade** has been created by an administrator. If the group **brocade** does not exist, you must log in to TEKA as officer and create the group, then assign the group to a manager.

1. From the **Encryption Center Devices** table, select a switch that needs to have a TEKA client, then select **Properties**.
2. Click **Key Vault User Name**.

   The **Key Vault User Information** dialog box displays. (Refer to Figure 285.)

FIGURE 285 TEKA Key Vault User Information

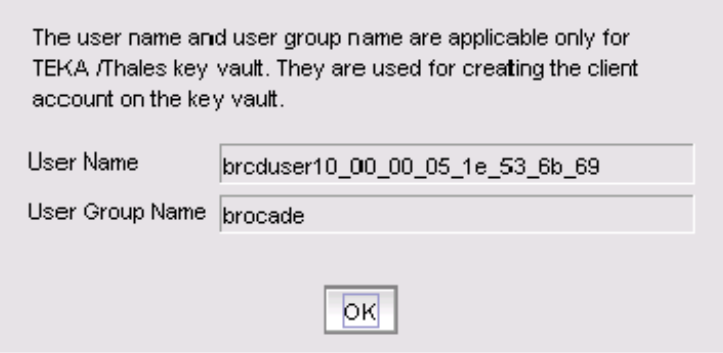

3. Copy the user name in the **User Name** field.
4. Log in to the Thales management program as a manager who has been assigned to the **brocade** group.
5. Select the **Clients** tab. (Refer to Figure 286.)

FIGURE 286 TEKA Clients tab

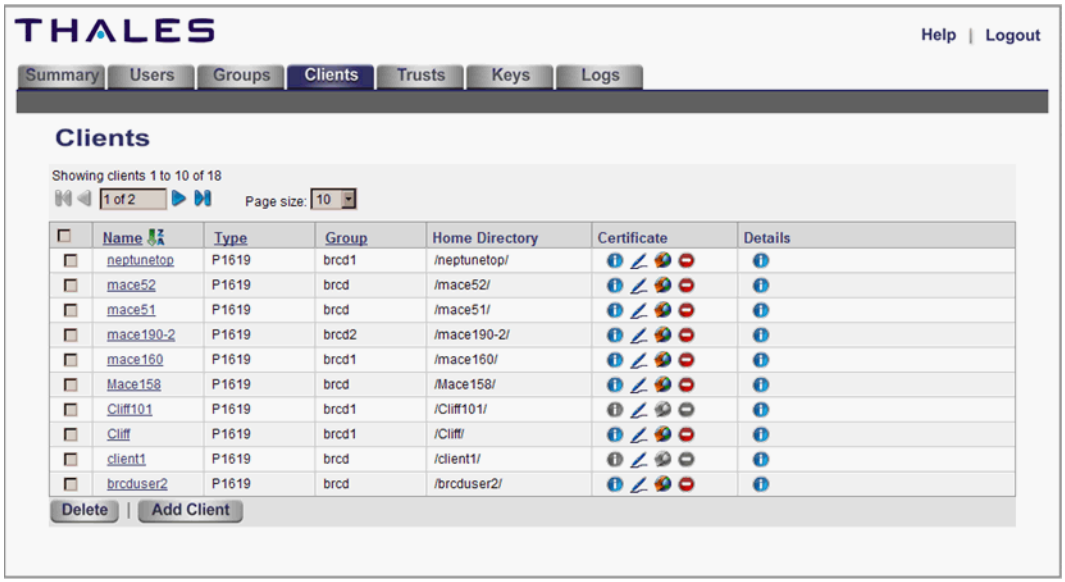

6. Click **Add Client**.
7. Enter the user name from step 3 in the **Name** field.
8. Enter a password in the **Password** and **Verify Password** fields.

9. Select the group **brocade** from the group pull-down menu, then click **Add Client**.

   A TEKA client user is created and is listed in the table.

Establishing TEKA key vault credentials on the switch

The credentials established for the TEKA client must be presented to TEKA by the Encryption switch. The primary and secondary TEKA key vaults must be installed and registered with the switch before you can configure CryptoTarget containers or LUNs.

1. From the **Encryption Center Devices** table, select a switch, then select **Switch > Key Vault Credentials** from the menu task bar.

   The **Key Vault Credentials** dialog box displays. (Refer to Figure 287.)

FIGURE 287 Key Vault Credentials dialog box

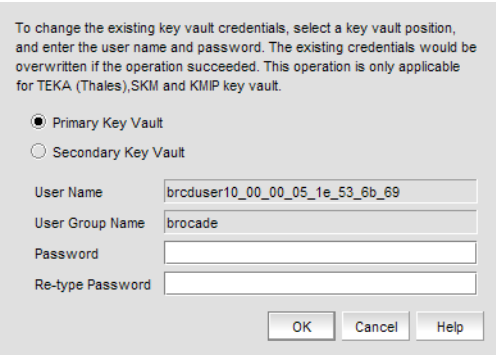

   The dialog box contains the following information:

   - **Primary Key Vault** selector**: Preselected.**
   - **Secondary Key Vault** selector**: A**ctive only if you are using a TEKA key vault.
   - **User Name: Used for creating the client account on the key vault.**
   - **User Group Name: Used for creating the client account on the key vault.**
   - **Password:** Enter a password for the Group Leader.
   - **Re-type Password:** Re-enter the password for verification.

2. Repeat the procedure for each node.
3. Copy the user name and password used when creating the TEKA client.

   You may create different credentials, but if you do, you must change the TEKA client credentials to match the new credentials.

4. Click **OK**.

   The following rules apply for TEKA:

   - The key vault user name and user group name are generated on the switch. To view those values, select **Switch > Properties**, then click **Key Vault User Name**.
   - The generated user name and user group name are registered with TEKA and are used for administering TEKA clients.
   - The password is established when the TEKA client is created.

Signing the encryption node KAC CSR on the TEKA appliance

The KAC certificate signing request (KAC CSR) generated when the encryption node is initialized must be exported for each encryption node and signed by the local CA on TEKA. The signed certificate must then be imported back into the encryption node.

1. From the **Encryption Center**, select **Switch > Export Certificate**.

   The **Export Switch Certificate** dialog box displays.

2. Select **Public Key Certificate Request (CSR)**, then click **OK**.

   A dialog box displays that allows you to save the CSR to your SAN Management Program client PC.

   Alternatively, you can select **Switch > Properties**, then click the **Export** button beside the **Public Key Certificate Request**, or you can copy the CSR for pasting in the **From Text** box on the Thales management program **Sign Certificate Request** page.

3. Log in to the Thales management program.
4. In the user table under the **Certificate** column, click the pen icon for the newly created user.

   The **Sign Certificate Request** page displays.

5. Enter the CSR file name exported from the switch in the **From File** box, or if you copied the CSR from **Switch > Properties**, paste the CSR file contents to the **From Text** box, then click **Sign**.
6. Under the **Certificate** column, click the export icon (globe with an arrow).

   A file save dialog displays.

7. Click **Save** and enter the destination location for this signed certificate. Save the certificate with a Privacy Enhanced Mail (.pem) extension.
8. Perform the above steps for both the primary and secondary key vaults using the same user name, password, and group.

Importing a signed KAC certificate into a switch

After a KAC CSR has been submitted and signed by a CA, the signed certificate must be imported into the switch.

1. From the Encryption Center, select **Switch > Import Certificate**.

   The **Import Signed Certificate** dialog box displays. (Refer to Figure 288.)

FIGURE 288 Import Signed Certificate dialog box

2. Browse to the location where the signed certificate is stored, then click **OK**.

   The signed certificate is stored on the switch.

Steps for connecting to a TKLM appliance

All switches you plan to include in an encryption group must have a secure connection to the Tivoli Key Lifecycle Manager (TKLM). A local LINUX host must be available to transfer certificates.

> NOTE
> Ensure that the time zone and clock time setting on the TKLM server and encryption nodes are the same. A difference of only a few minutes can cause the TLS connectivity to fail.

Repeat the same steps for configuring both the primary and secondary key vaults.

> NOTE
> The primary and secondary key vaults should be registered ***before*** you export the master key or encrypting LUNs. If the secondary key vault is registered ***after*** encryption is done for some of the LUNs, then the key database should be backed up and restored on the secondary TKLM from the registered primary TKLM before registering the secondary TKLM.

The following is a suggested order for the steps needed to create a secure connection to TKLM:

1. Initialize all encryption nodes to generate KAC certificates.
2. Export the signed KAC certificates to a local LINUX host. Refer to "Exporting the Fabric OS node self-signed KAC certificates" on page 660.
3. Obtain the necessary user credentials and log in to the TKLM server appliance from the TKLM management web console.
4. Create a default key store on TKLM. Refer to "Establishing a default key store and device group on TKLM" on page 661.
5. Create a device group named BRCD_ENCRYPTOR with device family LTO.
6. Add devices to the group. Refer to "Adding a device to the device group" on page 661.
7. Create a certificate for the TKLM server. Refer to "Creating a self-signed certificate for TKLM" on page 661.
8. Import the node KAC certificates. Refer to "Importing the Fabric OS encryption node KAC certificates to TKLM" on page 662.
9. Export the server CA certificate to a LINUX or Windows host. Refer to "Exporting the TKLM self-signed server certificate" on page 662.
10. Add encryption group members as needed. The first node added to an encryption group functions as the Group Leader. It is valid to have only one node in an encryption group.
11. Import the server CA certificate and register TKLM on the encryption Group Leader nodes. Refer to "Importing the TKLM certificate into the group leader" on page 663.
12. Enable the encryption engines.

Exporting the Fabric OS node self-signed KAC certificates

Each Fabric OS node generates a self-signed KAC certificate as part of the node initialization process as described under "Encryption node initialization and certificate generation". These certificates must be exported from each switch and stored on a local LINUX host to make them available for importing to TKLM.

1. Select a switch from the **Encryption Center Devices** table, then select **Switch > Export Certificate** from the menu task bar**.**

   The **Export Signed Certificate** dialog box displays.

2. Select **Signed switch certificate**, then click **OK**.

   A dialog box displays allowing you to save the signed certificate in .pem format in My Documents on your work station. Make a note of this location.

Converting the KAC certificate format

The KAC certificate exported from the encryption switch is in .pem format. It is automatically converted to a .der format during the export process; however, if you need to manually convert the file before importing it to the TKLM server, you can do so by completing the following steps:

1. Go to openssl utility.
2. Run `openssl x509 -outform der -in KAC_Certificate_Name.pem -out KAC_Certificate_Name.der`.

Establishing a default key store and device group on TKLM

To establish a default key store and Fabric OS device group on TKLM, complete the following steps:

1. Obtain the necessary user credentials, then log in to the TKLM user interface.
2. Select **Advanced Configuration > Keystore**.

   The **Keystore** page displays.
3. Click **OK** to accept the default keystore settings.

Adding a device to the device group

After you have established a default key store and Fabric OS device group on TKLM, add a Fabric OS device to the device group.

1. Select **Tivoli Key Lifecycle Manager > Welcome**.

   The device group **BRCD_ENCRYPTOR** you just created is displayed in the **Administration** panel.
2. Click **Go**.

   The **Configure Keys** page displays. This page identifies this step as **Step Two: Identify Drives**.
3. Click **Add** on the **Devices** table menu task bar, which adds the entry to the table.
4. Under **Device Serial Number**, enter **t**he serial number that is displayed for each node that you are adding to the device group.

Creating a self-signed certificate for TKLM

You must create a self-signed certificate for TKLM that can be downloaded to the Fabric OS encryption engines to verify the authenticity of TKLM.

1. **Select Tivoli Key Lifecycle Manager > Configuration**.

   The **Configuration** page displays.
2. Select **Create self-signed certificate**.
3. Under **Certificate label in key store**, enter a certificate label.
4. Under **Certificate description (common name)**, enter a descriptive name.

5. Under **Validity period of new certificate**, enter the desired life time for the certificate.
6. Select **Tivoli Key Lifecycle Manager > Advanced Configuration > Server Certificates** to verify that the certificate label is listed on **Administer Server Certificates** under **Certificates**.
7. Reboot the TKLM server.

Importing the Fabric OS encryption node KAC certificates to TKLM

The KAC certificates previously exported from the Fabric OS encryption nodes to an external LINUX host must now be imported into the TKLM server file system. You must import the KAC certificate in .der format. To do this, refer to "Converting the KAC certificate format" on page 661.

1. Import the KAC certificate from the external host into the TKLM server file system using a binary file transfer mechanism using FTP, USB, or SCP.
2. Select **Tivoli Key Lifecycle Manager > Advanced Configuration > Client Certificates.**

   The **Client Certificates** page displays.
3. Select **Import > SSL Certificate**.

   The **Import SSL Certificates for Clients** page displays.
4. Enter the Fabric OS KAC certificate name in the **Certificate** field.
5. Under **File name and location**, enter or browse to the location where the imported KAC certificate is stored, then select **Trust**.
6. Click **Import**.
7. Verify that the imported certificate is valid and active.

Exporting the TKLM self-signed server certificate

The TKLM self-signed server certificate must be exported in preparation for importing and registering the certificate on a Fabric OS encryption Group Leader node.

1. Enter the TKLM server wsadmin CLI.

   For Linux (in ./wsadmin.sh):

   ```
   <installed directory>/IBM/tivoli/tiptklmV2/bin/wsadmin.sh -username TKLMAdmin -password <password> -lang
   jython
   ```

   For Windows:

   ```
   <installed directory>\ibm\tivoli\tiptklmV2\bin\wsadmin.bat -username TKLMAdmin -password <password> -lang
   jython
   ```

2. Check the certificate list using the following command:

   ```
   print AdminTask.tklmCertList('[]')
   ```

   The listing will contain the UUID for all certificates. Use the UUID of the server certificate to export the server certificate from the database to the file system.

   ```
   print AdminTask.tklmCertExport('[
   -uuid <UUID of the certificate>
   -fileName <filename> -format DER]')
   ```

3. Exit the wsadmin CLI

   After export, the TKLM server certificate is at the following location:

   For LINUX:

   ```
   <installed directory>/ibm/tivoli/tiptklmV2/products/tklm/
   ```

   For Windows:

   ```
   <installed directory>\ibm\tivoli\tiptklmV2\products\tklm\
   ```

4. Transfer the TKLM certificate that was previously exported into the TKLM server file system to the Management application host using any binary file transfer mechanism via SCP, USB, or FTP.

Importing the TKLM certificate into the group leader

The TKLM certificate must be imported from the location on the host to the encryption Group Leader node. The encryption Group Leader exports the certificate to group member switches.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a switch from the **Encryption Center Devices** table, then select **Switch > Import Certificate** from the menu task bar.

   The **Import Signed Certificate** dialog box displays. (Refer to Figure 289.)

   FIGURE 289 Import Signed Certificate dialog box

   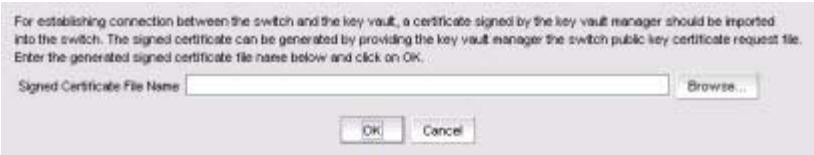

3. Browse to the location where the signed certificate is stored, then click **OK**.

   The signed certificate is stored on the switch.

Steps for connecting to a KMIP-compliant SafeNet KeySecure

With the introduction of Fabric OS 7.1.0, the Key Management Interoperability Protocol (KMIP) KeySecure Management Console can be used on the Encryption switch. Any KMIP-compliant server can be re-registered as a KMIP key vault on the Encryption switch after setting the key vault type to KMIP.

Currently, KMIP with SafeNet KeySecure 6.1 in native KMIP mode with the Brocade Encryption Switch in KMIP mode is supported. All nodes in an encryption group should be running Fabric OS 7.1.0 and later for the key vault type to be set to KMIP.

After installing the SafeNet KeySecure appliance (also referred to as the KeySecure), you must complete the following steps before the Encryption switch can be configured with the KeySecure. These steps must be performed only once, in preparation for first-time configuration.

> NOTE
> If you are configuring two KeySecure nodes, you must complete step 1 through step 6 on the primary node, then complete step 7 on the secondary node. If only a single node is being configured, step 7 is not needed.

The following suggested order of steps must be completed to create a secure connection to the SafeNet KeySecure.

1. Set FIPS compliance. (Refer to "Setting FIPS compliance" on page 664.)
2. Create a local CA. (Refer to "Creating a local CA" on page 664.)
3. Create a server certificate. (Refer to "Creating a server certificate" on page 664.)
4. Create a cluster. (Refer to "Creating a cluster" on page 665.)
5. Create a Brocade group on the KeySecure appliance. (Refer to "Configuring a Brocade group on the KeySecure" on page 665.)
6. Register the user name and password. (Refer to "Registering the KeySecure Brocade group user name and password" on page 666.)
7. Export and sign the encryption node certificate signing requests. (Refer to "Signing the encryption node KAC CSR on KMIP" on page 667.)
8. Import the signed certificates into the encryption node. (Refer to "Importing a signed KAC certificate into a switch" on page 668.)
9. Back up the certificates (Refer to "Backing up the certificates" on page 669.)
10. Configure the KMIP server. (Refer to "Configuring the KMIP server" on page 670.)
11. Add a secondary node to the cluster. (Refer to "Adding a node to the cluster" on page 670.)

Setting FIPS compliance

1. From the KeySecure Management Console, select the **Security** tab, then select **Advanced Security**, > **High Security**.

   The **High Security Configuration** page displays.

2. Under **FIPS Compliance**, set **FIPS Compliance** to **Yes**.

   This ensures that only TLS 1.0 connections are supported between the Encryption switch and the KeySecure.

Creating a local CA

1. From the KeySecure Management Console, select the **Security** tab, then select **CAs & SSL Certificates > Local CAs**.

   The **Certificate and CA Configuration** page displays.

2. Under **Create Local Certificate Authority**, enter the organization information in the fields provided, then click **Create**. The example is using SafeNetCA as the Local CA name.

   The new Local CA is listed in the **Local Certificate Authority List** table.

3. Verify the Local CA status is shown as **Active**.

Creating a server certificate

1. From the **Security** tab, select **CAs & SSL Certificates > SSL Certificates**.

   The **Certificate and CA Configuration** page displays.

2. Under **Create Certificate Request**, enter your organization information in the fields provided, then click **Create Certificate Request**. (The example is using "Safenet75ServerCert" as the server certificate name.)

After the page refreshes, the new certificate information is displayed in the **Certificate List** table.

3. Verify the server certificate status is shown as **Request Pending**.
4. Click on the server certificate name that you just created (Safenet75ServerCert), which displays the certificate contents.
5. Copy the certificate contents.
6. From the **Security** tab, select **CAs & SSL Certificates > Local CAs**.

   The **Certificate and CA Configuration** page displays.
7. Under **Local Certificate Authority List**, select the CA certificate you just created (SafeNetCA), then click **Sign Request**.

   The **Sign Certificate Request** dialog box displays.
8. Select **Server** as the **Certificate Purpose** and verify the **Certificate Duration** length. The default is 3649 days.
9. Paste the server certificate contents that you copied (refer to step 5) in the **Certificate Request** text box, then click **Sign Request**.

   The **Certificate and CA Configuration** page refreshes and the certificate information is displayed under **Certificate Request Information.**
10. Click **Download** after the request has been signed, and save the certificate to a local location.
11. Click **Install Certificate**.
12. Open the downloaded certificate and copy the certificate data from -----BEGIN CERTIFICATE REQUEST----- to -----END CERTIFICATE REQUEST----- lines. Be careful to exclude extra carriage returns or spaces after the data
13. Paste the server certificate request contents in the **Certificate Installation** text box, then click **Save.**
14. After the page refreshes, the new certificate information is displayed in the **Certificate List** table.
15. Verify the server certificate status is shown as **Active**.

Creating a cluster

1. From the KeySecure Management Console, select the **Device** tab, then select **Device Configuration > Cluster**.

   The **Cluster Configuration** page displays.
2. Under **Create Cluster**, enter a user-defined password in the fields provided, then click **Create**.

   The **Cluster Configuration** page refreshes; the new cluster information is listed in the **Cluster Members** table.
3. Verify the cluster status is shown as **Active**.
4. Under **Cluster Settings**, click **Download Cluster Key.**

   You are prompted to enter a local file name.

Configuring a Brocade group on the KeySecure

A Brocade group is configured on the KeySecure for all keys created by encryption switches and blades. This needs to be done only once for each key vault.

1. Log in to the KeySecure web management console using the admin password.
2. Select the **Security** tab.

3. Select **Local Users & Groups** under **Users & Groups**.
4. Select **Add** under **Local Users**.
5. Create a Brocade user name and password.
6. Select the **User Administration Permission** and **Change Password Permission** check boxes, then click **Save**.
7. Select **Add** under **Local Groups**.
8. Add a Brocade group under **Group**, then click **Save**.
9. Select the new Brocade group name, then select **Properties**.

   The **Local Group Properties** and a **User List** are displayed. (Refer to Figure 290.)

FIGURE 290 User & Group Configuration page - Local Group Properties and User List

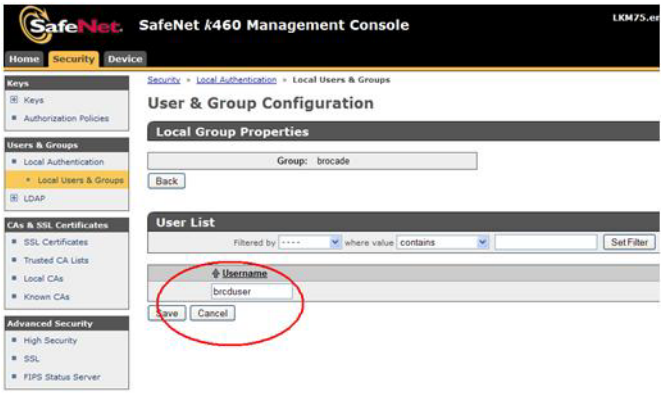

10. Under **User List**, select or type the Brocade user name under **Username**, then click **Save**.

    The Brocade user name and password are now configured on the KeySecure.

    NOTE
    The user name and password must also be registered on the Management application. Proceed to "Registering the KeySecure Brocade group user name and password".

Registering the KeySecure Brocade group user name and password

The Brocade group user name and password you created when configuring a Brocade group on the KeySecure must also be registered on each encryption node.

NOTE
This operation can be performed during or after the creation of the encryption group. During the creation of an encryption group, the key vault step will prompt for a user name and password.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select the group leader switch from the **Encryption Center Devices** table, then select **Switch > Key Vault Credentials** from the menu task bar.

   The **Key Vault Credentials** dialog box displays. (Refer to Figure 291).

FIGURE 291 Key Vault Credentials dialog box

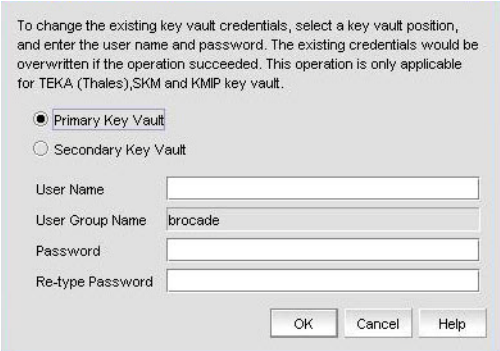

The dialog box contains the following information:

- **Primary Key Vault: Primary Key Vault** is preselected. KMIP key vaults are clustered, so only one set of credentials is needed.
- **Secondary Key Vault:** *(TEKA key vault only)*. Shown as inactive.
- **User Name**: Enter a user name for the group leader.
- **User Group Name: Displays the selected User Group Name.**
- **Password:** Enter a password for the group leader.
- **Re-type Password:** Re-enter the password for verification.

3. Enter the Brocade user name and password, then re-enter the password for verification.
4. Click **OK**.

Signing the encryption node KAC CSR on KMIP

The KAC certificate signing request generated when the encryption node is initialized must be exported for each encryption node and signed by the Brocade local CA on KMIP. The signed certificate must then be imported back into the encryption node.

1. **Select Configure > Encryption from the menu task bar to display the** The **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. **Select a switch from the Encryption Center Devices table, then select Switch > Export Certificate, from the menu task bar.**

   The **Export Switch Certificate** dialog box displays.
3. Select **Public Key Certificate Request (CSR)**, then click **OK**.

   You are prompted to save the CSR, which can be saved to your SAN Management Program client PC, or an external host of your choosing.

   Alternatively, you may select a switch, then select **Switch > Properties**. Click the **Export** button beside the **Public Key Certificate Request**, or copy the CSR for pasting into the **Certificate Request Copy** area on the KMIP **Sign Certificate Request** page.
4. Launch the KMIP administration console in a web browser and log in.
5. From the KeySecure Management Console, select the **Security** tab, then select **CAs & SSL Certificates > Local CAs**.
6. The **Certificate and CA Configuration** page displays.
7. Under **Local Certificate Authority List**, select the local CA name, and verify that its **CA Status** is shown as **Active**.
8. Click **Sign Request**.

The **Sign Certificate Request** page displays. (Refer to Figure 292.)

FIGURE 292 Certificate and CA Configuration page - Sign Certificate Request

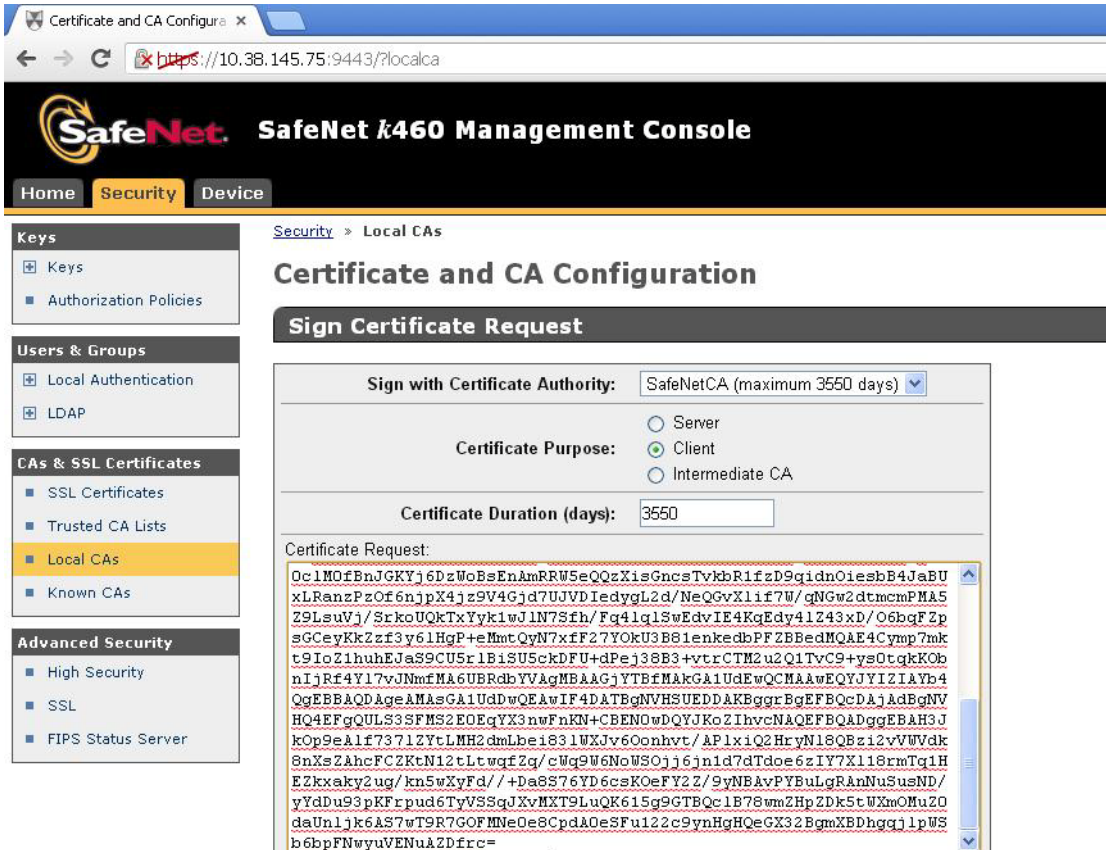

9. Select the local CA from the **Sign with Certificate Authority** drop-down list. The example is using “SafeNetCA”.
10. Select **Client** as **Certificate Purpose**.
11. Set **Certificate Duration.** (Default is 3649 days.)
12. Paste the file contents that you copied in step 3 in the **Certificate Request** area.
13. Click **Sign Request**.
14. Download the signed certificate to your local system as signed_kac_kmip_cert.pem.

    This file is ready to be imported to the encryption switch or blade.

Importing a signed KAC certificate into a switch

After a KAC CSR has been submitted and signed by a CA, the signed certificate must be imported into the switch.

NOTE
This operation can be performed only after the switch is added to the encryption group.

1. **Select Configure > Encryption from the menu task bar to display** the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a switch from the **Encryption Center Devices** table, then select **Switch > Import Certificate** from the menu task bar.

The **Import Signed Certificate** dialog box displays. (Refer to Figure 293.)

FIGURE 293 Import Signed Certificate dialog box

3. Browse to the location where the signed certificate is stored, then click **OK**.

   The signed certificate is stored on the switch.

Backing up the certificates

1. From the KeySecure Management Console, select the **Device** tab, then select **Maintenance > Backup & Restore > Create Backup**.

   The **Backup and Restore** page displays.

2. Select the server certificate from the list. The example is using **Safenet75ServerReq**.
3. Select the local CA from the list. The example is using **SafeNetCA**.
4. Select the **High Security** and **FIPS Status Server** check boxes, then click **Continue**.

   A list of backup device items displays. (Refer to Figure 294.)

FIGURE 294 Backup and Restore page - Device items

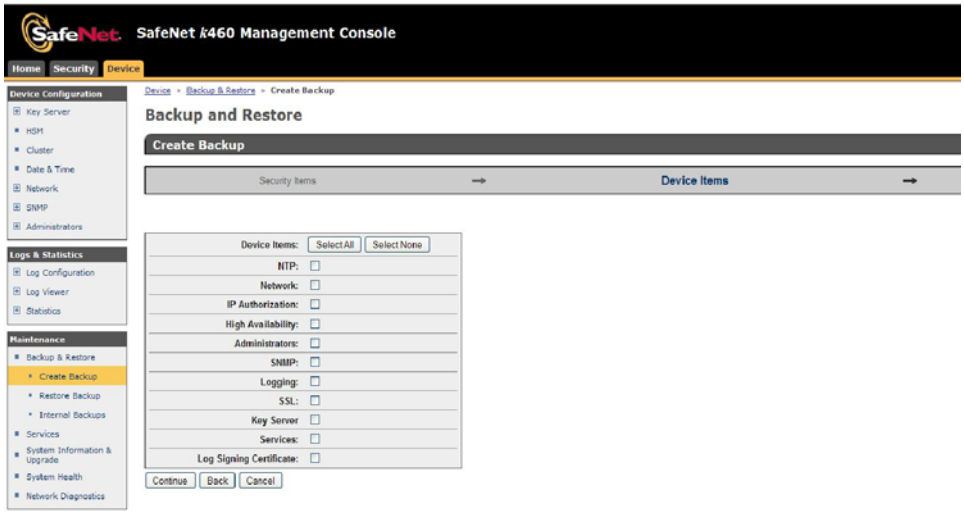

5. Select the items for backup, then click **Continue**.

   The **Create Backup** page displays, which is used for setting backup details. (Refer to Figure 295.)

FIGURE 295Backup and Restore page - Backup details

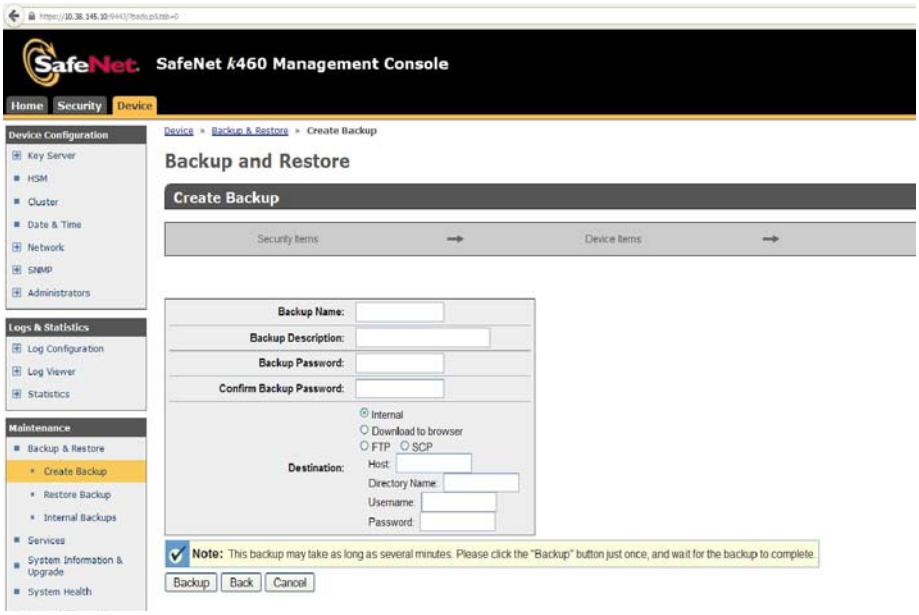

6. Enter backup details in the fields provided, then click **Backup** to initiate the backup process.
7. Restore this backup file on the Secondary clustered KeySecure server.

Configuring the KMIP server

1. From the KeySecure Management Console, select the **Device** tab, then select **Device Configuration > Key Server > Key Server**.

   The **Cryptographic Key Server Configuration** page displays.

2. Under **Cryptographic Key Server Settings**, select **KMIP** as the protocol.
3. Ensure that the **Use SSL** check box is selected.
4. Click **Edit** to open a dialog box for changing **IP**, **Port**, and **Server Certificate** settings.
5. After changing/adding your settings, save your settings.

   You are returned to the **Cryptographic Key Server Configuration** page. The settings are displayed in the table.

Adding a node to the cluster

Perform the following steps on the secondary KeySecure node when adding it to the cluster.

1. From the KeySecure Management Console, select the **Device** tab, then select **Device Configuration > Cluster**.

   The **Cluster Configuration** page displays.

2. Under **Join Cluster**, enter the cluster information that you configured for the primary KeySecure node. (Refer to "Creating a cluster" on page 665.)
3. Enter the primary KeySecure node IP address and port number in the respective **Cluster Member IP** and **Port** fields.
4. Enter the **Cluster Key File** or browse to the file location.

5. Enter the **Cluster Password**, then click **Join**.

   You are returned to the **Cluster Configuration** page with the cluster information listed in the **Cluster Members** table.

6. Verify that both KeySecure nodes are shown as **Active**.
7. From the **Devices** tab, select **Maintenance > Backup and Restore > Restore Backup**.

   The **Backup and Restore** page displays.

8. Under **Restore Backup**, select **Upload from browser**, then enter a file name or browse to the file location.
9. Enter the **Backup Password** in the field provided, then click **Restore**.
10. After the certificate is restored to the secondary node from the previously backed-up primary node, select **Maintenance > Services**.

    The **Services Configuration** page displays.

    NOTE
    A message displays, advising that the secondary node requires a restart.

11. Under **Restart/Halt**, select **Restart**, then click **Commit** and wait until the restart is completed.

    The primary and second KeySecure nodes are now in a cluster and active for use.

Steps for connecting to a KMIP-compliant keyAuthority

If you are using a TEKA KMIP-compliant server, only Thales e-Security keyAuthority running version 4.0 is supported; however, before selecting KMIP as the key vault type, all nodes in an encryption group must be running Fabric OS 7.2.0 or later.

For more information about configuration instructions, refer to Chapter 3 of the *Fabric OS Encryption Administrator's Guide Supporting Key Management Interoperability Protocol (KMIP) Key-Compliant Environments*.

Encryption preparation

Before you use the encryption setup wizard for the first time, you should have a detailed configuration plan in place and available for reference. The encryption setup wizard assumes the following:

- You have a plan in place to organize encryption devices into encryption groups.
- If you want redundancy and high availability in your implementation, you have a plan to create high availability (HA) clusters of two encryption switches or blades to provide failover support.
- All switches in the planned encryption group are interconnected on an I/O synch LAN.
- The management ports on all encryption switches and 8-slot Backbone Chassis CPs that have encryption blades installed, have a LAN connection to the SAN management program and are available for discovery.
- A supported key management appliance is connected on the same LAN as the encryption switches, 8-slot Backbone Chassis CPs, and the SAN Management program.
- An external host is available on the LAN to facilitate certificate exchange.
- Switch KAC certificates have been signed by a CA and stored in a known location.
- Key management system (key vault) certificates have been obtained and stored in a known location.

Creating a new encryption group

The following steps describe how to start and run the encryption setup wizard and create a new encryption group.

NOTE
When a new encryption group is created, any existing tape pools in the switch are removed.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 296.)

FIGURE 296 Encryption Center dialog box - No group defined

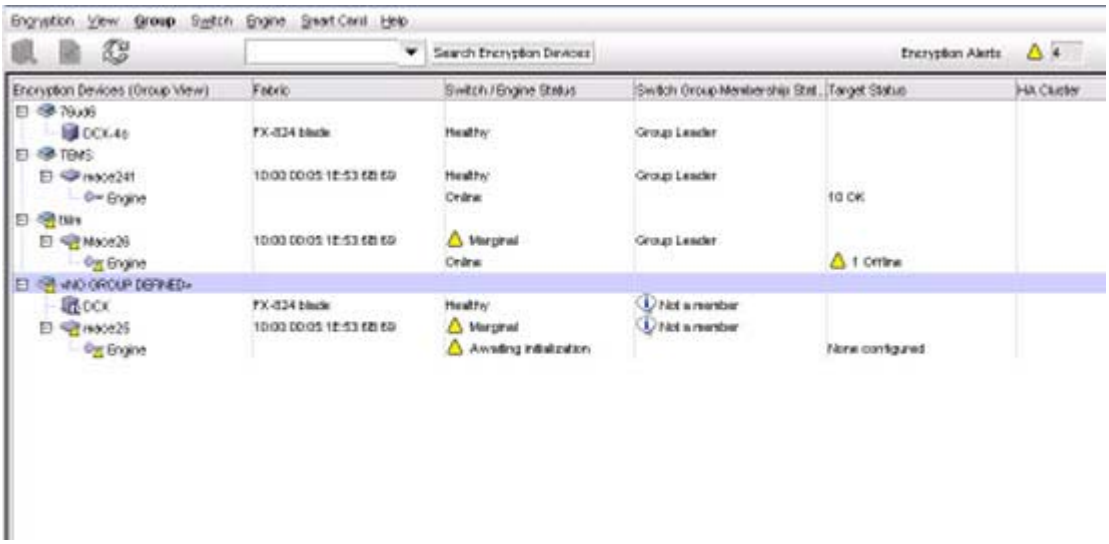

2. Select a switch from the **<NO GROUP DEFINED>** encryption group. (The switch must not be assigned to an encryption group.)
3. Select **Encryption > Create/Add to Group**, from the menu task bar.

   The **Configure Switch Encryption** wizard welcome screen displays. (Refer to Figure 297.) The wizard enables you to create a new encryption group, or add an encryption switch to an existing encryption group. The wizard also enables you to configure switch encryption.

   Click **Next** on each screen to advance to the next step in the wizard. Steps might vary slightly depending on the key vault type selected, but the basic wizard steps are as follows.

   a. Designate Switch Membership.
   b. Create a new encryption group or add a switch to an existing encryption group.
   c. Select the key vault.
   d. Specify the public key filename.
   e. Select Security Settings.
   f. Confirm the configuration.
   g. Configuration Status.
   h. Read Instructions.

FIGURE 297Configure Switch Encryption wizard - welcome screen

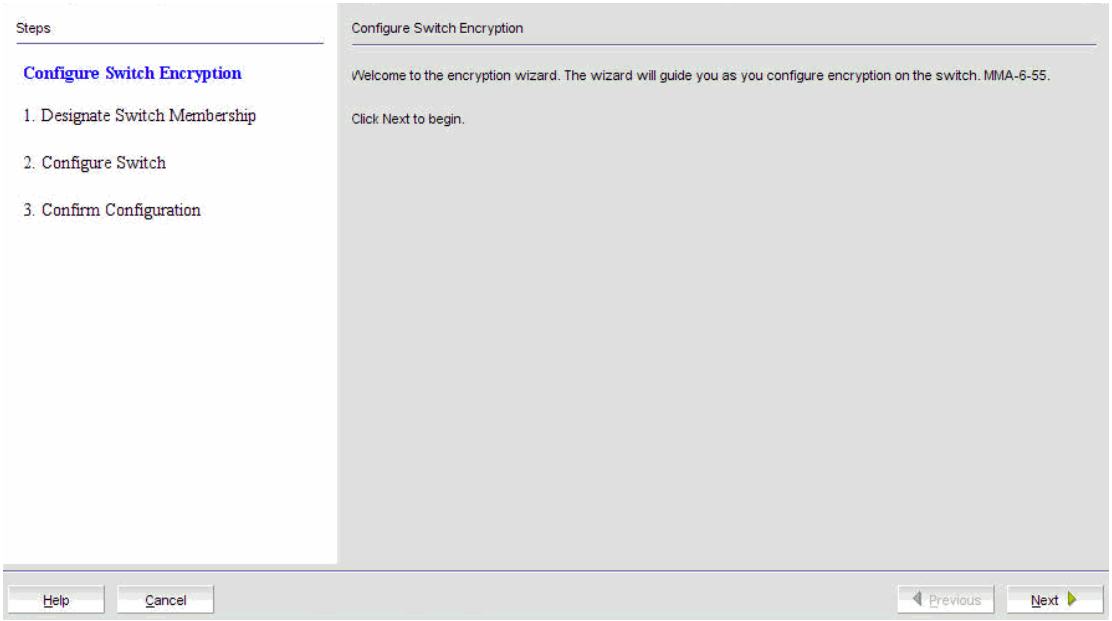

4. From the **Configure Switch Encryption** welcome screen, click **Next** to begin.

   The **Designate Switch Membership** dialog box displays (Figure 298). The dialog box contains the following options:

   - **Create a new encryption group containing just the switch**: Creates an encryption group for the selected switch
   - **Add this switch to an existing encryption group**: Adds the selected switch to an encryption group that already exists

FIGURE 298**Designate Switch Membership** dialog box

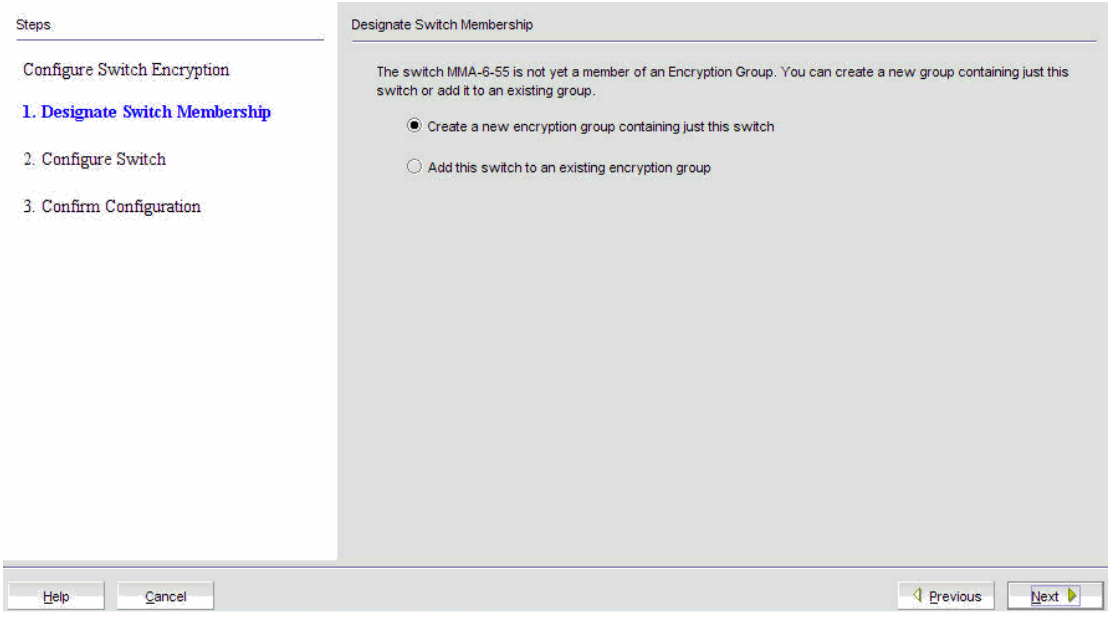

5. For this procedure, verify that **Create a new encryption group containing just this switch** is selected, then click **Next**.

   NOTE
   If you are adding a switch to an encryption, refer to "Adding a switch to an encryption group" on page 707.

The **Create a New Encryption Group** dialog box displays. (Refer to Figure 299.)

FIGURE 299 Create a New Encryption Group dialog box

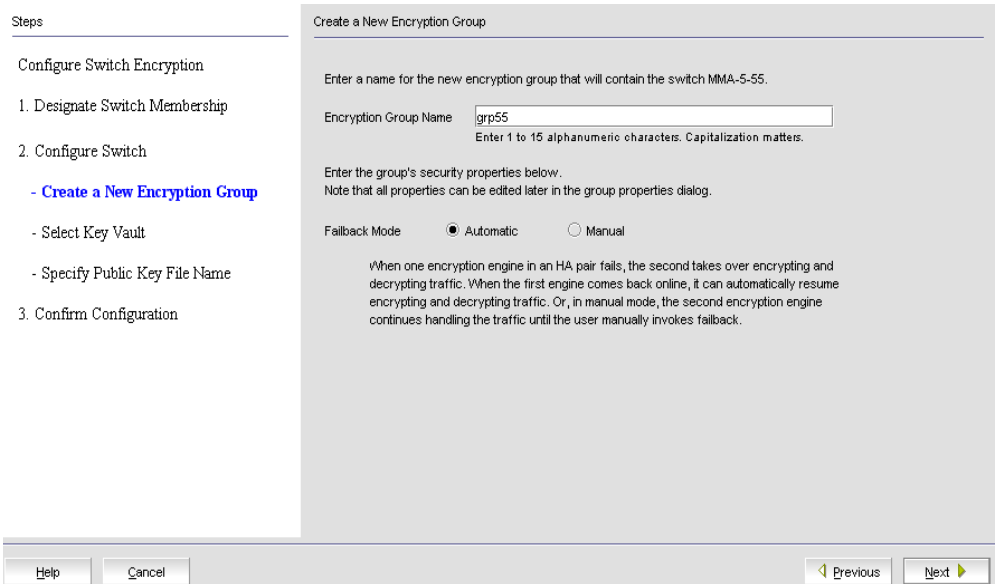

The dialog box contains the following information:

- **Encryption Group Name** text box: Encryption group names can have up to 15 characters. Letters, digits, and underscores are allowed. The group name is case-sensitive.
- **Failback mode**: Selects whether or not storage targets should be automatically transferred back to an encryption engine that comes online after being unavailable. Options are **Automatic** or **Manual**.

  NOTE
  When one encryption engine in the HA cluster fails, the second encryption engine in the HA cluster takes over the encryption and decryption of traffic to all encryption targets in the first encryption engine (failover). When the first encryption engine comes back online, the encryption group's failback setting (auto or manual) determines whether the first encryption engine automatically resumes encrypting and decrypting traffic to its encryption targets. In manual mode, the second encryption engine continues to handle the traffic until you manually invoke failback by way of the **Encryption Targets** dialog box.

6. Enter an **Encryption Group Name** for the encryption group and select **Automatic** as the Failback mode.

   If the name for the encryption group already exists, a pop-up warning message displays. Although unique group names avoid confusion while managing multiple groups, you are not prevented from using duplicate group names. Click **Yes** to use the same name for the new encryption group, or click **No** to enter another name.

7. Click **Next**.

   The **Select Key Vault**. dialog box displays. (Refer to Figure 300.)

FIGURE 300Select Key Vault dialog box

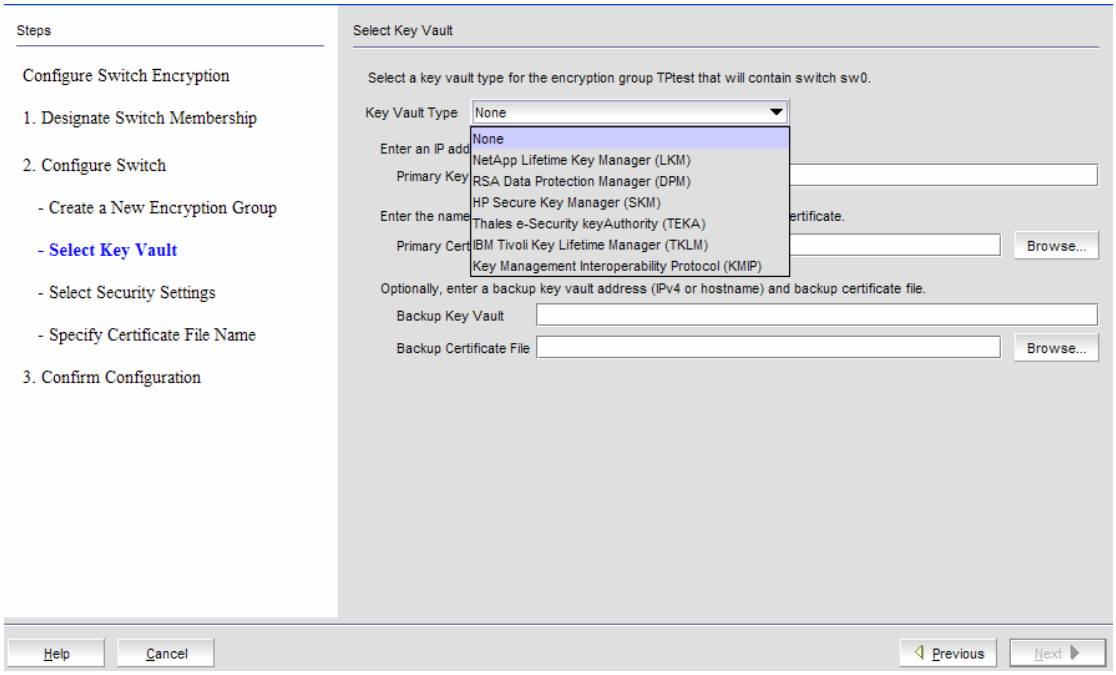

Using this dialog box, you can select a key vault for the encryption group that contains the selected switch. Prior to selecting your Key Vault Type, the selection is shown as **None**. The dialog box contains the following information:

- **Key Vault Type**:

  If an encryption group contains mixed firmware nodes, the Encryption Group Properties Key Vault Type name is based on the firmware version of the Group Leader. Options are:

  - **NetApp Lifetime Key Manager (LKM):** The NetApp Key Vault Type name is shown as NetApp Lifetime Key Manager (LKM) for both NetApp Lifetime Key Manager (LKM) and SafeNet KeySecure for key management (SSKM) Key Vault Types.
  - **RSA Data Protection Manager (DPM)**: If an encryption group contains mixed firmware nodes, the Encryption Group Properties Key Vault Type name is based on the firmware version of the Group Leader. For example, If a switch is running Fabric OS 7.1.0 or later, the Key Vault Type is displayed as "RSA Data Protection Manager (DPM)."If a switch is running a Fabric OS version prior to v7.1.0, Key Vault Type is displayed as "RSA Key Manager (RKM)".
  - **HP Secure Key Manager (SKM):** The HP Key Vault Type name is shown as HP Secure Key Manager (SKM) for both SKM and **Enterprise Secure Key Management (**ESKM) Key Vault Types.
  - **Thales e-Security keyAuthority (TEKA)**: If an encryption group contains mixed firmware nodes, the Encryption Group Properties Key Vault Type name is based on the firmware version of the Group Leader. For example, If a switch is running Fabric OS 7.1.0 or later, the Key Vault Type is displayed as "Thales e-Security keyAuthority (TEKA)."If a switch is running a Fabric OS version prior to v7.1.0, Key Vault Type is displayed as "Thales Key Manager (TEMS)".
  - **Tivoli Key Lifecycle Manager (TKLM)**
  - **Key Management Interoperability Protocol (KMIP)**: Any KMIP-compliant server can be registered as a key vault on the Encryption switch after setting the key vault type to KMIP.

    If you are using a SafeNet KeySecure server, only SafeNet KeySecure for key management (SSKM) native hosting LKM is supported from the Management application; however, before selecting KMIP as the key vault type, all nodes in an encryption group must be running Fabric OS 7.1.0 or later.

If you are using a TEKA KMIP-compliant server, only Thales e-Security keyAuthority running version 4.0 is supported (from the CLI); however, all nodes in an encryption group must be running Fabric OS 7.2.0 or later. For more information about supported platforms and configuration instructions, refer to the *Fabric OS Encryption Administrator's Guide Supporting Key Management Interoperability Protocol (KMIP) Key-Compliant Environments*.

- **Primary Key Vault**: The primary key vault name, either an IPv4 address or Host name.
- **Primary Certificate File**: The name of a file containing the key vault's public key certificate. This file can be generated from the key vault's administrative console.
- **User Name**: The key vault user name. This field is active for ESKM/SKM and TEKA key vaults. For ESKM/SKM, it is needed only for the primary key vault. For TEKA, it is needed only for the secondary key vault.
- **Password**: The key vault password. This field is active for ESKM/SKM and TEKA key vaults. For ESKM/SKM, it is needed only for the primary key vault. For TEKA, it is needed for both the primary and secondary key vaults.
- **Re-type Password**: Re-enter the password for verification.
- **Backup Key Vault**: *(Optional.)* The secondary key vault, either an IPv4 address or Host name. The backup address can be left blank.
- **Backup Certificate File**: *(Optional.)* If a backup key vault is entered, the backup certificate file must also be entered. Navigate to and select the secondary public key certificate from your desktop, if applicable.
- **Serial Number:** (*TKLM only*.) Serial number of the switch, which is required for registering the switch on the key vault.
- **Device Group: (***TKLM only*.) The name of the device group of which the switch is a member. This information is required for registering the switch on the key vault.

8. Select the **Key Vault Type**. Configuration options vary based on the key vault type you choose.

Configuring key vault settings for RSA Data Protection Manager (DPM)

The following procedure assumes you have already configured the initial steps in the **Configure Switch Encryption** wizard. If you have not already done so, go to "Creating a new encryption group" on page 672.

Figure 301 shows the key vault selection dialog box for DPM.

FIGURE 301 Select Key Vault dialog box for DPM

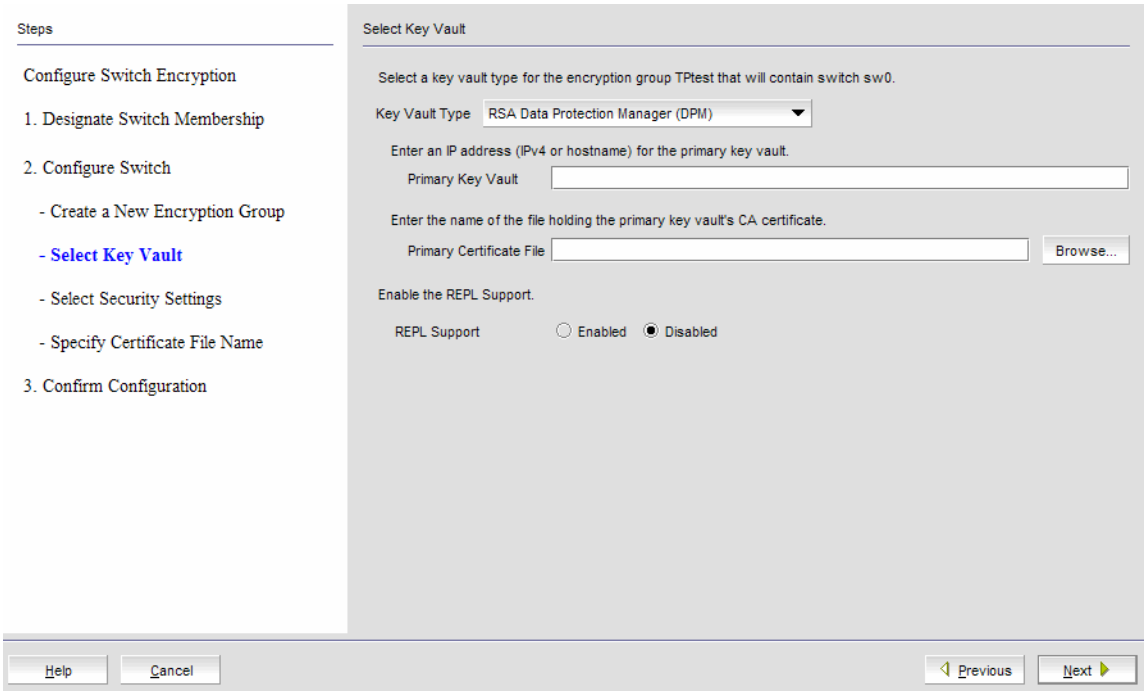

1. Enter the IP address or host name for the primary key vault. If you are clustering DPM appliances for high availability, IP load balancers are used to direct traffic to the appliances. Use the IP address of the load balancer.
2. Enter the name of the file that holds the Primary Key Vault's CA Key Certificate or browse to the desired location. This file can be generated from the key vault's administrative console.
3. If you are implementing encryption on data replication LUNs used by the EMC Symmetrix Remote Data Facility (SRDF), you must select **Enabled** for **REPL Support**.
4. Click **Next**.

   The **Specify Certificate Signing Request File Name** dialog box displays. (Refer to Figure 302.)

FIGURE 302 Specify Certificate Signing Request File Name dialog box

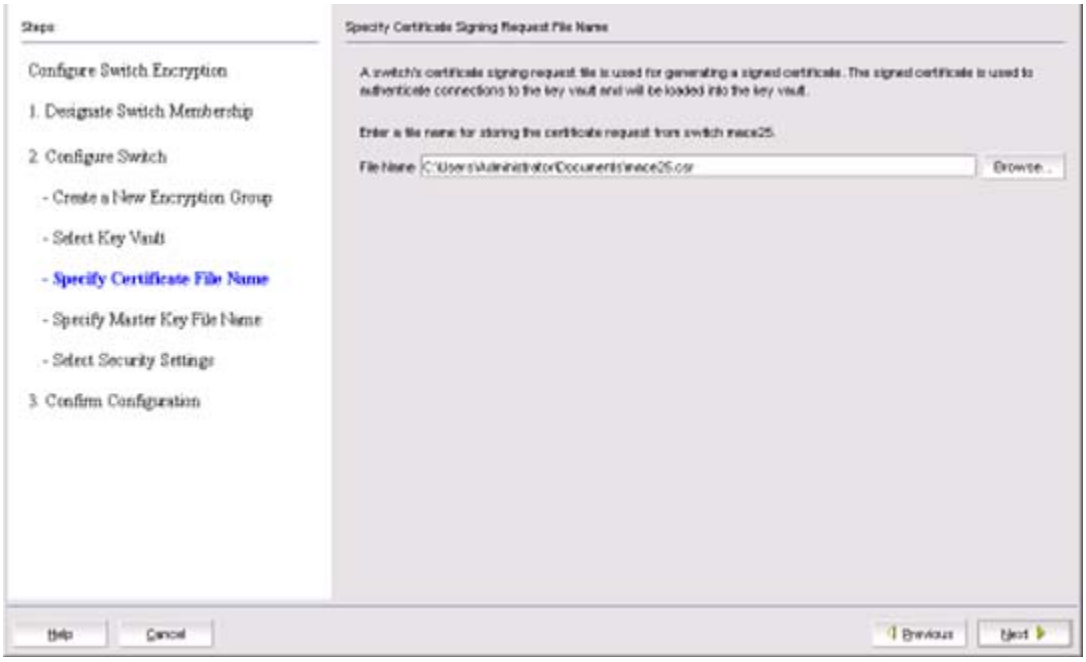

5. Enter the filename in which you want to store the certificate information, or browse to the file location.

   The certificate stored in this file is the switch's Switch Certificate Signing file. You will need to know this path and file name to install the switch's Switch Certificate Signing file on the key management appliance.

6. Click **Next**.

   The **Specify Master Key File Name dialog box** displays. (Refer to Figure 303.)

FIGURE 303 Specify Master Key File Name dialog box

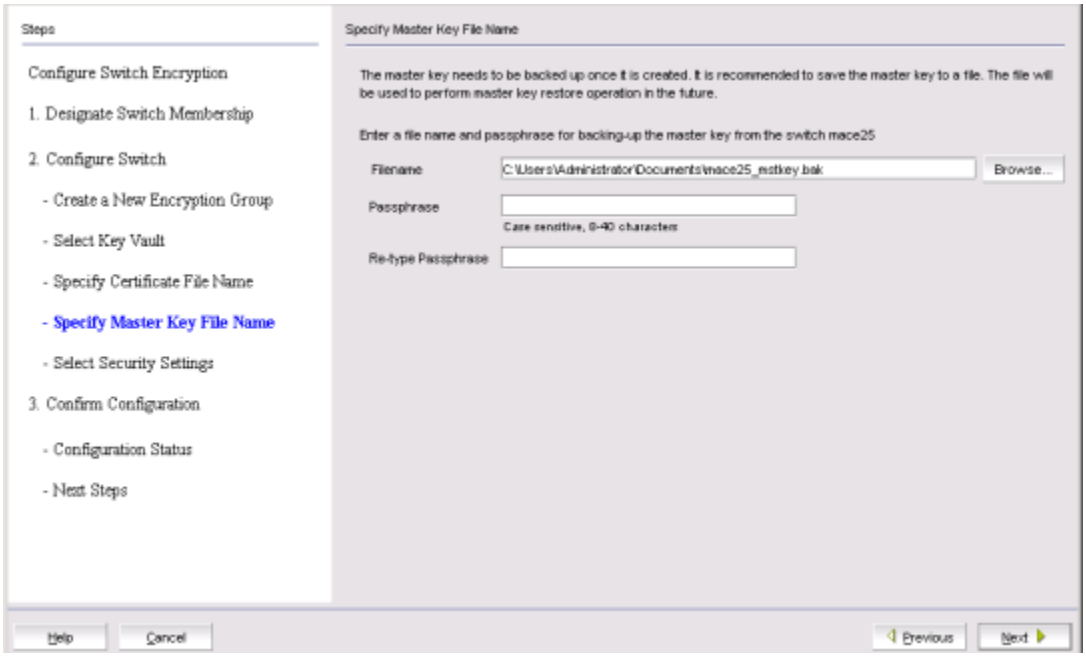

7. Enter the location of the file where you want to store back up master key information, or browse to the desired location.
8. Enter the passphrase, which is required for restoring the master key. The passphrase can be between eight and 40 characters, and any character is allowed.
9. Re-enter the passphrase for verification, then click **Next**.

   The **Select Security Settings** dialog box displays. (Refer to Figure 304.)

FIGURE 304 Select Security Settings dialog box

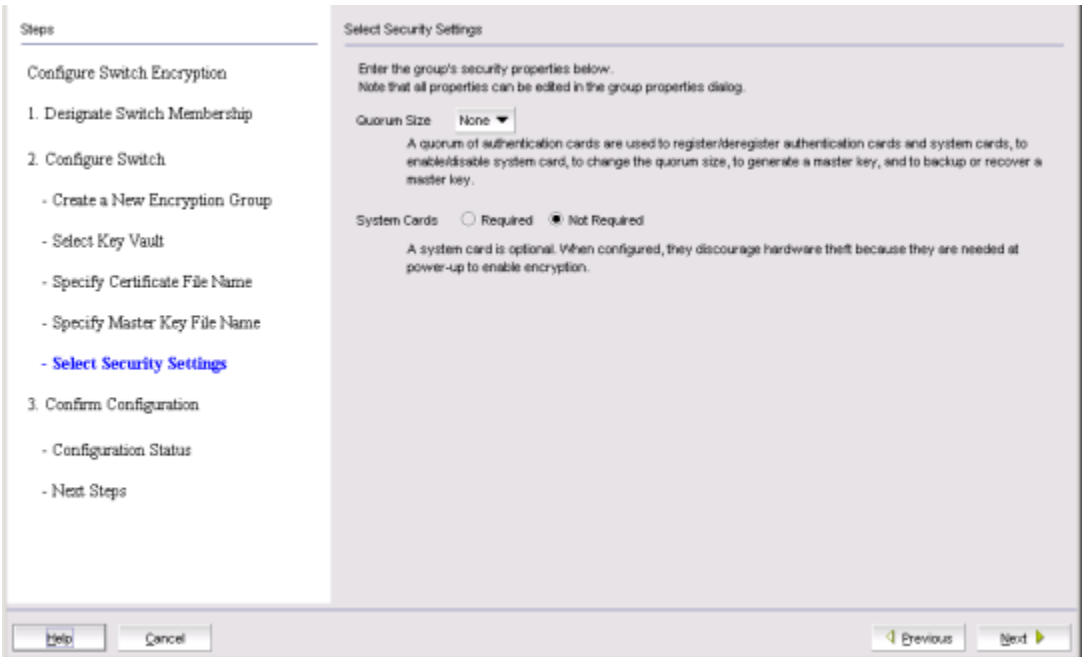

10. Set quorum size and system card requirements.

    The quorum size is the minimum number of cards necessary to enable the card holders to perform the security sensitive operations listed above. The maximum quorum size is five cards. The actual number of authentication cards registered is always more than the quorum size, so if you set the quorum size to five, for example, you will need to register at least six cards in the subsequent steps.

    Setting quorum size to a value greater than zero and/or setting system cards to **Required** launches additional wizard dialog boxes.

11. Click **Next**.

    The **Confirm Configuration** dialog box displays. (Refer to Figure 305.) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

FIGURE 305Confirm Configuration dialog box

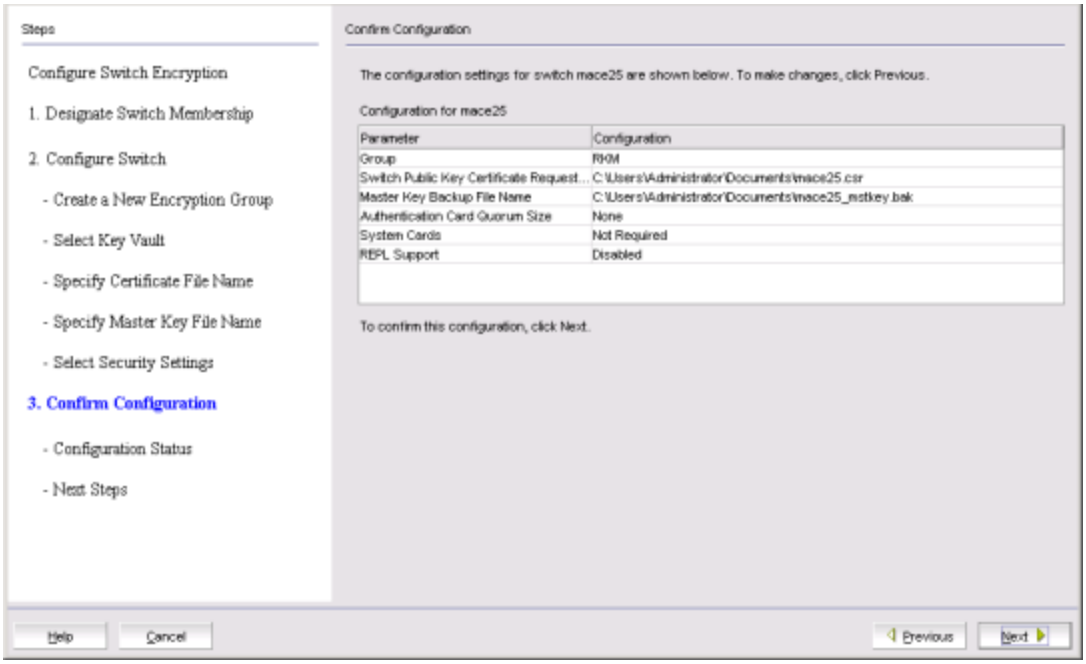

The **Configuration Status** dialog box displays. (Refer to Figure 306.)

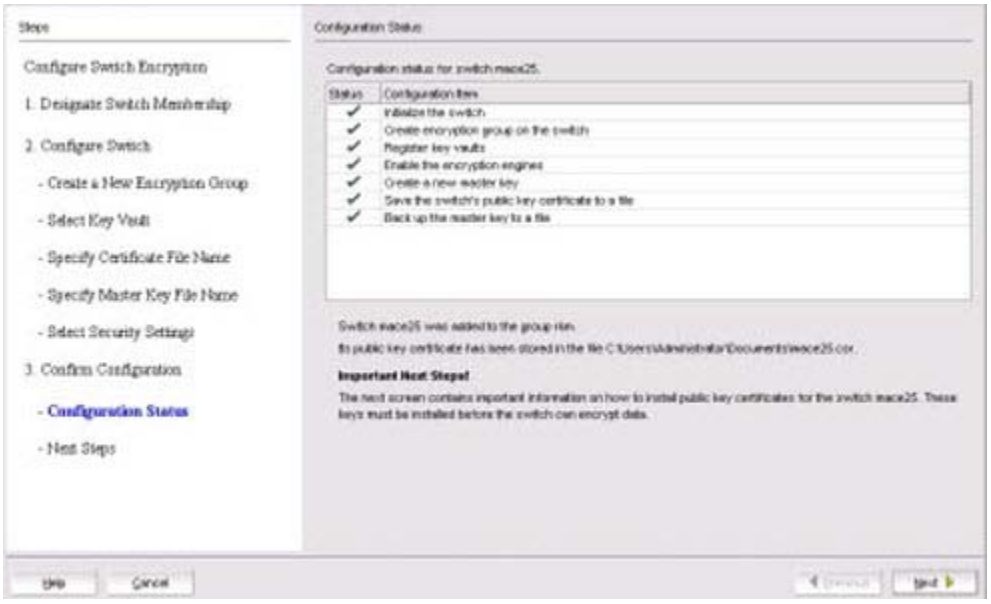

FIGURE 306Configuration Status dialog box

12. Review the post-configuration instructions, which you can copy to a clipboard or print for later, then click **Next**.

    The **Next Steps dialog box** displays. (Refer to Figure 307.) Instructions for installing public key certificates for the encryption switch are displayed.

FIGURE 307 Next Steps dialog box

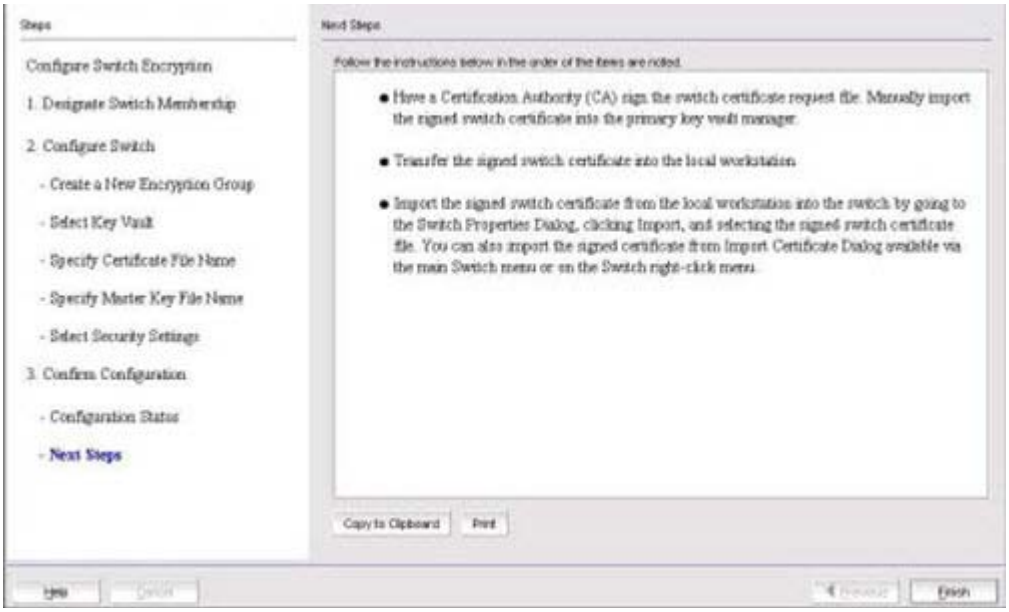

13. Review the post-configuration instructions, which you can copy to a clipboard or print for later, then click **Finish** to exit the wizard.

Configuring key vault settings for NetApp Link Key Manager (LKM/SSKM)

The following procedure assumes you have already configured the initial steps in the **Configure Switch Encryption** wizard. If you have not already done so, go to "Creating a new encryption group" on page 672.

Figure 308 shows the key vault selection dialog box for LKM/SSKM.

FIGURE 308 Select Key Vault dialog box for LKM/SSKM

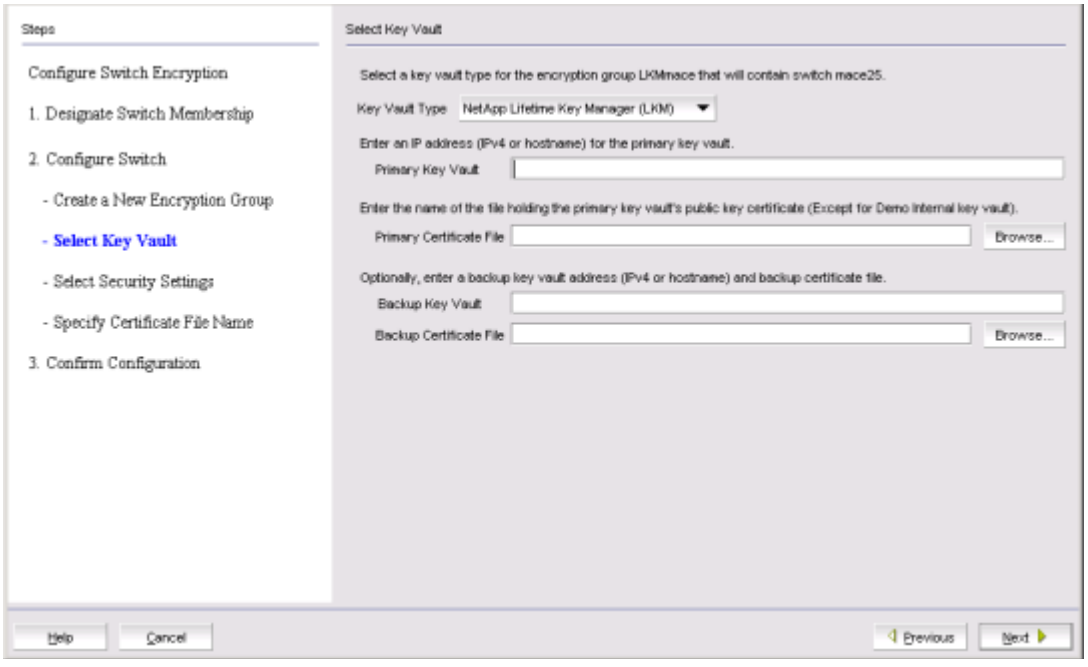

1. Enter the IP address or host name for the primary key vault.
2. Enter the name of the file that holds the primary key vault's public key certificate, or browse to the desired location.
3. If you are using a backup key vault, enter the IP address or host name, and the name of the file holding the backup key vault's public key certificate, then click **Next**.

   The **Specify Public Key Certificate (KAC) File Name** dialog box displays. (Refer to Figure 309.)

FIGURE 309 Specify Public Key Certificate (KAC) File Name dialog box

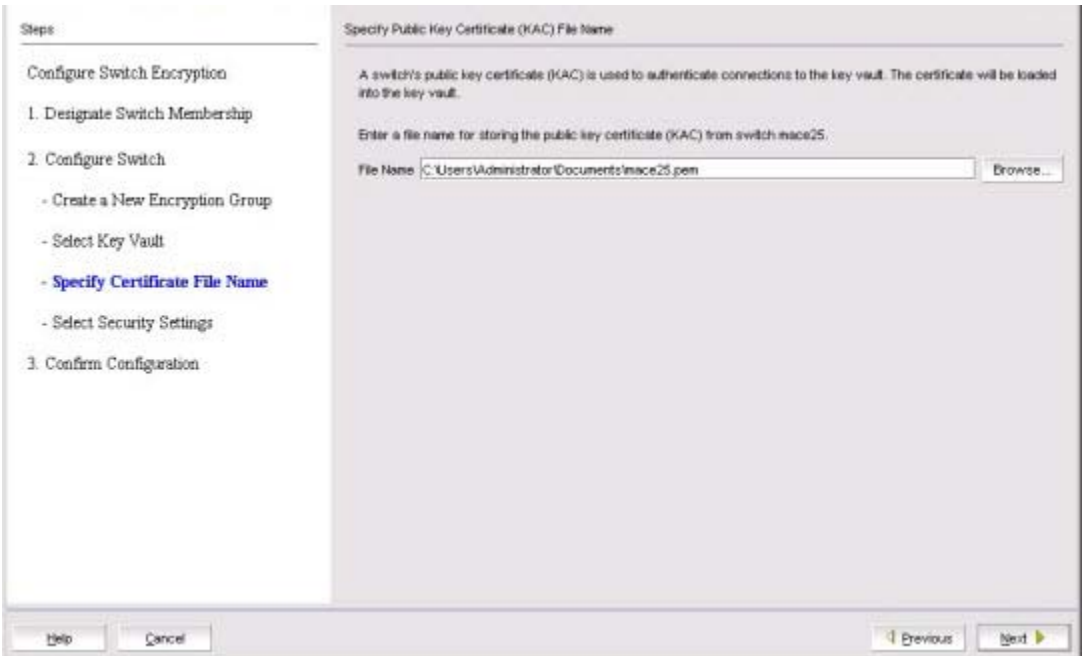

4. Specify the location of the file where you want to store the public key certificate that is used to authenticate connections to the key vault.

   The certificate stored in this file is the switch's public key certificate. You will need to know this path and file name to install the switch's public key certificate on the key management appliance.

5. Click **Next**.

   The **Select Security Settings** dialog box displays. (Refer to Figure 310.)

FIGURE 310Select Security Settings dialog box

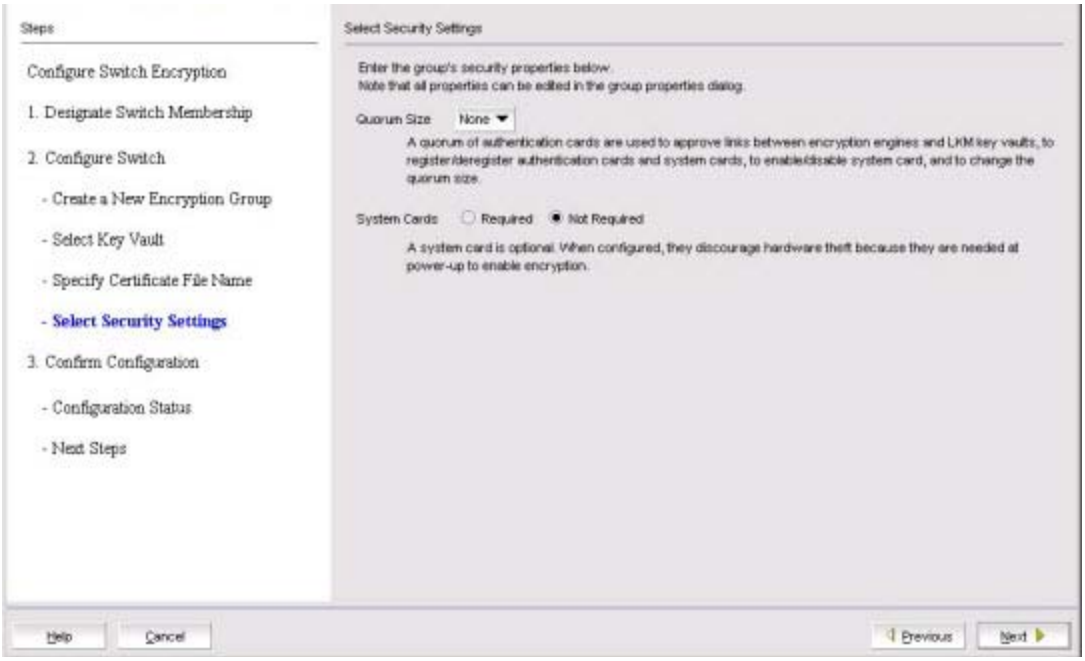

6. Set quorum size and system card requirements.

   The quorum size is the minimum number of cards necessary to enable the card holders to perform the security sensitive operations listed above. The maximum quorum size is five cards. The actual number of authentication cards registered is always more than the quorum size, so if you set the quorum size to five, for example, you will need to register at least six cards in the subsequent steps.

   Setting quorum size to a value greater than zero and/or setting system cards to **Required** launches additional wizard dialog boxes.

7. Click **Next**.

   The **Confirm Configuration** dialog box displays. (Refer to Figure 311.) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

FIGURE 311 Confirm Configuration dialog box

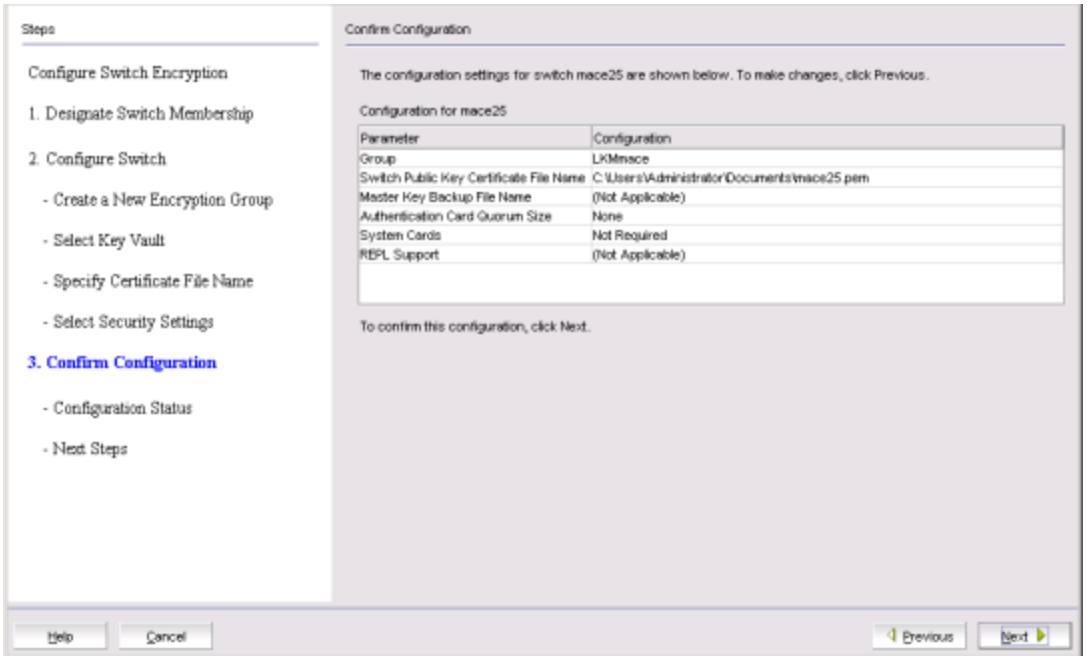

The **Configuration Status** dialog box displays. (Refer to Figure 312.)

FIGURE 312 Configuration Status dialog box

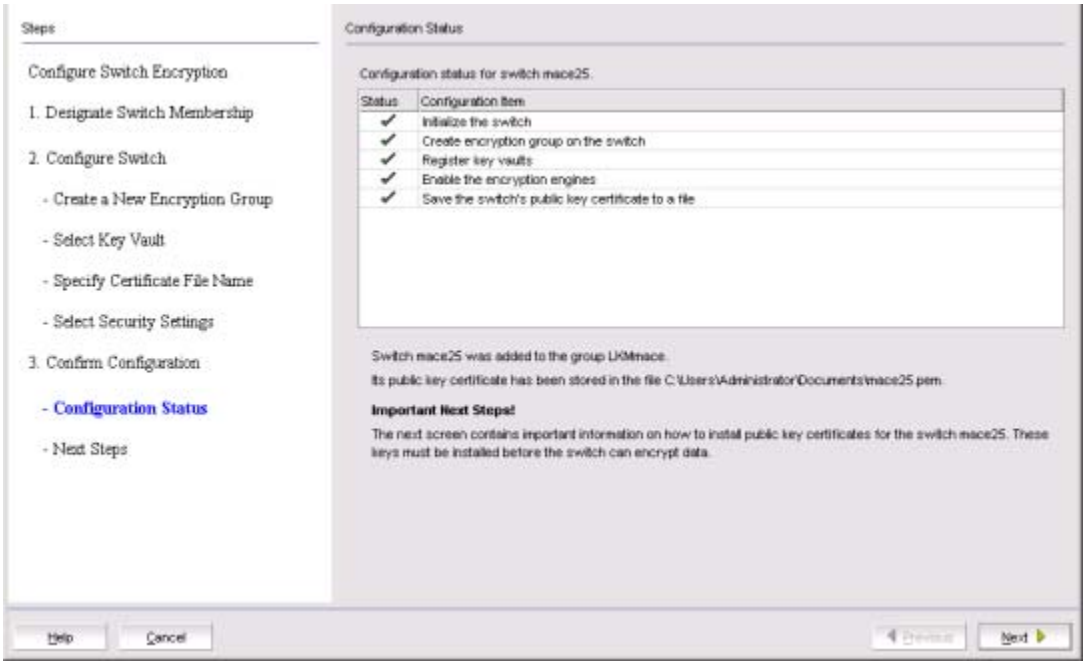

All configuration items have green check marks if the configuration is successful. A red stop sign indicates a failed step. A message displays below the table, indicating the encryption switch was added to the group you named, and the public key certificate is stored in the location you specified.

After configuration of the encryption group is completed, the Management application sends API commands to verify the switch configuration. See "Understanding configuration status results" on page 706 for more information.

8. Verify the information is correct, then click **Next**.

   The **Next Steps** dialog box displays. (Refer to Figure 313.) Instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

FIGURE 313 Next Steps dialog box

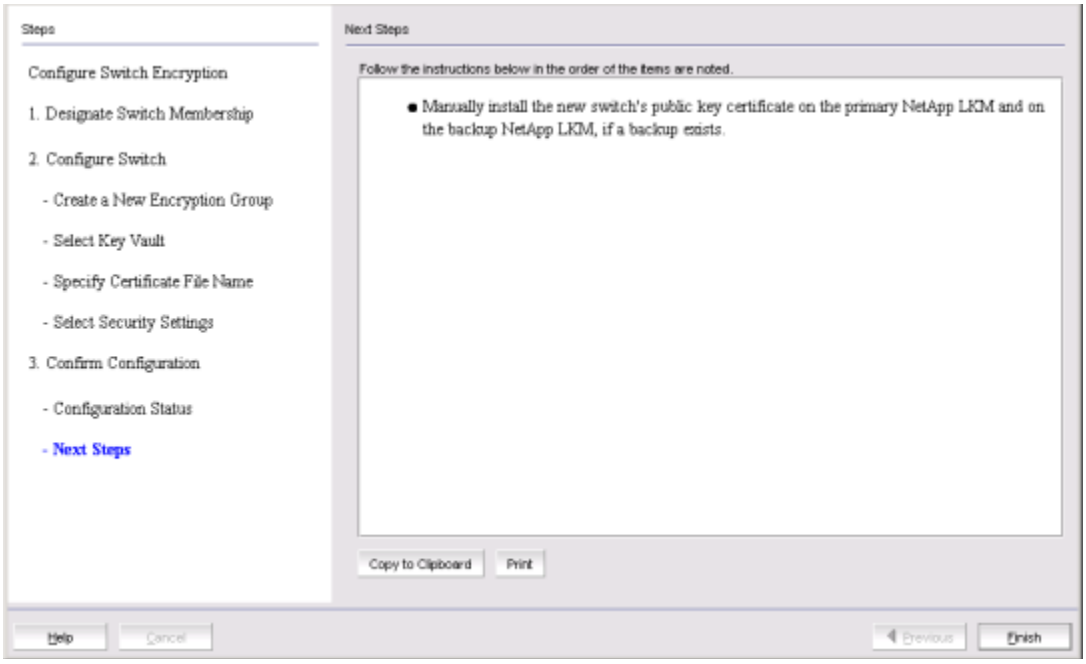

9. Review the post-configuration instructions, which you can copy to a clipboard or print for later, then click **Finish** to exit the **Configure Switch Encryption** wizard.
10. Refer to "Understanding configuration status results" on page 706.

Configuring key vault settings for HP Enterprise Secure Key Manager (ESKM/SKM)

The following procedure assumes you have already configured the initial steps in the **Configure Switch Encryption** wizard. If you have not already done so, go to "Creating a new encryption group" on page 672.

Figure 314 shows the key vault selection dialog box for ESKM/SKM.

FIGURE 314 Select Key Vault dialog box for ESKM/SKM

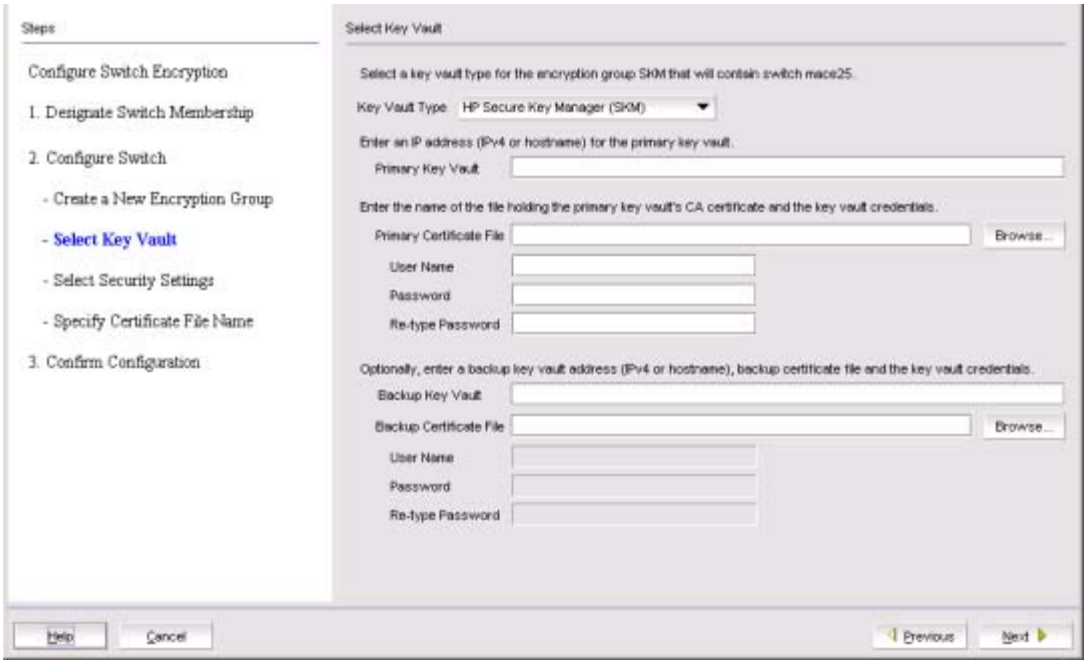

1. Enter the IP address or host name for the primary key vault.
2. Enter the name of the file that holds the primary key vault's CA key certificate, or browse to the desired location.
3. Enter the password you established for the Brocade user group.
4. If you are using a backup key vault, enter the IP address or host name, and the name of the file holding the backup key vault's public key certificate in the fields provided. The same user name and password used for the primary key vault are automatically applied to the backup key vault.
5. Click **Next**.

   The **Specify Certificate Signing Request File Name** dialog box displays. (Refer to Figure 315.)

FIGURE 315Specify Certificate Signing Request File Name dialog box

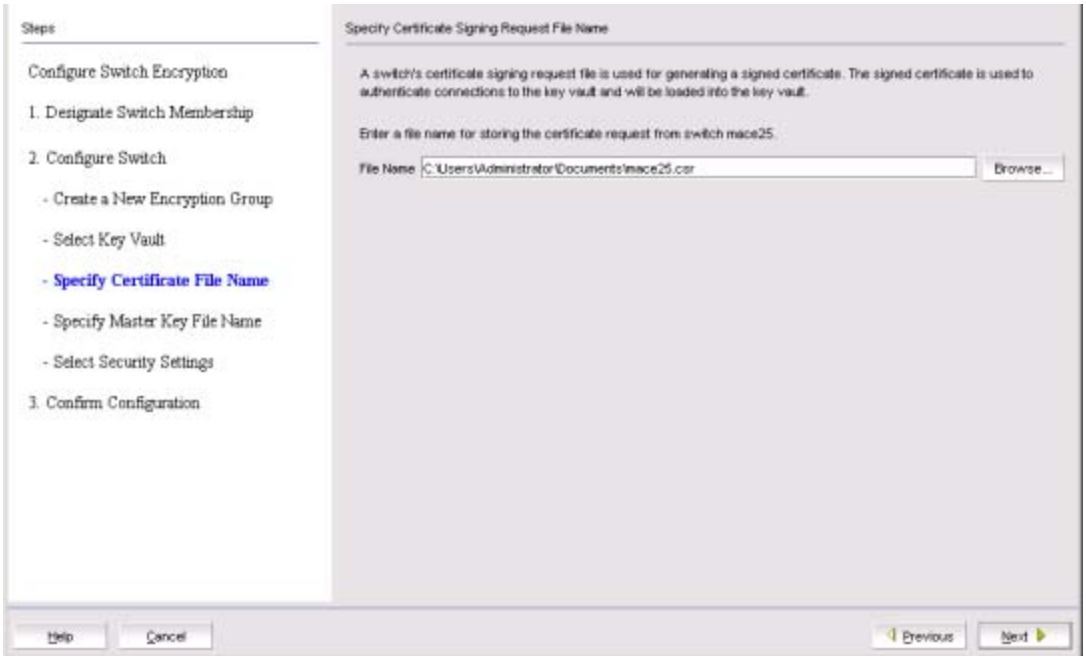

6. Enter the location of the file where you want to store the certificate information, or browse to the desired location, then click **Next**.

   The **Specify Master Key File Name** dialog box displays. (Refer to Figure 316.)

FIGURE 316Specify Master Key File Name dialog box

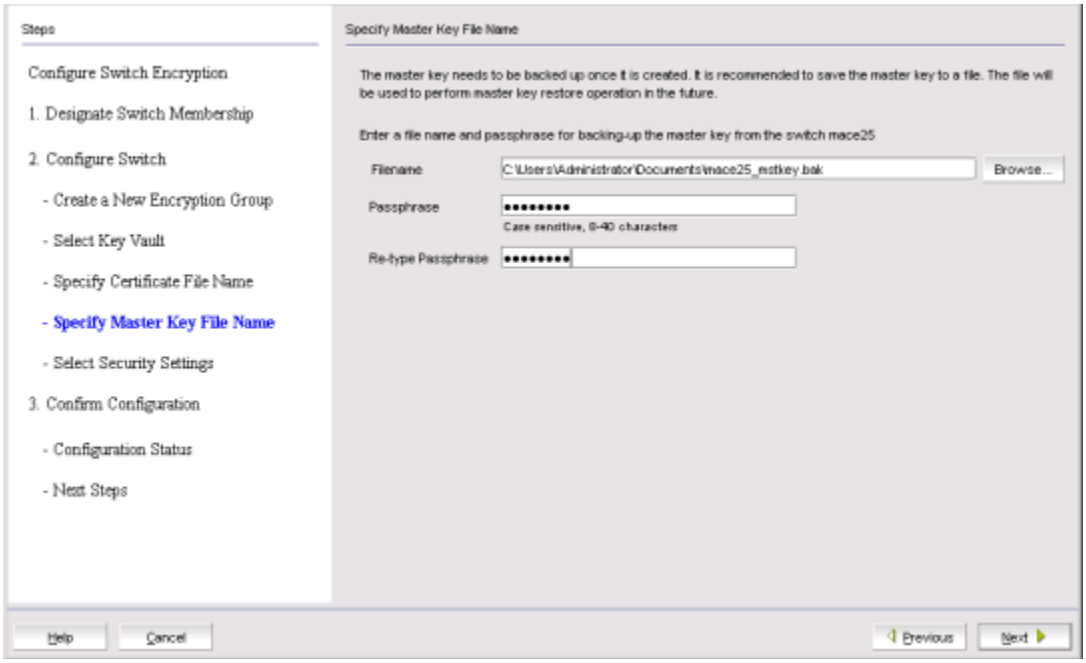

7. Enter the passphrase, which is required for restoring the master key. The passphrase can be between eight and 40 characters, and any character is allowed.
8. Re-enter the passphrase for verification, then click **Next**.

The **Select Security Settings** dialog box displays. (Refer to Figure 317.)

FIGURE 317 Select Security Settings dialog box

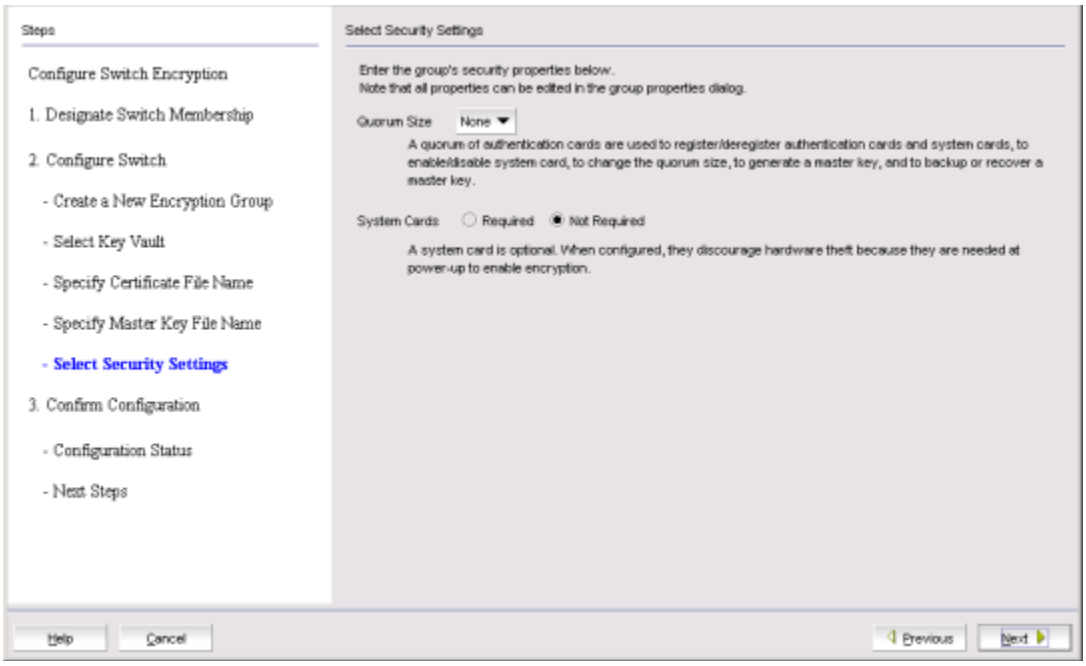

9. Set quorum size and system card requirements.

   The quorum size is the minimum number of cards necessary to enable the card holders to perform the security sensitive operations listed above. The maximum quorum size is five cards. The actual number of authentication cards registered is always more than the quorum size, so if you set the quorum size to five, for example, you will need to register at least six cards in the subsequent steps.

   Setting quorum size to a value greater than zero and/or setting system cards to **Required** launches additional wizard dialog boxes.

10. Click **Next**.

    The **Confirm Configuration** dialog box displays. (Refer to Figure 318.) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

FIGURE 318 Confirm Configuration dialog box

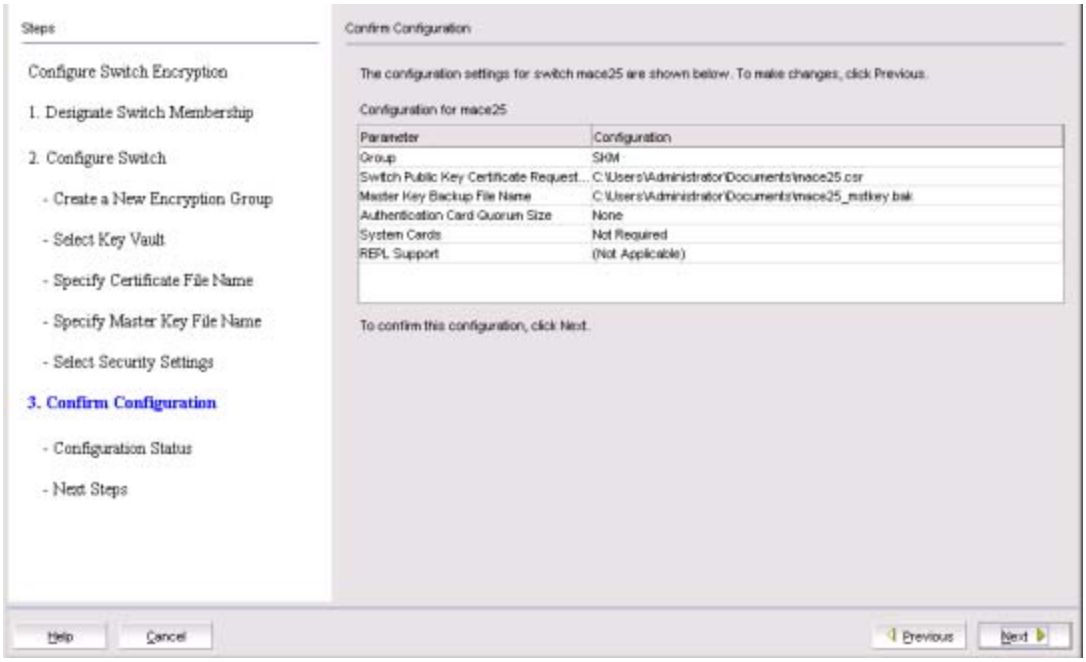

The **Configuration Status** dialog box displays. (Refer to Figure 319.)

FIGURE 319 Configuration Status dialog box

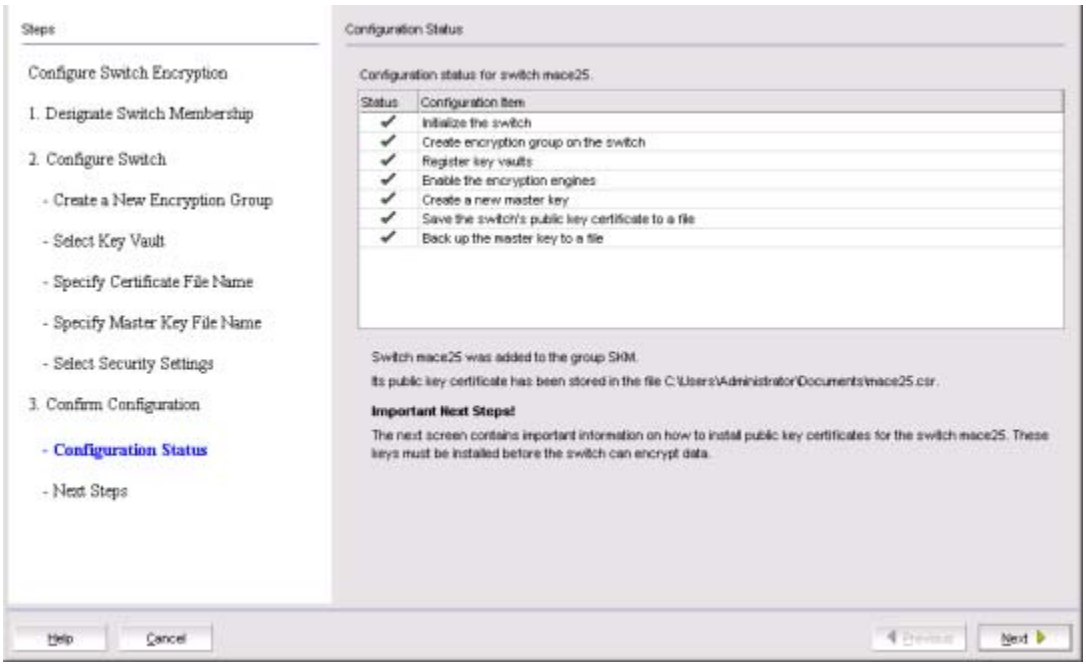

All configuration items have green check marks if the configuration is successful. A red stop sign indicates a failed step. A message displays below the table, indicating the encryption switch was added to the group you named, and the public key certificate is stored in the location you specified.

After configuration of the encryption group is completed, the Management application sends API commands to verify the switch configuration. See "Understanding configuration status results" on page 706 for more information.

11. Review important messages, then click **Next**.

    The **Next Steps** dialog box displays. (Refer to Figure 320.) Instructions for installing public key certificates for the encryption switch are displayed.

FIGURE 320 Next Steps dialog box

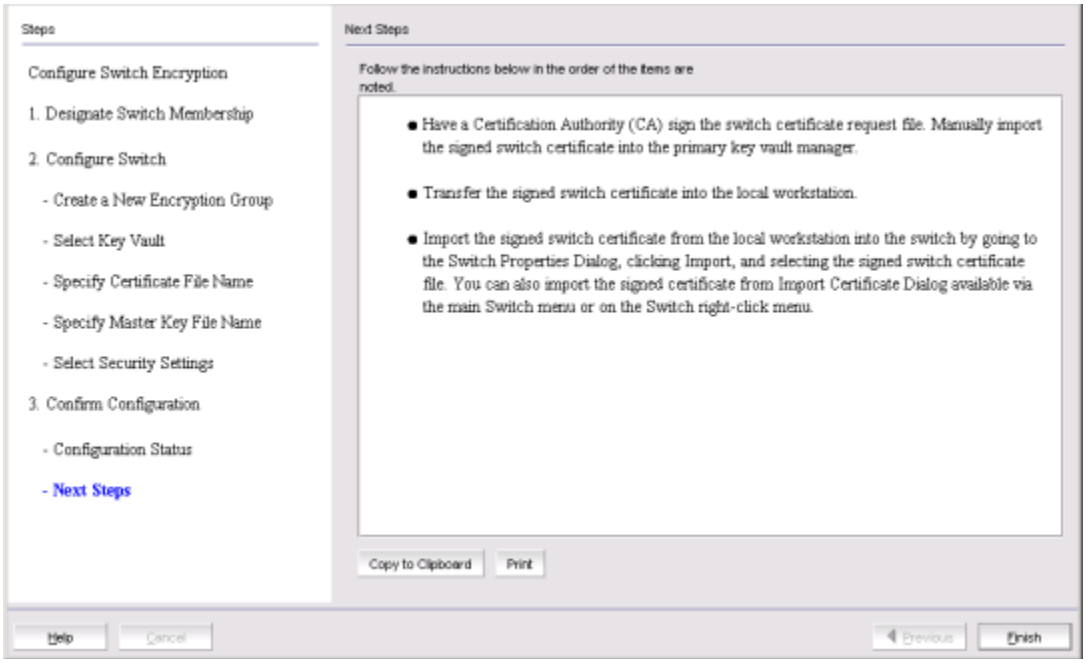

12. Review post-configuration instructions, which you can copy to a clipboard or print for later.
13. Click **Finish** to exit the **Configure Switch Encryption** wizard.
14. Refer to "Understanding configuration status results" on page 706.

Configuring key vault settings for Thales e_Security keyAuthority (TEKA)

The following procedure assumes you have already configured the initial steps in the **Configure Switch Encryption** wizard. If you have not already done so, go to "Creating a new encryption group" on page 672.

Figure 321 shows the key vault selection dialog box for TEKA.

FIGURE 321 Select Key Vault dialog box for TEKA

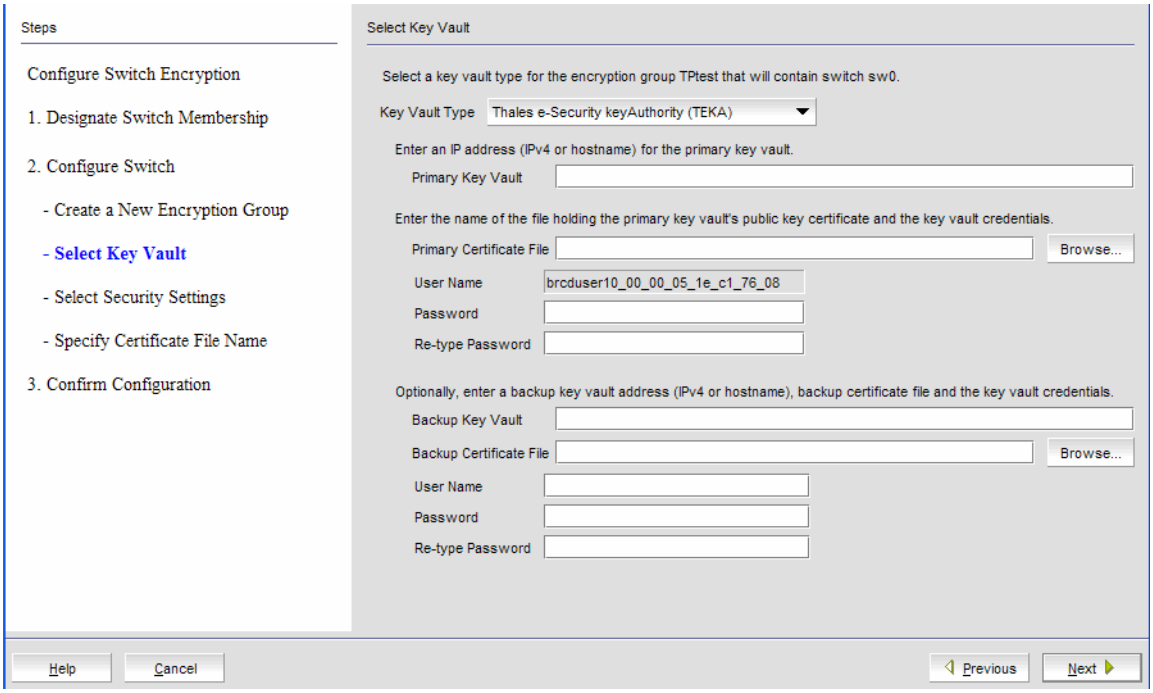

1. Enter the IP address or host name for the primary key vault.
2. Enter the name of the file that holds the primary key vault's public key certificate, or browse to the desired location.
3. Enter the password you created for the Brocade group TEKA client.
4. If you are using a backup key vault, enter the IP address or host name, the name of the file holding the backup key vault's public key certificate in the fields provided, and the user name and password for the backup key vault.
5. Click **Next**.

   The **Specify Master Key File Name** dialog box displays. (Refer to Figure 322.)

FIGURE 322Specify Master Key File Name dialog box

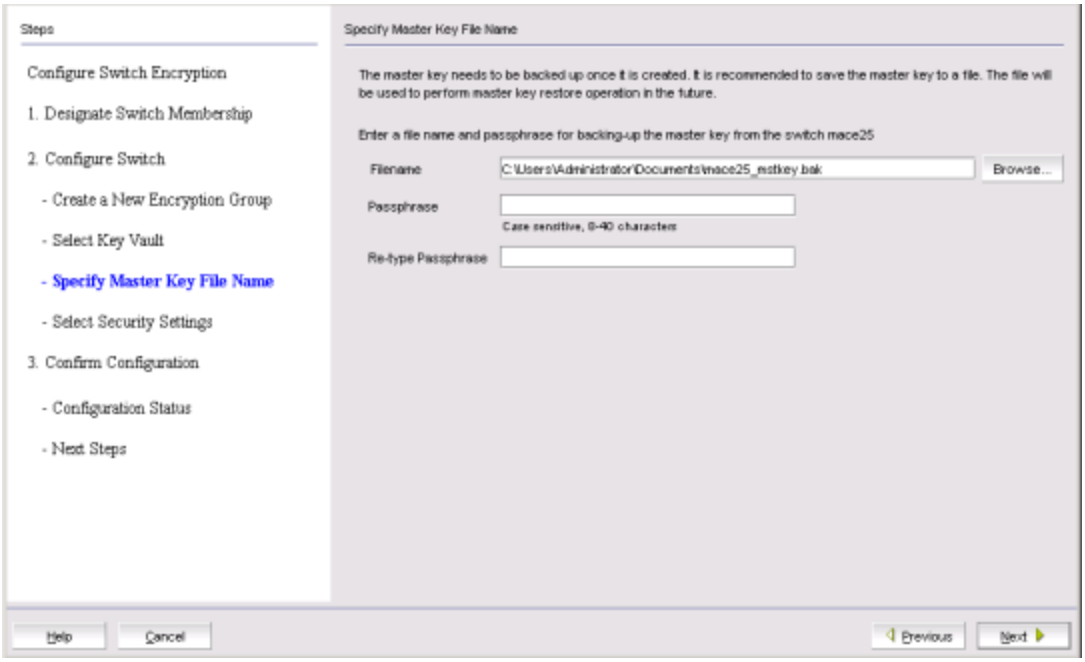

6. Enter the name of the file used for backing up the master key or browse to the desired location.
7. Enter the passphrase, which is required for restoring the master key. The passphrase can be between eight and 40 characters, and any character is allowed.
8. Re-enter the passphrase for verification, then click **Next**.

   The **Select Security Settings** dialog box displays. (Refer to Figure 323.)

FIGURE 323Select Security Settings dialog box

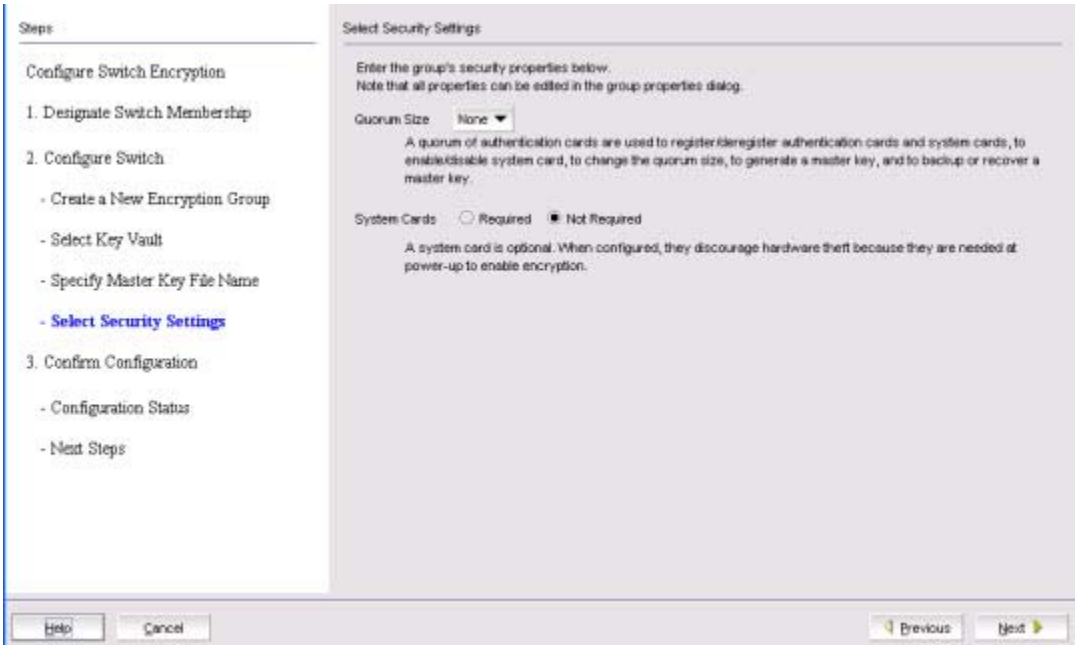

9. Set quorum size and system card requirements.

   The quorum size is the minimum number of cards necessary to enable the card holders to perform the security sensitive operations listed above. The maximum quorum size is five cards. The actual number of authentication cards registered is always more than the quorum size, so if you set the quorum size to five, for example, you will need to register at least six cards in the subsequent steps.

   Setting quorum size to a value greater than zero and/or setting system cards to **Required** launches additional wizard dialog boxes.

10. Click **Next**.

    The **Confirm Configuration** dialog box displays. (Refer to Figure 324.) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

FIGURE 324 Confirm Configuration dialog box

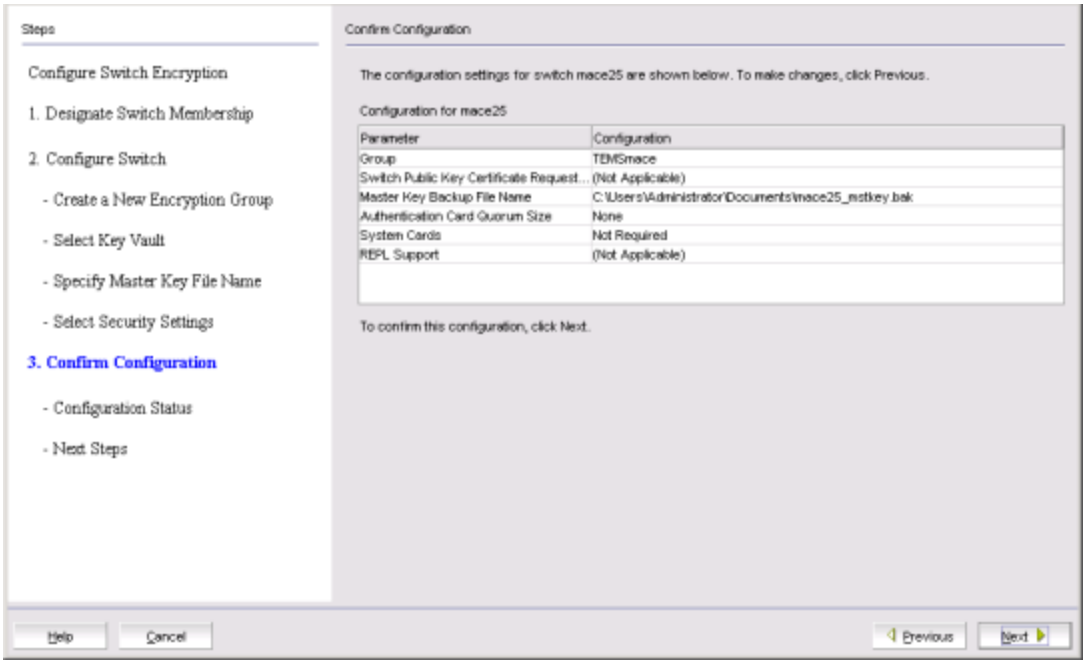

The **Configuration Status** dialog box displays. (Refer to Figure 325.)

FIGURE 325 Configuration Status dialog box

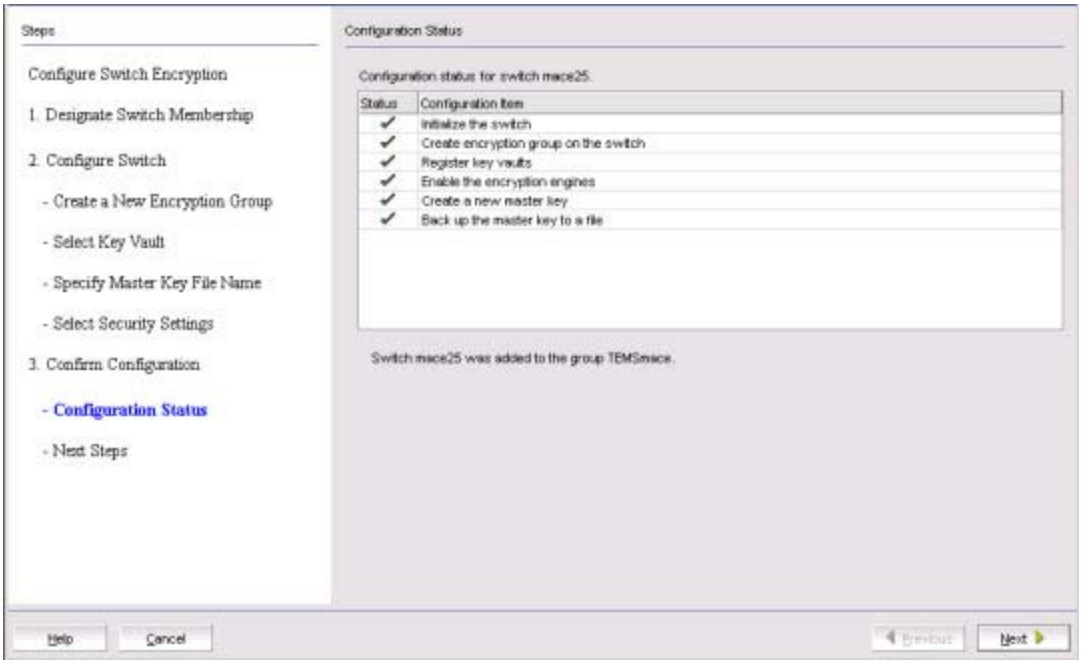

All configuration items have green check marks if the configuration is successful. A red stop sign indicates a failed step. A message displays below the table, indicating the encryption switch was added to the group you named, and the public key certificate is stored in the location you specified.

After configuration of the encryption group is completed, the Management application sends API commands to verify the switch configuration. See "Understanding configuration status results" on page 706 for more information.

11. Click **Next**.

    The **Next Steps** dialog box displays. (Refer to Figure 326.) Instructions for installing public key certificates for the encryption switch are displayed.

FIGURE 326Next Steps dialog box

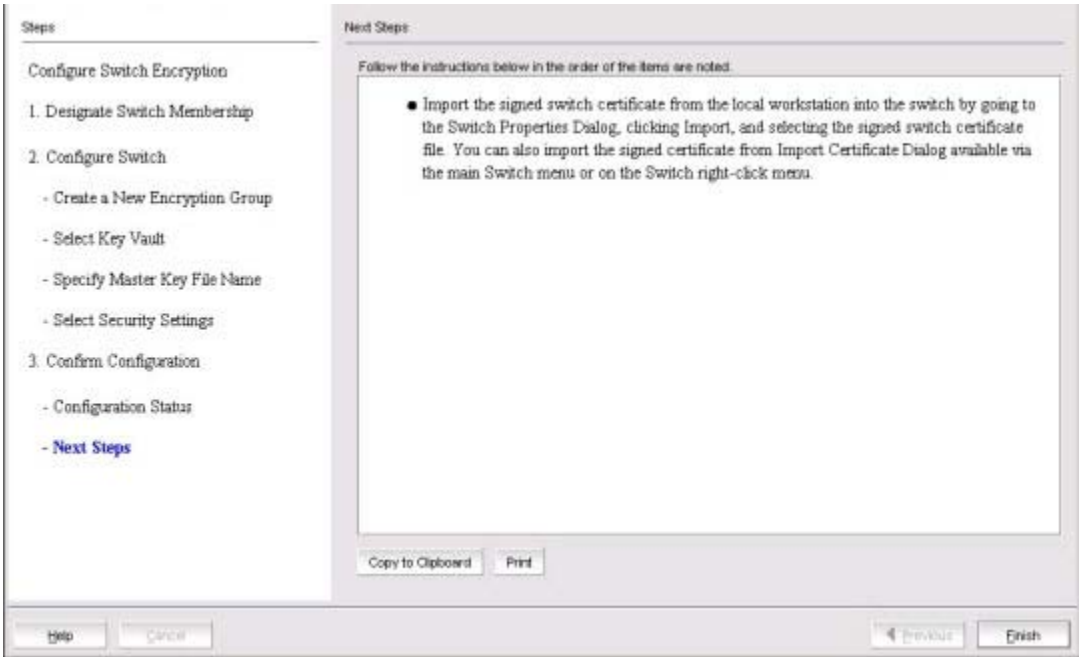

12. Review the post-configuration instructions, which you can copy to a clipboard or print for later.
13. Click **Finish** to exit the **Configure Switch Encryption** wizard.
14. Refer to "Understanding configuration status results" on page 706.

Configuring key vault settings for IBM Tivoli Key Lifetime Manager (TKLM)

The following procedure assumes you have already configured the initial steps in the **Configure Switch Encryption** wizard. If you have not already done so, go to "Creating a new encryption group" on page 672.

Figure 327 shows the key vault selection dialog box for TKLM.

FIGURE 327 Select Key Vault dialog box for TKLM

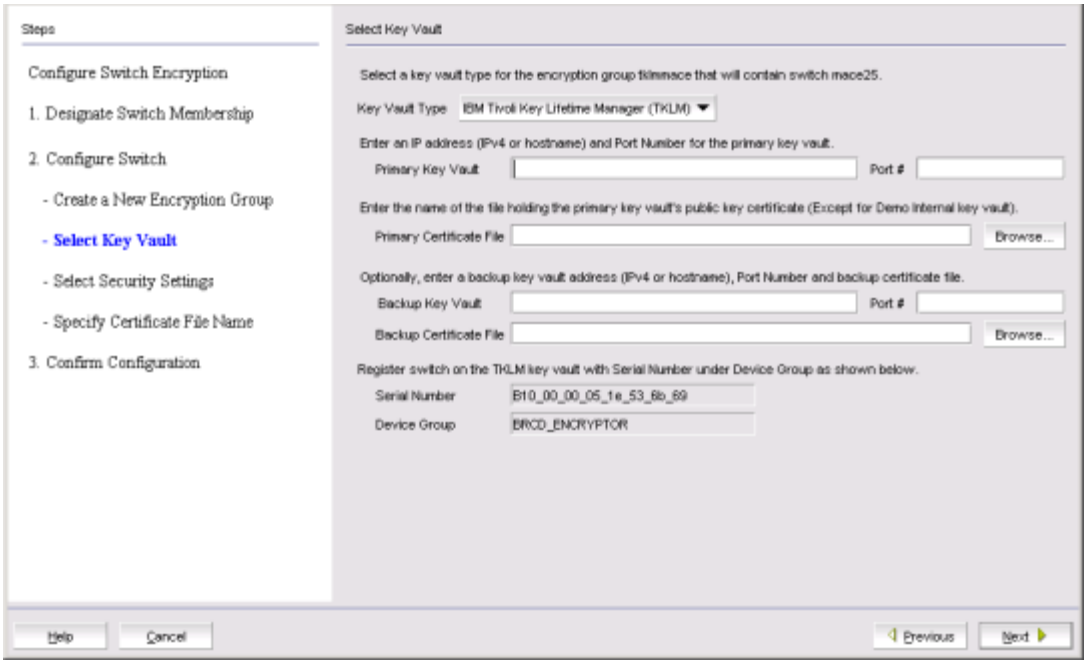

1. Enter the IP address or host name for the primary key vault.
2. Enter the name of the file that holds the primary key vault's public key certificate or browse to the desired location.
3. If you are using a backup key vault, enter the IP address or host name, and the name of the file holding the backup key vault's public key certificate in the fields provided.
4. Click **Next**.

   The **Specify Master Key Certificate File Name** dialog box displays. (Refer to Figure 329.)

FIGURE 328Specify Public Key Certificate (KAC) File Name dialog box

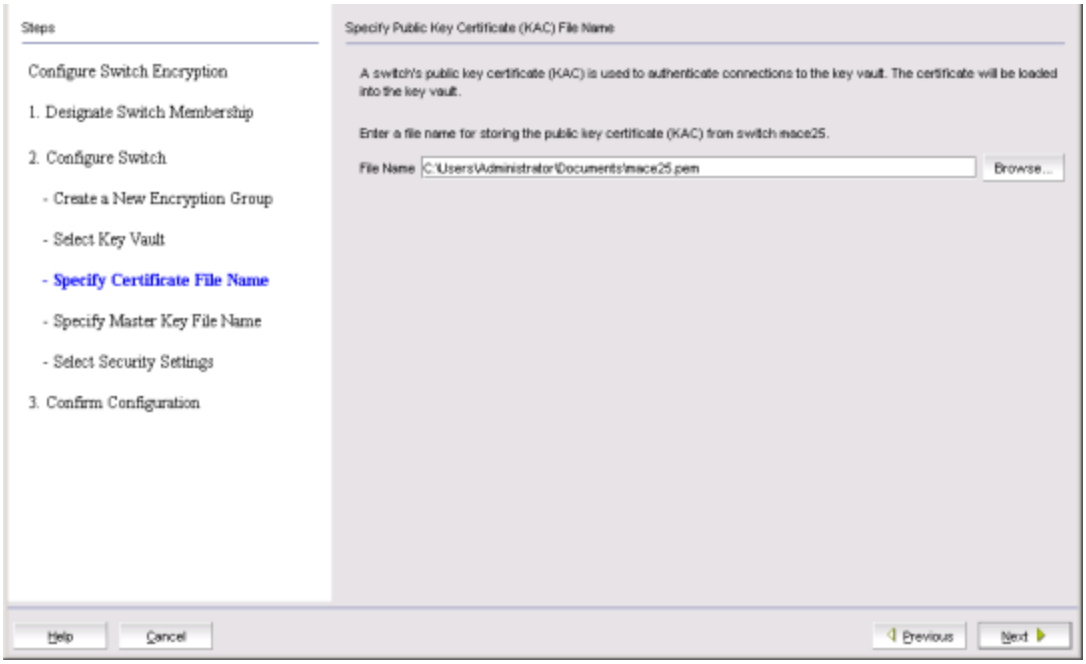

5. Enter the name of the file where the switch's public key certificate is stored, or browse to the desired location, then click **Next**.

   The **Specify Master Key File Name** dialog box displays. (Refer to Figure 329.)

FIGURE 329Specify Master Key File Name dialog box

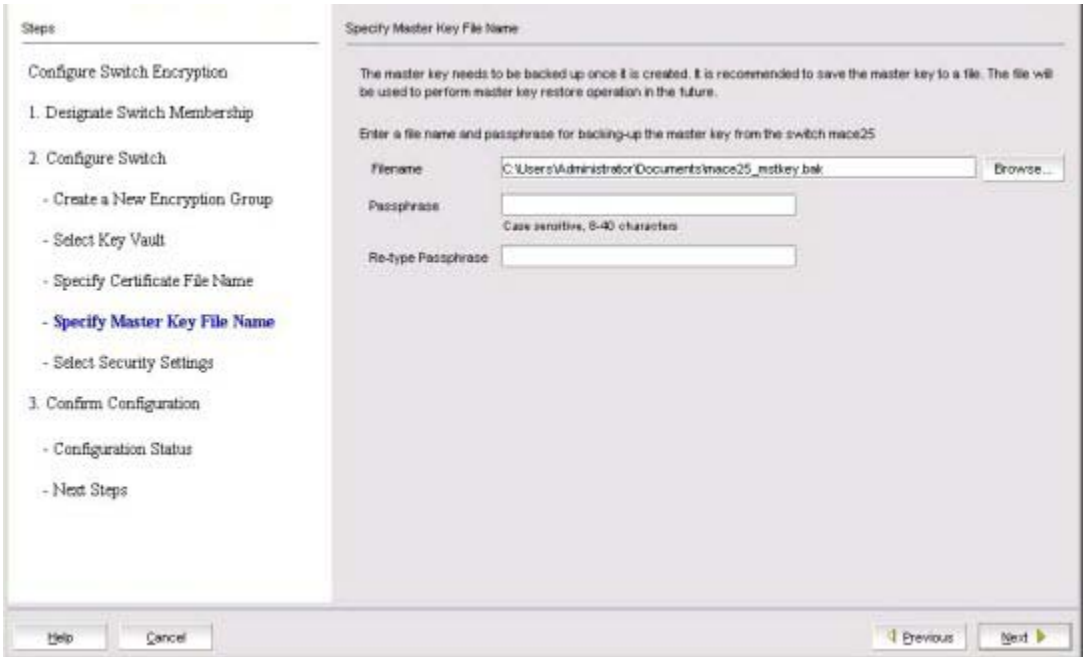

6. Enter the name of the file used for backing up the master key, or browse to the desired location.
7. Enter the passphrase, which is required for restoring the master key. The passphrase can be between eight and 40 characters, and any character is allowed.

8. Re-enter the passphrase for verification, then click **Next**.

   The **Select Security Settings** dialog box displays. (Refer to Figure 330.)

FIGURE 330 Select Security Settings dialog box

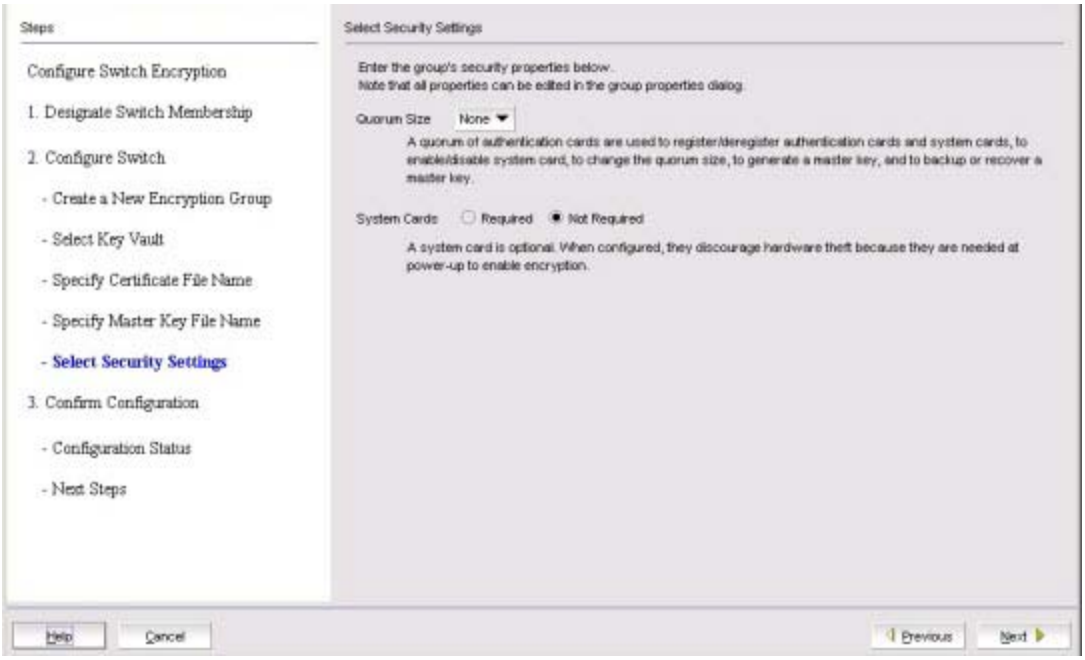

9. Set quorum size and system card requirements.

   The quorum size is the minimum number of cards necessary to enable the card holders to perform the security sensitive operations listed above. The maximum quorum size is five cards. The actual number of authentication cards registered is always more than the quorum size, so if you set the quorum size to five, for example, you will need to register at least six cards in the subsequent steps.

   Setting quorum size to a value greater than zero and/or setting system cards to **Required** launches additional wizard dialog boxes.

10. Click **Next**.

    The **Confirm Configuration** dialog box displays. (Refer to Figure 331.) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

FIGURE 331 Confirm Configuration dialog box

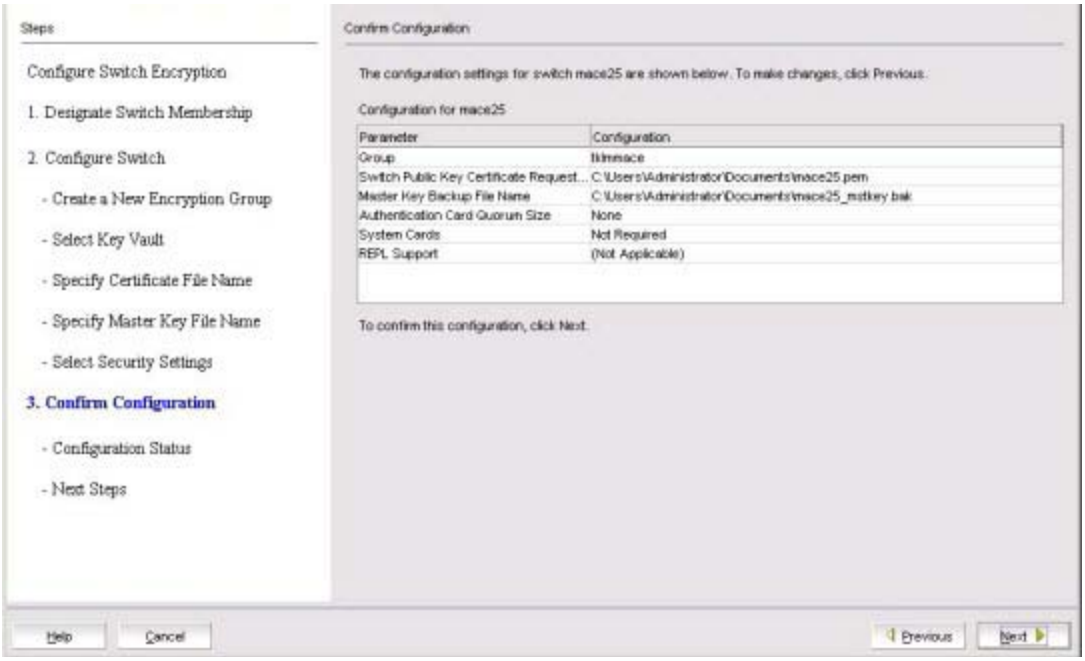

The **Configuration Status** dialog box displays. (Refer to Figure 332.)

FIGURE 332 Configuration Status dialog box

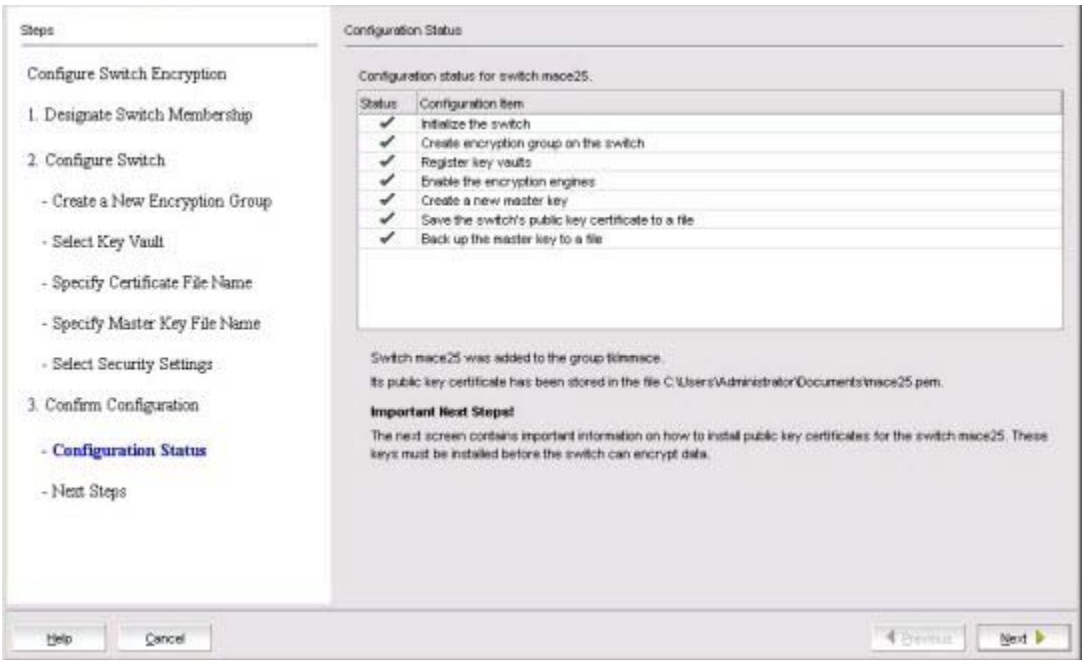

All configuration items have green check marks if the configuration is successful. A red stop sign indicates a failed step. A message displays below the table, indicating the encryption switch was added to the group you named, and the public key certificate is stored in the location you specified.

After configuration of the encryption group is completed, the Management application sends API commands to verify the switch configuration.

11. Click **Next**.

    The **Next Steps** dialog box displays. (Refer to Figure 333.) Instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

FIGURE 333 Next Steps dialog box

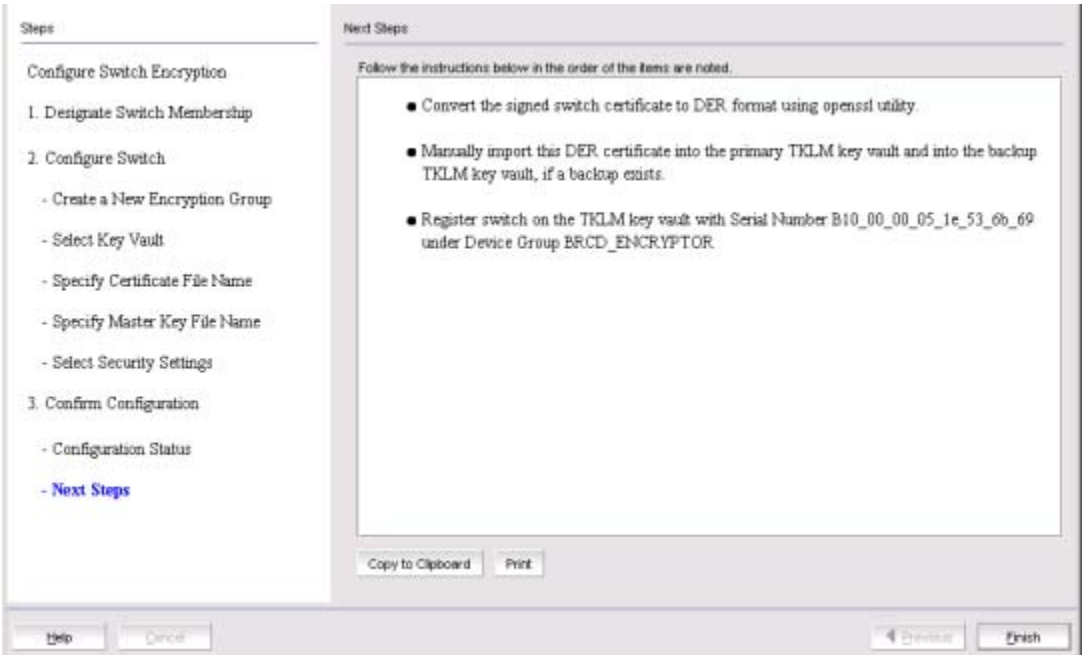

12. Review the post-configuration instructions, which you can copy to a clipboard or print for later.
13. Click **Finish** to exit the **Configure Switch Encryption** wizard.
14. Refer to "Understanding configuration status results" on page 706.

Configuring key vault settings for Key Management Interoperability Protocol

The following procedure assumes you have already configured the initial steps in the **Configure Switch Encryption** wizard. If you have not already done so, go to "Creating a new encryption group" on page 672.

NOTE:

- With the introduction of Fabric OS 7.1.0, KMIP with SafeNet KeySecure for key management (SSKM) native hosting LKM is supported. Before selecting KMIP as the key vault type, all nodes in a KeySecure encryption group must be running Fabric OS 7.1.0 or later.
- With the introduction of Fabric OS 7.2.0, KMIP with TEKA 4.0 is also supported, but must be configured using the CLI. All nodes in a keyAuthority encryption group must be running Fabric OS 7.2.0 or later. For configuration instructions, refer to the *Fabric OS Encryption Administrator's Guide Supporting Key Management Interoperability Protocol (KMIP) Key-Compliant Environments*.

Figure 334 shows the key vault selection dialog box for KMIP.

FIGURE 334 Select Key Vault dialog box for KMIP

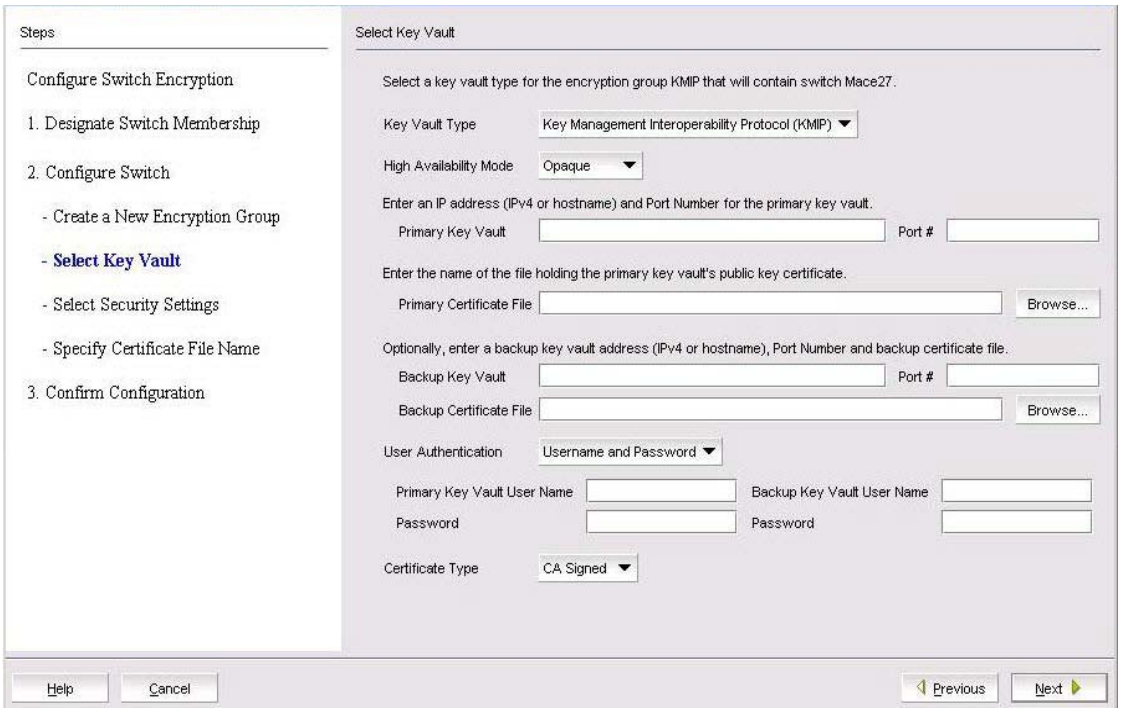

1. Select the High Availability mode. Options are:
    - **Opaque:** Both the primary and secondary key vaults are registered on the BES. The client archives the key to a single (primary) key vault. For disk operations, an additional hardening check is done on the secondary key vault before the key is used for encryption.
    - **Transparent:** A single key vault should be registered on the BES. The client assumes the entire HA is implemented on the key vault. Key archival and retrieval is done to the KMIP without any additional hardening checks.
    - **No HA:** Both the primary and secondary key vaults are registered on the BES. The client archives keys to both key vaults and ensures that the archival is successful before the key is used for encryption.
2. Enter the **Primary Key Vault** IP address or hostname, and port number.
3. Enter the **Primary Certificate** file name, or browse to the file location.
4. (Optional) Enter a **Backup Key Vault** IP address or hostname, and port number, and **Backup Certificate File**, or browse to the desired location.
5. Select the method for user authentication. Options are:
    - **Username and Password:** Activates the Primary and Backup Key Vault User Names and password fields for completion.
    - **Username:** Activates the Primary and Backup Key Vault User Names for completion.
    - **None:** Deactivates Primary and Backup Key Vault User Names and password fields.
6. Select the Certificate Type. Options are:

- **CA Signed:** The Encryption switch KAC certificate is signed by a CA, imported back on the Encryption switch and registered as a KAC certificate. The CA will be registered as a key vault certificate on the Encryption switch. If you selected **CA Signed**, the wizard opens the **Specify Public Key Certificate (KAC) File Name** dialog box (Figure 335). Go to Step 7.
- **Self Signed:** The self-signed certificates are exchanged and registered on both ends. The key vault certificate is registered on the Encryption switch and the Encryption switch KAC certificate is registered on the key vault. If you selected **Self Signed**, the wizard opens the **Specify Master Key File Name** dialog box. (Refer to Figure 336.) Go to Step 8.

7. Click **Next**.

   The **Specify Public Key Certificate (KAC) File Name** dialog box displays. (Refer to Figure 335.)

FIGURE 335 Specify Public Key Certificate (KAC) File Name dialog box

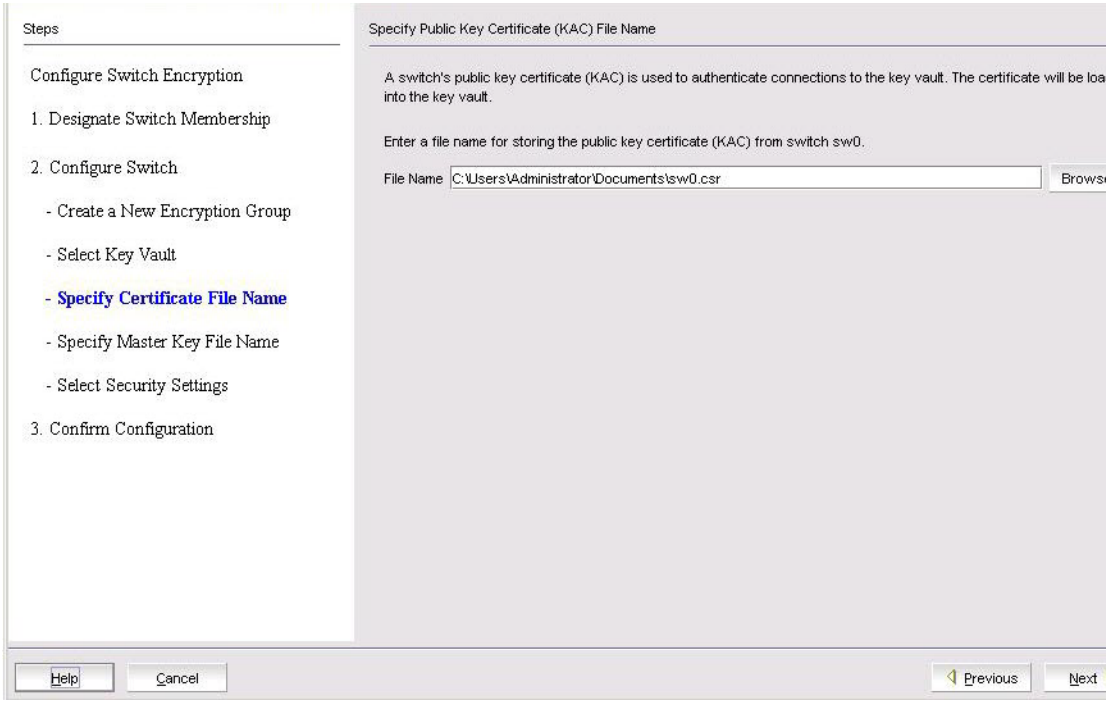

8. Enter the name of the file where the switch's public key certificate is stored, or browse to the desired location, then click **Next**.

   The **Specify Master Key File Name** dialog box displays. (Refer to Figure 336.)

FIGURE 336 Specify Master Key File Name dialog box

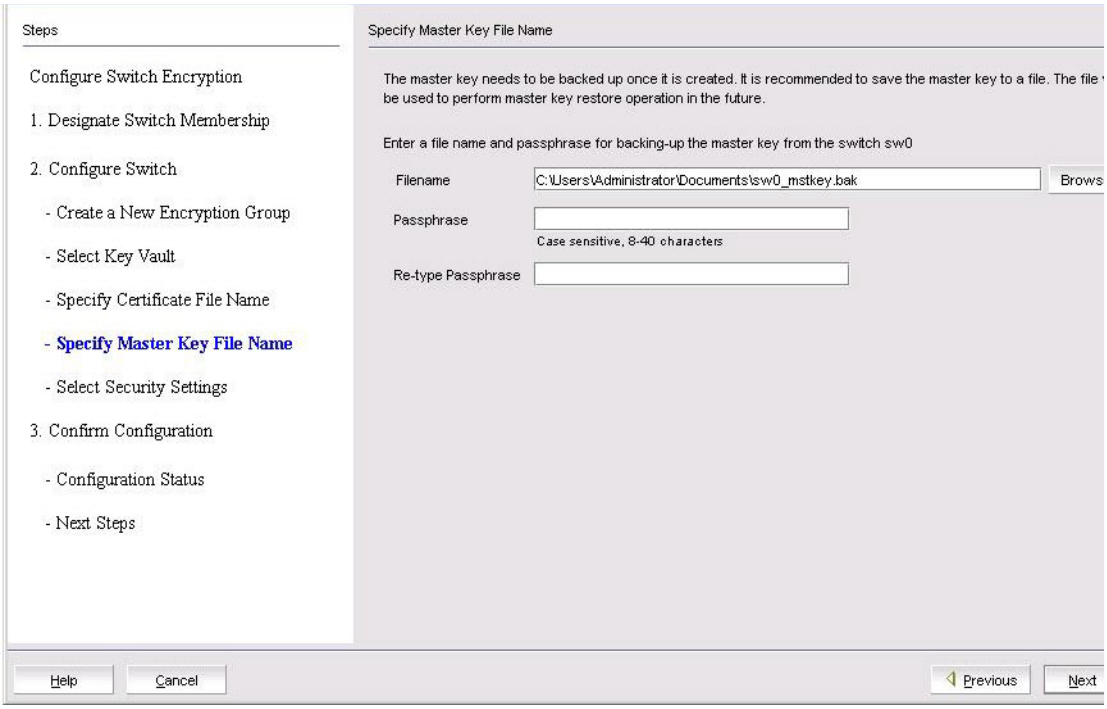

9. Enter the name of the file used for backing up the master key, or browse to the desired location.
10. Enter the passphrase, which is required for restoring the master key. The passphrase can be between eight and 40 characters, and any character is allowed.
11. Re-enter the passphrase for verification, then click **Next**.

    The **Select Security Settings** dialog box displays. (Refer to Figure 337.)

FIGURE 337 Select Security Settings dialog box

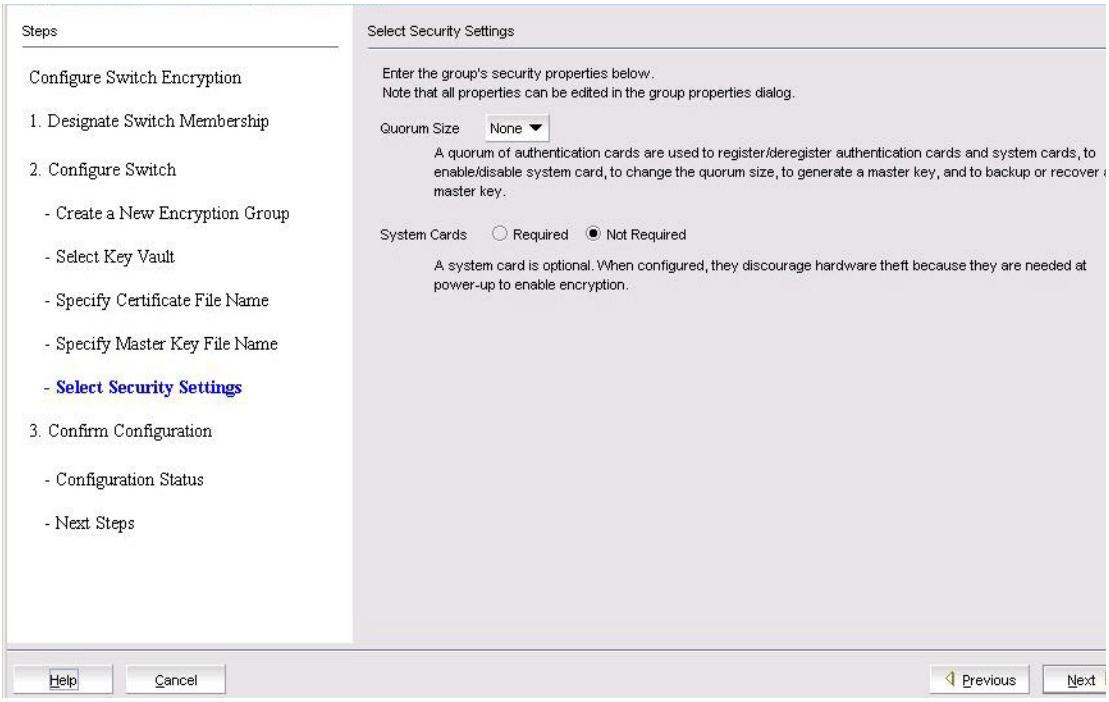

12. Set quorum size and system card requirements.

    The quorum size is the minimum number of cards necessary to enable the card holders to perform the security sensitive operations listed above. The maximum quorum size is five cards. The actual number of authentication cards registered is always more than the quorum size, so if you set the quorum size to five, for example, you will need to register at least six cards in the subsequent steps.

    Setting quorum size to a value greater than zero and/or setting system cards to **Required** launches additional wizard dialog boxes.

13. Click **Next**.

    The **Confirm Configuration** dialog box displays. (Refer to Figure 338.)

FIGURE 338 Confirm Configuration dialog box

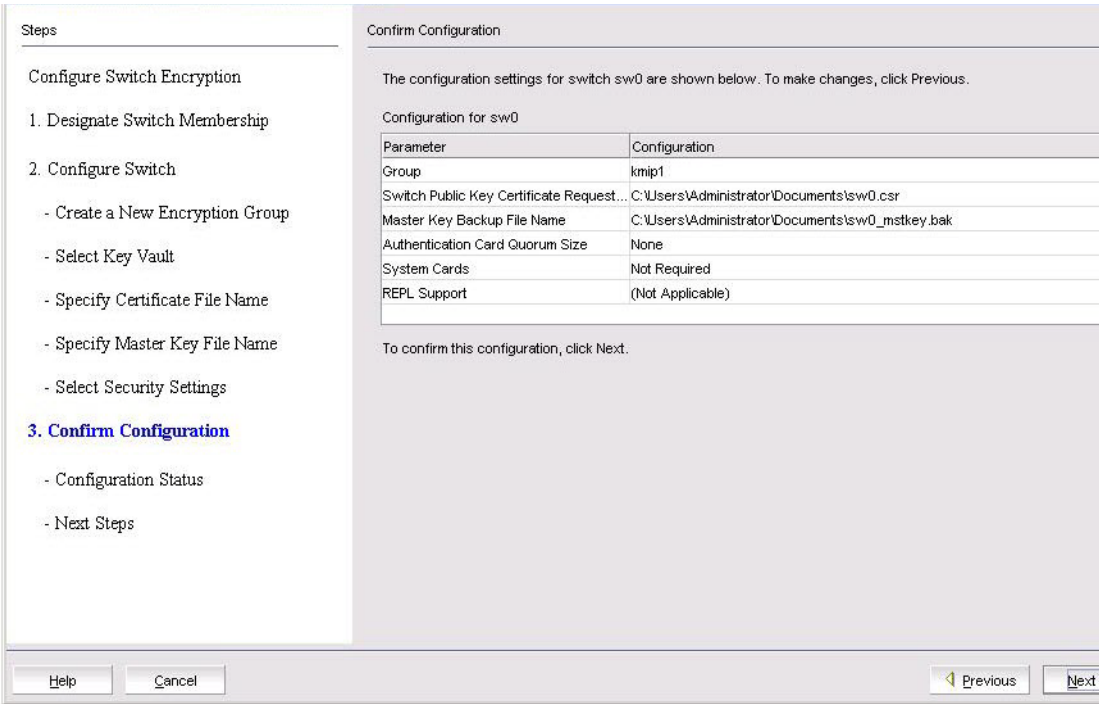

14. Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

    The **Configuration Status** dialog box displays. (Refer to Figure 339.)

FIGURE 339 Configuration Status dialog box

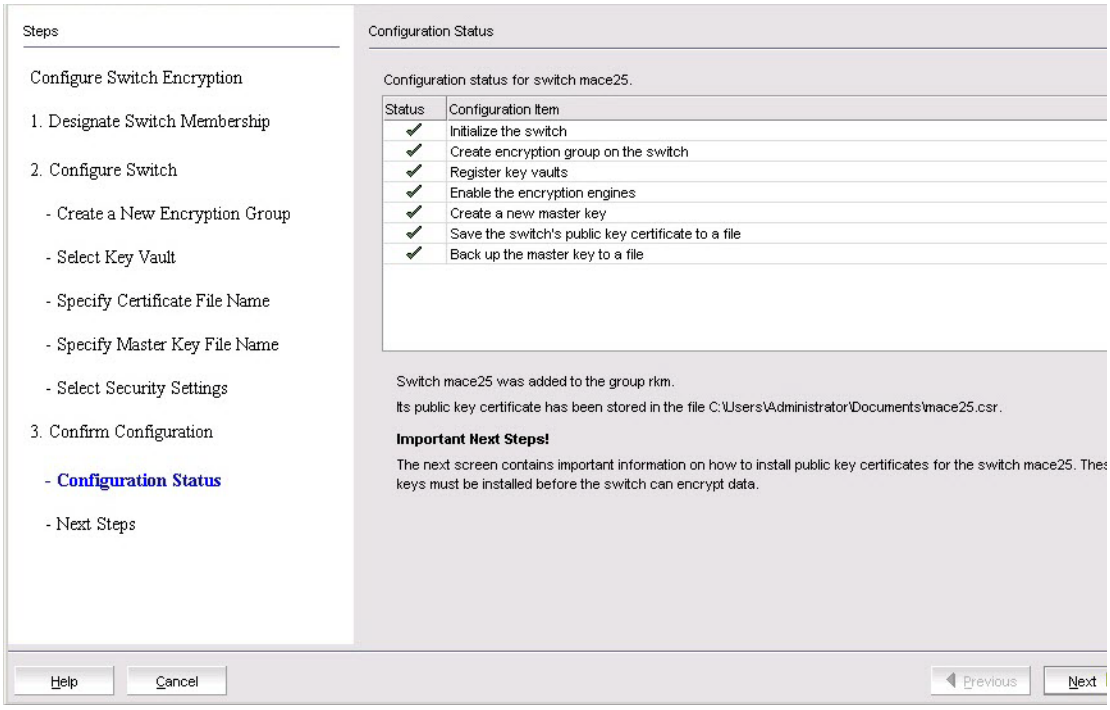

All configuration items have green check marks if the configuration is successful. A red stop sign indicates a failed step. A message displays below the table, indicating the encryption switch was added to the group you named, and the public key certificate is stored in the location you specified.

After configuration of the encryption group is completed, the Management application sends API commands to verify the switch configuration.

15. Click **Next**.

    The **Next Steps** dialog box displays. (Refer to Figure 340.) Instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

FIGURE 340 Next Steps dialog box

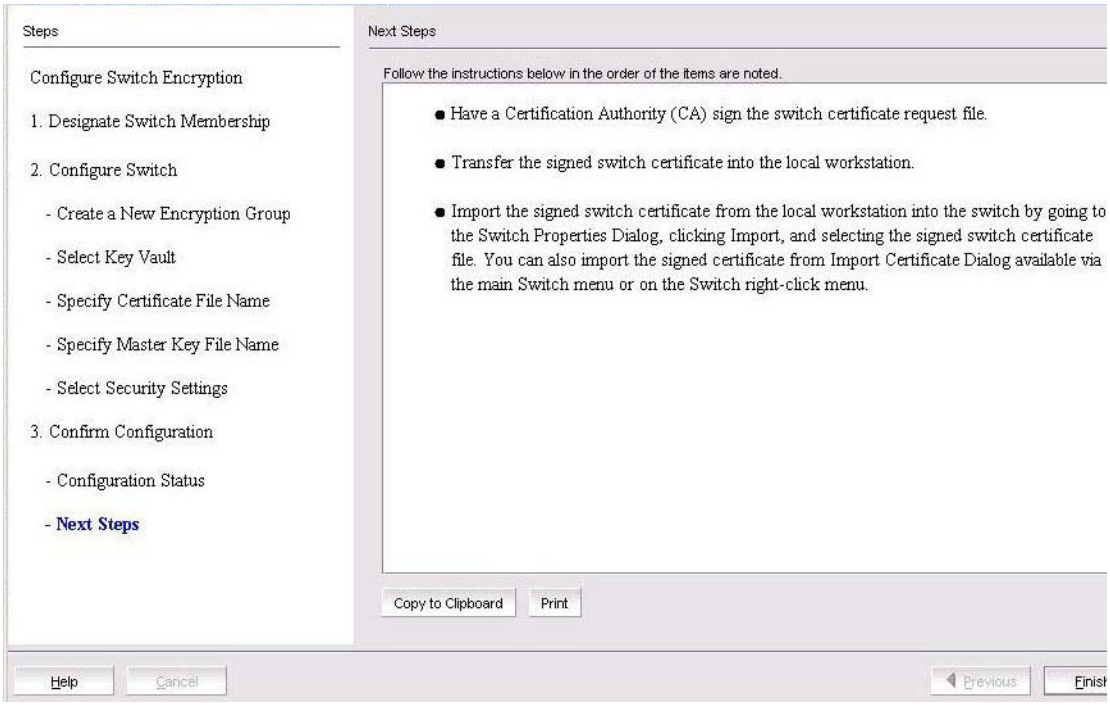

16. Review the post-configuration instructions, which you can copy to a clipboard or print for later, then click **Finish** to exit the **Configure Switch Encryption** wizard.

Refer to "Understanding configuration status results".

Understanding configuration status results

After configuration of the encryption group is completed, the Management application sends API commands to verify the switch configuration. The CLI commands are detailed in the encryption administrator's guide for your key vault management system.

1. Initialize the switch. If the switch is not already in the initiated state, the Management application performs the **cryptocfg --initnode** command.
2. Create an encryption group on the switch. the Management application creates a new group using the **cryptocfg --create -encgroup** command, and sets the key vault type using the **cryptocfg --set -keyvault** command.
3. Register the key vault. the Management application registers the key vault using the **cryptocfg --reg keyvault** command.

4. Enable the encryption engines. the Management application initializes an encryption switch using the **cryptocfg --initEE [<slotnumber>] and cryptocfg --regEE [<slotnumber>]** commands.
5. Create a new master key. **(Opaque key vaults only).** the Management application checks for a new master key. New master keys are generated from the **Security** tab located in the **Encryption Group Properties** dialog box.

   NOTE
   A master key is not generated if the key vault type is LKM/SSKM. LKM/SSKM manages DEK exchanges through a trusted link, and the LKM/SSKM appliance uses its own master key to encrypt DEKs.

6. Save the switch's public key certificate to a file. the Management application saves the KAC certificate in the specified file.
7. Back up the master key to a file. **(Opaque key vaults only).** the Management application saves the master key in the specified file.

Adding a switch to an encryption group

The setup wizard allows you to either create a new encryption group, or add an encryption switch to an existing encryption group. Use the following procedure to add a switch to an encryption group:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a switch to add from the **Encryption Center Devices** table, then select **Switch > Create/Add to Group from the menu task bar**.

   NOTE
   The switch must not already be in an encryption group.

   The **Configure Switch Encryption** wizard welcome screen displays. (Refer to Figure 341.)

FIGURE 341 Configure Switch Encryption wizard - welcome screen

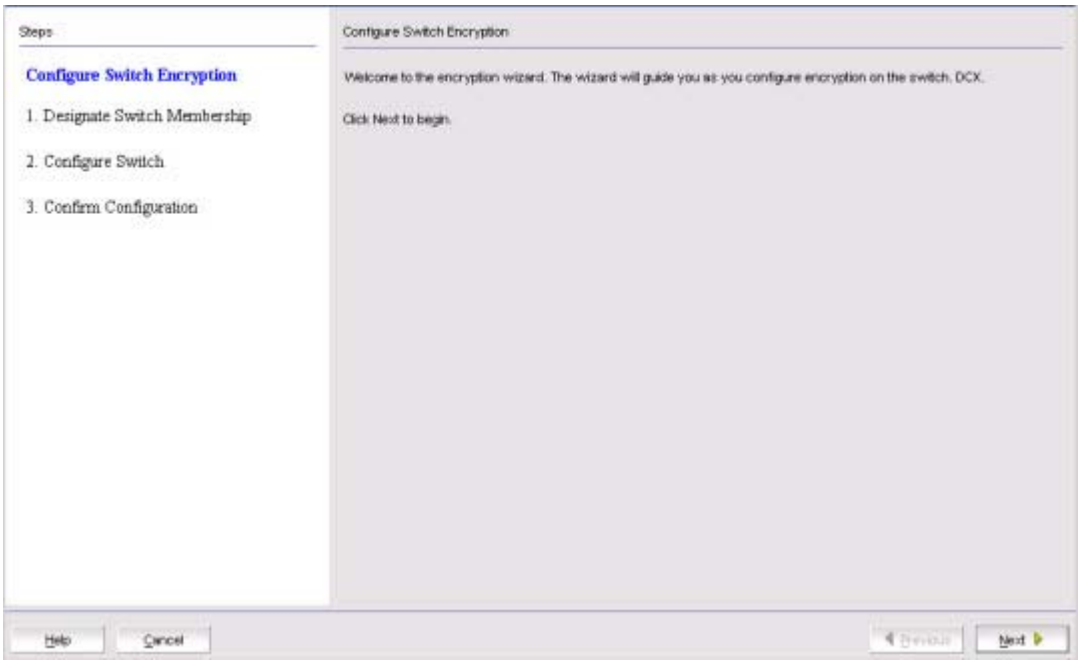

3. Click **Next**.

   The **Designate Switch Membership** dialog box displays. (Refer to Figure 342.)

FIGURE 342 Designate Switch Membership dialog box

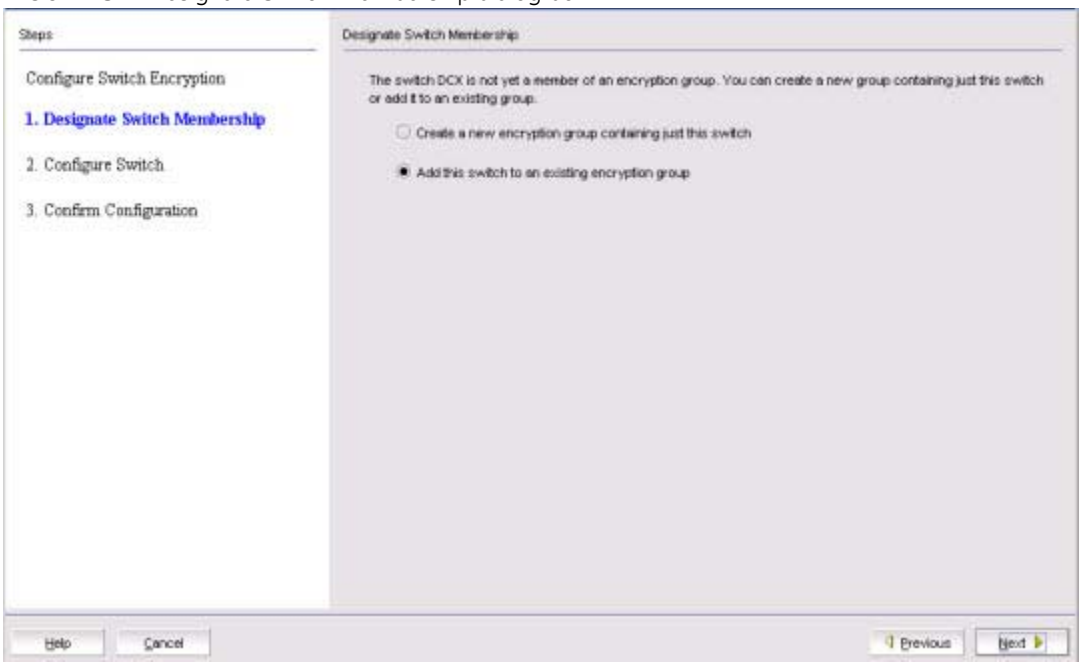

4. For this procedure, select **Add this switch to an existing encryption group**, then click **Next**.

   The **Add Switch to Existing Encryption Group** dialog box displays. (Refer to Figure 343.)

   The dialog box contains the following information:

   - **Encryption Groups** table: Enables you to select an encryption group in which to add a switch.
   - **Member Switches** table: Lists the switches in the selected encryption group.

     NOTE
     If you are creating a new encryption group, refer to "Creating a new encryption group" on page 672.

FIGURE 343 Add Switch to Existing Encryption Group dialog box

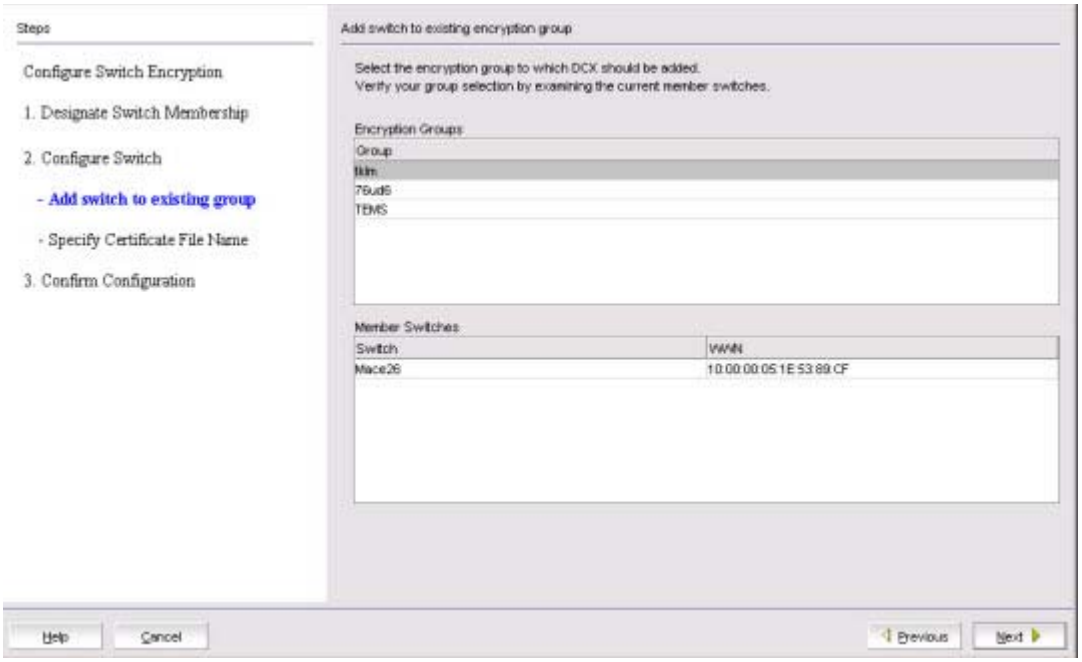

5. Select the group in which to add the switch, then click **Next**.

   The **Specify Public Key Certificate (KAC) File Name** dialog box displays. (Refer to Figure 344.)

FIGURE 344 Specify Public Key Certificate (KAC) File Name dialog box

6. Enter the location where you want to store the public key certificate that is used to authenticate connections to the key vault, or browse to the desired location, then click **Next**.

The **Confirm Configuration** dialog box displays. (Refer to Figure 345.) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

FIGURE 345 Confirm Configuration dialog box

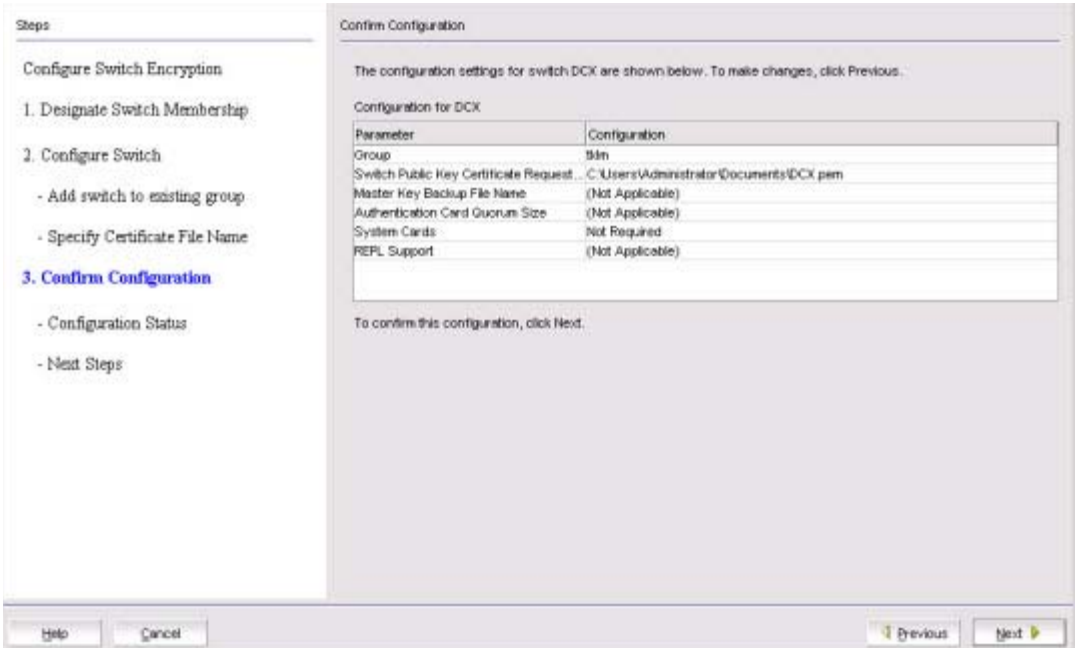

The **Configuration Status** dialog box displays. (Refer to Figure 346.)

FIGURE 346 Configuration Status dialog box

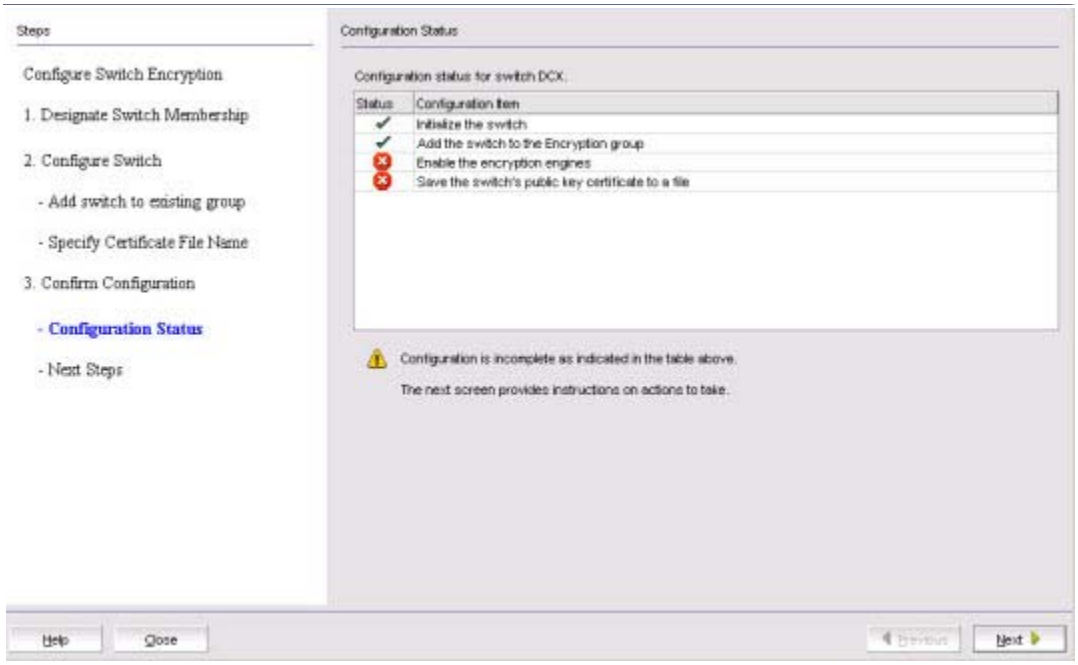

All configuration items have green check marks if the configuration is successful. A red stop sign indicates a failed step. A message displays below the table, indicating the encryption switch was added to the group you named, and the public key certificate is stored in the location you specified.

7. Review important messages, then click **Next**.

   The **Error Instructions** dialog box displays. (Refer to Figure 347.) Instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

FIGURE 347 Error Instructions dialog box

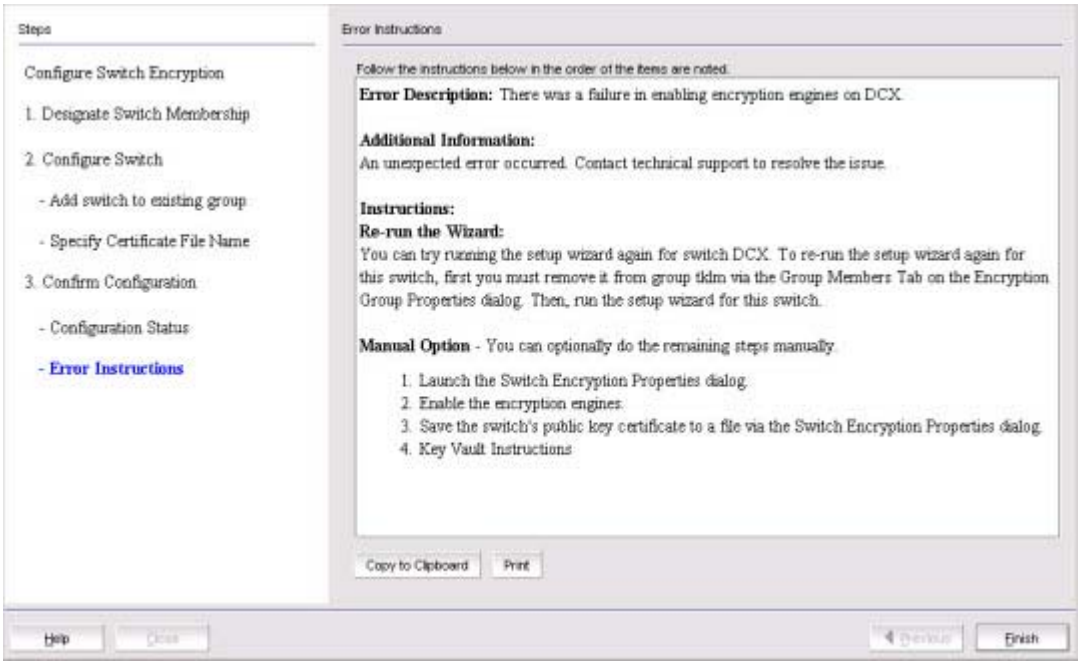

8. Review the post-configuration instructions, which you can copy to a clipboard or print for later.
9. Click **Finish** to exit the **Configure Switch Encryption** wizard.

Replacing an encryption engine in an encryption group

To replace an encryption engine in an encryption group with another encryption engine within the same DEK Cluster, complete the following steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption engine from the **Encryption Center Devices table**, then select **Engine > Replace** from the menu task bar.

   The **Encryption Group Properties** dialog box displays with the **Engine Operations** tab selected. (Refer to Figure 348.)

   You can also display the **Engine Operations** tab by selecting an encryption group from the **Encryption Center Devices** table, selecting **Group > Properties** from the menu task bar, then selecting the **Engine Operations** tab.

FIGURE 348Engine Operations tab

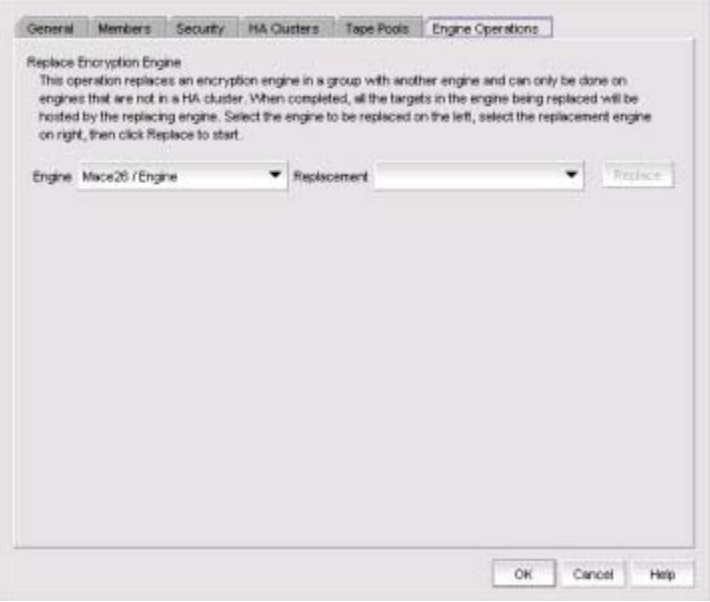

3. Select the engine to replace from the **Engine** list.
4. Select the engine to use as the replacement from the **Replacement** list, then click **Replace**.

   All containers hosted by the current engine (**Engine** list) are replaced by the new engine (**Replacement** list).

High availability clusters

A high availability (HA) cluster consists of exactly two encryption engines configured to host the same CryptoTargets and to provide Active/Standby failover and failback capabilities in a single fabric. One encryption engine can take over encryption and decryption tasks for the other encryption engine if that member fails or becomes unreachable.

> NOTE
> High availability clusters between two encryption engines (EEs) should not be confused with High Availability opaque mode that is supported in KMIP.

When creating a new HA cluster, add one engine to create the cluster, then add the second engine. You can make multiple changes to the HA clusters list; the changes are not applied to the switch until you click **OK**.

> NOTE
> An IP address is required for the management port for any cluster-related operations.

HA cluster configuration rules

The following rules apply when configuring an HA cluster:

- The encryption engines that are part of an HA cluster must belong to the same encryption group and be part of the same fabric.
- An HA cluster cannot span fabrics and it cannot provide failover/failback capability within a fabric transparent to host MPIO software.
- HA cluster configuration and related operations must be performed on the group leader.

- HA clusters of FS8-18 blades should not include blades in the same DCX Backbone chassis.

  NOTE
  HA cluster creation is blocked when encryption engines belonging to FS8-18 blades in the same DCX Backbone chassis are specified.

- Cluster links must be configured before creating an HA cluster.
- It is recommended that the HA cluster configuration be completed before you configure storage devices for encryption.
- It is mandatory that the two encryption engines in the HA cluster belong to two different nodes for true redundancy. This is always true for Encryption switches, but is not true if two FS8-18 blades in the same DCX Backbone chassis are configured in the same HA cluster.

Creating HA clusters

For the initial encryption node, perform the following procedure.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption group from the **Encryption Center Devices** table, then select **Group > HA Cluster** from the menu task bar.

   NOTE
   If groups are not visible in the **Encryption Center Devices** table, select **View > Groups** from the menu task bar.

   The Encryption Group Properties dialog box displays, with the **HA Clusters** tab selected. (Refer to Figure 349.)
3. Select an available encryption engine from the **Non HA Encryption Engines** table and a destination HA cluster from the **High Availability Clusters** table**.** Select **New HA Cluster** if you are creating a new cluster.

   NOTE
   If you are creating a new HA cluster, a dialog box displays requesting a name for the new HA cluster. HA cluster names can have up to 31 characters. Letters, digits, and underscores are allowed.
4. Click the right arrow to add the encryption engine to the selected HA cluster.

FIGURE 349 Encryption Group Properties dialog box - HA Clusters tab

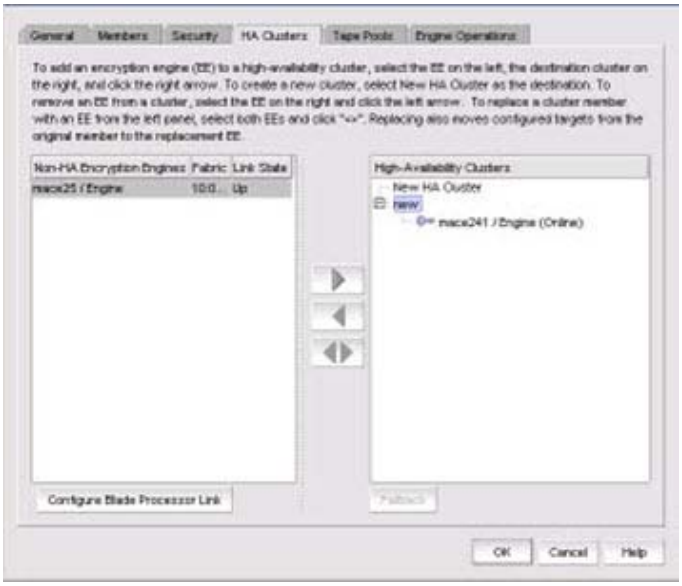

To add the second encryption node to the HA cluster, perform the following procedure.

1. Select the desired HA cluster from the right panel.
2. Select the desired encryption engine to be added from the left panel.
3. Click the right arrow to add the encryption engine to the selected HA cluster.
4. Click **OK**.

Removing engines from an HA cluster

Removing the last engine from an HA cluster also removes the HA cluster.

If only one engine is removed from a two-engine cluster, you must either add another engine to the cluster, or remove the other engine.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption group from the **Encryption Center Devices** table, then select **Group > HA Cluster** from the menu task bar.

   The Encryption Group Properties dialog box displays, with the **HA Clusters** tab selected.
3. Select an engine from the **High Availability Clusters** table, then click the left arrow. (Refer to Figure 349.)
4. Either remove the second engine or add a replacement second engine, making sure all HA clusters have exactly two engines.
5. Click **OK**.

Swapping engines in an HA cluster

Swapping engines is useful when replacing hardware. Swapping engines is different from removing an engine and adding another because when you swap engines, the configured targets on the former HA cluster member are moved to the new HA cluster member.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box.
2. Select an encryption group from the **Encryption Center Devices** table, then select **Group > HA Cluster** from the menu task bar

   The Encryption Group Properties dialog box displays, with the **HA Clusters** tab selected. (Refer to Figure 349.)

To swap engines, select one engine from the **High Availability Clusters** table and one unclustered engine from encryption engine from the **Non HA Encryption Engines** table, then click the dual arrow.

NOTE
The two engines being swapped must be in the same fabric.

Failback option

The **Failback** option determines the behavior when a failed encryption engine is restarted. When the first encryption engine comes back online, the encryption group's failback setting (auto or manual) determines how the encryption engine resumes encrypting and decrypting traffic to its encryption targets.

- In auto mode, when the first encryption engine restarts, it automatically resumes encrypting and decrypting traffic to its encryption targets.
- In manual mode, the second encryption engine continues handling the traffic until you manually invoke failback using the CLI or the Management application, or until the second encryption engine fails. When the encryption engine recovers, it can automatically fail back its CryptoTarget containers if the second encryption engine is not hosting them.

Invoking failback

To invoke failback to the restarted encryption engine from the Management application, complete the following steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption group from the **Encryption Center Devices** table to which the encryption engine belongs, then click **Group > HA Clusters**.

   The Encryption Group Properties dialog box displays, with the **HA Clusters** tab selected (Refer to Figure 349.)
3. Select the online encryption engine, then click **Failback**.
4. Click **OK**, then close the **Encryption Center** dialog box.

Configuring encryption storage targets

Adding an encryption target maps storage devices and hosts to virtual targets and virtual initiators within the encryption switch. The storage encryption wizard enables you to configure encryption for a storage device (target).

> NOTE
> It is recommended that you configure the host and target in the same zone before configuring them for encryption. If the host and target are not already in the same zone, you can still configure them for encryption, but you will need to configure them in the same zone before you can commit the changes. If you attempt to close the **Encryption Targets** dialog box without committing the changes, you are reminded of uncommitted changes in the Management application.

The wizard steps are as follows:

1. Select Encryption Engine
2. Select Target
3. Select Hosts
4. Name Container
5. Confirmation
6. Configuration Status
7. Important Instructions

Adding an encryption target

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table to which to add the target, then select **Group/Switch/Engine > Targets** from the menu task bar.

   > NOTE
   > You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   The **Encryption Targets** dialog box displays. (Refer to Figure 350.)

FIGURE 350 Encryption Targets dialog box

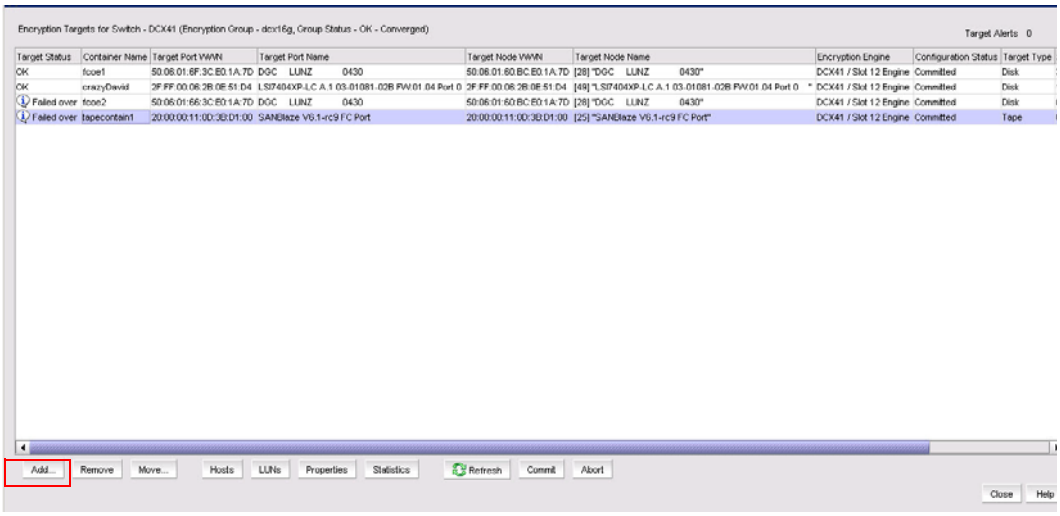

3. Click **Add**.

   The **Configure Storage Encryption** wizard welcome screen displays. (Refer to Figure 351.)

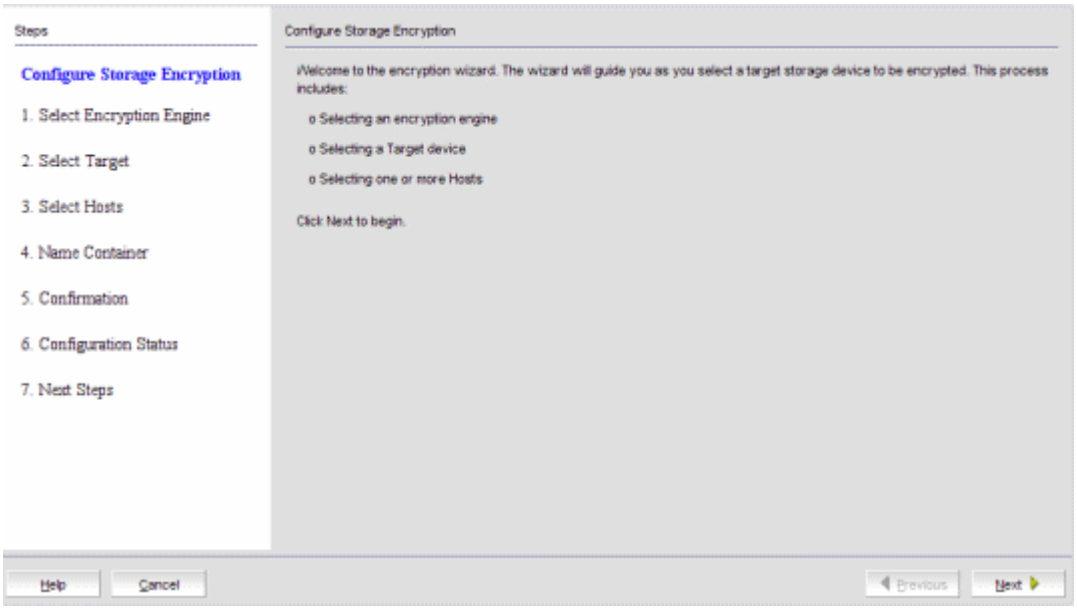

FIGURE 351 Configure Storage Encryption wizard - welcome screen

4. Click **Next**.

   The **Select Encryption Engine** dialog box displays. (Refer to Figure 352.)

FIGURE 352Select Encryption Engine dialog box

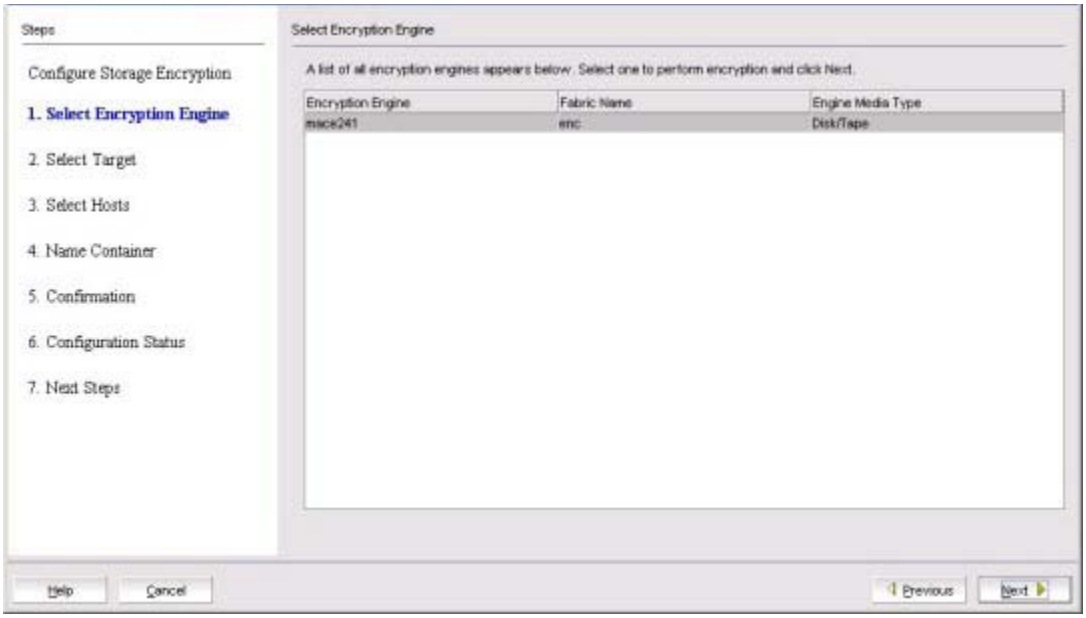

The dialog box contains the following information:

- **Encryption engine:** The name of the encryption engine. The list of engines depends on the scope being viewed:
  - If an encryption group was selected, the list includes all engines in the group.
  - If a switch was selected, the list includes all encryption engines for the switch.
  - If a single encryption engine was selected, the list contains only that engine.
- **Fabric Name:** The name of the fabric to which the selected encryption engine (blade or switch) is configured.
- **Engine Media Type:** The media type of the encryption engine. Options are: **Tape** and **Disk**.

5. Select the encryption engine (blade or switch) to configure, then click **Next.**

   The **Select Target** dialog box displays. (Refer to Figure 353.) The dialog box lists all target ports and target nodes in the same fabric as the encryption engine. The **Targets in Fabric table** does *not* show targets that are already configured in an encryption group.

FIGURE 353Select Target dialog box

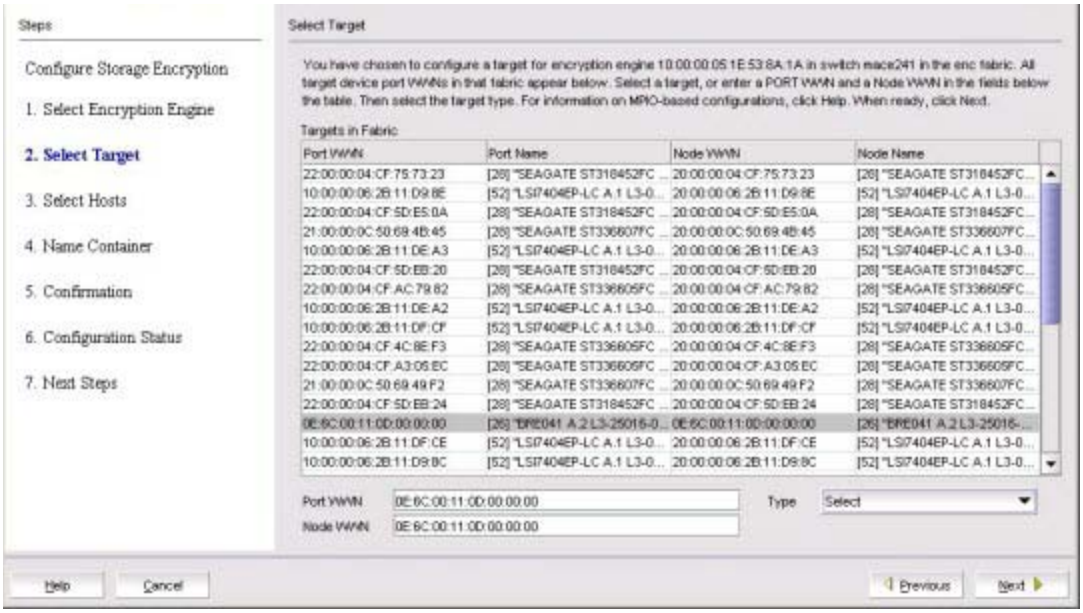

The dialog box contains the following information:

- **Target Port WWN:** The world wide name of the target port in the same fabric as the encryption engine.
- **Target Port Name:** The name of the target port in the same fabric as the encryption engine.
- **Target Node WWN:** The world wide name of the target node in the same fabric as the encryption engine.
- **Target Node Name:** The name of the target device.
- **Targets** list: Options are: **Tape** and **Disk**.

NOTE
The **Targets** list does not show targets that are already configured in the encryption group.

6. Select a target from the list. (The **Target Port WWN** and **Target Node WWN** fields contain all target information that displays when using the **nsShow** command.) You can also enter WWNs manually, for example, to specify a target that is not on the list.
7. Select a target type from the **Type** list, then click **Next**.

The **Select Hosts** dialog box displays. (Refer to Figure 354.) You can configure hosts for selected target device ports. All hosts that are in the same fabric as the encryption engine are listed.

NOTE
The selected target and initiator port must be in the same zone, or an error will result.

FIGURE 354 Select Hosts dialog box

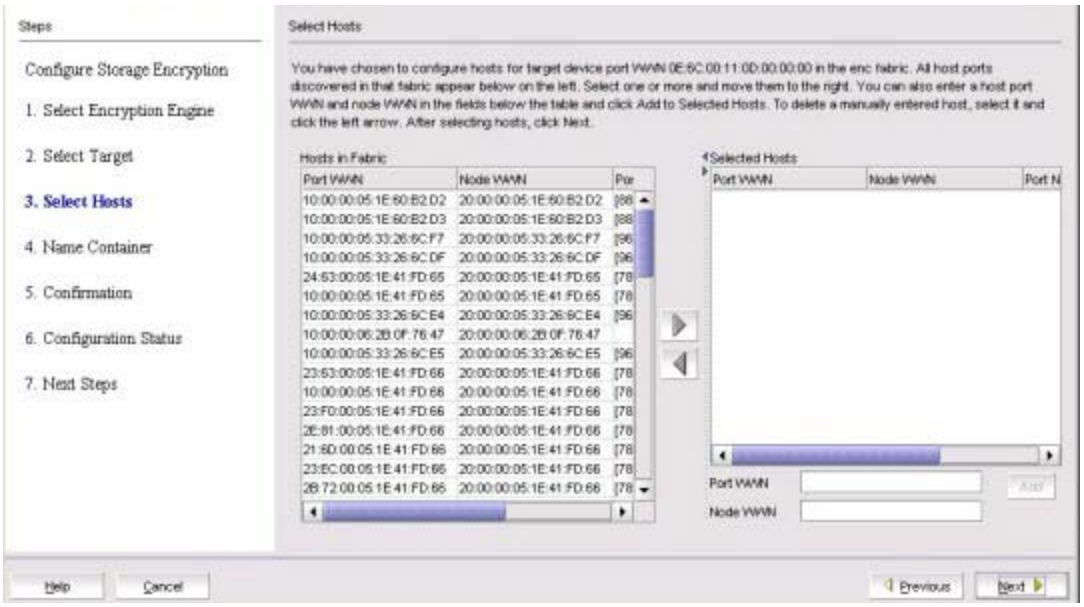

The dialog box contains the following information:

- **Hosts in Fabric** table: Lists the available hosts in the fabric.
- **Selected Hosts** table: Lists the hosts that have been selected to access the target.
- **Port WWN:** The world wide name of the host ports that are in the same fabric as the encryption engine.
- **Node WWN:** The world wide name of the host nodes that are in the same fabric as the encryption engine.
- **Port Name:** The user-assigned port name, if one exists; otherwise, the symbolic port name from the device.
- **Port ID:** The 24-bit Port ID of the host port.
- **VI Port WWN:** The world wide name of the virtual initiator port.
- **VI Node WWN:** The world wide name of the virtual initiator node.
- **Host Name:** The name of the hosts that are in the same fabric as the encryption engine.
- **Port WWN** text box: Type a world wide name for a host port.

  NOTE
  You must enter the host node world wide name before clicking **Add**, to add the WWN to the **Selected Hosts** table.

- **Node WWN** text box: Type a world wide name for a host node.

  NOTE
  You must also enter the host port world wide name before clicking **Add** to add the node WWN to the **Selected Hosts** table.

- **Device Type:** The device type indicated by the fabric's name service. The value is either **Initiator** or **Initiator + Target**.
- Right arrow button: Moves a host from the **Host in Fabric** table to the **Selected Hosts** table.
- Left arrow button: Removes a host from the **Selected Hosts** table.
- **Add** button: Click to manually add host port world wide names or host node world wide names to the **Selected Hosts** table.

8. Select hosts using either of the following methods:

a. Select a maximum of 1024 hosts from the **Hosts in Fabric** table, then click the right arrow to move the hosts to the **Selected Hosts** table. (The **Port WWN** column contains all target information that displays when using the **nsshow** command.)

b. Manually enter world wide names in the **Port WWN** and **Node WWN** text boxes if the hosts are not included in the table. You must fill in both the Port WWN and the Node WWN. Click **Add** to move the host to the **Selected Hosts** table.

9. Click **Next**.

   The **Name Container** dialog box displays. (Refer to Figure 355.) You can specify a name for the target container that is created in the encryption engine to hold the target configuration data. The name is only needed when configuring the storage using the command line interface (CLI).

   The container name defaults to the target WWPN. You can, however, rename the container name. Target container names can have up to 31 characters. Letters, digits, and underscores are allowed.

FIGURE 355 Name Container dialog box

10. Enter the container name. The container name is a logical encryption name to specify a name other than the default. You can use a maximum of 31 characters. Letters, digits, and underscores are allowed.

11. Click **Next**.

    The **Confirmation** screen displays. (Refer to Figure 356.) The confirmation screen confirms and completes configuration of encryption engines, targets, and hosts.

FIGURE 356Confirmation screen

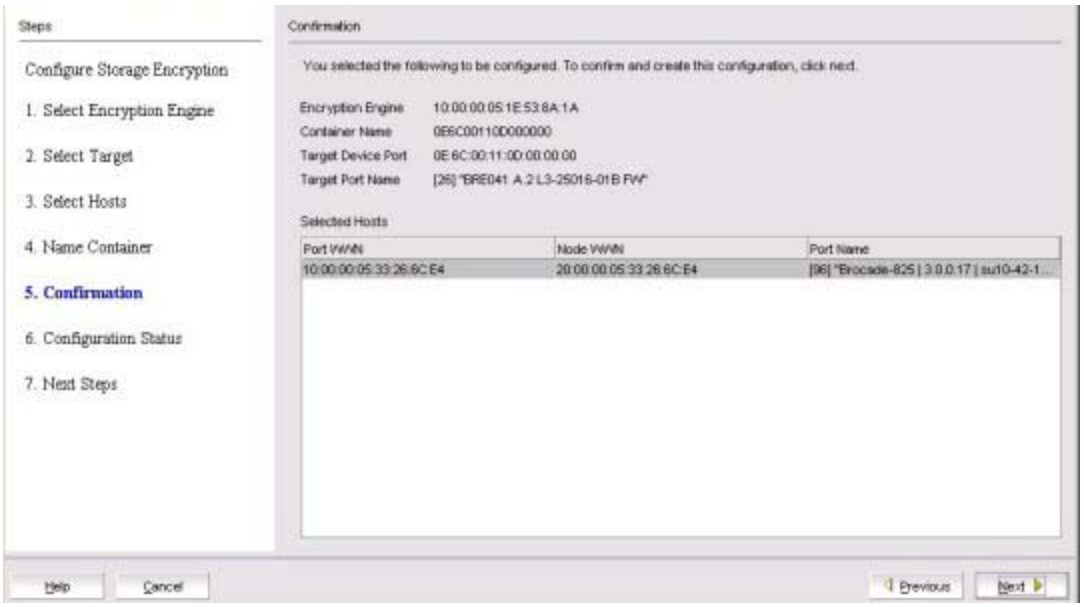

The **Confirmation** screen contains the following information:

- **Encryption Engine:** The slot location of the encryption engine.
- **Container Name:** The logical encryption name used to map storage targets and hosts to virtual targets and virtual initiators.
- **Target Device Port:** The world wide name of the target device port.
- **Host Node WWN:** The world wide name of the host node.
- **Host Port WWN:** The world wide name of the host port.
- **Host Name:** The name of the host.

12. Verify the information is correct, then click **Next**, which creates the configuration.

    The **Configuration Status** screen displays, which shows the status of the new container configuration. (Refer to Figure 357.) The target and host that are configured in the target container are listed, as well as the virtual targets (VT) and virtual initiators (VI).

    NOTE
    If you can view the VI/VT Port WWNs and VI/VT Node WWNs, the container has been successfully added to the switch.

FIGURE 357Configuration Status screen

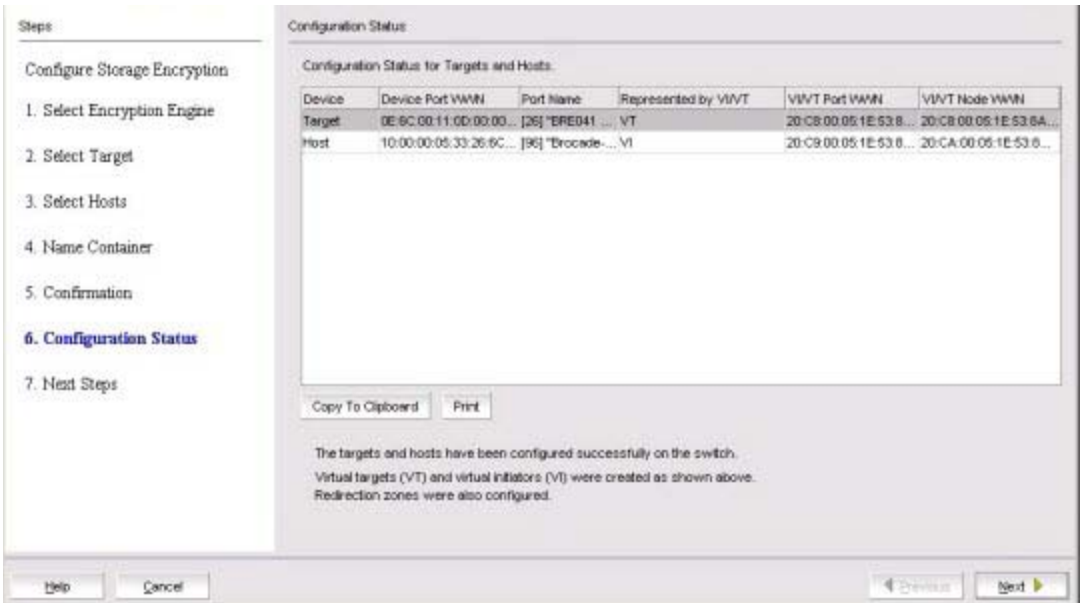

The screen contains the following information:

- **Device:** The device type (target or host).
- **Device Port WWN:** The port world wide name.
- **Represented by VI/VT:** The virtual target (VT) mapped to the physical target or virtual initiator (VI) representing the host.
- **VI/VT Port WWN:** The port world wide name of the virtual target or virtual initiator.
- **VI/VT Node WWN:** The node world wide name of the virtual target or virtual initiator.

13. Review any post-configuration instructions or messages, which you can copy to a clipboard or print for later, then click **Next**.

    The **Next Steps screen** displays. (Refer to Figure 358.) Post-configuration instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

FIGURE 358Next Steps screen

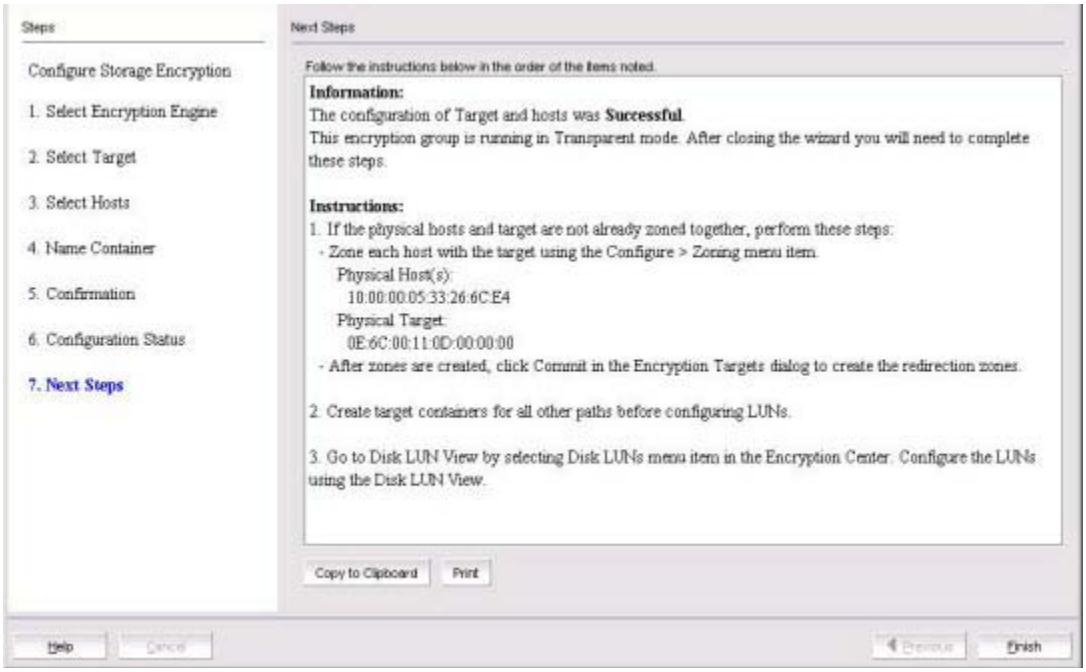

The **Next Steps** screen contains the following information:

- **Important Instructions**: Instructions about post-configuration tasks you must complete after you close the wizard. For example, you must zone the physical hosts and the target together and then you encrypt the LUNs using the **Storage Device LUNs** dialog box.
- **Copy to Clipboard** button: Saves a copy of the instructions.
- **Print** button: Prints the configuration.

14. Review the post-configuration instructions, which you can copy to a clipboard or print for later, then click **Finish** to exit the **Configure Switch Encryption** wizard.

Configuring hosts for encryption targets

Use the **Encryption Target Hosts** dialog box to edit (add or remove) hosts for an encrypted target.

NOTE
Hosts are normally selected as part of the **Configure Switch Encryption** wizard, but you can also edit hosts later using the **Encryption Target Hosts** dialog box.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table that contains the storage device to be configured, then select **Group/Switch/Engine > Targets** from the menu task bar.

NOTE
You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

The **Encryption Targets** dialog box displays. (Refer to Figure 359.)

FIGURE 359Encryption Targets dialog box

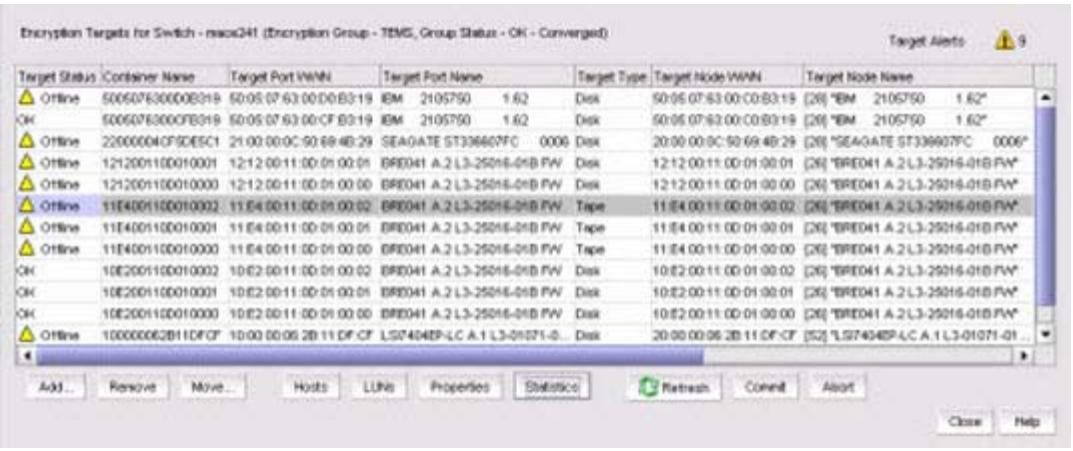

3. Select a target storage device from the list, then click **Hosts**.

   The **Encryption Target Hosts** dialog box displays. The **Hosts in Fabric** table lists the configured hosts in a fabric.

   The table displays the following information:

   - **Port WWN**: The world wide name of the host ports that are in the same fabric as the encryption engine.
   - **Node WWN:** The world wide name of the host nodes that are in the same fabric as the encryption engine.
   - **Port Name**: The name of the hosts that are in the same fabric as the encryption engine.
   - **Port ID**: Displays the 24-bit port ID (PID) of the host port in both the **Host Ports in Fabric** table and the **Selected Hosts** table.

     NOTE
     Both the **Hosts in Fabric** table and the **Selected Hosts** table now contain a **Port ID** column to display the 24-bit PID of the host port.

4. Select one or more hosts in a fabric using either of the following methods:

   a. Select a maximum of 1024 hosts from the **Hosts in Fabric** table, then click the right arrow to move the hosts to the **Selected Hosts** table. (The **Port WWN** column contains all target information that displays when using the **nsShow** command.)

   b. Manually enter world wide names in the **Port WWN** and **Node WWN** text boxes if the hosts are not included in the table. You must fill in both the Port WWN and the Node WWN. Click **the** right arrow button to move the host to the **Selected Hosts** table.

      NOTE
      The selected host and target must be in the same zone, or an error will result.

   The Selected Hosts table lists the following:

   - **Port WWN**: The selected host port's world wide name.
   - **Node WWN**: The selected host node's world wide name.
   - **Port Name**: The name of the host selected to access the encryption target.
   - **Port WWN** text box: Type a world wide name for a host port, and click the Add to Selected Hosts button to add to the Selected Hosts table.
   - **Port ID**: Displays the 24-bit port ID (PID) of the host port in both the Host Ports in Fabric table and the Selected Hosts table.

- **VI Port WWN**: The world wide name of the virtual initiator port.
- **VI Node WWN**: The world wide name of the virtual initiator node.

  NOTE
  To remove an encryption engine from the **Selected Hosts** table, select the engine(s), then click the left arrow button.

5. Click **OK** or **Apply** to apply your changes.

Adding target disk LUNs for encryption

You can add a new path to an existing disk LUN or add a new LUN and path by launching the **Add New Path** wizard.

NOTE
Before you can add a target disk LUN for encryption, you must first configure the Storage Arrays. For more information, see "Configuring storage arrays" on page 730.

Complete the following steps to add a target disk LUN:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box.
2. Select a group, switch, or engine from the **Encryption Center Devices** table, then select **Group/Switch/Engine > Disk LUNs** from the menu task bar.

   The **Encryption Disk LUN View** dialog box displays. (Refer to Figure 360.)

FIGURE 360 Encryption Disk LUN View dialog box

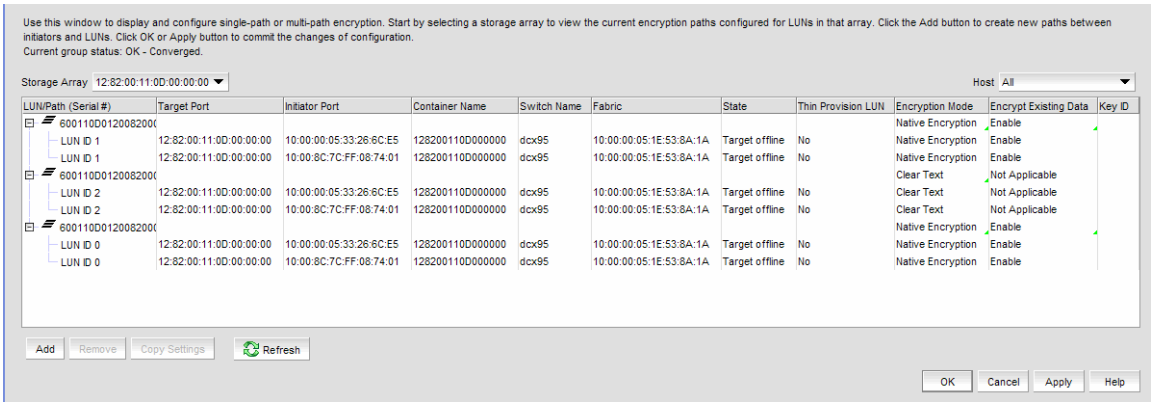

The dialog box provides a convenient way to view and manage disk LUNs that are provisioned from different hosts, identify conflicts between configuration policies on storage systems, and to provide a launching point for the **Add New Path** wizard for configuring multiple I/O paths to the LUN.

The dialog box contains the following information:

- **Storage Array** selector: Determines which LUN paths are displayed in the table. Enables you to select a storage array from the LUN view prior to launching the **Add New Path** wizard. Only ports that belong to at least one target container are listed.
- **Host** selector: Used to select a host from the LUN view prior to launching the **Add New Path** wizard. Only ports that belong to at least one target container are listed.

- **Encryption path** table: Should be LUN/Path identified by the following:
  - **LUN Path Serial #**
  - **Target Port**
  - **Initiator Port**
  - **Container Name**
  - **Switch Name**
  - **Fabric**
  - **State**
  - **Thin Provision LUN**
  - **Encryption Mode**
  - **Encrypt Existing Data**
  - **Key ID**
- **Remove button:** Removes a selected entry from the table.

3. Click **Add** to launch the **Add New Path** wizard.

   The **Select Target Port** dialog box displays. (Refer to Figure 361.)

FIGURE 361 Select Target Port dialog box

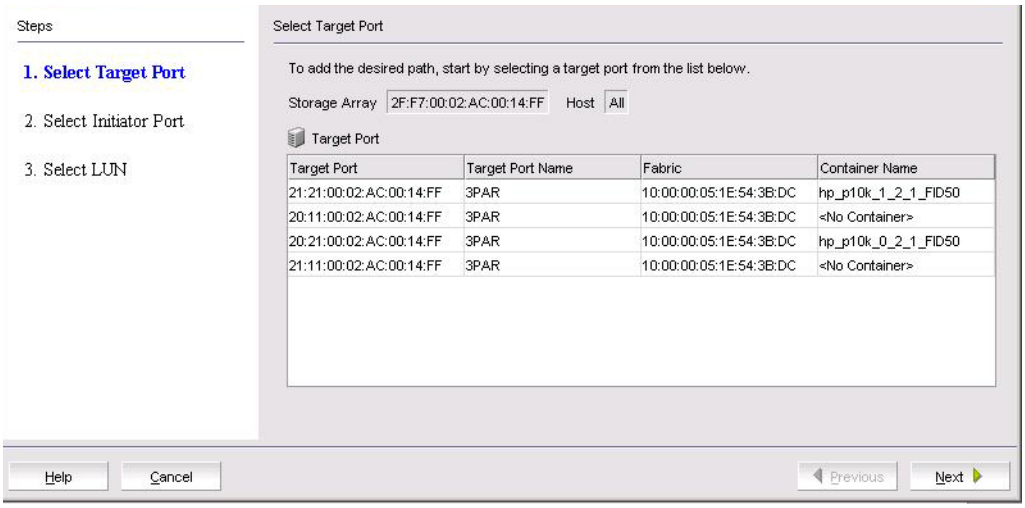

The dialog box is used to select a target port when configuring multiple I/O paths to a disk LUN. The dialog box contains the following information:

- **Storage Array** The storage array selected from the LUN view prior to launching the **Add New Path** wizard.
- **Host:** The host selected from the LUN view prior to launching the **Add New Path** wizard.
- **Target Port** table: Lists target ports using the following identifiers:
  - **Target Port**
  - **Target Port Name**
  - **Fabric**
  - **Container Name**

4. Select the target port from the **Target Port** table, then click **Next**.

   The **Select Initiator Port** dialog box displays. (Refer to Figure 362.)

FIGURE 362 Select Initiator Port dialog box

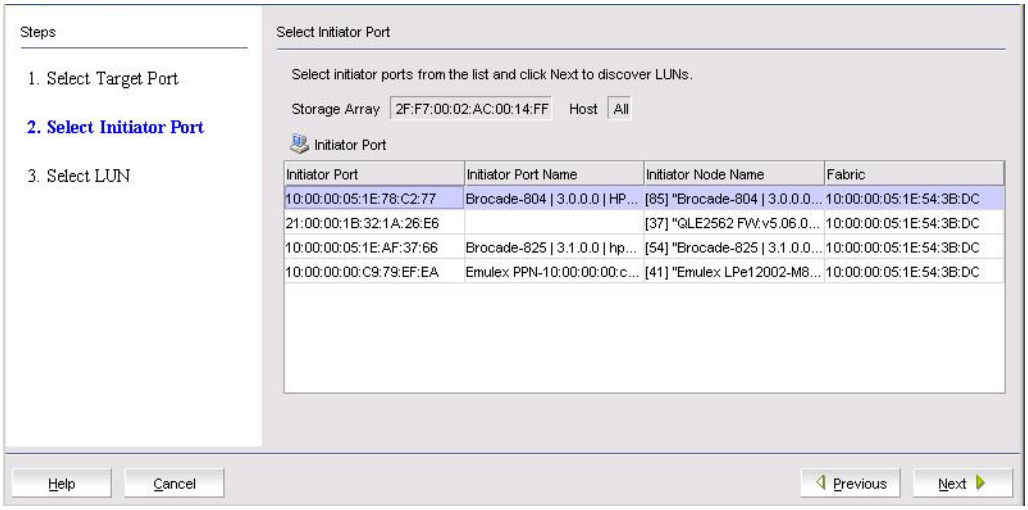

The dialog box is used to select an initiator port when configuring multiple I/O paths to a disk LUN. The dialog box contains the following information:

- **Storage Array**: Displays the storage array that was selected from the LUN view prior to launching the wizard.
- **Host**: The host selected from the LUN view prior to launching the wizard.
- **Initiator Port** table: Lists initiator ports using the following identifiers:
  - **Initiator Port**
  - **Initiator Port Name**
  - **Initiator Node Name**
  - **Fabric**

5. Select the initiator port from the **Initiator Port** table, then click **Next**.

   LUN discovery is launched and a progress bar displays. There are four possible outcomes:

   - A message displays indicating no LUNs were discovered. Click **OK** to dismiss the message and exit the wizard.
   - A message displays indicating LUNs have been discovered, but are already configured. Click **OK** to dismiss the message and exit the wizard.
   - A message displays indicating that the target is not in the right state for discovering LUNs. Click **OK** to dismiss the message and exit the wizard.
   - The **Select LUN** dialog box displays, which lists discovered LUNs that are available. (Refer to Figure 363.)

FIGURE 363 Select LUN dialog box

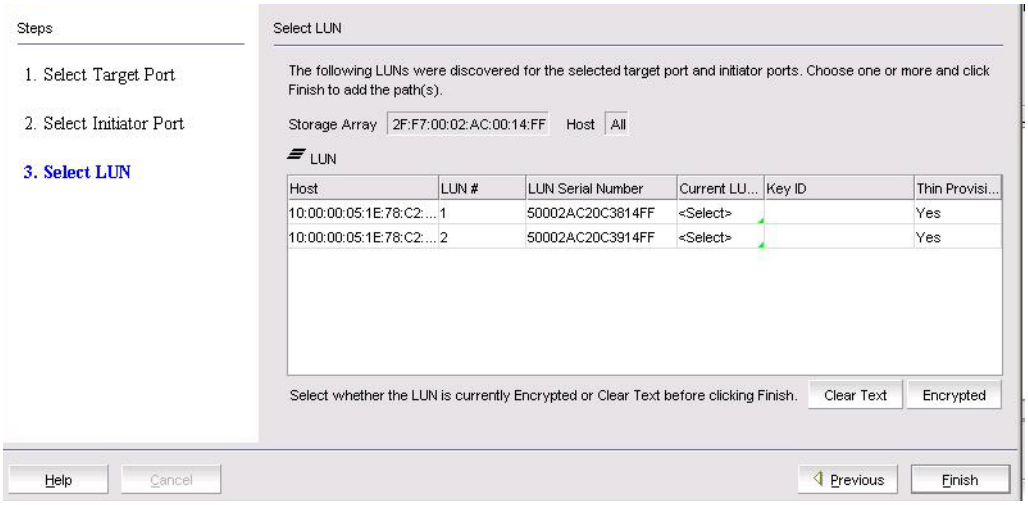

The dialog box is used to select a LUN when configuring multiple I/O paths to a disk LUN. The dialog box contains the following information:

- **Storage Array**: The storage array selected from the LUN view prior to launching the **Add New Path** wizard.
- **Host:** The host elected from the LUN view prior to launching the **Add New Path** wizard.
- **LUN** table: Available LUNs identified by the following:
  - **Host**
  - **LUN Number**
  - **LUN Serial Number**
  - **Current LUN State: Options are Encrypted, which is automatically selected if the LUN has a key ID; Clear Text, and <select>** for LUNs without a key ID. User selection is required.
- **Key ID**: Identifies the key ID for discovered LUNs.
- **Thin Provision LUN**: Identifies if the new LUN is a thin provisioned LUN.Options are **Yes**, **No**, **Unknown**., or **Not Applicable**.

  NOTE
  Thin provision support is limited to Brocade-tested storage arrays. The thin provisioned LUN status will be displayed as **Yes** for supported storage arrays only.

- **New LUN:** Displayed only if remote replication is enabled.

6. Select the LUN from **LUN** list.
7. Set the **Current LUN State** as required. If the LUN already has an existing key ID, the **Current LUN State** field is automatically set to **Encrypted**. You can accept the automatically assigned state or change this value if desired.
8. If **REPL Support** was enabled by the **Configure Switch Encryption wizard**, a **New LUN** check box is presented and enabled by default. If this LUN is to be paired with another LUN for SRDF data replication, the **New LUN** option must be enabled. Refer to "Metadata requirements and remote replication" for information about how this option works. If **REPL support** was not enabled, this check box is not displayed.
9. Click **Finish**.

   The new LUN path is added to the **Encryption Disk LUN View** table.
10. Click **OK** on the LUN view to commit the operation.

NOTE
With the introduction of Fabric OS v7.1.0, the maximum number of uncommitted configuration changes per disk LUN (or maximum paths to a LUN) is 512 transactions. The 512 LUN operations can be for the same LUN or be subjected to 25 distinct LUNs. This change of restriction in commit limit is applicable when using the Management application only. Earlier Fabric OS versions allowed a maximum of 25 uncommitted changes per disk LUN. Adding or modifying more than 25 paths on the same LUN is not recommended unless the LUN is encrypted.

In environments where there are multiple paths to the same LUNs, it is critical that the same LUN policies are configured on all instances of the LUN. Be sure to return to the **Encryption Disk LUN View** dialog box to determine if there are configuration mismatches. Check under **Encryption Mode** for any entries showing **Mismatch**. To correct the mismatch, click the incorrect mode to display the options, then select the correct mode. (Refer to Figure 364.)

FIGURE 364 Correcting an encryption mode mismatch

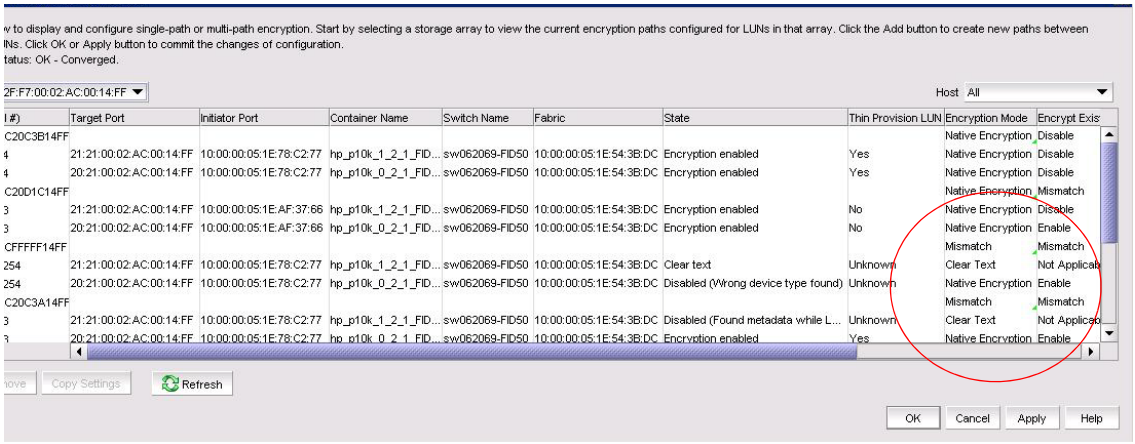

When you correct a policy on a LUN, it is automatically selected for all paths to the selected LUN. When you modify LUN policies, a **Modify** icon displays to identify the modified LUN entry.

11. Click **OK** or **Apply** to apply the changes.

Configuring storage arrays

The storage array contains a list of storage ports that will be used later in the LUN centric view. You must assign storage ports from the same storage array for multi-path I/O purposes. On the LUN centric view, storage ports in the same storage array are used to get the associated CryptoTarget containers and initiators from the database. Storage ports that are not assigned to any storage array but are within the fabrics of the encryption group will be listed as a single target port on the LUN centric view. Storage Arrays are configured using the **Storage Port Mapping** dialog box. You will need to:

1. Configure target and zone initiator ports in the same zone in order for the target container to come online and discover LUNs in the storage system.
2. Create CryptoTarget containers for each target port in the storage array from the Target Container dialog box. Add initiator ports to the container. You must create target containers for those target ports in the configured storage arrays or unassigned target ports before mapping any LUN on the LUN centric view. If you do not create the container, LUN discovery will not function.

For more detailed information on creating a CryptoTarget container, refer to the chapter describing storage arrays in this administrator's guide.

Remote replication LUNs

The Symmetrix Remote Data Facility (SRDF) transmits data that is being written to both a local Symmetrix array and a remote symmetrix array. The replicated data facilitates a fast switchover to the remote site for data recovery.

SRDF supports the following methods of data replication:

- Synchronous Replication provides real-time mirroring of data between the source Symmetrix and the target Symmetrix systems. Data is written simultaneously to the cache of both systems in real time before the application I/O is completed, thus ensuring the highest possible data availability.
- Semi-Synchronous Replication writes data to the source system, completes the I/O, then synchronizes the data with the target system. Since the I/O is completed prior to synchronizing data with the target system, this method provides an added performance advantage. A second write will not be accepted on a Symmetrix source device until its target device has been synchronized.
- Adaptive Copy Replication transfers data from the source devices to the remote devices without waiting for an acknowledgment. This is especially useful when transferring large amounts of data during data center migrations, consolidations, and in data mobility environments.
- Asynchronous Replication places host writes into chunks and then transfers an entire chunk to the target system. When a complete chunk is received on the target system, the copy cycle is committed. If the SRDF links are lost during data transfer, any partial chunk is discarded, preserving consistency on the target system. This method provides a consistent point-in-time remote image that is not far behind the source system and results in minimal data loss if there is a disaster at the source site.

SRDF pairs

Remote replication is implemented by establishing a synchronized pair of SRDF devices connected by FC or IP links. A local source device is paired with a remote target device while data replication is taking place. While the SRDF devices are paired, the remote target device is not locally accessible for read or write operations. When the data replication operation completes, the pair may be split to enable normal read/write access to both devices. The pair may be restored to restore the data on the local source device.

Figure 365 shows the placement of encryption switches in an SRDF configuration. When encryption is enabled for the primary LUN, encrypted data written by the local application server to the primary LUN is replicated on the secondary LUN. The data is encrypted using a DEK that was generated on the local encryption switch and stored on the local DPM key vault. When each site has an independent key vault, as shown in Figure 365, the key vaults must be synchronized to ensure the availability of the DEK at the remote site. Refer to DPM user documentation for information about how to synchronize the key vaults. Both sites may share the same key vault, which eliminates the need for synchronization across sites. Depending on distance between sites, sharing a key vault may add some latency when retrieving a key.

FIGURE 365Basic SRDF configuration with encryption switches

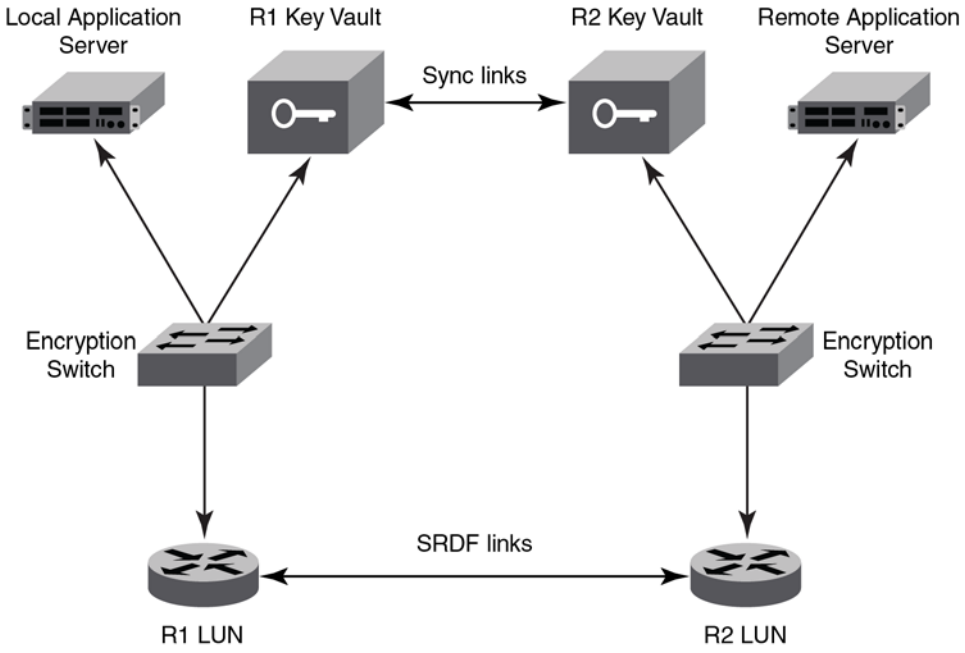

Metadata requirements and remote replication

When the metadata and key ID are written, the primary metadata on blocks 1–16 is compressed and encrypted. However, there are scenarios whereby these blocks cannot be compressed, and the metadata is not written to the media. If blocks 1–16 are not compressible on the local source device and metadata is not written, obtaining the correct DEK for the remote target device becomes problematic. This problem is avoided by reserving the last three blocks of the LUN for a copy of the metadata. These blocks are not exposed to the host initiator. When a host reads the capacity of the LUN, the size reported is always three blocks less than the actual size. The behavior is enforced by selecting the **New LUN** check box on the **Select LUN** screen of the **Add New Path** wizard when adding LUNs for an SRDF pair (for example, R1 and R2 in Figure 365).

Note the following when using the **New LUN** option:

- Both LUNs that form an SRDF pair must be added to their containers using the **New LUN** option.
- For any site, all paths to a given SRDF device must be configured with the **New LUN** option.
- All LUNs configured with the **New LUN** option will report three blocks less than the actual size when host performs READ CAPACITY 10/READ CAPACITY 16.
- If a LUN is added with the **New LUN** option and with encryption enabled, it will always have valid metadata even if blocks 1–16 of the LUN is not compressible.
- LUNs configured as cleartext must also be added with the **New LUN** option if they are part of an SRDF pair. This is to handle scenarios whereby the LUN policy is changed to encrypted at some later time, and to verify formation of DEK clusters and LUN accessibility prior to enabling encryption for the LUN. When cleartext LUNs are configured with the **New LUN** option, no metadata is written to the last three blocks, but will still report three blocks less than the actual size when host performs READ CAPACITY 10/READ CAPACITY 16.
- The **New LUN** option is used only if a DPM key vault is configured for the encryption group.
- The **New LUN** option can be used only if replication is enabled for the encryption group.

- If the local LUN contains host data, configuring it with the **New LUN** option will cause the data on the last three blocks of the LUN to be lost. Before using the **New LUN** option, you must migrate the contents of the LUN to another LUN that is larger by at least three blocks. The new, larger LUN can then be used when creating the SRDF pair. The remote LUN of the SRDF pair must be of the same size. The original smaller LUN with user data can be decommissioned.

Adding target tape LUNs for encryption

You can configure a Crypto LUN by adding the LUN to the CryptoTarget container and enabling the encryption property on the Crypto LUN. You must add LUNs manually. After you add the LUNs, you must specify the encryption settings.

When configuring a LUN with multiple paths, the same LUN policies must be configured on all paths to the LUN. If there are multiple paths to the same physical LUNs, then the LUNs are added to multiple target containers (one target per storage device port).

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table that contains the storage device to be configured, then select **Group/Switch/Engine > Targets** from the menu task bar.

   NOTE
   You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   The **Encryption Targets** dialog box displays. (Refer to Figure 366.) Initially, the table is empty. You must add LUNs manually.

FIGURE 366 Encryption Targets dialog box

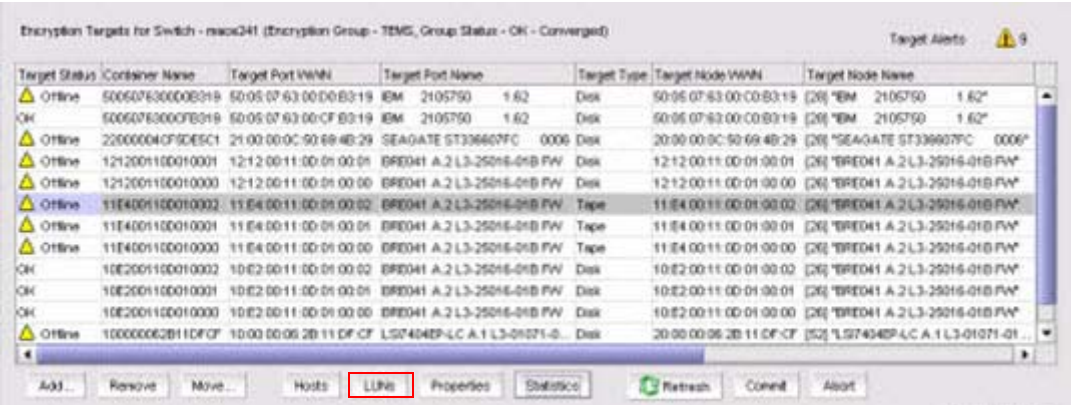

3. Select a target tape storage device from the **Encryption Targets** table, then click **LUNs**.

   The **Encryption Target Tape LUNs** dialog box displays. (Refer to Figure 367.)

FIGURE 367 Encryption Target Tape LUNs dialog box

4. Click **Add**.

   The **Add Encryption Target Tape LUNs** dialog box displays. (Refer to Figure 368.) A table of all LUNs in the storage device that are visible to hosts is displayed. LUNs are identified by the **Host** world wide name, **LUN** number, **Volume Label Prefix** number, and **Enable Write Early ACK** and **Enable Read Ahead** status. The LUN numbers may be different for different hosts.

FIGURE 368 Add Encryption Target Tape LUNs dialog box

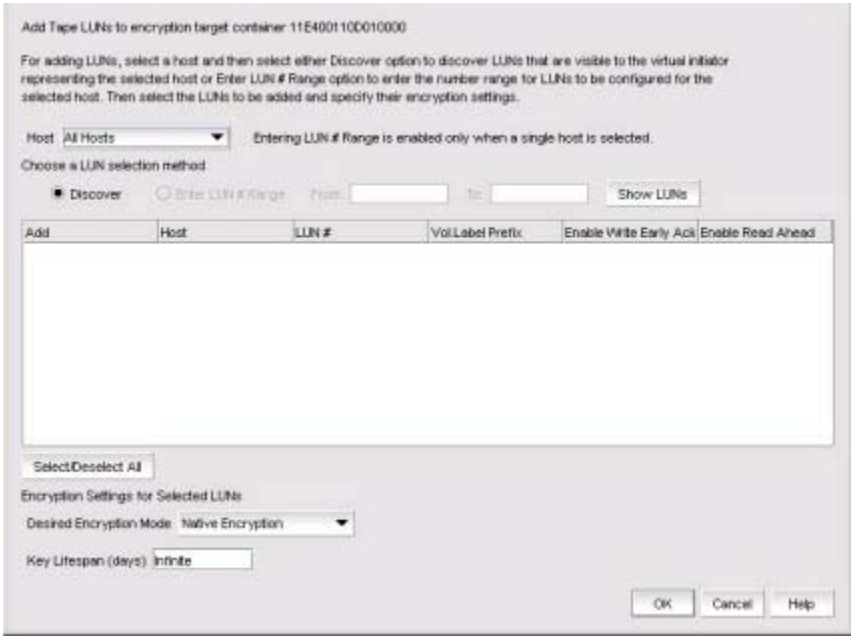

5. Select a host from the **Host** list.

   Before you encrypt a LUN, you must select a host, then either discover LUNs that are visible to the virtual initiator representing the selected host, or enter a range of LUN numbers to be configured for the selected host.

   When you select a specific host, only the LUNs visible to that host are displayed. If you select **All Hosts**, LUNs visible to all configured hosts are displayed. If a LUN is visible to multiple hosts, it is listed once for each host.

6. Choose a LUN to be added to an encryption target container using one of the two following methods:

- **Discover**: Identifies the exposed logical unit number for a specified initiator. If you already know the exposed LUNs for the various initiators accessing the LUN, you can enter the range of LUNs using the alternative method.
- **Enter a LUN number range**: Allows you to enter a **From** value and a **To** value to manually enter the logical unit numbers for the selected host(s).

7. Click **Show LUNs**.

   The LUN needed for configuring a Crypto LUN is the LUN that is exposed to a particular initiator.

   The table displays the following information:

   - **Host**: The host on which the LUN is visible.
   - **LUN #**: The logical unit's number.
   - **Vol. Label Prefix**: *Optional*. The user-supplied tape volume label prefix to be included in tape volume labels generated b the switch for encrypted tapes.
   - **Enable Write Early Ack**: When selected, enables tape write pipelining on this tape LUN. Use this option to speed long serial writes to tape, especially for remote backup operations.
   - **Enable Read Ahead**: When selected, enables read pre-fetching on this tape LUN. Use this option to speed long serial read operations from tape, especially for remote restore operations.

     NOTE
     The **Select/Deselect All** button allows you to select or deselect all available LUNs.

8. Select the desired encryption mode. Options are: **Native Encryption**, **DF-Compatible Encryption, and Cleartext**.

   - If you change a LUN policy from **Native Encryption** or **DF-Compatible Encryption** to **Clear Text**, you disable encryption.
   - The LUNs of the target that are not enabled for encryption must still be added to the CryptoTarget container with the **Clear Text** encryption mode option.

     NOTE
     The rekeying interval can only be changed for disk LUNs. For tape LUNs, expiration of the rekeying interval simply triggers the generation of a new key to be used on future tape volumes. Tapes that are already made are not rekeyed. To rekey a tape, you need to read the tape contents using a host application that decrypts the tape contents using the old key, then rewrite the tape, which re-encrypts the data with the new key.

9. Set the **Key Lifespan** setting, then click **OK**.

   The selected tape LUNs are added to the encryption target container.

Moving targets

The **Move Targets** dialog box is used to redistribute which engine encrypts which targets. It is also useful for transferring all targets to another engine before replacing or removing engine hardware. Moving targets to another engine may be done while traffic is flowing between the host and target. Traffic is interrupted for a short time but resumes before the host applications are affected.

1. Select **Configure > Encryption**.

   The **Encryption Center** dialog box displays.

2. Select one or more encryption engines from the **Encryption Center Devices** table, then select **Engine > Targets** from the menu task bar. The encryption engine must be in the same group and same fabric.

   The **Encryption Targets** dialog box displays.

3. Select one or more targets in the Encryption Targets dialog and click **Move**.

   The **Move Targets** dialog box is displayed.

4. Select an encryption engine, then click **OK** to close the dialog and start the move operation.

Configuring encrypted tape storage in a multi-path environment

This example assumes one host is accessing one storage device using two paths:

- The first path is from Host Port A to Target Port A, using Encryption Engine A for encryption.
- The second path is from Host Port B to Target Port B, using Encryption Engine B for encryption.

Encryption Engines A and B are in switches that are already part of Encryption Group *X*.

The following procedure is used to configure this scenario using the Management application.

1. Configure Host Port A and Target Port A in the same zone by selecting **Configure > Zoning** from the BNA main menu.
2. Configure Host Port B and Target Port B in the same zone by selecting **Configure > Zoning** from the BNA main menu.
3. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box (Refer to Figure 269 on page 620).
4. Click **View Groups** to display the encryption groups if groups are not already displayed.
5. Select Encryption Group *X*, then click **the Targets icon**.
6. From the **Encryption Targets** dialog box, click **Add** to open the **Configure Storage Encryption** wizard. Use the wizard to create a target container for Encryption Engine A with Target Port A and Host Port A.
7. Repeat Step 6 to create a target container for Encryption Engine B with Target Port B and Host Port B.

   Up to this point, the Management application has been automatically committing changes as they are made. The targets and hosts are now fully configured; only the LUN configuration remains.

8. In the **Encryption Targets** dialog box, select Target Port A, click **LUNs**, then click **Add**. Select the LUNs to be encrypted and the encryption policies for the LUNs.
9. In the **Encryption Targets** dialog box, select Target Port B, click **LUNs**, then click **Add**. Select the LUNs to be encrypted and the encryption policies for the LUNs, making sure that the encryption policies match the policies specified in the other path.
10. Click **Commit** to make the LUN configuration changes effective in both paths simultaneously.

the Management application does not automatically commit LUN configuration changes. You must manually commit any LUN configuration changes, even in non-multi–path environments. Committing LUN configuration changes manually allows the matching changes made in a multi-path environment to be committed together, preventing cases where one path may be encrypting and another path is not, thus causing corrupted data.

NOTE
There is a limit of 16 uncommitted tape LUN configuration changes. When adding more than 8 LUNs in a multi-path environment, repeat step 8 and step 9 above, adding only 8 LUNs to each target container at a time. Each commit operation will commit 16 LUNs, 8 in each path.

Tape LUN write early and read ahead

The tape LUN write early and read ahead feature uses tape pipelining and prefetch to speed serial access to tape storage. These features are particularly useful when performing backup and restore operations, especially over long distances.

You can enable tape LUN write early and read ahead while adding the tape LUN for encryption, or you can enable or disable these features after the tape LUN has been added for encryption.

Enabling and disabling tape LUN write early and read ahead

To enable or disable tape LUN write early and read ahead, follow these steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table, then select **Group/Switch/Engine > Targets** from the menu task bar.

   NOTE
   You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   The **Encryption Targets** dialog box displays. (Refer to Figure 369.)

FIGURE 369 Encryption Targets dialog box

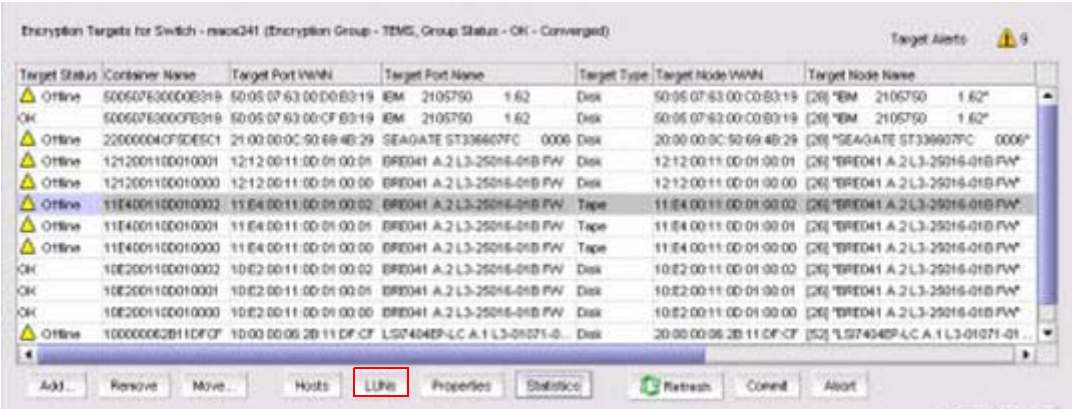

3. Select a target tape storage device from the table, then click **LUNs**.

   The **Encryption Target Tape LUNs** dialog box displays. (Refer to Figure 370.)

FIGURE 370Encryption Target Tape LUNs dialog box - Setting tape LUN read ahead and write early

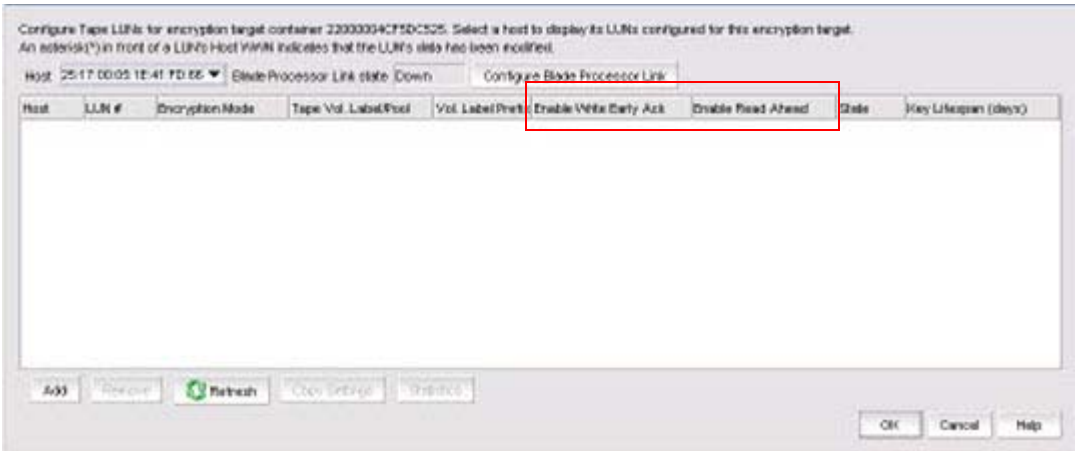

4. In the **Enable Write EarlyAck** and **Enable Read Ahead** columns, when the table is populated, you can set these features as desired for each LUN:
   - To enable write early for a specific tape LUN, select **Enable Write Early Ack** for that LUN.
   - To enable read ahead for a specific LUN, select **Enable Read Ahead** for that LUN.
   - To disable write early for a specific tape LUN, deselect **Enable Write Early Ack** for that LUN.
   - To disable read ahead for a specific LUN, deselect **Enable Read Ahead** for that LUN.
5. Click **OK**.
6. Commit the changes on the related crypto target container:
   a. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
   b. Select a group, switch, or engine from the **Encryption Center Devices** table that contains the storage device to be configured, then select **Group/Switch/Engine > Targets** from the menu task bar.

      NOTE
      You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   c. Select the appropriate crypto target container, then click **Commit**.

Tape LUN statistics

This feature enables you to view and clear statistics for tape LUNs. These statistics include the number of compressed blocks, uncompressed blocks, compressed bytes and uncompressed bytes written to a tape LUN.

The tape LUN statistics are cumulative and change as the host writes more data on tape. You can clear the statistics to monitor compression ratio of ongoing host I/Os.

The encryption management application allows you to select tape LUN from either a tape LUN container through the **Encryption Targets** dialog box, or from the **Target Tape LUNs** dialog box.

Viewing and clearing tape container statistics

You can view LUN statistics for an entire crypto tape container or for specific LUNs.

To view or clear statistics for tape LUNs in a container, follow these steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices** table, then select **Group > Targets** from the menu task bar.

   The **Encryption Targets** dialog box displays. (Refer to Figure 371.) A list of the configured CryptoTarget containers is displayed.

FIGURE 371 Encryption Targets dialog box

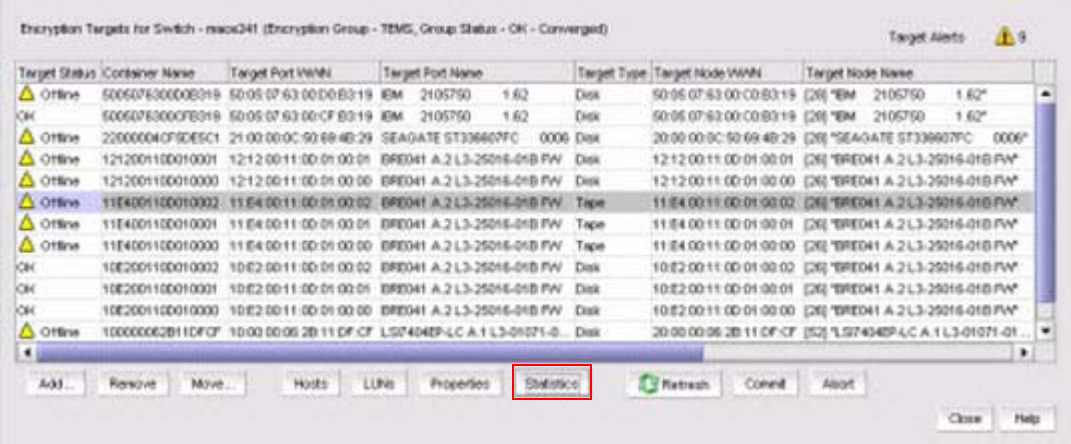

3. Select **Tape as** the container of type for which to display or clear statistics, then click **Statistics**.

   The **Tape LUN Statistics** dialog box displays. (Refer to Figure 372.) A list of the statistics for all LUNs that are members of the selected tape container is displayed.

FIGURE 372 Tape LUN Statistics dialog box

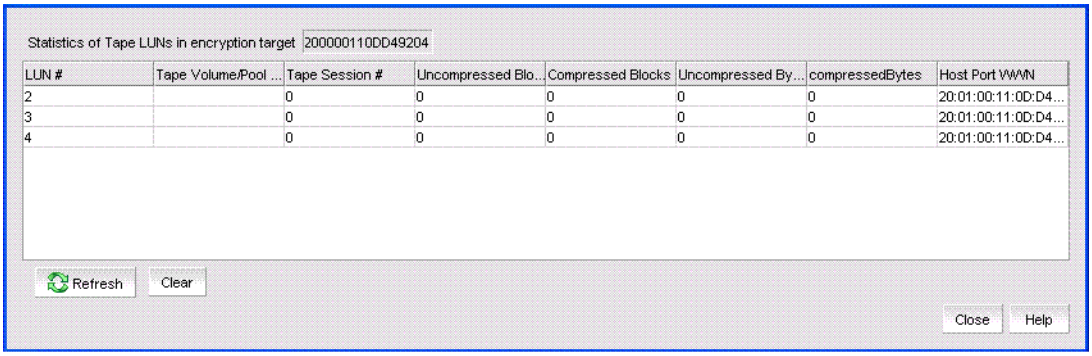

The dialog box contains the following information:

- **LUN #:** The number of the logical unit for which statics are displayed.
- **Tape Volume/Pool:** The tape volume label of the currently-mounted tape, if a tape session is currently in progress.
- **Tape Session #:** The number of the ongoing tape session.
- **Uncompressed blocks:** The number of uncompressed blocks written to tape.
- **Compressed blocks:** The number of compressed blocks written to tape.

- **Uncompressed Bytes:** The number of uncompressed bytes written to tape.
- **Compressed Bytes:** The number of compressed bytes written to tape.
- **Host Port WWN:** The WWN of the host port that is being used for the write operation.
- A **Refresh** button updates the statistics on the display since the last reset.
- A **Clear** button resets all statistics in the display.

4. To clear the tape LUN statistics for all member LUNs for the container, click **Clear**, then click **Yes** to confirm.

To view statistics for specific LUNs:

1. Select a tape container, then click **LUNs**.
2. From the **Target Tape LUNs** dialog box, select the LUNs you want to monitor.

Viewing and clearing tape LUN statistics for specific tape LUNs

To view or clear statistics for tape LUNs in a container, complete these steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table that contains the storage device to be configured, then select **Group/Switch/Engine > Targets** from the menu task bar.

   NOTE
   You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   The **Encryption Targets** dialog box displays. (Refer to Figure 369.)

3. Select a tape target storage device, then click **LUNs**.

   The **Target Tape LUNs** dialog box displays. (Refer to Figure 373.) A list of the configured tape LUNs is displayed.

FIGURE 373 Target Tape LUNs dialog box

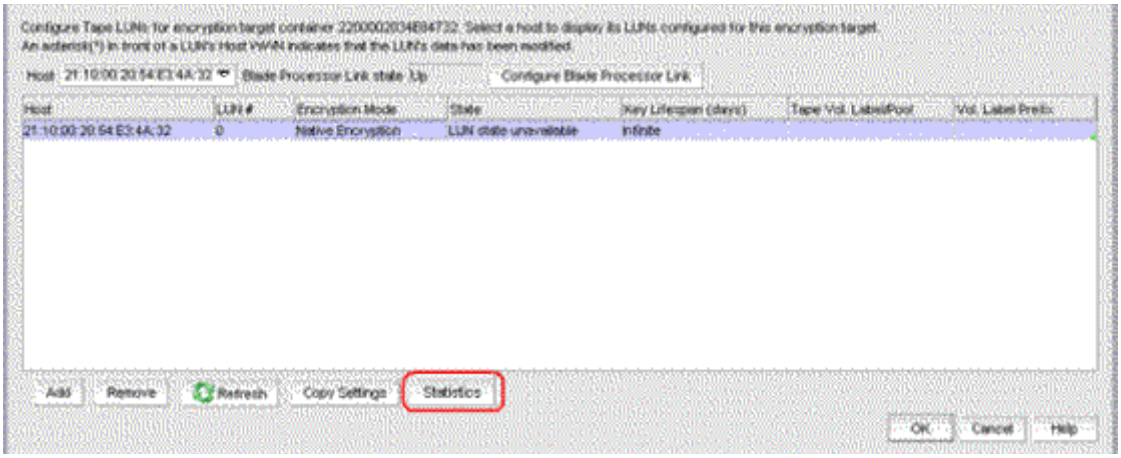

4. Select the LUN or LUNs for which to display or clear statistics, then click **Statistics**.

   The **Tape LUN Statistics** dialog box displays. (Refer to Figure 374.) The statistic results based on the LUN or LUNs you selected is displayed. Tape LUN statistics are cumulative.

FIGURE 374 Tape LUN Statistics dialog box

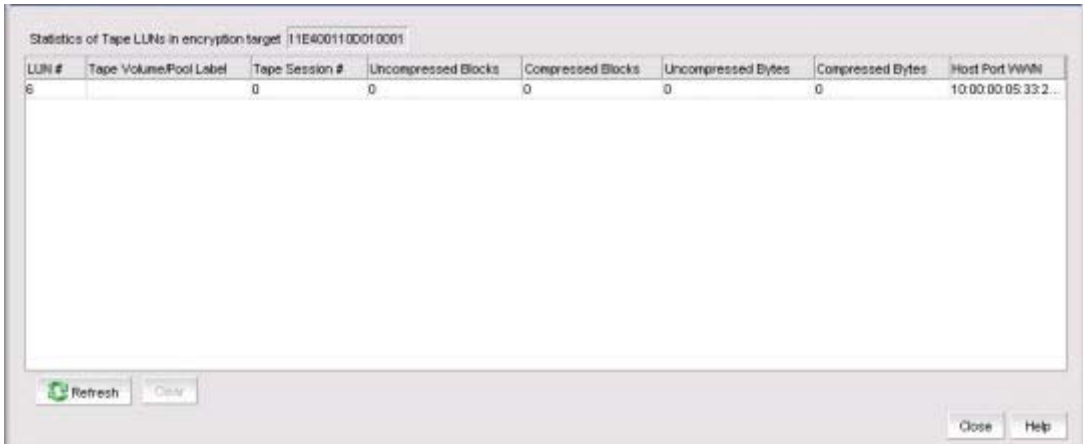

The dialog box contains the following information:

- **LUN #:** The number of the logical unit for which statics are displayed.
- **Tape Volume/Pool:** The tape volume label of the currently-mounted tape, if a tape session is currently in progress.
- **Tape Session #:** The number of the ongoing tape session.
- **Uncompressed blocks:** The number of uncompressed blocks written to tape.
- **Compressed blocks:** The number of compressed blocks written to tape.
- **Uncompressed Bytes:** The number of uncompressed bytes written to tape.
- **Compressed Bytes:** The number of compressed bytes written to tape.
- **Host Port WWN:** The WWN of the host port that is being used for the write operation.
- A **Refresh** button updates the statistics on the display since the last reset.
- A **Clear** button resets all statistics in the display.

5. Do either of the following:

   a. Click **Clear** to clear the tape LUN statistics, then click **Yes** to confirm.

   b. Click **Refresh** to view the current statistics cumulative since the last reset.

Viewing and clearing statistics for tape LUNs in a container

To view or clear statistics for tape LUNs in a container, follow these steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table that contains the storage device to be configured, then select **Group/Switch/Engine > Targets** from the menu task bar.

   NOTE
   You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   The **Encryption Targets** dialog box displays. (Refer to Figure 375.) A list of configured CryptoTarget containers is displayed.

FIGURE 375 Encryption Targets dialog box

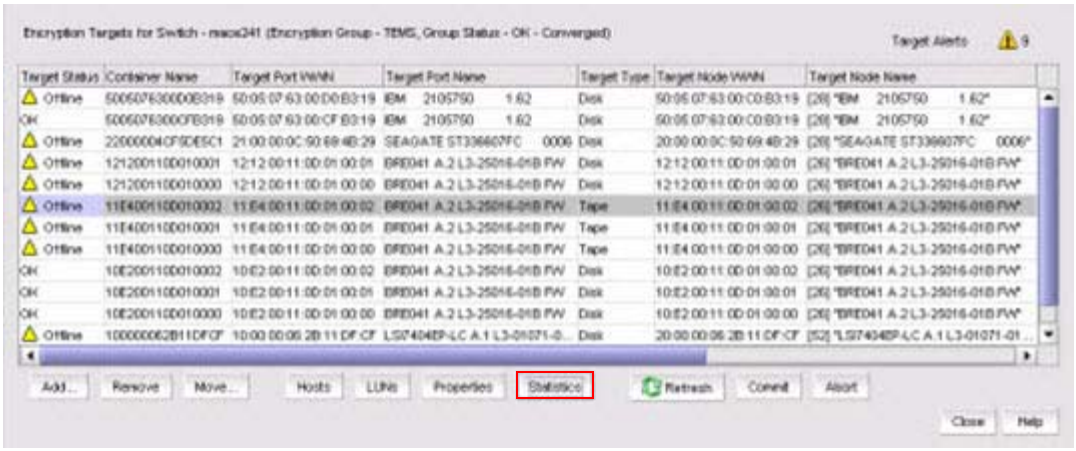

3. Select **Tape as** the container of type for which to display or clear statistics, then click **Statistics**.

   The **Tape LUN Statistics** dialog box displays. (Refer to Figure 376.) The statistics for all LUNs that are members of the selected tape container are displayed.

FIGURE 376 Tape LUN Statistics dialog box

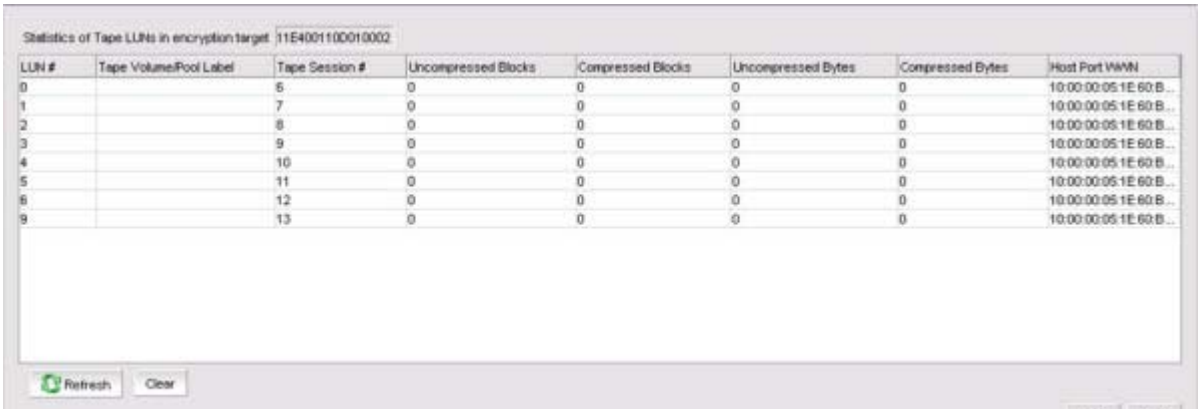

The dialog box contains the following information:

- **LUN #:** The number of the logical unit for which statics are displayed.
- **Tape Volume/Pool:** The tape volume label of the currently-mounted tape, if a tape session is currently in progress.
- **Tape Session #:** The number of the ongoing tape session.
- **Uncompressed blocks:** The number of uncompressed blocks written to tape.
- **Compressed blocks:** The number of compressed blocks written to tape.
- **Uncompressed Bytes:** The number of uncompressed bytes written to tape.
- **Compressed Bytes:** The number of compressed bytes written to tape.
- **Host Port WWN:** The WWN of the host port that is being used for the write operation.

4. Do either of the following:

   - Click **Clear** to clear the tape LUN statistics for member LUNs in the container, then click **Yes** to confirm.
   - Click **Refresh** to update the tape LUN statistics on the display.

Encryption engine rebalancing

If you are currently using encryption and running Fabric OS 6.3.x or earlier, you are hosting tape and disk target containers on different encryption switches or blades. Beginning with Fabric OS 6.4, disk and tape target containers can be hosted on the same switch or blade. Hosting both disk and tape target containers on the same switch or blade might result in a drop in throughput, but it can reduce cost by reducing the number of switches or blades needed to support encrypted I/O in environments that use both disk and tape.

The throughput drop can be mitigated by re-balancing the tape and disk target containers across the encryption engine. This ensures that the tape and disk target containers are distributed within the encryption engine for maximum throughput.

All nodes within an encryption group must be upgraded to Fabric OS 6.4 or later to support hosting disk and tape target containers on the same encryption engine. If any node within an encryption group is running an earlier release, disk and tape containers must continue to be hosted on separate encryption engines.

During rebalancing operations, be aware of the following:

- You might notice a slight disruption in Disk I/O. In some cases, manual intervention may be needed.
- Backup jobs to tapes might need to be restarted after rebalancing is completed.

To determine if rebalancing is recommended for an encryption engine, check the encryption engine properties. Beginning with Fabric OS 6.4, a field is added that indicates whether or not rebalancing is recommended.

You might be prompted to rebalance during the following operations:

- When adding a new disk or tape target container.
- When removing an existing disk or tape target container.
- After failover to a backup encryption engine in an HA cluster.
- After a failed encryption engine in an HA cluster is recovered, and failback processing has occurred.

Rebalancing an encryption engine

To re-balance an encryption engine, complete the following steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an engine, then select **Engine > Re-Balance** from the menu task bar.

   A warning message displays, noting the potential disruption of disk and tape I/O, and that the operation may take several minutes.
3. Click **Yes** to begin rebalancing.

Master keys

Master keys belong to the group and are managed from **Group Properties**.

When an opaque key vault is used, a master key is used to encrypt the data encryption keys. The master key status indicates whether a master key is used and whether it has been backed up. Encryption is not allowed until the master key has been backed up.

Only the active master key can be backed up, and multiple backups are recommended. You can back up or restore the master key to the key vault, to a file, or to a recovery card set. A recovery card set is set of smart cards. Each recovery card holds a portion of the master key. The cards must be gathered and read together from a card reader attached to a PC running the Management application to restore the master key.

Although it is generally not necessary to create a new master key, you might be required to create one due to the following:

- The previous master key has been compromised.
- Corporate policy might require a new master key every year for security purposes.

With regard to DPM, any DEK in the key vault that is either compromised, or needs to be deactivated or destroyed, must first undergo the decommissioning procedure. For more information, refer to "Disk device decommissioning" on page 751.

When you create a new master key, the former active master key automatically becomes the alternate master key.

The new master key cannot be used (no new data encryption keys can be created, so no new encrypted LUNs can be configured), until you back up the new master key. After you have backed up the new master key, it is strongly recommended that all encrypted disk LUNs be rekeyed. Rekeying causes a new data encryption key to be created and encrypted using the new active master key, thereby removing any dependency on the old master key. Refer to "Creating a new master key" on page 749 for more information.

Master key actions are disabled if they are unavailable. For example:

- The user does not have Storage Encryption Security permissions.
- The Group Leader is not discovered or managed by the Management application.

NOTE
It is important to back up the master key because if the master key is lost, none of the data encryption keys can be restored and none of the encrypted data can be decrypted.

Active master key

The active master key is used to encrypt newly created data encryption keys (DEKs) prior to sending them to a key vault to be stored. You can restore the active master key under the following conditions:

- The active master key has been lost, which happens if all encryption engines in the group have been zeroized or replaced with new hardware at the same time.
- You want multiple encryption groups to share the same active master key. Groups should share the same master key if the groups share the same key vault and if tapes (or disks) are going to be exchanged regularly between the groups.

Alternate master key

The alternate master key is used to decrypt data encryption keys that were not encrypted with the active master key. Restore the alternate master key for the following reasons:

- To read an old tape that was created when the group used a different active master key.
- To read a tape (or disk) from a different encryption group that uses a different active master key.

Master key actions

NOTE
Master keys belong to the group and are managed from Group Properties.

Master key actions are as follows:

- **Backup master key**: Enabled any time a master key exists. Selecting this option launches the **Backup Master Key for Encryption Group** dialog box.

  You can back up the master key to a file, to a key vault, or to a smart card. You can back up the master key multiple times to any of these media in case you forget the passphrase you originally used to back up the master key, or if multiple administrators each needs a passphrase for recovery. Refer to the following procedures for more information:

  - "Saving the master key to a file" on page 745
  - "Saving a master key to a key vault" on page 746
  - "Saving a master key to a smart card set" on page 746

  You must back up the master key when the status is **Created but not backed up**.

- **Restore master key**: Enabled when no master key exists or the previous master key has been backed up. This option is also enabled when using a DPM key vault.

  When this option is selected, the **Restore Master Key for Encryption Group** dialog box displays, from which you can restore a master key from a file, key vault, or smart card set. Refer to the following procedures for more information:

  - "Restoring a master key from a file" on page 747
  - "Restoring a master key from a key vault" on page 748
  - "Restoring a master key from a smart card set" on page 748

- **Create new master key**: Enabled when no master key exists, or the previous master key has been backed up. Refer to "Creating a new master key" on page 749.

  You must create a new master key when the status is **Required but not created**.

  NOTE
  If a master key was not created, **Not Used** is displayed as the status and the **Master Key Actions** list is unavailable. In this case, you must create a new master key. Additional master key statuses are **Backed up but not propagated** and **Created and backed up**.

Saving the master key to a file

Use the following procedure to save the master key to a file.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices table**, then select **Group > Security** from the menu task bar.

   The **Encryption Group Properties** dialog box displays with the **Security** tab selected.

3. Select **Backup Master Key** as the **Master Key Action**.

   The **Master Key Backup** dialog box displays, but only if the master key has already been generated.

4. Select **File** as the **Backup Destination**.
5. Enter a file name, or browse to the desired location.

6. Enter the passphrase, which is required for restoring the master key. The passphrase can be between eight and 40 characters, and any character is allowed.
7. Re-enter the passphrase for verification, then click **OK**.

   ATTENTION
   Save the passphrase. This passphrase is required if you ever need to restore the master key from the file.

Saving a master key to a key vault

Use the following procedure to save the master key to a key vault.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices table**, then select **Group > Security** from the menu task bar.

   The **Encryption Group Properties** dialog box displays with the **Security** tab selected.
3. Select **Backup Master Key** as the **Master Key Action**.

   The **Backup Master Key for Encryption Group** dialog box displays.
4. Select **Key Vault** as the **Backup Destination**.
5. Enter the passphrase, which is required for restoring the master key. The passphrase can be between eight and 40 characters, and any character is allowed.
6. Re-enter the passphrase for verification, then click **OK**.

   A dialog box displays that shows the **Key ID**. The **Key ID** identifies the storage location in the key vault.
7. Store both the **Key ID** and the passphrase in a secure place. Both will be required to restore the master key in the future.
8. Click **OK**. after you have copied the **Key ID**.

Saving a master key to a smart card set

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices table**, then select **Group > Security** from the menu task bar.

   The **Encryption Group Properties** dialog box displays with the **Security** tab selected.
3. Select **Backup Master Key** as the **Master Key Action**.

   The **Backup Master Key for Encryption Group** dialog box displays.
4. Select **A Recovery Set of Smart Cards** as the **Backup Destination**.
5. Enter the recovery card set size.
6. Insert the first blank card and wait for the card serial number to appear.
7. Run the additional cards through the reader that are needed for the set. As you read each card, the card ID displays in the **Card Serial#** field. Be sure to wait for the ID to appear.
8. Enter the mandatory last name and first name of the person to whom the card is assigned.

9. Enter a Card **Password**.
10. Re-enter the password for verification.
11. Record and store the password in a secure location.
12. Click **Write Card**.

    You are prompted to insert the next card, up to the number of cards specified in step 5.
13. Repeat step 6 through step 12 for each card in the set.
14. After the last card is written, click **OK** in the **Master Key Backup** dialog box to finish the operation.

Overview of saving a master key to a smart card set

A card reader must be attached to the SAN Management application PC to save a master key to a recovery card. Recovery cards can only be written once to back up a single master key. Each master key backup operation requires a new set of previously unused smart cards.

> NOTE
> Windows operating systems do not require smart card drivers to be installed separately; the driver is bundled with the operating system. However, you must install a smart card driver for UNIX operating systems. For instructions, refer to the Installation Guide that comes with your system.

The key is divided among the cards in the card set, up to 10. The quorum of cards required to restore the master key must be less than the total number of cards in the set, and no greater than five. For example, when the master key is backed up to a set of three cards, a quorum of any two cards can be used together to restore the master key. When the master key is backed up to a set of 10 cards, a quorum size of up to five cards can be configured for restoring the master key. Backing up the master key to multiple recovery cards is the recommended and most secure option.

> NOTE
> When you write the key to the card set, be sure you write the full set without canceling. If you cancel, all previously written cards become unusable; you will need to discard them and create a new set.

Restoring a master key from a file

Use the following procedure to restore the master key from a file.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices table**, then select **Group > Security** from the menu task bar.

   The **Encryption Group Properties** dialog box displays with the **Security** tab selected.
3. Select **Restore Master Key** as the **Master Key Action**.

   The **Restore Master Key for Encryption Group** dialog box displays.
4. Choose the active or alternate master key for restoration, as appropriate.
5. Select **File** as the **Restore From** location.
6. Enter a file name, or browse to the desired location.
7. Enter the passphrase. The passphrase that was used to back up the master key must be used to restore the master key.
8. Click **OK**.

Restoring a master key from a key vault

Use the following procedure to restore the master key from a key vault:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices table**, then select **Group > Security** from the menu task bar.

   The **Encryption Group Properties** dialog box displays with the **Security** tab selected.
3. Select **Restore Master Key** as the **Master Key Action**.

   The **Restore Master Key for Encryption Group** dialog box displays.
4. Choose the active or alternate master key for restoration, as appropriate.
5. Select **Key Vault** as the **Restore From** location.
6. Enter the key ID of the master key that was backed up to the key vault.
7. Enter the passphrase. The passphrase that was used to back up the master key must be used to restore the master key.
8. Click **OK**.

Restoring a master key from a smart card set

A card reader must be attached to the SAN Management application PC to complete this procedure.

Use the following procedure to restore the master key from a set of smart cards.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices table**, then select **Group > Security** from the menu task bar.

   The **Encryption Group Properties** dialog box displays with the **Security** tab selected.
3. Select **Restore Master Key** as the **Master Key Action**.

   The **Restore Master Key for Encryption Group** dialog box displays.
4. Choose the active or alternate master key for restoration, as appropriate.
5. Select **A Recovery Set of Smart Cards** as the **Restore From** location.
6. Insert the recovery card containing a share of the master key that was backed up earlier, and wait for the card serial number to appear.
7. Enter the password that was used to create the card. After five unsuccessful attempts to enter the correct password, the card becomes locked and unusable.
8. Click **Restore**.

   You are prompted to insert the next card, if needed.
9. Repeat step 6 through step 8 until all cards in the set have been read.
10. Click **OK**.

Creating a new master key

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices table**, then select **Group > Security** from the menu task bar.

   The **Encryption Group Properties** dialog box displays with the **Security** tab selected.
3. Select **Create a New Master Key** from the list.

   A warning displays.
4. Click **Yes** to proceed.

Security settings

Security settings help you identify if system cards are required to initialize an encryption engine and also determine the number of authentication cards needed for a quorum.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices table**, then select **Group > Security** from the menu task bar.

   The **Select Security Settings** dialog box displays. The dialog box contains the following information:

   - **Quorum Cards**: Select the number of authentication cards needed for a quorum. The quorum is always set to one card less than the number of cards registered. For example, if you register three cards, the quorum needed for authentication is two.
   - **System Cards:** Determine whether or not a system card is required to initialize the encryption engine

   NOTE
   The **Select Security Settings** dialog box only sets a quorum number for authentication cards. To register authentication cards, click **Next** to display the **Authentication Cards** dialog box.

Zeroizing an encryption engine

Zeroizing is the process of erasing all data encryption keys and other sensitive encryption information in an encryption engine. You can zeroize an encryption engine manually to protect encryption keys. No data is lost because the data encryption keys for the encryption targets are stored in the key vault.

Zeroizing has the following effects:

- All copies of data encryption keys (DEKs) kept in the encryption switch or blade are erased.
- Internal public and private key pairs that identify the encryption engine are erased and the encryption switch or blade is in the FAULTY state.
- All encryption operations on this engine are stopped and all virtual initiators (VI) and virtual targets (VT) are removed from the fabric's name service.
- The key vault link key (for NetApp LKM/SSKM key vaults) or the master key (for other key vaults) is erased from the encryption engine.

  Once enabled, the encryption engine is able to restore the necessary data encryption keys from the key vault when the link key (for the NetApp Lifetime Key Management application) or the master key (for other key vaults) is restored.

- If the encryption engine was part of an HA cluster, targets fail over to the peer, which assumes the encryption of all storage targets. Data flow will continue to be encrypted.
- If there is no HA backup, host traffic to the target will fail as if the target has gone offline. The host will not have unencrypted access to the target. There will be no data flow at all because the encryption virtual targets will be offline.

NOTE
Zeroizing an engine affects the I/Os, but all target and LUN configurations remain intact. Encryption target configuration data is not deleted.

You can zeroize an encryption engine only if it is enabled (running), or disabled but ready to be enabled. If the encryption engine is not in one of these states, an error message results.

When using a **NetApp LKM/SSKM key vault,** if all encryption engines in a switch are zeroized, the switch loses the link key required to communicate with the LKM/SSKM vault. After the encryption engines are rebooted and re-enabled, you must use the CLI to create new link keys for the switch.

When using an **opaque key vault**, if all encryption engines in an encryption group are zeroized, the encryption group loses the master key required to read data encryption keys from the key vault. After the encryption engines are rebooted and re-enabled, you must restore the master key from a backup copy, or alternatively, you can generate a new master key and back it up. Restoring the master key from a backup copy or generating a new master key and backing it up indicates that all previously generated DEKs will not be decryptable unless the original master key used to encrypt them is restored.

Setting zeroization

Use the **Restore Master key** wizard from the **Encryption Group Properties** dialog box to restore the master key from a backup copy.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption engine from the **Encryption Center Devices** table, then select **Engine > Zeroize** from the menu task bar.

   A warning dialog box describes consequences and actions required to recover.
3. Click **Yes** to zeroize the encryption engine.
   - For an encryption blade: After the zeroize operation is successful, a message displays noting that the encryption blade will be powered off and powered on to make it operational again. Click **OK** to close the message. After the encryption blade is powered on, click **Refresh** in the **Encryption Center** dialog box to update the status of the encryption blade and perform any operations.
   - For an encryption switch: After the zeroization operation is successful, you are instructed to reboot the encryption switch. Click **OK** to close the message, then reboot the encryption switch. After the encryption switch is rebooted, click **Refresh** in the **Encryption Center** dialog box to update the status of the encryption switch and perform any operations.

Using the Encryption Targets dialog box

The **Encryption Targets** dialog box enables you to send outbound data that you want to store as ciphertext to an encryption device. The encryption target acts as a virtual target when receiving data from a host, and as a virtual initiator when writing the encrypted data to storage.

NOTE
The **Encryption Targets** dialog box enables you to launch a variety of wizards and other related dialog boxes.

To access the Encryption Targets dialog box, complete the following steps.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table, then select **Group/Switch/Engine > Targets** from the menu task bar.

   NOTE
   You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   The **Encryption Targets** dialog box displays. (Refer to Figure 377.) The targets currently being encrypted by the selected group, switch, or encryption engine are listed. If a group is selected, all configured targets in the group are displayed. If a switch is selected, all configured targets for the switch are displayed.

FIGURE 377 Encryption Targets dialog box

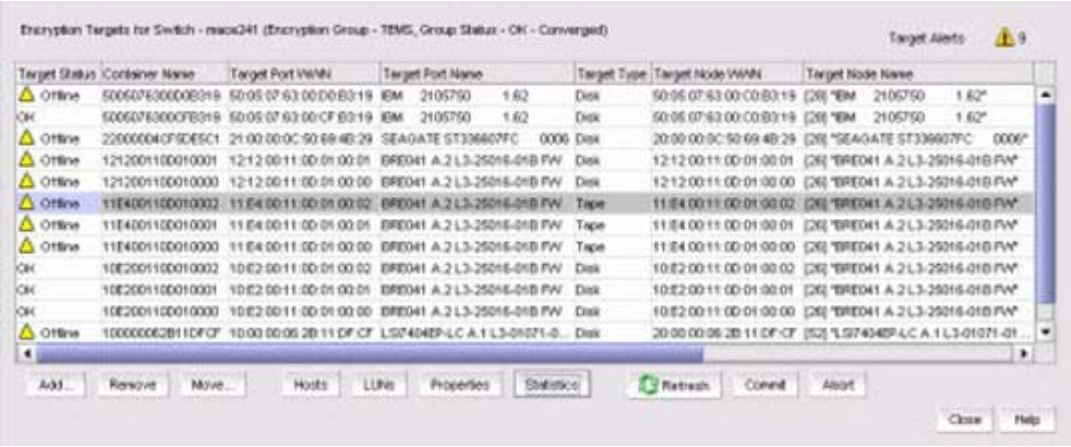

Redirection zones

It is recommended that you configure the host and target in the same zone ***before*** you configure them for encryption. Doing so creates a redirection zone to redirect the host/target traffic through the encryption engine; however, a redirection zone can only be created if the host and target are in the same zone. If the host and target are not already configured in the same zone, you can configure them for encryption, but you will still need to configure them in the same zone, which will then enable you to create the redirection zone as a separate step.

NOTE
If the encryption group is busy when you click **Commit**, you are given the option to either force the commit, or abort the changes. Click **Commit** to re-create the redirection zone.

Disk device decommissioning

A disk device needs to be decommissioned when any of the following occurs:

- The storage lease expires for an array, and devices must be returned or exchanged.
- Storage is reprovisioned for movement between departments.
- An array or device is removed from service.

In all cases, all data on the disk media must be rendered inaccessible. Device decommissioning deletes all information that could be used to recover the data, for example, information related to master key IDs and cache files.

NOTE
With regard to DPM, any DEK in the key vault that is either compromised, or needs to be deactivated or destroyed, must first undergo the decommissioning procedure.

After device decommissioning is performed, the following actions occur:

- Metadata on the LUN is erased and the reference is removed from cache on the Encryption switch.
- The LUN state is shown as decommissioned in the key vault.
- The LUN is removed from the container.

NOTE
The key IDs that were used for encrypting the data are returned.

When disk LUNs are decommissioned, the decommissioned keys are still stored on the switch. In order to delete them from the switch, you must view them from the **Decommissioned Key IDs** dialog box. (Refer to Figure 378.)

When a device decommission operation fails on the encryption Group Leader for any reason, the crypto configuration remains uncommitted until a user-initiated commit or a subsequent device decommission operation issued on the encryption Group Leader completes successfully. Device decommission operations should always be issued from a committed configuration. If not, the operation will fail with the error message **An outstanding transaction is pending in Switch/EG**. If this occurs, you can resolve the problems by committing the configuration from the encryption Group Leader.

Provided that the crypto configuration is not left uncommitted because of any crypto configuration changes or a failed device decommission operation issued on a encryption Group Leader node, this error message will not be seen for any device decommission operation issued serially on an encryption group member node. If more than one device decommission operation is attempted in an encryption group from member nodes simultaneously, this error message is transient and will go away after device decommission operation is complete. If the device decommissioning operation fails, retry the operation after some time has passed.

With the introduction of Fabric OS 7.1.0, all key vault types support the ability to decommission disk LUNs. For earlier Fabric OS versions, (for example, Fabric OS 7.0.x) the command that is used to decommission LUNs is only recognized on DPM (formerly RKM) and LKM/SSKM key vault types.

Decommissioning disk LUNs

Use the following procedure to decommission a disk LUN.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table that contains the storage device to be configured, then select **Group/Switch/Engine > Targets** from the menu task bar.

   NOTE
   You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   The **Encryption Targets** dialog box displays. (Refer to Figure 371 on page 739.)
3. Select a Target storage device from the list, then click **LUNs**.

   The **Encryption Target Disk LUNs** dialog box displays.

4. Select the LUNs associated with the device, then click **Decommission**.

   A warning message displays.

5. Click **Yes** to proceed with the decommissioning process.

   A **LUN Decommission Status** dialog box is displayed while the LUNs are being decommissioned. Click **OK** to close the dialog box.

   If a rekey operation is currently in progress on a selected LUN, a message is displayed that gives you a choice of doing a **Forced Decommission**, or to **Cancel** and try later after the rekey operation is complete.

6. To check on the progress of the decommissioning operation, click **Refresh**. When decommissioning is complete, the LUNs are removed from the **Encryption Target LUNs** table.

Displaying and deleting decommissioned key IDs

With the introduction of Fabric OS 7.1.0, the ability to decommission disk LUNs is supported on all key vault platforms. Earlier releases restricted this functionality to DPM (formerly RKM) and LKM/SSKM key vaults only.

When disk LUNs are decommissioned, the process includes the disabling of the key record in the key vault and indication that the key has been decommissioned. These decommissioned keys are still stored on the switch. You can display, copy, and delete them as an additional security measure.

The **Decommissioned Key IDs** dialog box lists Key IDs that have been decommissioned at the key vault. They should also be deleted from the switch for added security, and to create room for new key IDs. Using this dialog box, you can delete key IDs that are decommissioned at the key vault, but still stored on the switch.

In order to delete keys from the key vault, you need to know the Universal ID (UUID). To display vendor-specific UUIDs of decommissioned key IDs, complete the following procedure:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a switch from the **Encryption Center Devices table**, then select **Switch > Decommissioned key IDs from the menu task bar**.

   The **Decommissioned Key IDs** dialog box displays. (Refer to Figure 378.)

   FIGURE 378 Decommissioned Key IDs dialog box

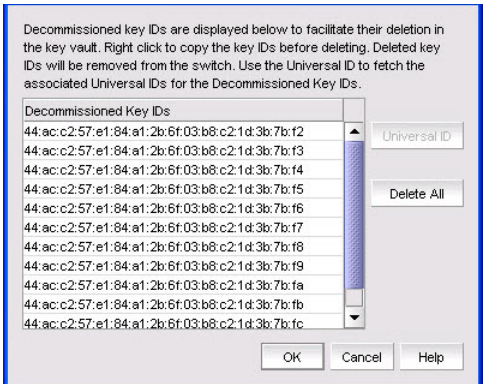

   The dialog box contains the following information:

   - **Decommissioned key IDs** that have been decommissioned at the key vault are listed in a table.

- **Universal ID** button: Launches the **Universal ID** dialog box to display the universal ID for each selected decommissioned key.

  You need to know the Universal ID (UUID) associated with the decommissioned disk LUN key IDs in order to delete keys from the key vault. You can display vendor-specific UUIDs of decommissioned key IDs. For more information, refer to "Displaying Universal IDs" on page 754.

- **Delete All** button: Deletes all of the listed decommissioned key IDs.

3. Click **Delete All** to delete the decommissioned keys from the switch. As a precaution, copy the keys to a secure location before deleting them from the switch. Right-click on an entry in the table to individually select a key ID. You may also copy or export a single row within the table or the entire table. To export the keys, right-click and select **Export,** which will export the key IDs.

Displaying Universal IDs

In order to delete keys from the key vaults, you need to know the Universal ID (UUID) associated with the decommissioned disk LUN key IDs. To display the Universal IDs, complete the following procedure:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a switch from the **Encryption Center Devices** table, then select **Switch > Decommissioned key IDs from the menu task bar**.

   The **Decommissioned Key IDs** dialog box displays. (Refer to Figure 378.)

3. Select the desired decommissioned key IDs from the **Decommissioned Key IDs** table, then click **Universal ID**.

   The **Universal IDs** dialog box displays the universal ID for each selected decommissioned key. (Refer to Figure 379.)

   FIGURE 379 Universal IDs dialog box

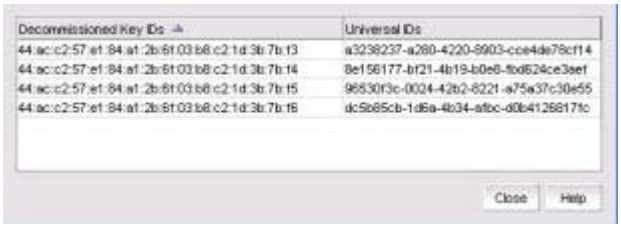

4. Click **Close**.

   NOTE
   You will need to export the decommissioned key ID to the key vault.

Rekeying all disk LUNs manually

The encryption management application allows you to perform a manual rekey operation on all encrypted primary disk LUNs and all non-replicated disk LUNs hosted on the encryption node that are in the read-write state.

Manual rekeying of all LUNs might take an extended period of time. The management application allows manual rekey of no more than 10 LUNs concurrently. If the node has more than 10 LUNs, additional LUN rekey operations will remain in the pending state until others have finished.

The following conditions must be satisfied for the manual rekey operation to run successfully:

- The node on which you perform the manual rekey operation must be a member of an encryption group, and that encryption group must have a key vault configured.
- The node must be running Fabric OS 7.0.0 or later.
- The encryption group must be in the converged state.
- The target container that hosts the LUN must be online.

In addition to providing the ability to launch manual rekey operations, the management application also enables you to monitor their progress.

Setting disk LUN Re-key All

To rekey all disk LUNs on an encryption node, complete these steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select the switch on which to perform a manual re-key from the **Encryption Center Devices** table, then select **Switch > Re-Key All** from the menu task bar. (Refer to Figure 380.)

FIGURE 380 Selecting the Re-Key All operation

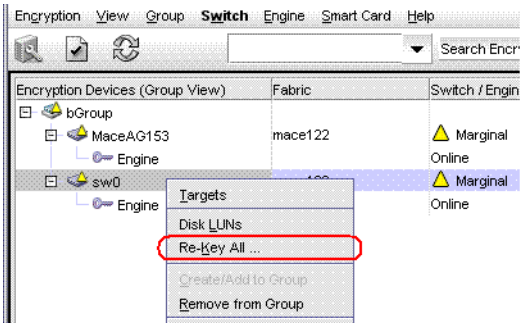

If REPL support is enabled on the encryption group, a confirmation dialog box displays, asking whether to rekey mirror LUNs.

3. Click **Yes** to includes mirror LUNs, or click **No** to exclude mirror LUNs.

   A warning message displays, requesting confirmation to proceed with the rekey operation.

4. Click **Yes**.

   Rekeying operations begin on up to 10 LUNs. If more than 10 LUNs are configured on the switch, the remaining rekey operations are held in the pending state.

5. Open the **Encryption Target Disk LUNs** dialog box to see LUNs being rekeyed and LUNs pending.
   a. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
   b. Select the encryption switch from the **Encryption Center Devices** table, then select **Targets** from the menu task bar.

      The **Encryption Targets** dialog box displays. (Refer to Figure 359.)

6. Select a disk LUN device from the table, then click **LUNs**.

   The **Encryption Targets Disk LUNs** dialog box displays. (Refer to Figure 381.) The dialog box lists the status of the rekey operation.

FIGURE 381 Pending manual rekey operations

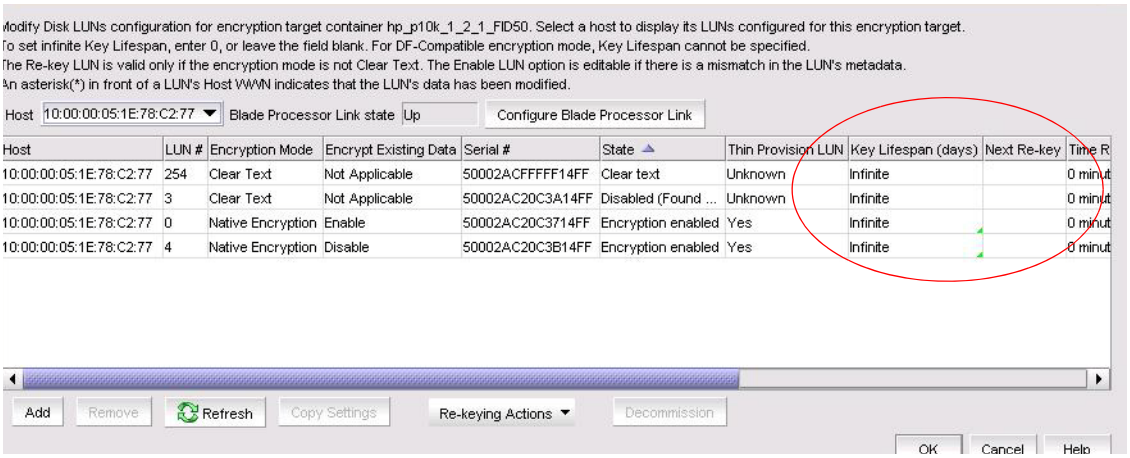

Viewing disk LUN rekeying details

You can view details related to the rekeying of a selected target disk LUN from the **LUN Re-keying Details** dialog box.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table, then select **Group/Switch/Engine > Targets**, or right-click the group, switch, or engine and select **Targets**.

   NOTE
   You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   The **Encryption Targets** dialog box displays.
3. Select a Target storage device, then **select Group/Switch/Engine > Disk LUNs**.

   The **Encryption Target Disk LUNs** dialog box displays. (Refer to Figure 382.) Initially the list is empty. You must add LUNs manually.

FIGURE 382 Encryption Target Disk LUNs dialog box

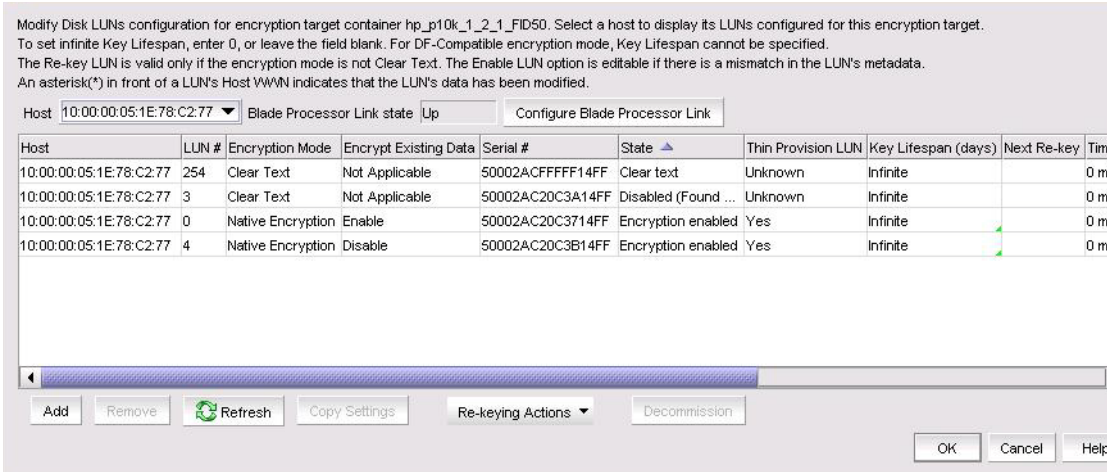

4. Click **Add**.

The **Add Disk LUNs** dialog box displays. This dialog box includes a table of all LUNs in the storage device that are visible to the hosts.

5. Click **Re-keying Details**.

   The **LUN Re-keying Details** dialog box displays. The dialog box contains the following information:

   - **Key ID:** The LUN key identifier.
   - **Key ID State:** The state of the LUN rekeying operation.
   - **Encryption Algorithm:** The algorithm of the LUN rekeying operation.
   - **Re-key Session Number:** The session number of the LUN rekeying operation.
   - **Re-key Role:** The role of the LUN rekeying operation.
   - **Re-key State**: The state of a manual LUN rekeying operation. Options are:
     - **Read Phase**
     - **Write Phase**
     - **Pending**
     - **Disabled**
   - **Block Size**: The block size used on the LUN.
   - **Number of Blocks**: The number of blocks written.
   - **Current LBA**: The Logical Block Address (LBA) of the block that is currently being written.
   - **Re-key Completion:** The status of the LUN rekeying operation's progress.

Viewing the progress of manual rekey operations

To monitor the progress of manual rekey operations, complete these steps:

1. Select **Configure > Encryption** from the menu task bar to display t**he Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption group from the **Encryption Center Devices** table, then select **Group > Re-Key Sessions** from the menu task bar.

   The **Re-Key Sessions Status** dialog box displays, which enables you to check on the status of each LUN that is being rekeyed within an encryption group. (Refer to Figure 383.)

FIGURE 383 Re-Key Sessions Status dialog box

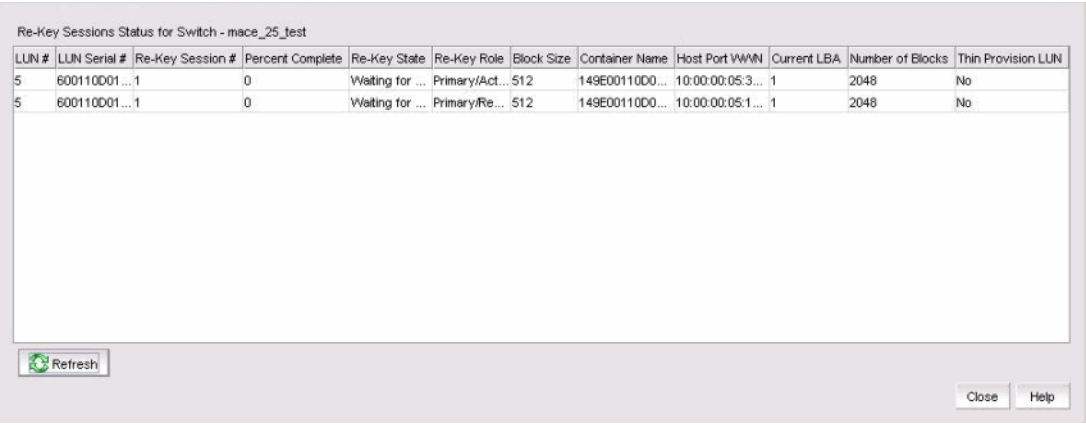

The dialog box contains the following information:

- **LUN #:** The LUN number.
- **LUN Serial #:** The LUN serial number.
- **Re-Key Session #:** The number assigned to the rekeying session.
- **Percent Complete:** The percentage of completion of the rekeying session.
- **Re-Key State:** Options are:
  - **Re-Key Setup**
  - **LUN Prep**
  - **LUN Clean-up**
  - **Key Update**
  - **Read Phase**
  - **Write Phase**
  - **HA Sync Phase**
- **Re-Key Role:** Options are:
  - **Primary/Active**
  - **Backup/Active**
- **Block Size:** The block size used on the LUN.
- **Container Name:** The CryptoTarget container name.
- **Host Port WWN:** The WWN of the host port that is being used for the write operation.
- **Current LBA:** The Logical Block Address (LBA) of the block that is currently being written.
- **Number of Blocks:** The number of blocks written.
- **Thin Provision LUN**: Identifies if the new LUN is a thin provisioned LUN. Options are:
  - **Yes:** Thin provision support is limited to Brocade-tested storage arrays. The thin provision LUN status will be displayed as **Yes** for supported storage arrays only.
  - **No:** Shown as No if the LUN is not a thin provisioned LUN.
  - **Unknown:** Shown if the LUN status cannot be determined.
  - **Not Applicable:** Applies to Encryption switches that are running a Fabric OS version earlier than v7.1.0.

3. Click **Refresh** periodically to update the display.

Thin provisioned LUNs

With the introduction of Fabric OS 7.1.0, the Encryption switch can discover if a disk LUN is a thin provisioned LUN. Support for a thin provisioned LUN is limited to disk containers only. Thin provisioned LUNs can be created with the new LUN option.

NOTE
Currently, thin provisioned LUN support is limited to Brocade-tested storage arrays running specific supported firmware releases. The thin provisioned LUN status will be displayed as **Yes** for supported storage arrays only.

Thin provisioned LUNs rely on on-demand allocation of blocks of data, instead of the traditional method of allocating all blocks up front. If a thin provisioned LUN status is shown as **Yes**, then first-time encryption and rekey are done on the allocated blocks only, which results in the provisioned region of the LUN to remain the same after the rekey is performed.

Thin provisioned LUN support requires no action by the user. The Encryption switch can automatically detect if a LUN is a thin provisioned LUN.

NOTE:

- For thin provisioned LUNs that were previously full provisioned then converted to thin, a **discoverLUN** command must be performed prior to any rekeying operations. Failure to do so results in the full capacity of the LUN to be encrypted as if it were not thin provisioned. Updated thin provisioned status can be verified using the **cryptocfg --show -container -all -stat** command and checking the output for "`Thin Provision LUN: Yes`". Similarly, if a thin- to full-LUN conversion has been performed, a **discoverLUN** command must be performed for this LUN change to reflect on the Encryption switch or FS8-18 blade.
- If a LUN is a thin provisioned LUN, LUN status is shown as **Yes**. (Thin provision support is limited to Brocade-tested storage arrays. The thin provisioned LUN status will be displayed as **Yes** for supported storage arrays only.)
- If a LUN is not a thin provisioned LUN or if thin provisioning is not supported with the LUN, LUN status is shown as **No**. (This can be a result of the array not supporting thin provisioning, or the Encryption switch or blade does not support the thin provisioning features of the array. Refer to the Fabric OS release notes for supported arrays.)
- If LUN status cannot be determined, LUN status is shown as **Unknown**.
- If you are running a Fabric OS version earlier than v7.1.0, LUN status is shown as **Not Applicable**.
- Zero detect with encryption is not supported.

Thin Provisioning support

Thin-provisioned logical unit numbers (LUNs) are increasingly used to support a pay-as-you-grow strategy for data storage capacity. Also known as dynamic provisioning, virtual LUNs, or thin LUNs, the same technology that allows storage administrators to allocate physical disk space to LUNs on an as-needed basis creates limitations around certain data-at-rest encryption operations that use the Encryption switch or blade. Performing first-time encryption (FTE) (conversion of cleartext to ciphertext) and data rekeying operations (applying new data encryption keys to ciphertext data) on thin-provisioned LUNs results in an attempt by the encryption switch to overwrite data up to the size of the logical size of the thin-provisioned LUN, rather than limiting FTE/rekeying to the size of the physically allocated LUN size or to the data that has been written. This generally triggers the allocation of additional blocks to the thin-provisioned LUN, using up the amount of physical disk space that is available to the LUN and defeating the objective of using thin provisioning.

Additionally, for thin-provision capable storage products that support space reclamation based on data pattern recognition (for example, 'string of zeros'), the encryption of such patterns will interfere with the space reclamation functionality of the storage and should be avoided.

Certain types of storage, including 3PAR, have been successfully tested by limiting the use of thin provisioning to "greenfield" LUNs, or LUNs that do not have any written data yet. Rekeying operations on these LUNs, like FTE, are also not permitted. As these limitations are not feasible for most environments, the recommendation from Brocade is that any encrypted LUNs be fully provisioned with disk.

Viewing time left for auto rekey

You can view the time remaining until auto rekey is no longer active for a disk LUN. The information is expressed as the difference between the next rekey date and the current date and time, and is measured in days, hours, and minutes.

Although you cannot make changes directly to the table, you can modify the time left using CLI. For more information, refer to the administrator's guide supporting your key vault management system.

To view the time left for auto rekey, follow these steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group, switch, or engine from the **Encryption Center Devices** table for which to view the auto rekey information, then select **Group/Switch/Engine > Targets** from the menu task bar.

   NOTE
   You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Targets** icon.

   The **Encryption Targets** dialog box displays. (Refer to Figure 359.)
3. Select a target disk device from the table, then click **LUNs**.

   The **Encryption Target Disk LUNs** dialog box displays. The time left for auto rekey information is listed in the table. (Refer to Figure 384.)

FIGURE 384 Encryption Targets Disk LUNs dialog box - Time left for auto rekey

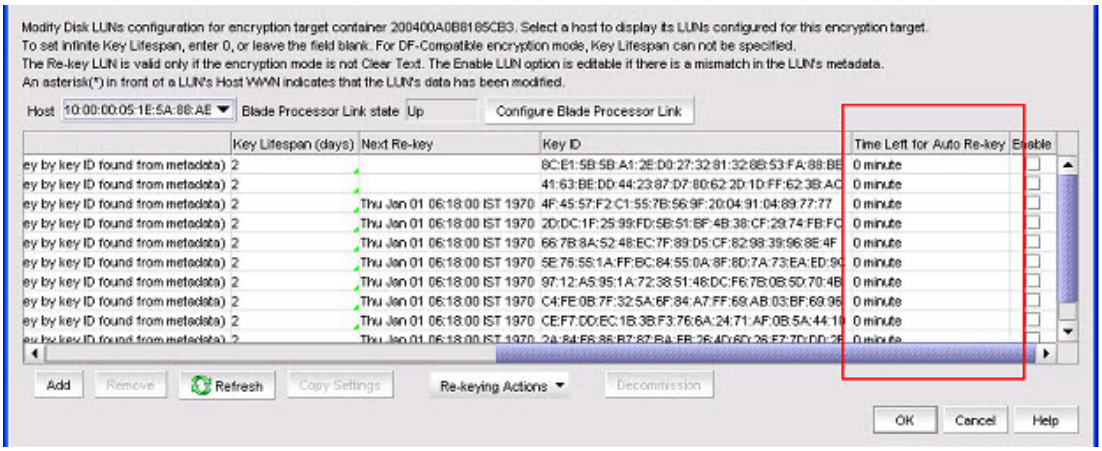

Viewing and editing switch encryption properties

To view switch encryption properties, complete the following steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a switch or encryption engine from the **Encryption Center Devices** table, then select **Switch/Engine > Properties** from the menu task bar, or right-click a switch or encryption engine and select **Properties**.

   NOTE
   You can also select a group, switch, or engine from the **Encryption Center Devices** table, then click the **Properties** icon.

   The **Encryption Switch Properties** dialog box displays. (Refer to Figure 385.)

FIGURE 385 Encryption Switch Properties dialog box

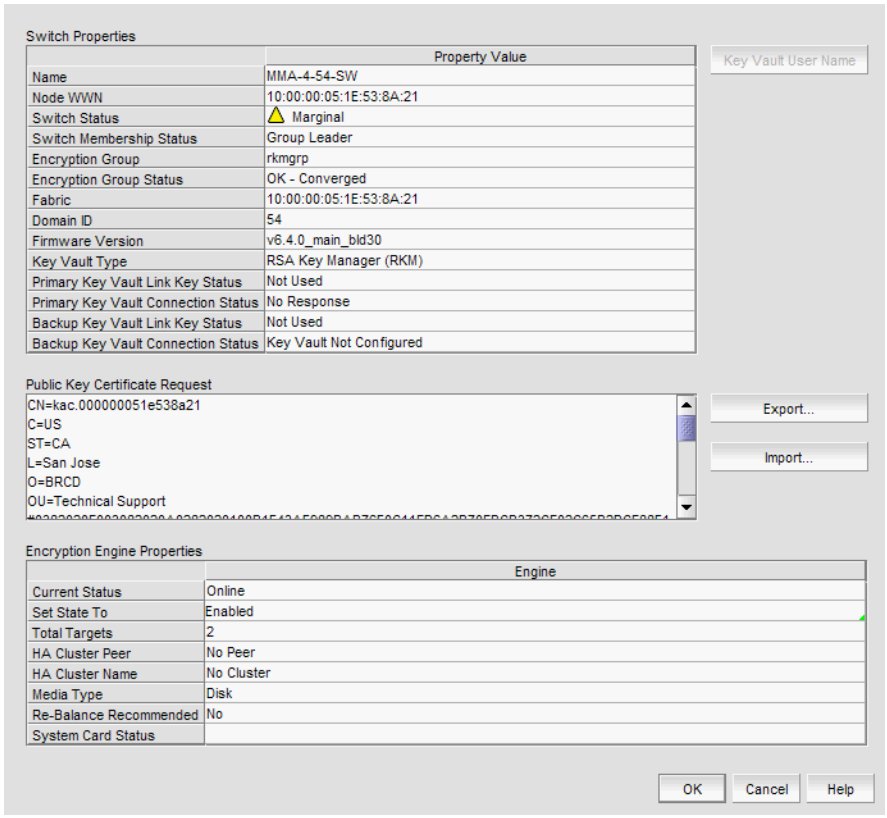

The dialog box contains the following information:

- **Switch Properties** table: A list of properties associated with the selected switch
- **Name**: The name of the selected switch
- **Node WWN:** The world wide name of the node
- **Switch Status:** The health status of the switch. Options are:
  - **Healthy**
  - **Marginal**
  - **Down**
  - **Unknown**
  - **Unmonitored**
  - **Unreachable**
- **Switch Membership Status:** The alert or informational message description, which details the health status of the switch. Options are:
  - **Group Member**
  - **Leader-Member Comm**
  - **Error**
  - **Discovering**
  - **Not a member**
- **Encryption Group:** The name of the encryption group to which the switch belongs

- **Encryption Group Status:** Status options are:
  - **OK/Converged**: the Group Leader can communicate with all members
  - **Degraded**: the Group Leader cannot communicate with one or more members. The following operations are not allowed: key vault changes, master key operations, enable/disable encryption engines, Failback mode changes, HA Cluster creation or addition (removal is allowed), tape pool changes, and any configuration changes for storage targets, hosts, and LUNs.
  - **Unknown**: The Group Leader is in an unmanaged fabric
- **Fabric:** The name of the fabric to which the switch belongs
- **Domain ID:** The domain ID of the selected switch
- **Firmware Version:** The current encryption firmware on the switch.
- **Key Vault type**:
- **Primary Key Vault Link Key Status/Backup Key Vault Link Key Status:** Status options are:
  - **Not Used**: The key vault type is not LKM/SSKM.
  - **No Link Keys, ready to establish**: No access request has been sent to an LKM/SSKM, or a previous request was not accepted.
  - **Link key requested, waiting for LKM approval**: A request has been sent to LKM/SSKM and is waiting for the LKM/SSKM administrator's approval.
  - **Created, not validated**: An interim state until first used **Link Key valid, online**: (LKM/SSKM only) a shared link key exists and has been successfully used.
- **Primary Key Vault Connection Status/Backup Key Vault Connection Status:** Whether the primary key vault link is connected. Options are:
  - **Unknown/Busy**
  - **Key Vault Not Configured**
  - **No Response**
  - **Failed authentication**
  - **Connected**
- **Key Vault User Name** button: (*TEKA only.*) Launches a dialog box to identify key vault user information. A user name is automatically generated on the switch side for use in defining a TEKA client for the switch.
- **Public Key Certificate Request** text box: The switch's KAC certificate signing request, which must be signed by a certificate authority (CA). The signed certificate must then be imported onto the switch and onto the primary and backup key vaults.
- **Export** button: Exports the public key certificate in CSR format to an external file for signing by a certificate authority (CA).
- **Import** button: Imports a signed public key certificate.
- **Encryption Engine Properties** table: The properties for the encryption engine. There may be 0 to 4 slots, one for each encryption engine in the switch.
- **Current Status:** The status of the encryption engine. Many possible values exist. Common options are:
  - **Not Available (the engine is not initialized)**
  - **Disabled**
  - **Operational**
  - **need master/link key**
  - **Online**
- **Set State To:** Identifies if the state is enabled or disabled. You can click the line item in the table to change the value, then click **OK** to apply the change.
- **Total Targets:** The number of encrypted target devices.

- **HA Cluster Peer:** The name and location of the high-availability (HA) cluster peer (another encryption engine in the same group), if in an HA configuration. If no peer is configured, **No Peer** is displayed.
- **HA Cluster Name:** The name of the HA cluster (for example, Cluster1), if in an HA configuration. HA Cluster names can have up to 31 characters. Letters, digits, and underscores are allowed.
- **Media Type:** The media type of the encryption engine. Options are **Disk and Tape**, or **Disk/Tape** when both are present.
- **Re-Balance Recommended:** Indicates if LUN rebalancing is recommended for an encryption engine that is hosting both disk and tape LUNs. Options are Yes and No.
- **System Card Status:** The current status of system card information for the encryption engine. Options are **Enabled and Disabled**.

Exporting the public key certificate signing request from properties

To export the certificate signing request (CSR) under Public Key Certificate Request, complete the following steps.

1. Click **Export**, then browse to the location where you want to save the certificate and click **Save**.

   Alternatively, you may also copy the CSR and paste it to a file.

2. Submit the CSR to a certificate authority (CA) for signing. CA signing requirements and procedures differ per key manager appliance.

Importing a signed public key certificate from properties

To import a signed public key certificate, complete the following steps.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption engine from the **Encryption Center Devices** table, then select **Engine > Properties** from the menu task bar, or right-click a switch or encryption engine and select **Properties**.

   NOTE
   You can also select a an engine from the **Encryption Center Devices** table, then click the **Targets** icon.

3. Click **Import**.

   The **Import Signed Certificate** dialog box displays. (Refer to Figure 386.)

   FIGURE 386 Import Signed Certificate dialog box

   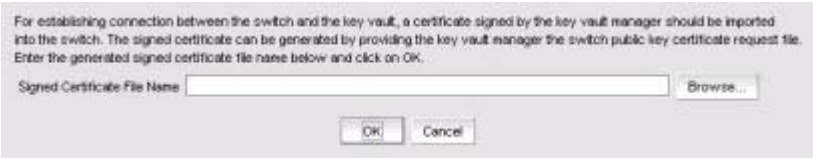

4. Enter or browse to the file containing the signed certificate, then click **OK**.

   The file is imported onto the switch.

Enabling and disabling the encryption engine state from Properties

To enable the encryption engine, complete the following steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select an encryption engine from the **Encryption Center Devices** table, then select **Engine > Properties** from the menu task bar, or right-click a switch or encryption engine and select **Properties**.

   NOTE
   You can also select a an engine from the **Encryption Center Devices** table, then click the **Targets** icon.

3. In the **Encryption Engine Properties** table, locate **Set State To**.
4. Click the adjacent **Engine** field and select **Enabled** or **Disabled** accordingly, then click **OK**.

Viewing and editing encryption group properties

Whenever you add or change a key vault address, you must also load the corresponding key vault certificate. When adding or changing a key vault, if the switches in the encryption group have not been previously registered with the new key vault, you must add the switch certificates to the key vault.

To view encryption group properties, complete the following steps.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices** table, then select **Group > Properties** from the menu task bar.
3. You can also select a group from the **Encryption Center Devices** table, then click the **Properties** icon.

   NOTE
   If groups are not visible in the **Encryption Center Devices** table, select **View > Groups** from the menu task bar.

   The **Encryption Group Properties** dialog box includes several tabs that are used to configure the various functions for encryption groups. All tabs are visible for all key vault types with one exception; the **Link Keys** tab is visible only if the key vault type is NetApp LKM/SSKM. Unless otherwise specified, the **Encryption Group Properties** dialog box opens with the **General** tab displayed. (Refer to Figure 387.)

FIGURE 387Encryption Group Properties dialog box

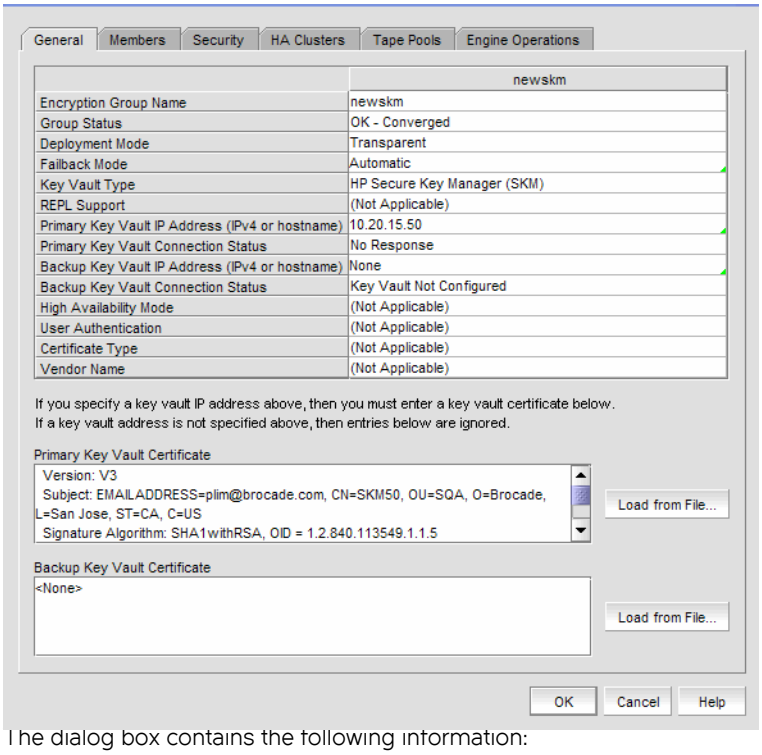

The dialog box contains the following information:

- **General** tab: For a description of the dialog box, refer to "General tab" on page 765.
- **Members** tab: For a description of the dialog box, refer to "Members tab" on page 768.
- **Security** tab: For a description of the dialog box, refer to "Security tab" on page 770.
- **HA Clusters** tab: For a description of the dialog box, refer to "HA Clusters tab" on page 772.
- **Tape Pools** tab: For a description of the dialog box, refer to "Tape Pools tab" on page 775.
- **Engine Operations** tab: For a description of the dialog box, refer to "Engine Operations tab" on page 777.

General tab

The General tab is viewed from the **Encryption Group Properties** dialog box. (Refer to Figure 388.) To access the **General** tab, select a group from the **Encryption Center Devices** table, then select **Group > Properties** from the menu task bar.

NOTE
You can also select a group from the **Encryption Center Devices** table, then click the **Properties** icon.

FIGURE 388 Encryption Group Properties dialog box - General tab

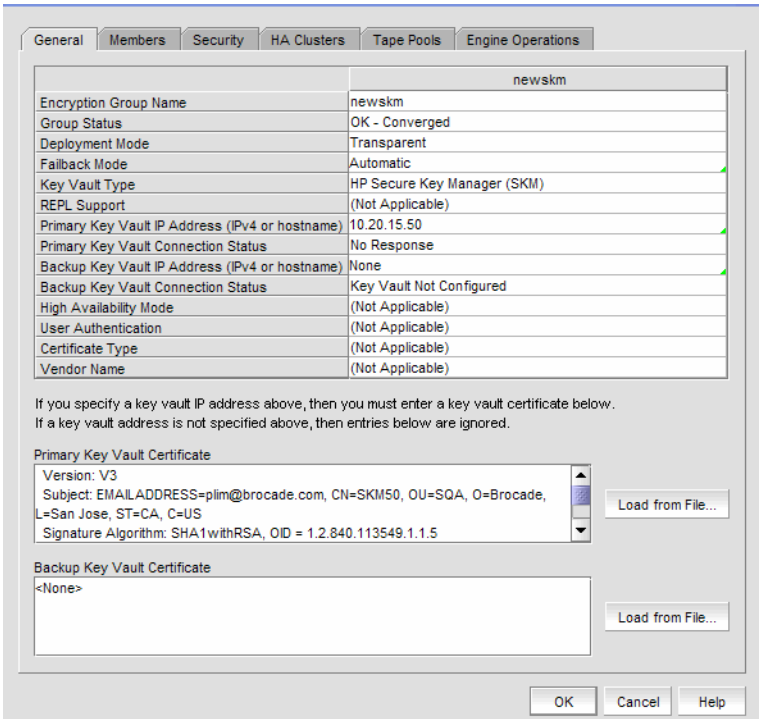

The dialog box contains the following information:

- **Encryption Group Name:** The name of the encryption group.
- **Group Status:** The status of the encryption group. Options are:
  - **OK-Converged**: The Group Leader can communicate with all members.
  - **Degraded**: The Group Leader cannot contact one or more of the configured group members. When the group is in a degraded state, many operations are not permitted, including configuring targets, hosts, LUNs, HA clusters, and tape pools.
- **Deployment Mode:** The group's deployment mode, which is transparent mode.
- **Failback Mode:** Identifies the group's failback mode. Options are: **Automatic** and **Manual**. Failback mode can be changed by clicking on the field and selecting the desired mode.

  The HA failback option determines the behavior when a failed encryption engine is restarted. When one encryption engine in an HA cluster fails, the second encryption engine in the HA cluster takes over the encryption and decryption of traffic to all encryption targets in the first encryption engine.

  When the first encryption engine comes back online, the encryption group's failback setting determines whether the first encryption engine automatically resumes encrypting and decrypting traffic to its encryption targets. In manual mode, the second encryption engine continues handling the traffic until you manually invoke failback using the CLI, or until the second encryption engine fails.

- **Key Vault Type:** Options are:
  - **RSA Data Protection Manager (DPM):** If an encryption group contains mixed firmware nodes, the Encryption Group Properties Key Vault Type name is based on the firmware version of the Group Leader. For example, If a switch is running Fabric OS 7.1.0 or later, the Key Vault Type is displayed as "RSA Data Protection Manager (DPM)."If a switch is running a Fabric OS version prior to v7.1.0, Key Vault Type is displayed as "RSA Key Manager (RKM)".
  - **NetApp Lifetime Key Manager (LKM):** The NetApp Key Vault Type name is shown as NetApp Lifetime Key Manager (LKM) for both NetApp Lifetime Key Manager (LKM) and SafeNet KeySecure for key management (SSKM) Key Vault Types.
  - **HP Secure Key Manager (SKM):** The HP Key Vault Type name is shown as HP Secure Key Manager (SKM) for both SKM and **Enterprise Secure Key Management (**ESKM) Key Vault Types.
  - **Thales e-Security keyAuthority (TEKA):** If an encryption group contains mixed firmware nodes, the Encryption Group Properties Key Vault Type name is based on the firmware version of the Group Leader. For example, If a switch is running Fabric OS 7.1.0 or later, the Key Vault Type is displayed as "Thales e-Security keyAuthority (TEKA)."If a switch is running a Fabric OS version prior to v7.1.0, Key Vault Type is displayed as "Thales Key Manager (TEMS)".
  - **Tivoli Key Lifetime Manager (TKLM):** (No other key vault name is used)
  - **Key Management Interoperability Protocol (KMIP)**: Any KMIP-compliant server can be registered as a key vault on the Encryption switch after setting the key vault type to KMIP.

    With the introduction of Fabric OS 7.1.0, KMIP with SafeNet KeySecure for key management (SSKM) native hosting LKM is supported.

    With the introduction of Fabric OS 7.2.0, KMIP with TEKA 4.0 is also supported (using the CLI only). For more information about supported platforms and configuration instructions, refer to the *Fabric OS Encryption Administrator's Guide Supporting Key Management Interoperability Protocol (KMIP) Key-Compliant Environments*.
- **REPL Support:** Identifies if the remote replication LUN support is enabled or disabled. You can change the current setting by clicking on the field and selecting the desired state.
- **Primary Key Vault IP Address:** The IP address of the primary key vault, either IPv4 or host name.
- **Primary Key Vault Connection Status:** The status of the primary key vault link. In an operating environment, the status should be **Connected**. Other options are:
  - **Unknown/Busy**
  - **Not configured**
  - **Not responding**
  - **Failed authentication**
- **Backup Key Vault IP Address:** *(Optional.)* The IP address of the backup key vault. This field can be left blank.
- **Backup Key Vault Connection Status:** The status of the backup key vault link. Options are:
  - **Connected**
  - **Unknown/Busy**
  - **Not configured**
  - **Not responding**
  - **Failed authentication**

- **High Availability Mode: (***For KMIP key vault type.***) Options are:**
  - **Opaque:** Both the primary and secondary key vaults are registered on the Encryption switch. The client archives the key to a single (primary) key vault. For disk operations, an additional key hardening check is done on the secondary key vault before the key is used for encryption.
  - **Transparent:** A single key vault should be registered on the Encryption switch. The client assumes the entire HA is implemented on the key vault. Key archival and retrieval is done to the KMIP without any additional key hardening checks.
  - **No HA:** Both the primary and secondary key vaults are registered on the Encryption switch. The client archives keys to both key vaults and ensures that the archival is successful before the key is used for encryption.
  - **None: High availability is not configured.**
  - **Not Applicable: Displayed if your selected key vault type is not KMIP.**
- **User Authentication: (***For KMIP key vault type.***)** The methods used to authenticate a user. Options are:
  - **Username and Password:** Activates the Primary and Backup Key Vault User Names and password fields for completion.
  - **Username:** Activates the Primary and Backup Key Vault User Names for completion.
  - **None:** Deactivates Primary and Backup Key Vault User Names and password fields.
  - **Not Applicable: Displayed if your selected key vault type is not KMIP.**
- **Certificate Type: (***For KMIP key vault type.***)** Displays the TLS certificate type used between the BES and the key vault. Options are:
  - **CA Signed:** The BES KAC certificate is signed by a CA, imported back on the Encryption switch and registered as a KAC certificate. The CA will be registered as a key vault certificate on the Encryption switch.
  - **Self Signed:** The self-signed certificates are exchanged and registered on both ends. The key vault certificate is registered on the BES and the BES KAC certificate is registered on the key vault.
- **Vendor Name: (*For KMIP key vault type*)** Displays the supported key vendor server. The vendor name will display the connected key vault through KMIP.
- **Primary Key Vault Certificate table:** Displays the details of the primary vault certificate; for example, version and signature information. The **Load from File** button allows you to locate and load a primary key vault certificate from a different location.
- **Backup Key Vault Certificate table:** Displays the details of the backup vault certificate; for example, version and signature information. The **Load from File** button allows you to locate and load a backup key vault certificate from a different location.

Members tab

The **Members** tab lists group switches, their role, and their connection status with the Group Leader. The table columns are not editable. The tab displays the configured membership for the group and includes the following:

- **Node WWN:** The member switch's world wide name.
- **IP Address:** The switch's IP address or host name.
- **Node Name:** The switch's node name, if known. If unknown, this field is blank.
- **Connection Status:** The switch's connection status. Possible values are:
  - **Group Leader:** The switch designated as the Group Leader, so there is no connection status.
  - **Trying to Contact:** The member is not responding to the Group Leader. This might occur if the member switch is not reachable by way of the management port, or if the member switch does not believe it is part of the encryption group.
  - **Configuring:** The member switch has responded and the Group Leader is exchanging information. This is a transient condition that exists for a short time after a switch is added or restored to a group.
  - **OK:** The member switch is responding to the Group Leader switch.
  - **Not Available:** The Group Leader is not a managed switch, so connection statuses are not being collected from the Group Leader.

The **Members** table might not match the list of members displayed in the **Encryption Center** dialog box if some configured members are unmanaged, missing, or in a different group.

NOTE
When the encryption group is in the Degraded state, the **Members** tab indicates the group member that the leader cannot contact. If the non-responding switch should no longer be included in the encryption group, it can be removed using the **Remove** button.

The **Members** tab is viewed from the **Encryption Group Properties** dialog box. (Refer to Figure 389.) To access the **Members** tab, select a group from the **Encryption Center Devices** table, then select **Group > Properties** from the menu task bar.

NOTE
You can also select a group from the **Encryption Center Devices** table, then click the **Properties** icon.

FIGURE 389 Encryption Group Properties dialog box - Members tab

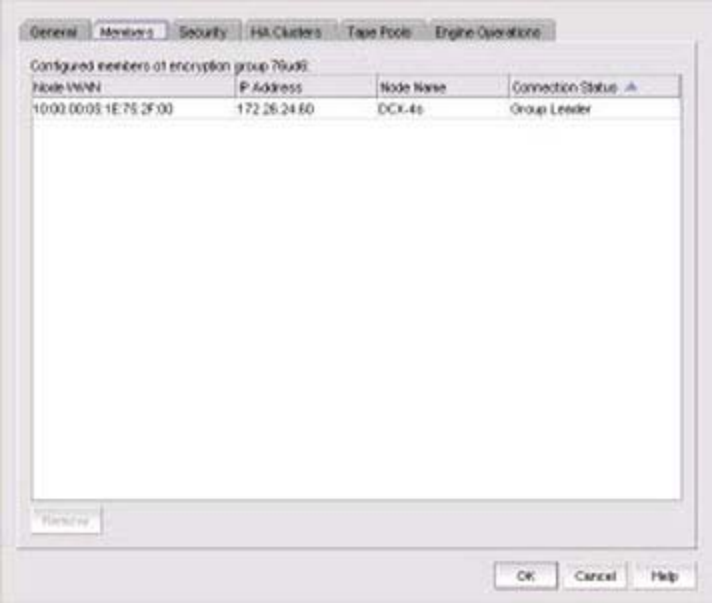

Members tab Remove button

You can click the **Remove** button to remove a selected switch or group from the encryption group table.

- You cannot remove the Group Leader unless it is the only switch in the group. If you remove the Group Leader, the Management application also removes the HA cluster, the target container, and the tape pool (if configured) that are associated with the switch.
- If you remove a switch from an encryption group, the Management application also removes the HA cluster and target container associated with the switch.

  NOTE
  If the encryption group is in a degraded state, the Management application does not remove the HA clusters or target containers associated with the switch. In this case, an error message displays.

- If you remove the last switch from a group, the Management application also deletes the group.

Consequences of removing an encryption switch

The consequences of removing a switch from an encryption group are as follows:

- All configured targets on the switch are deleted from the switch's configuration.
- Any encryption being performed by the switch is halted.
- If the removed switch was in an HA cluster, the switch can no longer provide HA support. HA clusters that contained the encryption engine from the removed switch are deleted.

The consequences of removing the last switch in a group (which will be the Group Leader) are all switch removal consequences noted above, plus the following:

- The encryption group is deleted.
- All configured tape pools are deleted.

Table 60 explains the impact of removing switches.

TABLE 60 Switch removal impact

Switch configuration	Impact of removal
The switch is the only switch in the encryption group.	The encryption group is also removed.
The switch has configured encryption targets on encryption engines.	• The switch is configured to encrypt traffic to one or more encryption targets. • The target container configuration is removed. • The encrypted data remains on the encryption target but is not usable until the encryption target is manually configured on another encryption switch.  **The encryption target data is visible in encrypted format to zoned hosts. It is strongly recommended that you remove the encryption targets from all zones before you disable encryption. Otherwise, hosts might corrupt the encrypted data by writing directly to the encryption target without encryption.**
The switch has encryption engines in HA clusters.	The HA clusters are removed. High availability is no longer provided to the other encryption engine in each HA cluster.

A warning message is displayed when you attempt to remove a switch or an encryption group. After you have read the warning, you must click **Yes** to proceed.

Security tab

The **Security** tab displays the status of the master key for the encryption group and whether smart cards are required. From here, you register smart cards for use.

The **Security** tab is viewed from the **Encryption Group Properties** dialog box. (Refer to Figure 390.) To access the **Security** tab, select a group from the **Encryption Center Devices** table, then select **Group > Security** from the menu task bar. The **Properties** dialog box displays with the **Security** tab selected.

NOTE
You can also select a group from the **Encryption Center Devices** table, then click the **Properties** icon.

FIGURE 390 Encryption Group Properties dialog box - Security tab

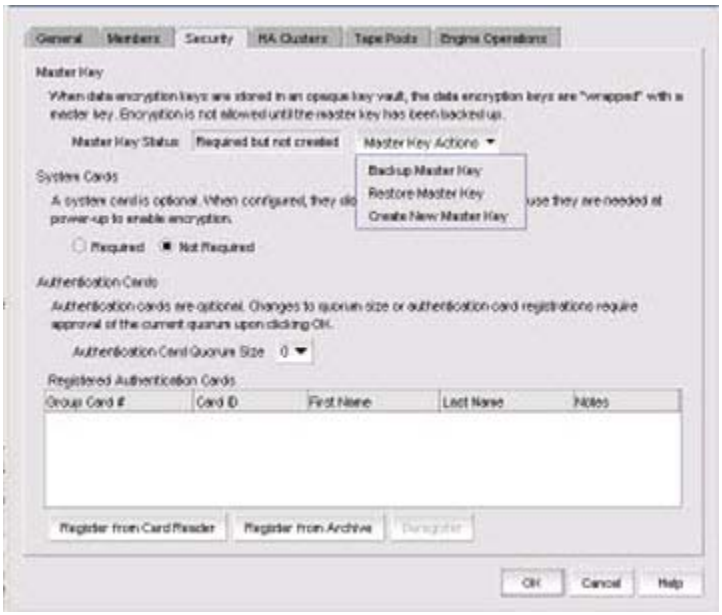

The dialog box contains the following information:

- **Master Key Status**: Displays the status of the master key. Possible values are:
  - **Not used:** Displays when LKM/SSKM is the key vault.
  - **Required but not created:** Displays when a master key needs to be created.
  - **Created but not backed up:** Displays when the master key needs to be backed up. For safety, the master key cannot be used until it is backed up.
  - **Created and backed up:** Indicates the master key is usable.
- **Master Key Actions** list: Master Key actions are disabled if the master key state is not correct. Master key actions are:
  - **Create a new master key:** Enabled when no master key exists or the previous master key has been backed up.
  - **Back up a master key:** Enabled any time a master key exists.
  - **Restore a master key:** Enabled when either no master key exists or the previous master key has been backed up.
- **System Cards:** Identifies if the use of a system card is required for controlling activation of the encryption engine. You must indicate if cards are required or not required. If a system card is required, it must be read by the card reader on the switch.
- **Authentication Cards**, which identifies if one or more authentication cards must be read by a card reader attached to a Management application PC to enable certain security-sensitive operations.
- **Authentication Cards quorum size** selector: Determines the number of registered authentication cards needed for a quorum. The number should always be one less than the actual number registered.

  NOTE
  When registering authentication cards, you must register the defined quorum size plus one.

- **Registered Authentication Cards** table: Lists the registered authentication cards.
  - **Group Card #**: The number of cards that are registered.
  - **Card ID:** The card serial number.
  - **First Name** and **Last Name**: The first and last name of the person assigned to the card. The names are identified when the authentication card is first registered.
  - **Notes**: An optional entry of information.
- **Register from Card Reader** button: Launches the **Add Authentication Card** dialog box.
- **Register from Archive** button: Launches the **Add Authentication Card** dialog box.
- **Deregister** button: Deregisters authentication cards, thus enabling them to be removed from the switch and the database.

Encryption is not allowed until the master key has been backed up. Master keys are needed for all key vaults except LKM/SSKM.

NOTE
You must enable encryption engines before you back up or restore master keys.

NOTE
If all encryption engines are otherwise operating normally but are missing the master key, the following message displays below the Master Key status:

```
"None of the encryption engines in this encryption group have a copy of the master key. The master key should be restored from a backup."
```

This situation can occur if all encryption engines in a group are zeroized and then re-enabled.

HA Clusters tab

The **HA Clusters** tab allows you to create and delete HA clusters, add encryption engines to and remove encryption engines from HA clusters, and failback an engine. Changes are not applied to the encryption group until you click **OK**.

Each HA cluster must have exactly two encryption engines. The two encryption engines in the cluster must be in the same fabric (they will always be in the same encryption group since only the engines in the group are listed for selection).

HA clusters are groups of encryption engines that provide high availability features. If one of the engines in the group fails or becomes unreachable, the other cluster member takes over the encryption and decryption tasks of the failed encryption engine. An HA cluster consists of exactly two encryption engines. Refer to "Creating HA clusters" on page 713.

The **HA Clusters** tab is viewed from the **Encryption Group Properties** dialog box. (Refer to Figure 391.) To access the **HA Clusters** tab, select a group from the **Encryption Center Devices** table, then select **Group > HA Clusters** from the menu task bar. The **Properties** dialog box displays with the **HA Clusters** tab selected.

NOTE
You can also select a group from the **Encryption Center Devices** table, then click the **Properties** icon.

The tab displays the includes the following information:

- **Non-HA Encryption Engines** table: Displays a list of encryption engines that are not configured for high-availability clustering
- **High-Availability Clusters** table: A list of encryption engines that have been selected for high-availability clustering.

- Right and left arrow buttons: You can select an encryption engine in the **Non-HA Encryption Engines** table and click the right arrow button to add the encryption engine to the **High-Availability Clusters**. (If you are creating a new HA cluster, a dialog box displays requesting a name for the new HA cluster.) Similarly, you can select an encryption engine in the **High-Availability Clusters** table and click the left arrow button to remove it from a cluster. The encryption engine is removed from the table and shown as available.
- Dual arrow button: After selecting an encryption engine in both the **Non-HA Encryption Engines** table and the **High-Availability Clusters** table, clicking the dual arrow button swaps the cluster members.

  NOTE
  Swapping engines using the dual arrow button is not the same as removing one engine and adding another. When swapping engines, all configured targets are moved from the former HA cluster member to the new HA cluster member. Swapping engines is useful when replacing hardware.

- **Configure Blade Processor Link** button: When active, clicking the button displays the **Configure Blade Processor Link** dialog box. Blade processor links must be configured and functioning to enable the failover/failback capabilities of a high availability cluster. For more information, refer to "Configuring blade processor links" on page 632.
- **Failback** button: After selecting an online encryption engine in the **High-Availability Clusters** table, you can click **Failback** to manually invoke failback. For more information, refer to "Invoking failback" on page 715.

FIGURE 391 Encryption Group Properties dialog box - HA Clusters tab

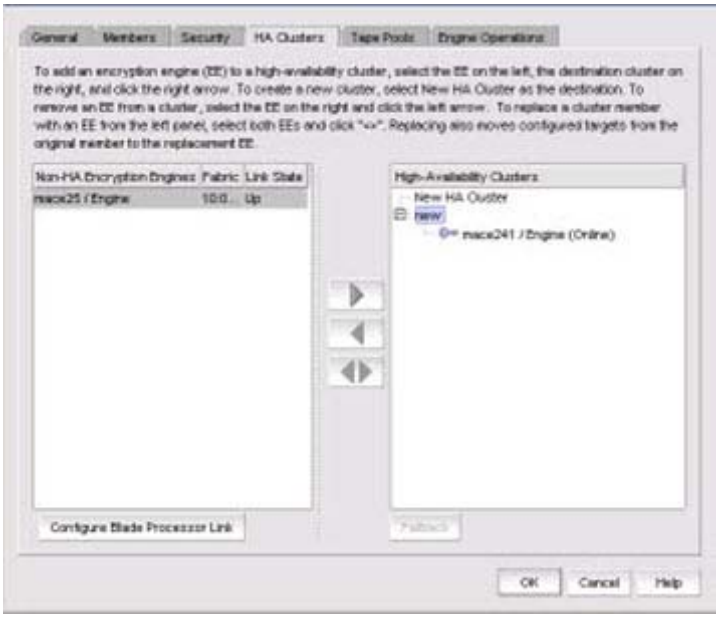

Link Keys tab

NOTE
The **Link Keys** tab displays only if the key vault type is NetApp LKM/SSKM.

Connections between a switch and an NetApp LKM/SSKM key vault require a shared link key. Link keys are used only with LKM/SSKM key vaults. Link keys are used to protect data encryption keys in transit to and from the key vault. There is a separate link key for each key vault for each switch. The link keys are configured for a switch but are stored in the encryption engines, and all of the encryption engines in a group share the same link keys. You must create link keys under the following circumstances:

- When a new encryption group is created.

- When a new switch is added to an encryption group.
- When a new key vault is added to an encryption group.
- After all encryption engines in a switch have been zeroized.
- When all of the encryption blades have been removed from a director and one or more new encryption blades have been added.

The **Link Keys** tab is viewed from the **Encryption Group Properties** dialog box. (Refer to Figure 392.) A table displays link key status for each switch in an encryption group, which includes the following information:

- **Switch**: The name of the selected switch in the encryption group.
- **Key Vault**: The type of key vault, either Primary or Secondary.
- **Link Key Status**: The link key status can be one of the following:
  - No Link Key: No access request was sent to LKM/SSKM yet, or a previous request was not accepted.
  - No Link Key, ready to establish: No link key exists, and no link key has been requested.
  - Link Key requested, waiting for LKM/SSKM approval: A request was sent to LKM/SSKM and is waiting for LKM/SSKM approval.
  - Waiting for local approval: A response was received from LKM/SSKM and needs local quorum of cards approval.
  - Created, not validated: The interim state until first used.
  - Link Key Valid, Online: A shared link key exists and has been successfully used.

Included on the Link Keys tab is the **Establish** button and the **Accept** button.

- If a switch shows a status of No Link Key, ready to establish, you may select the switch and click **Establish** to send a Trust Establishment Package (TEP) message to LKM/SSKM.
- If a switch shows a status of Link Key requested, waiting for LKM/SSKM approval, you may click **Accept** to accept the Trust Acceptance Package (TAP) that was sent in response to the TEP that was sent when you clicked **Establish**.

To access the **Link Keys** tab, select an LKM/SSKM group from the **Encryption Center Devices** table, then select **Group > Link Keys** from the menu task bar. The **Properties** dialog box displays with the **Link Keys** tab selected.

NOTE
You can also select a group from the **Encryption Center Devices** table, then click the **Properties** icon.

FIGURE 392Encryption Group Properties dialog box - Link Keys tab

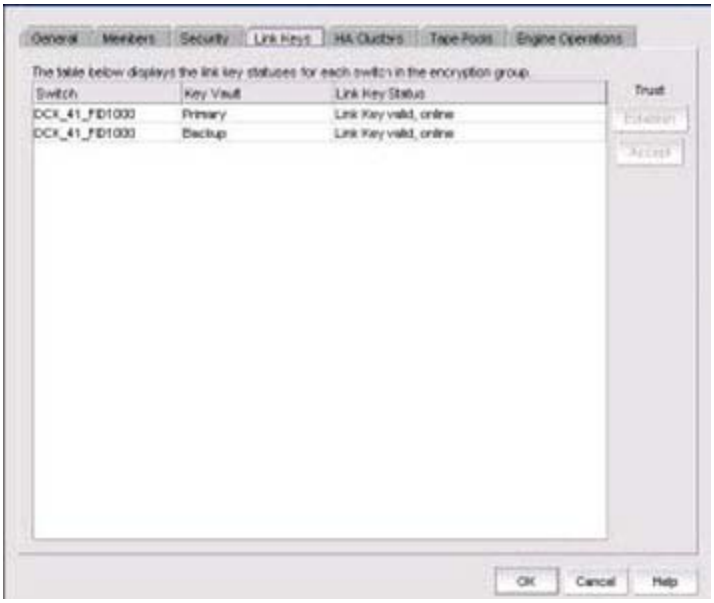

Tape Pools tab

Tape pools are managed from the **Tape Pools** tab. From the **Tape Pools** tab, you can add, modify, and remove tape pools.

- To add a tape pool, click **Add**, then complete the **Add Tape Pool** dialog box.
- To remove an encryption switch or engine from a tape pool, select one or more tape pools listed in the table, then click **Remove**.
- To modify a tape pool, you must remove the entry, then add a new tape pool.

The **Tape Pools** tab is viewed from the **Encryption Group Properties** dialog box. (Refer to Figure 393.) To access the **Tape Pools** tab, select a group from the **Encryption Center Devices** table, then select **Group > Tape Pools** from the menu task bar. The **Properties** dialog box displays with the **Tape Pools** tab selected.

NOTE
You can also select a group from the **Encryption Center Devices** table, then click the **Properties** icon.

FIGURE 393 Encryption Group Properties dialog box - Tape Pools tab

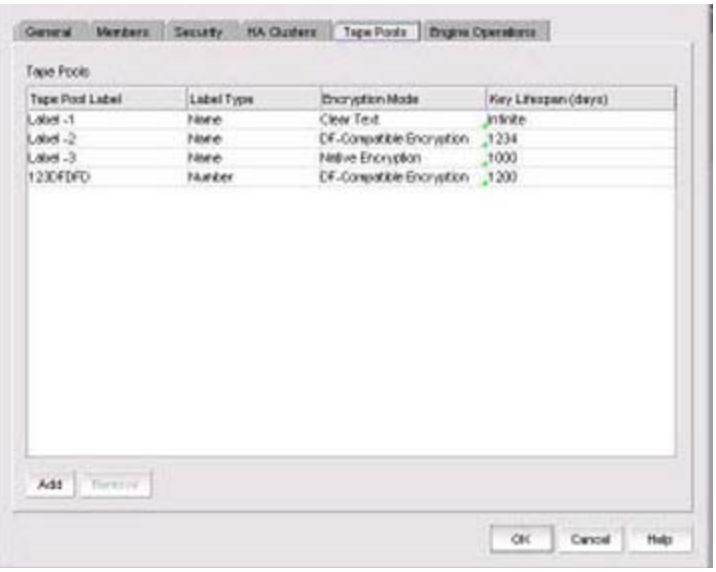

Tape pools overview

Tape cartridges and volumes can be organized into a tape pool (a collection of tape media). The same data encryption keys are used for all cartridges and volumes in the pool. Tape pools are used by backup application programs to group all tape volumes used in a single backup or in a backup plan. The tape pool name or number used must be the same name or number used by the host backup application. If the same tape pool name or number is configured for an encryption group, tapes in that tape pool are encrypted according to the tape pool settings instead of the tape LUN settings.

Encryption switches and encryption blades support tape encryption at the tape pool level (for most backup applications) and at the LUN (tape drive) level. Since Tape Pool policies override the LUN (tape drive) policies, the LUN pool policies are used only if no tape pools exist or if the tape media/volume does not belong to any configured tape pools.

All encryption engines in the encryption group share the tape pool definitions. Tapes can be encrypted by any encryption engine in the group where the container for the tape target LUN is hosted. The tape media is mounted on the tape target LUN.

Tape pool definitions are not needed to read a tape. The tape contains enough information (encryption method and key ID) to read the tape. Tape pool definitions are only used when writing to tape. Tape pool names and numbers must be unique within the encryption group.

Adding tape pools

A tape pool can be identified by either a name or a number, but not both. Tape pool names and numbers must be unique within the encryption group. When a new encryption group is created, any existing tape pools in the switch are removed and must be added.

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to Figure 269 on page 620.)
2. Select a group from the **Encryption Center Devices** table, then select **Group > Tape Pools** from the menu task bar.

   NOTE
   If groups are not visible in the **Encryption Center Devices** table, select **View > Groups** from the menu task bar.

3. Click **Add**.

The **Add Tape Pool** dialog box displays. The **Name** tape pool label type is the default; however, you can change the tape pool label type to **Number**

4. Based on your selection, do one of the following:
   - If you selected **Name** as the **Tape Pool Label Type**, enter a name for the tape pool. This name must match the tape pool label or tape ID that is configured on the tape backup/restore application.
   - If you selected **Number** as the **Tape Pool Label Type**, enter a (hex) number for the tape pool. This number must match the tape pool label or tape number that is configured on the tape backup/restore application.
5. Select the **Encryption Mode**. Options are Clear Text, DF-Compatible Encryption, and Native Encryption. Note the following:
   - DF-Compatible Encryption is valid only when LKM/SSKM is the key vault.
   - The **Key Lifespan (days)** field is editable only if the tape pool is encrypted.
   - If **Clear Text** is selected as the encryption mode, the key lifespan is disabled.

   NOTE
   You cannot change the encryption mode after the tape pool I/O begins. DF-compatible encryption requires a DF-compatible encryption license to be present on the switch. If the license is not present, a warning message displays.
6. Enter the number of days to use a key before obtaining a new one, if you choose to enforce a key lifespan. The default is Infinite (a blank field or a value of 0), which is the recommended setting.

   NOTE
   The key lifespan interval represents the key expiry timeout period for tapes or tape pools. You can only enter the **Key Lifespan** field if the tape pool is encrypted. If **Clear Text** is selected as the encryption mode, the **Key Lifespan** field is disabled.
7. Click **OK**.

Engine Operations tab

The **Engine Operations** tab enables you to replace an encryption engine in a switch with another encryption engine in another switch within a DEK Cluster environment. A DEK Cluster is a set of encryption engines that encrypt the same target storage device. DEK Clusters do not display in the Management application; they are an internal implementation feature and have no user-configurable properties. Refer to "Replacing an encryption engine in an encryption group" on page 711.

The **Engine Operations** tab is viewed from the **Encryption Group Properties** dialog box. (Refer to Figure 394.) To access the **Engine Operations** tab, select a group from the **Encryption Center Devices** table, then select **Group > Engine Operations** from the menu task bar**.** The **Properties** dialog box displays with the **Engine Operations** tab selected.

NOTE
You can also select a group from the **Encryption Center Devices** table, then click the **Properties** icon.

You simply select the encryption engine you want to replace from the Engine list, select the encryption engine to use for the group from the **Replacement** list, then click **Replace**.

FIGURE 394 Encryption Group Properties Dialog Box - Engine Operations Tab

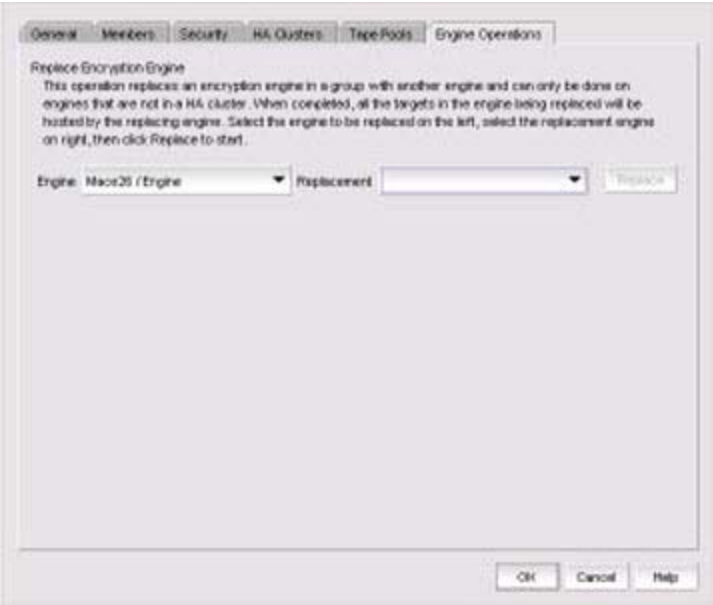

NOTE
You cannot replace an encryption engine if it is part of an HA cluster.

Encryption-related acronyms in log messages

Fabric OS log messages related to encryption components and features may have acronyms embedded that require interpretation. Table 61 lists some of those acronyms.

TABLE 61 Encryption acronyms

Acronym	Name
EE	Encryption Engine
EG	Encryption Group
HAC	High Availability Cluster

Zoning

In this chapter

Zoning overview

NOTE
Zoning is supported on Fabric OS devices.

Zoning is a fabric-based service that enables you to partition your network into logical groups of devices that can access each other and prevent access from outside the group. Grouping devices into zones in this manner not only provides security, but also relieves the network from Registered State Change Notification (RSCN) storms that occur when too many native FCoE devices attempt to communicate with one another.

You can use zoning to partition your network in many ways. For example, you can partition your network into two zones, ***winzone*** and ***unixzone***, so that your Windows servers and storage do not interact with your UNIX servers and storage. You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions; for example, you can create a temporary zone to back up nonmember devices.

A device in a zone can communicate only with other devices connected to the fabric within the same zone. A device not included in the zone is not available to members of that zone. When zoning is enabled, devices that are not included in ***any*** zone configuration are inaccessible to all other devices in the fabric.

Zones can be configured dynamically. They can vary in size, depending on the number of fabric-connected devices, and devices can belong to more than one zone.

Consider Figure 395, which shows configured zones, Red, Green, and Blue.

- Server 1 can communicate only with the Storage 1 device.
- Server 2 can communicate only with the RAID and Storage 2 devices.
- Server 3 can communicate with the RAID and Storage 1 devices.
- The Storage 3 device is not assigned to a zone; no other zoned fabric device can access it.

FIGURE 395Zoning

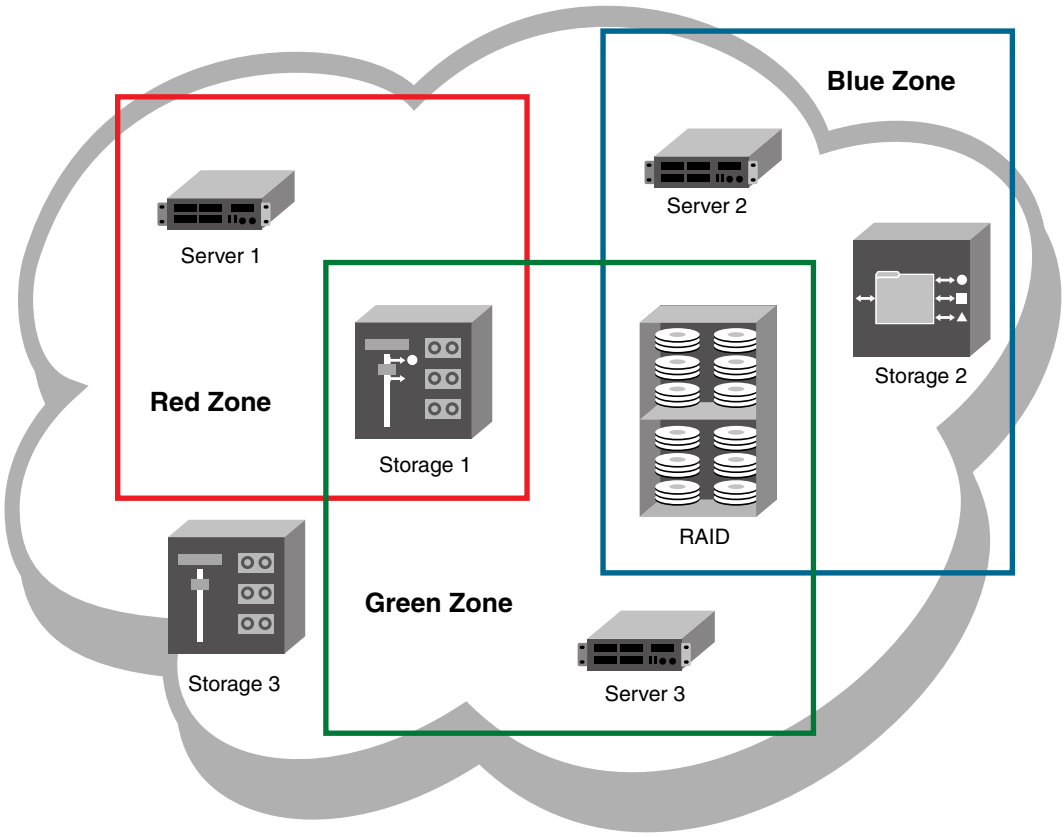

NOTE
Zone objects based on physical port number or port ID (D,I ports) are not supported in Network OS fabrics.

Types of zones

The following types of zones are supported:

- Standard zones

  Enable you to partition your fabric into logical groups of devices that can access each other. These are “regular” or “standard” zones. Unless otherwise specified, all references to zones refer to these standard zones.

- Frame redirection zones

  Reroute frames between an initiator and target through a Virtual Initiator and Virtual Target for special processing or functionality, such as for storage virtualization or encryption. Refer to “Redirection zones” on page 751 for more information.

- LSAN zones

  Provide device connectivity between fabrics without merging the fabrics. Refer to “LSAN zones” on page 805 for more information.

- LSAN Peer zones

  Provide device connectivity between LSAN Peer zone fabrics without merging the fabrics. Refer to “Creating an LSAN Peer zone” on page 831 for more information.

- QoS zones

  Assign high or low priority to designated traffic flows. Quality of Service (QoS) zones are standard zones with additional QoS attributes that you select when you create the zone. Beginning with the Management application 14.0.1, if the vTap/QoSH compatibility mode is enabled, the user will not be able to create or rename zones with "_QoSH5" prefixes.

- Traffic Isolation zones (TI zones)

  Isolate inter-switch traffic to a specific, dedicated path through the fabric. Refer to "Traffic Isolation zones" on page 810 for more information.

- Peer zones

  Allows communication between principal members and peer or non-principal members present within that zone. Refer to "Creating a Peer zone" on page 825 for more information.

- Target Driven Peer zones

  A Target Driven Peer zone is a special zone that allows you to read, delete, activate or deactivate its members present in the zone. Refer to "Target Driven Peer zones" on page 833 for more information.

Zoning best practices

The following are recommendations for using zoning:

- When using a mixed fabric — that is, a fabric containing two or more switches running different release levels of Fabric OS — you should use the switch with the latest Fabric OS level to perform zoning tasks.
- Switches with earlier versions of Fabric OS do not have the same capability to view all the functionality that more recent versions of Fabric OS provide, as functionality is backwards-compatible but not forward-compatible.
- Zone using the core switch in preference to using an edge switch.
- Zone using a Backbone rather than a switch. A Backbone has more resources to handle zoning changes and implementations.
- When you are adding a switch to an existing fabric, prior to joining the fabric you must set the defzone policy of the switch being added as follows:
  - If the joining switch has locally-attached devices that are online, the defzone policy of the switch being added should be set to "No Access".
  - If the joining switch has no online locally-attached devices the defzone policy of the switch being added can be set to "All Access".

  You can set the Zoning Policies in the Management application by clicking the Zoning Policies button in the Zoning dialog box. This is done to avoid a transitional state where the "All Access" policy might lead to excessive RSCN activity; with extreme cases having the potential for additional adverse effects. This is especially important for fabrics having a very high device count.
- Initial WWNs starting with '00' are for engineering use only and are auto-created. Starting with Fabric OS 7.3.0, zones having initial WWNs starting with '00' are treated as peer zones. This means that when you attempt one of the following,
  - A firmware upgrade from firmware earlier than Fabric OS 7.3.0
  - A merge with a switch running firmware earlier than Fabric OS 7.3.0
  - A configuration download using firmware earlier than Fabric OS 7.3.0

  and the WWN of the first zone member starts with '00', Fabric OS will treat the zone as a peer zone. You must delete such invalid zones before continuing (refer to "Removing a zone from a zone configuration" on page 796.

Online zoning

Online zoning allows you to do the following:

- View both defined and active zone information in the fabric.
- Create and modify zones and zone configurations in the software zone database.
- Activate a zone configuration in order to publish the zone information in the selected fabric.
- Deactivate the current active zone configuration.
- Configure zoning policies in the selected fabric.
- Generate zoning reports for the fabric.

Offline zoning

NOTE
Offline zoning is available only for Enterprise and Professional Plus editions.

Offline zoning enables you to copy a fabric zone database and edit it offline. The benefits to offline zoning include the following:

- You want to make changes to the zone database now, but apply them later.

  For example:

  - If you make incremental changes to zoning on an ongoing basis, but want to apply the changes to the fabric during scheduled downtime.
  - If you are expecting new servers to be delivered, but want to make changes to zoning now and apply the changes after the servers are delivered and ready to go online.
- You want to keep multiple copies of the zone database and switch between them.

  For example, if you want to allow specific servers access to tape drives for backup during specific time windows, you can have multiple zone databases (one or more for backup and one for normal operation) and switch between them easily.
- You want to analyze the impact of changes to storage access before applying the changes.

  For example, if you deploy a new server and want to ensure that the zoning changes result in only the new server gaining access to specific storage devices and nothing else. Refer to "Comparing zone databases" on page 817.

Zoning naming conventions

The naming rules for zone names, zone aliases, and zone configuration names vary with the type of fabric. The following conventions apply:

- Zone names, Zone configuration names, and Alias name can begin with an alphanumeric character or can have one or more special characters, such as " _" , "-", "$" or "^".
- Names are not case-sensitive.
- Zone, alias, and configuration names cannot begin with "bfa_","red_", "lsan_red_", or "d__efault__". Zone configuration names cannot begin with "r_e_d_i_r_c__fg". These prefixes are reserved.
- Names cannot begin with a special character.
- Normal zone names cannot begin with "bfa_" prefixes or end with "blun_" suffixes.
- The recommended character limit is 64 characters.
- Duplicate names are not allowed between zones, zone aliases, and zone configurations within a zone database.

If you enter an invalid zone or zone configuration name, an error or warning message displays depending on the type of fabric you are trying to zone.

Zoning and FICON

Session-based hardware enforcement is in effect if the zone has a mix of WWN and Domain,Port members.

Session-based hardware enforcement is also in effect if a port is in multiple zones, and is defined by WWN in one zone and by Domain,Port in another.

Session-based zoning enforcement is not recommended in a FICON environment.

When configuring a zone for FICON, ensure that all members of the zone, including members of any zone aliases, are either all WWN or all Domain,Port members, but not a mix of both.

Zone database size

The supported maximum zone database size is 1 MB for fabric with minimum one pizza box and 2MB for director-only fabrics.

If the fabric contains only Backbone Chassis platforms, the supported maximum zone database size is 2 MB.

Virtual Fabric considerations: If Virtual Fabrics is enabled, the sum of the zone database sizes on all of the logical fabrics must not exceed the maximum size allowed for the chassis (2MB) and pizza box (1MB).

The Professional Edition does not support large zone databases. In the Professional Edition, the maximum size of the zone database without zone aliases is 32 KB. If the zone database contains aliases, the maximum size is less than 32 KB.

Zoning configuration

At a minimum, zoning configuration entails creating zones and zone members. However, you can also create zone aliases, zone configurations, and zone databases. You can define multiple zone configurations, deactivating and activating individual configurations as your needs change. Zoning configuration can also involve enabling or disabling the default zone.

Configuring zoning

The following procedure provides an overview of the steps you must perform to configure zoning.

Note that for any zoning-related procedure, changes to a zone database are not saved until you click **OK** or **Apply** on the **Zoning** dialog box. If you click **Cancel** or the close button (X), no changes are saved.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays, as shown in Figure 396.

FIGURE 396 Zoning dialog box

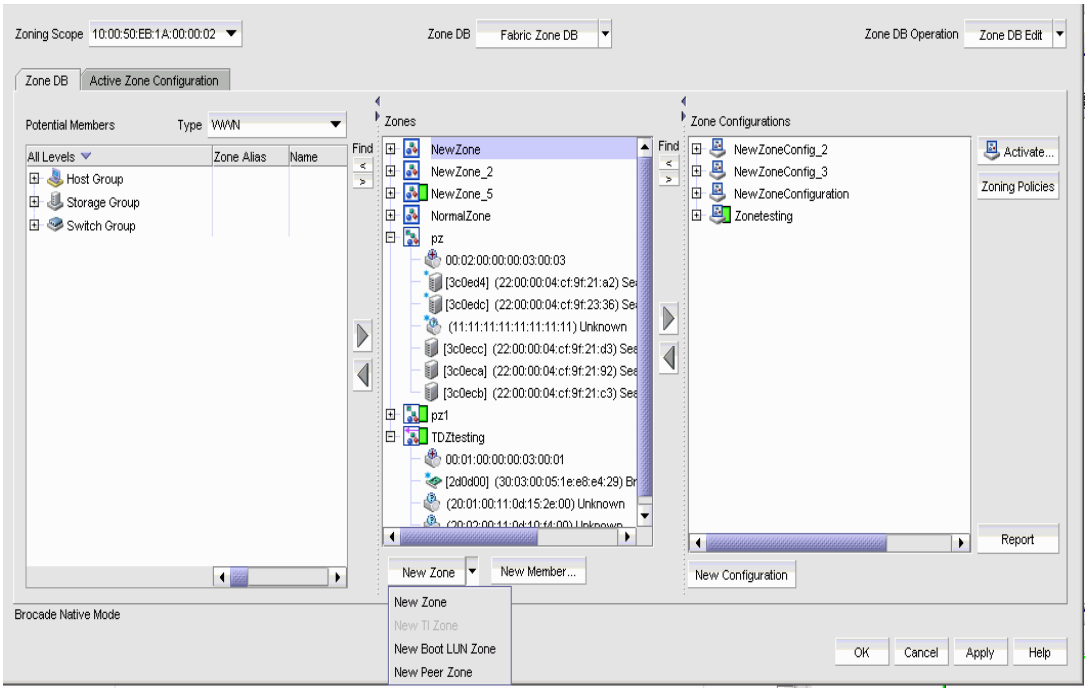

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.
4. If you want to show all the discovered fabrics in the **Potential Members** list, right-click in the **Potential Members** list and select **Display All**.
5. Create the zones.

   For specific instructions, refer to "Creating a zone" on page 785.
6. Add members to each zone.

   For specific instructions, refer to "Adding members to a zone" on page 786 and "Creating a new member in an LSAN zone" on page 808.
7. Create a zone configuration.

   For specific instructions, refer to "Creating a zone configuration" on page 794.
8. Activate the zone configuration.

   For specific instructions, refer to "Activating a zone configuration" on page 796.
9. Set zoning policies, if necessary.

   For specific instructions, refer to "Enabling or disabling the default zone for fabrics" on page 790.
10. Click **OK** or **Apply** to save your changes.

    Any zones or zone configurations you have changed are saved in the zone database.

Creating a zone

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Click **New Zone**.

   A new zone displays in the **Zones** list.

5. Type the name for the zone.

   For zone name requirements and limitations, refer to "Zoning naming conventions" on page 782.

6. (*Optional — Fabric OS only*) Set the QoS for the zone by right-clicking the zone and selecting **QoS >** ***Priority_Level*** (High, Medium, or Low).

   NOTE
   QoS priority support is available for zones with WWN or Domain,Index (D,I) members.

   QoS zones using D,I notation cannot be created if any of the switches in the fabric are running Fabric OS versions earlier than 6.3.0.

   The zone name is automatically renamed to QoS***X_Zone_Name***, where ***X*** is the priority level (H — High, M — Medium, or L — Low) and ***Zone_Name*** is the name you entered for the zone.

The new, empty zone is created. You cannot save an empty zone. Refer to "Adding members to a zone" on page 786 for instructions on adding members and saving the zone.

Viewing zone properties

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Right-click the zone you want to review in the **Zones** list and select **Properties**.

   The **Zone Properties** dialog box displays.

5. Review the zone properties.

   Note that when any modifications are made to an active zone, the **Zone Properties** dialog box continues to show the status as Active until the changes are saved to the zone database.

   You can change the zone name by double-clicking the name and then modifying the name in the editable field.

6. Click **OK** to close the **Zone Properties** dialog box.

Adding members to a zone

Use this procedure to add a member to a zone when the member is listed in the **Potential Members** list of the **Zone DB** tab.

Enterprise and Professional Plus editions: For instructions to add a member to a zone when the member is not listed in the **Potential Members** list, refer to the procedure "Creating a member in a zone" on page 787.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

   If you want to show all the discovered fabrics in your fabric group in the **Potential Members** list, right-click in the **Potential Members** list and select **Display All**.

4. Select one or more zones to which you want to add members in the **Zones** list. (Press **SHIFT** or **CTRL** and click each zone name to select more than one zone.)
5. Select an option from the **Type** list.

   By default, the first time you launch the **Zoning** dialog box for a zoning scope, the **Potential Members** list displays valid members using the following rules:

   - If you select the **WWN** type, the valid members display by the Attached Ports.
   - If you select the **WWN-Fabric Assigned** type, the valid members display by the ports on which FA-PWWN is configured.
   - If you select the **Domain,Port Index** type, the valid members display by ALL Product Ports (both occupied and unoccupied). This option is available for FC fabrics only.
   - If you select the **Alias** type, the valid members display by the device Alias.

6. Select one or more members to add to the zone in the **Potential Members** list. (Press **SHIFT** or **CTRL** and click each member to select more than one member. To add all ports on a device, select the device.)

   You cannot add duplicate members to the same zone.

7. Click the right arrow between the **Potential Members** list and **Zones** list to add the selected members to the zone.

   A message is displayed if unsupported potential members are moved to the **Zones** list. Click **OK** to close the message box. Reconsider your selections and make corrections as appropriate.

8. For offline zone databases only, complete the following steps to save the zone configuration into the switch from the offline zone database:

   a. Select **Save to Switch** from the **Zone DB Operation** list.

   b. Click **Yes** on the confirmation message.

      The selected zone database is saved to the fabric without enabling a specific zone configuration.

9. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Creating a member in a zone

Use this procedure to add a member to a zone when the member is not listed in the **Potential Members** list of the **Zone DB** tab.

For instructions to add a member to a zone when the member is listed in the **Potential Members** list, refer to the procedure "Adding members to a zone" on page 786.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select one or more zones to which you want to add members in the **Zones** list. (Press **SHIFT** or **CTRL** and click each zone name to select more than one zone.)
5. Click **New Member**.

   The **Add Zone Member** dialog box displays.

6. Select an option from the **Member Type** list.

   The fields in the dialog box vary based on the **Member Type** option you select.

7. Fill in the remaining fields in the dialog box.

   Click the **Help** button for additional information on each field.

8. Click **OK** to save your changes and close the **Add Zone Member** dialog box.

   OR

   Click **Apply** to save your changes and keep the **Add Zone Member** dialog box open so you can add more new members. Repeat step 5 through step 8 as many times as needed, and proceed to step 9 when appropriate.

9. For offline zone databases only, complete the following steps to save the zone configuration into the switch from the offline zone database:

   a. Select **Save to Switch** from the **Zone DB Operation** list.

   b. Click **Yes** on the confirmation message.

      The selected zone database is saved to the fabric without enabling a specific zone configuration.

10. Click **OK** or **Apply** to save your changes.

    Any zones or zone configurations you have changed are saved in the zone database.

Removing a member from a zone

Use the following procedure to remove one or more members from a zone or zones. Note that the member is not deleted; it is only removed from the zone.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.
4. Click the plus sign (+) by the appropriate zone in the **Zones** list to expand the listing and show the zone's members.
5. Perform one of the following actions:
   - Right-click the name of the zone member you want to remove in the **Zones** list and select one of the following options from the shortcut menu that displays:
     - **Remove** - To remove the zone member from the selected zone.
     - **Remove All** - To remove the zone member from all zones to which it belongs.
   - To remove multiple zone members, select the members to be removed from the zone, and click the left arrow between the **Potential Members** list and the **Zones** list.

   When successful, the zone member is removed from the **Zones** list.
6. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Renaming a zone

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.
4. Right-click the name of the zone you want to change in the **Zones** list and select **Rename**.
5. Type the new name for the zone.

   For zone name requirements and limitations, refer to "Zoning naming conventions" on page 782.
6. Press **Enter** to save the new name.

   For FC fabrics, if an invalid name is entered for a zone or zone configuration, the application displays a warning message. If there is a naming violation according to the vendor, the switch returns the error message for the exact information along with the zone configuration activation failure message.
7. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

   NOTE
   You can rename Zone names, Zone configuration names, and Alias name. Names can begin with an alphanumeric character and can have one or more special characters, such as " _" , "-", "$" or "^".

Deleting a zone

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select one or more zones in the **Zones** list that you want to delete, then right-click and select **Delete**.

   A message displays asking you to confirm the deletion.

5. Click **Yes** to delete the selected zones.

   The message closes and the zone or zones are removed from the **Zones** list.

   NOTE
   If you delete something in error, click **Cancel** on the **Zoning** dialog box to exit without saving changes. When you reopen the dialog box, the zone is restored.

6. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Duplicating a zone

When you duplicate a zone, you make a copy of it in the same zone database. The first time a zone is duplicated, the duplicate is automatically given the name <*zonelabel*>*_copy*. On subsequent duplications, a sequential number is assigned to the zone name, such as <*zonelabel*>*_copy_1*, <*zonelabel*>*_copy_2,* and <*zonelabel*>*_copy_3*.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select one or more zones in the **Zones** list that you want to duplicate, then right-click and select **Duplicate**.

   The duplicated zone or zones display in the **Zones** list.

5. (*Optional*) Type a new name for the zone and press **Enter** to save the name.

   Depending on the characters included in the name you enter, a message may display informing you the name contains characters that are not accepted by some switch vendors. Click **OK** and enter a different name or accept the default name assigned to the zone. (For zone name requirements and limitations, refer to "Zoning naming conventions" on page 782.)

6. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Customizing the zone member display

In the **Zoning** dialog box, you can customize which properties are displayed and in what order.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays, based on the **Configure > Zoning** menu selection.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Click the plus sign (+) by the appropriate zone in the **Zones** list to expand the listing and show the zone members.
5. Right-click the name of any zone member and select **Member Display**.

   The **Zone Member Display** dialog box displays, as shown in Figure 397.

   FIGURE 397 Zone Member Display dialog box

6. Select or clear the check boxes for the properties you want to display or hide.

   All of the options are selected by default. You cannot clear the **WWN/Domain,Port Index** check box. It is always selected.

7. Select a property and click the **Up** or **Down** buttons to rearrange the order in which the properties are displayed.
8. Click **OK**.

   The display is changed for all zone members in the **Zones** list.

Enabling or disabling the default zone for fabrics

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select the zoning database you want from the **Zone DB** list.
5. Click **Zoning Policies**.

   The **Zoning Policies** dialog box displays.

6. Make sure the appropriate fabric is named on the **Zoning Policies** dialog box.
7. Perform one of the following actions based on the task you want to complete:
   - To enable the default zone, click **Enable**, and then click **OK**.
   - To disable the default zone, click **Disable**, and then

When you are adding a switch to an existing fabric, prior to joining the fabric you must set the defzone policy of the switch being added as follows:

- If the joining switch has locally-attached devices that are online, the defzone policy of the switch being added should be set to "No Access".
- If the joining switch has no online locally-attached devices the defzone policy of the switch being added can be set to "All Access".

You can set the Zoning Policies in the Management application by clicking the Zoning Policies button in the Zoning dialog box. This is done to avoid a transitional state where the "All Access" policy might lead to excessive RSCN activity; with extreme cases having the potential for additional adverse effects. This is especially important for fabrics having a very high device count.

8. Click **OK** on the **Zoning Policies** dialog box.

   The **Zoning Policies** dialog box closes and the **Zone DB** tab displays.

9. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Creating a zone alias

An alias is a logical group of port index numbers and WWNs. Specifying groups of ports or devices as an alias makes zone configuration easier by enabling you to configure zones using an alias rather than inputting a long string of individual members. You can specify members of an alias using the following methods:

- Identifying members by switch domain and port index (D,I) number pair.
- Identifying members by device node and device port WWNs.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.
4. Select **Alias** from the **Type** list.
5. Click **New Alias**.

   The **New Alias** dialog box displays.

6. Type a name for the alias in the **Alias Name** field.

   Refer to "Zoning naming conventions" on page 782 for rules about zone alias names.

7. (*Optional*) Select an option from the **Type** list to choose how to display the objects in the **Potential Members** list.
8. (*Optional*) Show all discovered fabrics in the **Potential Members** list by right-clicking in the **Potential Members** list and selecting **Display All**.

   This right-click option is not available if you selected **WWN-Fabric Assigned** in the **Type** list.

9. Select one or more members that you want to add to the alias in the **Potential Members** list. (Press **SHIFT** or **CTRL** and click each member to select more than one member.)

You can also add WWNs not listed in the **Potential Members** list by entering the WWN in the **Detached WWN** field and clicking **Add**.

10. Click the right arrow between the **Potential Members** list and the **Selected Member(s)** list to add the selected members to the alias.

    NOTE
    Beginning with Fabric OS 8.1.0 and later, you can create zone alias with device Port WWN by default. Right-click in the P**otential Members** and select **Move Node WWN** check box. Only the device Port WWN is added to the **Selected Member(s)** list.

11. Click **OK** or **Apply** on the **New Alias** dialog box to save your changes.
12. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

Editing a zone alias

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.
4. Select **Alias** from the **Type** list.
5. Select the alias you want to edit in the **Alias** list and click **Edit**.

   The **Edit Alias** dialog box displays.

6. Add members to the alias by completing the following steps.
   a. Select an option from the **Type** list to choose how to display the objects in the **Potential Members** list.
   b. Show all discovered fabrics in the **Potential Members** list by right-clicking in the **Potential Members** list and selecting **Expand All**.

      This right-click option is not available if you selected **WWN-Fabric Assigned** in the **Type** list.

   c. Select one or more members that you want to add to the alias in the **Potential Members** list. (Press **SHIFT** or **CTRL** and click each member to select more than one member.)

      You can also add WWNs not listed in the **Potential Members** list by entering the WWN in the **Detached WWN** field and clicking **Add**.

   d. Click the right arrow between the **Potential Members** list and the **Selected Member(s)** list to add the selected members to the alias.

7. Remove members from the alias by completing the following steps.
   a. Select one or more members that you want to remove from the alias in the **Selected Member(s)** list. (Press **SHIFT** or **CTRL** and click each member to select more than one member.)
   b. Click the left arrow between the **Potential Members** list and the **Selected Member(s)** list to remove the selected members from the alias.

8. Click **OK** or **Apply** on the **Edit Alias** dialog box to save your changes.
9. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

Removing an object from a zone alias

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.
4. Select **Alias** from the **Type** list.
5. Show all objects in the **Alias** list by right-clicking an object and selecting **Tree > Expand All**.
6. Select one or more objects that you want to remove from the alias in the **Alias** list. (Press **SHIFT** or **CTRL** and click each member to select more than one member.)

   You can select objects from different zone aliases.

7. Right-click one of the selected objects and select **Remove.**

   The selected objects are removed from the associated zone aliases.

8. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

Exporting zone aliases

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.
4. Select **Alias** from the **Type** list.
5. Click **Export.**

   The **Export Alias** dialog box displays.

6. Browse to the location to which you want to export the zone alias data.
7. Enter a name for the export file in the **File Name** field.
8. Click **Export Alias**.
9. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

Renaming a zone alias

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.
4. Select **Alias** from the **Type** list.

5. Right-click the zone alias you want to rename and select **Rename.**
6. Edit the name and press **Enter**.

   Refer to "Zoning naming conventions" on page 782 for rules about zone alias names.
7. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

Deleting a zone alias

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.
4. Select **Alias** from the **Type** list.
5. Right-click the zone alias you want to delete and select **Delete.**
6. Click **Yes** on the confirmation message.

   The selected zone alias is deleted from the **Alias** list.
7. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

Duplicating a zone alias

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.
4. Select **Alias** from the **Type** list.
5. Right-click the zone alias you want to duplicate and select **Duplicate.**

   The duplicated zone alias displays in the **Alias** list (for example, <*Zone_Alias*>_Copy).
6. Edit the name.

   To edit the name, refer to "Renaming a zone alias" on page 793.
7. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

Creating a zone configuration

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Click **New Configuration**.

   A new configuration displays in the **Zone Configurations** list.

5. Enter a name for the zone configuration.

   For zone name requirements and limitations, refer to "Zoning naming conventions" on page 782.

6. Press **Enter**.

   Depending on the characters included in the name you enter, a message may display informing you the name contains characters that are not accepted by some switch vendors. Click **OK** and enter a different name or accept the default name assigned to the zone. (For zone name requirements and limitations, refer to "Zoning naming conventions" on page 782.)

7. Add zones to the zone configuration.

   For step-by-step instructions, refer to "Adding zones to a zone configuration" on page 795.

8. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Viewing zone configuration properties

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Potential Members** list.
4. Right-click the zone configuration you want to review in the **Zone Configurations** list and select **Properties**.

   The **Zone Configuration Properties** dialog box displays.

5. Review the zone configuration properties.
6. Click **OK** to close the **Zone Configuration Properties** dialog box.

Adding zones to a zone configuration

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select one or more zone configurations to which you want to add zones in the **Zone Configurations** list. (Press **SHIFT** or **CTRL** and click each zone configuration name to select more than one zone configuration.)
5. Select one or more zones to add to the zone configurations in the **Zones** list. (Press **SHIFT** or **CTRL** and click each zone name to select more than one zone.)

6. Click the right arrow between the **Zones** list and the **Zone Configurations** list to add the zones to the zone configurations.
7. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Removing a zone from a zone configuration

Use the following procedure to remove a zone from a zone configuration. Note that the zone is not deleted; it is only removed from the zone configuration.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.
4. Click the plus sign (+) by the appropriate zone configuration in the **Zone Configurations** list to expand the listing and show the zone configuration members.
5. Perform one of the following actions:
   - Right-click the name of the zone you want to remove in the **Zone Configurations** list and select **Remove**.
   - To remove multiple zones, select the zones to be removed from the zone configuration, and click the left arrow between the **Zones** list and the **Zone Configurations** list.

   When successful, the zone is removed from the **Zone Configurations** list.
6. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Activating a zone configuration

When a zone configuration is active, its members can communicate with one another. Only one zone configuration can be active at any given time.

> NOTE
> Only one server should be run at a time (actual servers performing discovery) or logon conflicts may occur. Also, activation speeds may differ depending on the hardware vendor and type of zoning used.

You cannot activate a zone configuration if any of the following is true:

- You do not have access privileges to activate zone configurations. You will not be able to activate a zone configuration unless your access privileges are redefined.
- The fabric is not manageable.
- You do not have Read/Write or Activate privileges for the selected fabric and the selected zone database (for FC fabrics only).
- The selected fabric is not supported by the Management application.
- The selected fabric is no longer discovered.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. (*Optional*) Select a zone database from the **Zone DB** list (Enterprise and Professional Plus editions only).
5. Select the zone configuration you want to activate in the **Zone Configurations** list.
6. Click **Activate**.
7. Review the information in the **Activate Zone Configuration** dialog box.
   a. Make sure the selected zone configuration is the one you want to activate.
   b. (*Optional*) Select the **Generate a report with the activation of new zone configuration** check box to generate the Zone Configuration Activation report.
   c. If you are activating a zone configuration from the offline zone database, select or clear the **Save only the selected zone configuration to the existing zone database in the fabric** check box.
      - If the check box is cleared (default), the entire offline zone database is saved to the switch and replaces the existing online zone database.
      - If the check box is selected, only the selected zone configuration and any TI zones in the offline zone database are saved to the switch and are added to the existing online zone database.
8. Click **OK** to activate the zone configuration.

   A message displays informing you that the zones and zone configurations you change will be saved in the zone database and asking whether you want to proceed. Click **Yes** to confirm the activation, or click **No** to cancel the activation.

   When you click **Yes**, a busy window displays indicating the activation is in progress. A status field informs you whether the activation succeeded or failed. When it succeeds, icons for the active zone configuration and its zones display green. When it fails, the message includes the reason for the failure.

9. Click **OK** to continue.

   The **Activate Zone Configuration** dialog box is closed and the **Zone DB** tab displays.

10. Click **OK**.

    The zone configuration is activated and the entire zone database is sent to the fabric. If the members of the active zone are a part of any alias, the alias name is shown in parentheses in the fabric tree. If the members belong to multiple aliases, then the alias names are displayed in comma-separated format.

Deactivating a zone configuration

Use this procedure to deactivate the active zone configuration.

There are several conditions that could cause the **Deactivate** button to be unavailable. They include the following:

- There is no active zone configuration in the selected fabric.
- The fabric is not manageable.
- You do not have Read/Write or Activate privileges for the selected fabric and the selected zone database (for FC fabrics only).

- The selected fabric is not supported by the Management application.
- The selected fabric is no longer discovered.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Active Zone Configuration** tab.
3. Select a fabric from the **Active Zone Configuration** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Click **Deactivate**.
5. Click **Yes** on the confirmation message.

   If the deactivation succeeded, the zone configuration no longer displays in the **Active Zone Configuration** tab.
   If the deactivation failed, the zone configuration still displays in the **Active Zone Configuration** tab.

6. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Renaming a zone configuration

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Right-click the name of the zone configuration you want to change in the **Zone Configurations** list and select **Rename**.
5. Type the new name for the zone configuration.

   For zone configuration name requirements and limitations, refer to "Zoning naming conventions" on page 782.

6. Press **Enter** to save the new name.
7. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Deleting a zone configuration

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select one or more zone configurations in the **Zone Configurations** list that you want to delete, then right-click and select **Delete**.

   A message displays asking you to confirm the deletion.

5. Click **Yes** to delete the selected zone configuration.

   The message closes and the selected zone configurations are removed from the **Zone Configurations** list.

   NOTE
   If you select "**Do not show me this again.**" on the confirmation message, the next time you delete a zone configuration, it will be deleted without requesting confirmation from you. If you delete something in error, click **Cancel** on the **Zoning** dialog box to exit without saving changes since the last operation (**Apply** or **Activate**). When you reopen the dialog box, the zone configuration is restored.

6. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Duplicating a zone configuration

When you duplicate a zone configuration, you make a copy of it in the same zone database. The first time a zone configuration is duplicated, the duplicate is automatically given the name <*zonesetlabel*>*_copy*. On subsequent duplications, a sequential number is assigned to the zone configuration name, such as <*zonesetlabel*>*_copy_1*, <*zonesetlabel*>*_copy_2,* and <*zonesetlabel*>*_copy_3*.

Note that these naming conventions apply to both duplicate and deep duplicate operations.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.

3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select one or more zone configurations in the **Zone Configurations** list that you want to duplicate, then right-click and select one of the following options:

   - **Duplicate** - To duplicate the zone configuration or configurations.
   - **Deep Duplicate** - To duplicate the zone configuration or configurations ***and*** all included zones.

   The duplicated zone configuration or sets display in the **Zone Configurations** list.

5. (*Optional*) Type a new name for the zone configuration and press **Enter** to save the name.

   For zone name requirements and limitations, refer to "Zoning naming conventions" on page 782.

6. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Creating an offline zone database

Offline zone databases are supported only in Enterprise and Professional Plus editions. Use this procedure to create a zone database and save it offline.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a zone database from the **Zone DB** list.
4. Select **Save As** from the **Zone DB Operation** list.

   The **Save Zone DB As** dialog box displays.

5. Enter a name for the database in the **Zone DB Name** field and click **OK**.
6. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

   If you want to show all discovered fabrics in the **Potential Members** list, right-click in the **Potential Members** list and select **Display All**.

7. Create the desired zones.

   For specific instructions, refer to "Creating a zone" on page 785.

8. Add members to each zone.

   For specific instructions, refer to "Adding members to a zone" on page 786 and "Creating a member in a zone" on page 787.

9. Create a zone configuration.

   For specific instructions, refer to "Creating a zone configuration" on page 794.

10. Activate the zone configuration.

    For specific instructions, refer to "Activating a zone configuration" on page 796.

11. Set zoning policies, if necessary.

    For specific instructions, refer to "Enabling or disabling the default zone for fabrics" on page 790.

12. Click **OK** or **Apply** to save your changes.

    Any zones or zone configurations you have changed are saved in the zone database.

Deleting an offline zone database

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning databases for the selected entity.

3. Select the offline zone database you want to delete in the **Zone DB** list.

   NOTE
   Only offline databases can be deleted.

4. Select **Delete** from the **Zone DB Operation** list.
5. Click **Yes** on the confirmation message.

   The message closes and the selected zone configurations are removed from the **Zone Configurations** list.

6. Click **OK** to save your work and close the **Zoning** dialog box.

   Any zones or zone configurations you have changed are saved in the zone database.

Refreshing a zone database

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a zone database from the **Zone DB** list.
4. Select **Refresh** from the **Zone DB Operation** list.

   A message displays informing you that refresh will overwrite the selected database. Click **Yes** to continue.

5. Click **OK**.

   Any zones or zone configurations you have changed are saved in the zone database.

Merging fabrics

When you merge fabrics, the defined and active zone configurations in both fabrics must match.

1. Compare and merge the two zone databases, and save the database to the offline repository.

   Refer to "Merging two zone databases" on page 802.

2. Ensure that the active zone configurations in each fabric are the same, including the same name.

   Refer to "Renaming a zone configuration" on page 798.

3. Load the newly merged zone database from the offline repository.
4. Activate the zone configuration.
5. If the active zone configuration names are the same in each fabric, then load the offline repository, and activate the zone configuration on each fabric.
6. If the active configuration names are different in each fabric, rename the zone configurations to be the same, and copy the zones.
7. Ensure that the active configurations are the same.
   a. Load the newly created offline zone database.
   b. Add the active zones to the zone configuration that is the active configuration on the other fabric.

c. Rename the inactive configuration.

Merging two zone databases

If a zone or zone configuration is merged, the resulting zone or zone configuration includes ***all*** members (including those selected for addition or removal).

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Select **Compare** from the **Zone DB Operation** list.

   The **Compare/Merge Zone DBs** dialog box displays, as shown in Figure 398.

   FIGURE 398 Compare/Merge Zone DBs dialog box

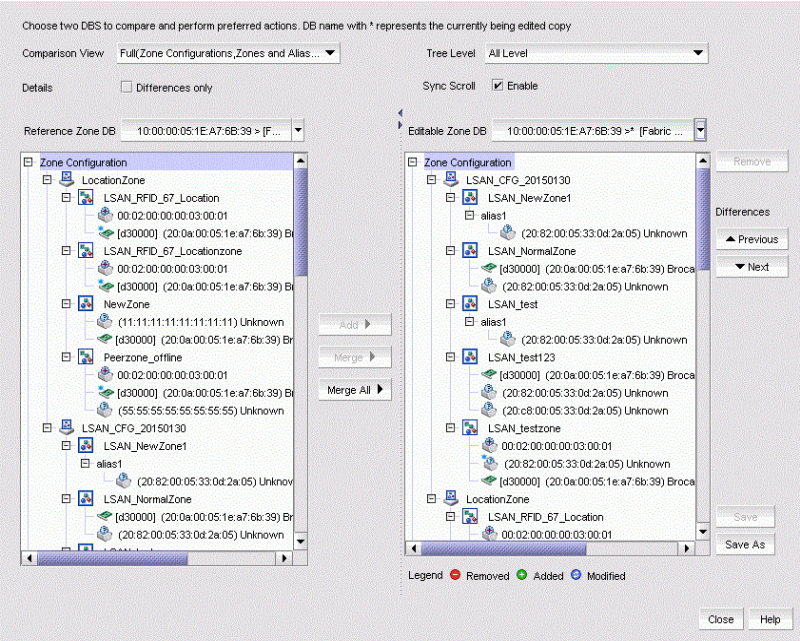

3. Select a database from the **Reference Zone DB** list.
4. Select a database from the **Editable Zone DB** list.

   The **Reference Zone DB** and **Editable Zone DB** areas display all available element types (zone configurations, zones, and aliases) for the two selected zone databases. In the **Editable Zone DB** area, each element type and element display with an icon indicator (Table 62) to show the differences between the two databases.

5. (*Optional*) Merge elements (zone configurations, zones, or aliases) by completing the following steps:

   a. Select one or more of the same element type from the **Reference Zone DB** area.

      You can select zone configurations/zones/aliases, but combination of element types are not allowed.

   b. Select the same type of element in the **Editable Zone DB** area.

      If you select a zone configuration in the **Reference Zone DB** area, you must select a zone configuration (same element type) in the **Editable Zone DB** area too.

c. Click **Merge**.

If the **Merge** button is inactive, make sure you have selected similar element types in both the **Reference Zone DB** area and the **Editable Zone DB** area. You can merge elements only with similar elements. For example, you cannot merge a zone with a zone configuration.

6. (*Optional*) Merge all elements by clicking **Merge All**.
7. (*Optional*) Add elements (zone configurations, zones, or aliases) to the editable database by completing the following steps.
   a. Select one or more of the same element type in the **Reference Zone DB** area.

   The selected elements are added to the editable zone database.

   b. Select an element in the **Editable Zone DB** area.

   You can add zone aliases and zone members to a zone, zones to a zone configuration, and zone configurations to the zone database.

   c. Click **Add**.

   If the **Add** button is inactive, make sure you have selected appropriate element types in both the **Reference Zone DB** area and the **Editable Zone DB** area.
8. (*Optional*) Remove elements from the editable zone database by selecting an available element from the **Editable Zone DB** area and click **Remove**.

   Note that if a zone is removed from a zone configuration, it is removed ***only*** from that single zone configuration. However, if the zone is removed from the list of zones, it is removed from ***all*** zone configurations.
9. Click **Save** to save the editable zone database to switch.
10. Click **Save As** to save the editable zone database in the offline repository (for Enterprise and Professional Plus editions only).

Creating a common active zone configuration in two fabrics

Before you can merge two fabrics, the defined and active zone configurations in both fabrics must match. Refer to "Merging two zone databases" on page 802 for instructions on how to merge the zone databases in two fabrics.

After you merge the two zone databases, you create a common active zone configuration before physically merging the fabrics.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Select **Compare** from the **Zone DB Operation** list.

   The **Compare/Merge Zone DBs** dialog box displays, as shown in Figure 398.
3. Select the database for the first fabric from the **Reference Zone DB** list.
4. Select the database for the second fabric from the **Editable Zone DB** list.
5. Set up a zone configuration that contains the active zones in both fabrics:
   a. Select the name of the active zone configuration from the **Reference Zone DB** area.
   b. Select the name of the active zone configuration in the **Editable Zone DB** area.
   c. Click **Merge**.

   All of the active zones from both fabrics are now in one zone configuration.

6. Click **Save As** to save the editable zone database in the offline repository for the second fabric.
7. Click **Save As** again, and select the name of the first fabric from the **Fabric** list to save the editable zone database in the offline repository for the first fabric.
8. Click **Close** to close the **Compare/Merge Zone DBs** dialog box and return to the **Zoning** dialog box.
9. In both fabrics, load the offline repository and activate the zone configuration from step b.

   Refer to "Activating a zone configuration" on page 796 for instructions.

Saving a zone database to a switch

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Select a zone database from the **Zone DB** list.
3. Select **Save to Switch** from the **Zone DB Operation** list.
4. Click **Yes** on the confirmation message.

   The selected zone database is saved to the fabric without enabling a specific zone configuration.
5. Click **OK** to save your work and close the **Zoning** dialog box.

Exporting an offline zone database

NOTE
You cannot export an online zone database.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Select an offline zone database from the **Zone DB** list.
3. Select **Export** from the **Zone DB Operation** list.

   The **Export Zone DB** dialog box displays.
4. Browse to the location where you want to export the zone database file (.xml format).
5. Click **Export Zone DB**.
6. Click **OK** to save your work and close the **Zoning** dialog box.

Importing an offline zone database

NOTE
You cannot import an online zone database. You cannot import a zone database that contains zones with duplicate members.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Select an offline zone database from the **Zone DB** list.
3. Select **Import** from the **Zone DB Operation** list.

   The **Import Zone DB** dialog box displays.
4. Browse to the zone database file (.xml format).
5. Click **Import Zone DB**.
6. Click **OK** to save your work and close the **Zoning** dialog box.

Rolling back changes to the offline zone database

Use this procedure to reverse changes made to an offline zone database.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Select the zone database you want to roll back from the **Zone DB** list.

   You must select an offline zone database that has a value in the **Last Saved to Fabric** column. You cannot roll back changes for zone databases that were never saved to the fabric.
3. Select **Roll Back** from the **Zone DB Operation** list.

   The selected zone database reverts back to what it was before the changes were applied.
4. Click **OK** to save your work and close the **Zoning** dialog box.

LSAN zones

Connecting to another network through a Fibre Channel (FC) router, you can create an LSAN zone to include zone objects on other fabrics. No merging takes place across the FC router when you create an LSAN zone.

Supported configurations for LSAN zoning

LSAN zoning is available only for backbone fabrics and any directly connected edge fabrics. A backbone fabric is a fabric that contains an FC router. All discovered backbone fabrics have the prefix LSAN_ in their fabric name, which is listed in the **Zoning Scope** list.

LSAN zones are supported between the following types of fabrics:

- Fabric OS and Fabric OS

  NOTE
  LSAN zoning is supported only in Enterprise and Professional Plus editions.

Configuring LSAN zoning

The following procedure provides an overview of the steps you must perform to configure LSAN zoning.

1. Select a backbone fabric from the Connectivity Map or Product List.
2. Select **Configure > Zoning > LSAN Zoning (Device Sharing)**.

   The **Zoning** dialog box displays, with the LSAN scope.

3. Click the **Zone DB** tab if that tab is not automatically displayed.
4. If you want to show all edge fabrics in your backbone fabric in the **Potential Members** list, right-click a device and select **Table > Expand All**.
5. Create the LSAN zones.

   For specific instructions, refer to "Creating an LSAN zone" on page 806.
6. Add members to each zone.

   For specific instructions, refer to "Adding members to the LSAN zone" on page 807.

   NOTE
   You cannot add an LSAN zone to a zone configuration. LSAN zones are automatically added to the active zone configuration. If the fabric does not have an active zone configuration, then a zone configuration with the name LSAN_CFG_*timestamp* is automatically created and the LSAN zone is added to it.
7. Click **Activate.**

   The **Activate LSAN Zones** dialog box displays.
8. Review the information in the **Activate LSAN Zones** dialog box.

   LSAN zones that contain online members are automatically included in the **Destination Fabrics** list. For LSAN zones that contain offline members, you can click the right arrow button to assign these zones to fabrics in the **Destination Fabrics** list.
9. Click **OK** to activate the LSAN zones and close the dialog box.

   A message displays informing you about the effects of LSAN zone activation and asking whether you want to proceed. Click **Yes** to confirm the activation, or click **No** to cancel the activation.

   You are prompted whether to activate the LSAN zone on the edge fabrics and backbone fabric. If the LSAN zone contains only online members, however, you are prompted only for the backbone fabric, and activation on the edge fabrics occurs automatically.
10. Click **OK** to close the **Zoning** dialog box.

Creating an LSAN zone

Create LSAN zones to enable communication between devices in different fabrics without merging the fabrics.

1. Select a backbone fabric from the Connectivity Map or Product List.
2. Select **Configure > Zoning > LSAN Zoning (Device Sharing)**.

   The **Zoning** dialog box displays, with the LSAN scope.
3. Click **New Zone**.

   The prefix LSAN_ is automatically added in the text field.
4. Enter a name for the zone.

   If LSAN tagging is configured, the zone name must match one of the configured tags.

   For zone name requirements and limitations, refer to "Zoning naming conventions" on page 782.
5. Press **Enter**.

6. Add members to the LSAN zone.
    a. Select one or more members to add to the zone in the **Potential Members** list. (Press **SHIFT** or **CTRL** and click each member to select more than one member.)
    b. Select an option from the **Type** list.

        For DCB-capable switches, you may need to change the port display options to see the ports. Right-click in the **Potential Members** list and select **Port Display** to change the options.
    c. Click the right arrow between the **Potential Members** list and the **Zones** list to add the selected members to the zone.
7. Click **Activate.**
8. Review the information in the **Activate LSAN Zones** dialog box.

    LSAN zones that contain online members are automatically included in the **Destination Fabrics** list. For LSAN zones that contain offline members, you can click the right arrow button to assign these zones to fabrics in the **Destination Fabrics** list.
9. Click **OK** to activate the LSAN zones.

    A message displays informing you about the effects of LSAN zone activation and asking whether you want to proceed. Click **Yes** to confirm the activation, or click **No** to cancel the activation.
10. Click **OK** to continue.

    All LSAN zones are activated on the selected fabrics and saved to their respective zone databases.
11. Click **OK** to close the **Zoning** dialog box.

Location Embedded LSAN zones

Beginning with Fabric OS v7.4.0 or later, you can read, edit, and delete the Location Embedded LSAN zones to enhance the LSAN zone and device mapping by the FC Routing. You can create the Location Embedded LSAN zones with the "LSAN_<name_part1>_RFID_<EdgeFabricId><name_part2>" naming convention, in which <name_partn> is optional. The Location Embedded LSAN zones must be pushed to their fabrics to be read by the Management application. You cannot move the Location Embedded LSAN zones to other fabrics.

Adding members to the LSAN zone

Use this procedure to add a member to an LSAN zone when the member is listed in the **Potential Members** list of the **Zone DB** tab.

LSAN zones do not support Domain,Port members.

1. Select a backbone fabric from the Connectivity Map or Product List.
2. Select **Configure > Zoning > LSAN Zoning (Device Sharing)**.

    The **Zoning** dialog box displays, with the LSAN scope.
3. Select the member type from the **Type** list.
    - If you select the **WWN** type, the valid members display by the Attached Ports.
    - If you select the **WWN-Fabric Assigned** type, the valid members display by the ports on which FA-PWWN is configured.

- If you select the **Alias** type, the valid members display by the device alias. Only aliases with WWN member types are displayed. Aliases that contain any Domain,Port members are not displayed.

For DCB-capable switches, you may need to change the port display options to see the ports. Right-click in the **Potential Members** list and select **Port Display** to change the options.

4. In the **Potential Members** list, select one or more members to add to the zone. (Press **Shift** or **Ctrl** and click each member to select more than one member.)

   If you want to show all discovered fabrics in the **Potential Members** list, right-click anywhere in the table and select **Display All**.

5. In the **Zones** list, select one or more LSAN zones to which you want to add members. (Press **Shift** or **Ctrl** and click each zone name to select more than one zone.)
6. Click the right arrow between the **Potential Members** list and the **Zones** list to add the selected members to the zone.
7. Click **Activate.**
8. Review the information in the **Activate LSAN Zones** dialog box.

   LSAN zones that contain online members are automatically included in the **Destination Fabrics** list. For LSAN zones that contain offline members, you can click the right arrow button to assign these zones to fabrics in the **Destination Fabrics** list.

9. Click **OK** to activate the LSAN zones.

   A message displays informing you about the effects of LSAN zone activation and asking whether you want to proceed. Click **Yes** to confirm the activation, or click **No** to cancel the activation.

10. Click **OK** to continue.

    All LSAN zones are activated on the selected fabrics and saved to their respective zone databases.

11. Click **OK** to close the **Zoning** dialog box.

Creating a new member in an LSAN zone

Use this procedure to add a member to an LSAN zone when the member is not listed in the **Potential Members** list of the **Zone DB** tab.

For instructions to add a member to a zone when the member is listed in the **Potential Members** list, refer to the procedure "Adding members to the LSAN zone" on page 807.

1. Select a backbone fabric from the Connectivity Map or Product List.
2. Select **Configure > Zoning > LSAN Zoning (Device Sharing)**.

   The **Zoning** dialog box displays, with the LSAN scope.

3. Select one or more zones to which you want to add members in the **Zones** list. (Press **SHIFT** or **CTRL** and click each zone name to select more than one zone.)
4. Click **New Member**.

   The **Add Zone Member** dialog box displays.

5. Select an option from the **Member Type** list.

   The fields in the dialog box vary based on the **Member Type** option you select.

6. Fill in the remaining fields in the dialog box.

   Click the **Help** button for additional information on each field.

7. Click **OK** to save your changes and close the **Add Zone Member** dialog box.

   OR

   Click **Apply** to save your changes and keep the **Add Zone Member** dialog box open so you can add more new members. Repeat step 3 through step 6 as many times as needed, and proceed to step 8 when you have finished adding members.

8. Click **Activate** and review the information in the **Activate LSAN Zones** dialog box.

   LSAN zones that contain online members are automatically included in the **Destination Fabrics** list. For LSAN zones that contain offline members, you can click the right arrow button to assign these zones to fabrics in the **Destination Fabrics** list.

9. Click **OK** to activate the LSAN zones.

   A message displays informing you about the effects of LSAN zone activation and asking whether you want to proceed. Click **Yes** to confirm the activation, or click **No** to cancel the activation.

10. Click **OK** to continue.

    All LSAN zones are activated on the selected fabrics and saved to their respective zone databases.

11. Click **OK** to close the **Zoning** dialog box.

Activating LSAN zones

1. Select a backbone fabric from the Connectivity Map or Product List.
2. Select **Configure > Zoning > LSAN Zoning (Device Sharing)**.

   The **Zoning** dialog box displays, with the LSAN scope.

3. Click **Activate**.
4. Review the information in the **Activate LSAN Zones** dialog box.

   LSAN zones that contain online members are automatically included in the **Destination Fabrics** list. For LSAN zones that contain offline members, you can click the right arrow button to assign these zones to fabrics in the **Destination Fabrics** list.

5. Click **OK** to commit the LSAN zones and activate them in the selected fabrics.

   A message displays informing you about the effects of LSAN zone activation and asking whether you want to proceed. Click **Yes** to confirm the activation, or click **No** to cancel the activation.

6. Click **OK** to close the **Zoning** dialog box.

LSAN tagging

You can configure two types of tags on an FC router:

- Enforce tag – Specifies which LSANs are to be enforced in an FC router.
- Speed tag – Specifies which LSANs are to be imported or exported faster than other LSANs.

You configure the tags using the command line interface. The Management application displays the tags.

If tags are configured, they are displayed in the **LSAN Zoning** dialog box. Note that although you can configure tags on FC routers running Fabric OS versions earlier than 7.2.0, the tags are displayed in the Management application only if the FC router is running Fabric OS 7.2.0 or later.

When these tags are configured, only LSAN zones that match the configured tags are retrieved from the edge fabrics and displayed in the LSAN Zoning dialog box.

Refer to the *Fabric OS Administrator's Guide* for information about configuring these tags.

Traffic Isolation zones

A Traffic Isolation zone (TI zone) is a special zone that isolates inter-switch traffic to a specific, dedicated path through the fabric. A TI zone contains a list of E_Ports, followed by a list of N_Ports. When the TI zone is activated, the fabric attempts to isolate all inter-switch traffic between N_Ports to only those E_Ports that have been included in the zone. The fabric also attempts to exclude traffic not in the TI zone from using E_Ports within that TI zone.

Traffic Isolation zoning is only supported with domain and port index number members.

To create a TI zone for a logical fabric that uses XISLs, you must create two TI zones: one in the logical fabric and one in the base fabric. The combination of TI zones in the base fabric and logical fabric sets the path through the base fabric for logical switches.

NOTE
TI zones are not supported with Network OS.

Failover options

A TI zone can have failover enabled or disabled.

Disable failover if you want to guarantee that TI zone traffic uses only the dedicated path, and that no other traffic can use the dedicated path.

Enable failover if you want traffic to have alternate routes if either the dedicated or non-dedicated paths cannot be used.

ATTENTION
If failover is disabled, use care when planning your TI zones so that non-TI zone devices are not isolated. If disabled failover is not used correctly, it can cause major fabric disruptions that are difficult to resolve.

For base switches, failover is always enabled, and you cannot change it.

Enhanced TI zones

Ports can be in multiple TI zones. Zones with overlapping port members are called ***enhanced TI zones*** (ETIZ).

Enhanced TI zones are supported only on the following platforms:

- 24-port, 8 Gbps FC Switch
- 40-port, 8 Gbps FC Switch
- 80-port, 8 Gbps FC Switch
- 48-port, 16 Gbps FC Switch
- 8 Gbps 12-port Embedded Switch
- 8 Gbps 24-port Embedded Switch
- 8 Gbps 16-port Embedded Switch
- 8 Gbps 24-port Embedded Switch

- 8 Gbps Extension Switch
- 8 Gbps 8-FC port, 10 GbE 24-DCB port Switch
- 8 Gbps 40-port Switch
- 16 Gbps 4-slot Backbone Chassis
- 16 Gbps 8-slot Backbone Chassis
- 8-slot Backbone Chassis
- 4-slot Backbone Chassis
- 8 Gbps Encryption Switch

Enhanced TI zones are supported only if the following conditions are met:

- Every switch must be one of the previously listed supported platforms.
- Every switch must be running Fabric OS 7.0 or later.

If the fabric contains a switch running an earlier version of Fabric OS, you cannot create an enhanced TI zone.

The failover mode must be the same for each enhanced TI zone to which a port belongs.

You cannot merge a down-level switch into a fabric containing enhanced TI zones, and you cannot merge a switch with enhanced TI zones defined into a fabric containing switches that do not support ETIZ.

Configuring Traffic Isolation zoning

The following procedure provides an overview of the steps you must perform to configure Traffic Isolation zoning.

Note that for any zoning-related procedure, changes to a zone database are not saved until you click **OK** or **Apply** on the **Zoning** dialog box. If you click **Cancel** or the close button (X), no changes are saved.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select **Domain, Port Index** from the **Type** list.
5. (*Optional*) If you want to show all discovered fabrics in the **Potential Members** list, right-click in the **Potential Members** list and select **Display All**.
6. Create the Traffic Isolation zones.

   For specific instructions, refer to "Creating a Traffic Isolation zone" on page 812.

7. Add members to each zone.

   For specific instructions, refer to "Adding members to a Traffic Isolation zone" on page 812.

   NOTE
   You cannot add a Traffic Isolation zone to a zone configuration.

8. Click **OK** or **Apply** to save your changes.

   The Traffic Isolation zones are saved, but are not activated. The Traffic Isolation zones are activated when you activate a zone configuration in the same zone database.

Creating a Traffic Isolation zone

Traffic Isolation zones are configurable only on a Fabric OS device. The seed switch must be running Fabric OS 7.0 or later.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select **Domain, Port Index** from the **Type** list.
5. Select **New TI Zone** from the **New Zone** list.
6. Enter a name for the zone.

   For zone name requirements and limitations, refer to "Zoning naming conventions" on page 782.

7. Press **Enter**.
8. Click **OK** or **Apply** to save your changes.

   The Traffic Isolation zone is saved, but is not activated.

Adding members to a Traffic Isolation zone

NOTE
Traffic Isolation zones are configurable only on a Fabric OS device.

Use this procedure to add a member to a zone when the member is listed in the **Potential Members** list of the **Zone DB** tab. Only ports can be added as members to a Traffic Isolation zone. You must add two or more N_Ports as well as all E_Ports on the path between the N_Ports.

NOTE
You cannot add a device as a member to a Traffic Isolation zone.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. (*Optional*) If you want to show all discovered fabrics in the **Potential Members** list, right-click in the **Potential Members** list and select **Display All**.

5. Select one or more Traffic Isolation zones to which you want to add members in the **Zones** list. (Press **SHIFT** or **CTRL** and click each zone name to select more than one zone.)
6. Select **Domain, Port Index** from the **Type** list.
7. Select two or more N_Ports (as well as all E_Ports on the path between the N_Ports) to add to the zone in the **Potential Members** list. (Press **SHIFT** or **CTRL** and click each port to select more than one port.)

   NOTE
   TI zones can be created in fabrics that contain logical switches; however, you can only select physical ports for TI zones.

   If you select a trunk port to add to the TI zone, all trunk ports in the trunk group are added to the TI zone automatically.
8. Click the right arrow between the **Potential Members** list and the **Zones** list to add the selected ports to the zone.
9. Click **OK** or **Apply** to save your changes.

   The TI zone is saved, but is not activated. Traffic Isolation zones are activated when you activate a zone configuration in the same zone database.

Enabling a Traffic Isolation zone

NOTE
Traffic Isolation zones are configurable only on a Fabric OS device.

Use this procedure to enable a Traffic Isolation zone. When a zone configuration in the same zone database is activated, the enabled TI zones are also activated at that time. Traffic Isolation zones are enabled by default when you create them.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.
4. Right-click the Traffic Isolation zone you want to enable in the **Zones** list and select **Configured Enabled**.
5. Click **OK** or **Apply** to save your changes.

   The Traffic Isolation zone is saved, but not activated. The Traffic Isolation zone is activated when you activate a zone configuration in the same zone database.

Disabling a Traffic Isolation zone

NOTE
Traffic Isolation zones are configurable only on a Fabric OS device.

Traffic Isolation zones are enabled by default when you create them. Use this procedure to disable a Traffic Isolation zone. To apply the settings and deactivate the zone, you must activate a zone configuration in the same zone database.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.
4. Right-click the Traffic Isolation zone you want to disable in the **Zones** list and clear the **Configured Enabled** check box.
5. Click **OK** or **Apply** to save your changes.

   The Traffic Isolation zone is not disabled until you activate a zone configuration in the same zone database.

Enabling failover on a Traffic Isolation zone

NOTE
Traffic Isolation zones are configurable only on a Fabric OS device.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.
4. Right-click the Traffic Isolation zone upon which you want to enable failover in the **Zones** list and select **Configured Failover**.
5. Click **OK** or **Apply** to save your changes.

Disabling failover on a Traffic Isolation zone

NOTE
Traffic Isolation zones are configurable only on a Fabric OS device.

If failover is disabled, be aware of the following considerations:

- Ensure that there are non-dedicated paths through the fabric for all devices that are not in a TI zone.
- If you create a TI zone with E_Ports only, failover must be enabled. If failover is disabled, the specified ISLs will not be able to route any traffic.
- Ensure that there are multiple paths between switches. Disabling failover locks the specified route so that only TI zone traffic can use it.

ATTENTION
If failover is disabled, use care when planning your TI zones so that non-TI zone devices are not isolated. If the disabled failover configuration is not correct, it can cause major fabric disruptions that are difficult to resolve.

You cannot disable failover if the TI zone was created in the base fabric or in a fabric in which a logical switch is configured to use XISLs (the **Base Fabric for Transport** check box is selected).

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Click the **Zone DB** tab if that tab is not automatically displayed.

3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Right-click the Traffic Isolation zone upon which you want to disable failover in the **Zones** list and clear the **Configured Failover** check box.

5. Click **OK** or **Apply** to save your changes.

Boot LUN zones

A Boot LUN zone is a special zone used to boot from SAN. Boot LUN zone names have the following format:

BFA_*HostPortWWN*_BLUN

After you create a Boot LUN zone, it is managed in the same way as standard zones.

You cannot add or remove members of a Boot LUN zone. Boot LUN zones cannot be merged.

Boot LUN zones are not supported for Network OS fabrics.

Creating a Boot LUN zone

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.

3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

   Boot LUN zones are not supported for Network OS fabrics.

4. Launch the **New Boot LUN Zone** dialog box by performing one of the following options:

   - Select **New Boot LUN Zone** from the **New Zone** list.
   - Right-click a zone in the **Zones** list and select **New Boot LUN Zone**.

   The **New Boot LUN Zone** dialog box displays. The scope of the dialog box is either the selected fabric or the selected zone, depending on how you launch it.

5. Select a host port WWN from the list or enter an offline WWN.

   You can click the ellipsis button to display and select the host port WWNs from a device tree with host group.

6. Select a storage port WWN from the list or enter an offline WWN.

   You can click the ellipsis button to display and select the storage port WWNs from a device tree with storage group.

7. Enter a 16-digit hexadecimal LUN number in the **LUN #** field.

   The hexadecimal LUN number needs to be in the second byte of the string.

8. Click **Generate**.

   The Boot LUN zone is generated and displayed in the **Boot LUN Zone details** area.

9. Click **OK** or **Apply** to save your changes.

The Boot LUN zone is saved to the Active Zone DB. To activate the Boot LUN zone, you must move it to a zone configuration and activate the configuration.

Modifying a Boot LUN zone

Only one Boot LUN zone can exist for a host port. If you want to change the target port or LUN number, you must create a new Boot LUN zone and overwrite the existing zone.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Right-click the Boot LUN zone you want to modify in the **Zones** list and select **New Boot LUN Zone**.

   The **New Boot LUN Zone** dialog box displays. You can modify the storage port WWN and LUN number.

5. Select a storage port WWN from the list or enter an offline WWN.

   You can click the ellipsis button to display and select the storage port WWNs from a device tree with storage group.

6. Enter a 16-digit hexadecimal LUN number in the **LUN #** field.

   The hexadecimal LUN number needs to be in the second byte of the string.

7. Click **Generate**.
8. Click **OK** or **Apply** to save your changes.

   A message displays that a Boot LUN zone already exists and asks whether you want to overwrite the existing zone.

9. Click **Yes**.

   The existing Boot LUN zone is replaced by the version you just created.

Deleting a Boot LUN zone

Boot LUN zones are deleted the same way that standard zones are deleted. Refer to "Deleting a zone" on page 789 for instructions.

Zoning administration

This section provides instructions for performing administrative functions with zoning. You can rename, duplicate, delete, and perform other tasks on zone members, zones, and zone configurations.

Comparing zone databases

You can compare zone databases against one another to identify any and all differences between their memberships prior to sending them to the switch. Once the two databases have been compared, icons display to show the differences between the two databases. These icons are illustrated and described in Table 62.

TABLE 62 Compare icon indicators

Icon	Description
	Added — Displays when an element is added to the editable database.
	Modified — Displays when an element is modified on the editable database.
	Removed — Displays when an element is removed from the editable database.

To compare two zone databases, complete the following steps.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Select **Compare** from the **Zone DB Operation** list.

   The **Compare/Merge Zone DBs** dialog box displays, as shown in Figure 399.

   FIGURE 399 Compare/Merge Zone DBs dialog box

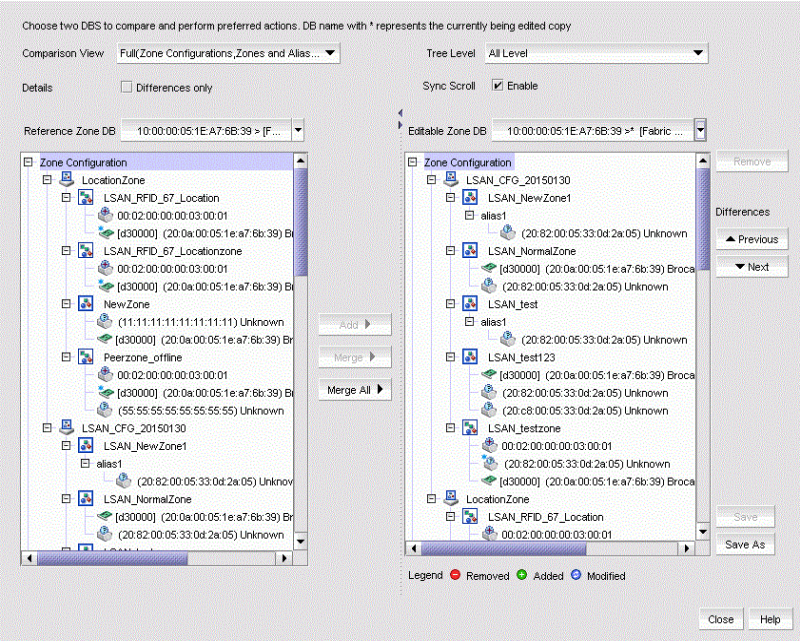

3. Select a database from the **Reference Zone DB** list.
4. Select a database from the **Editable Zone DB** list.

   The **Reference Zone DB** and **Editable Zone DB** areas display all available element types (zone configurations, zones, and aliases) for the two selected zone databases. In the **Editable Zone DB** area, each element type and element display with an icon indicator (Table 62) to show the differences between the two databases.

5. Set the display for the database areas by selecting one of the following from the **Comparison View** list:
   - **Storage-to-Host Connectivity** — Displays only storage and host devices.
   - **Host-to-Storage Connectivity** — Displays only host and storage devices.
   - **Full (Zone Configurations, Zones, and Aliases)** — Displays all zone configurations, zones, and aliases.
6. Set the level of detail for the database areas by selecting one of the following options from the **Tree Level** list:

   NOTE
   This list is only available when you set the **Comparison View** to **Full (Zone Configurations, Zones, and Aliases)**.

   - **All Level** — Displays all zone configurations, zones, and aliases.
   - **Zone Configurations** — Displays only zone configurations.
   - **Zones** — Displays only zones.
7. Select the **Differences only** check box to display only the differences between the selected databases.
8. Select the **Sync Scroll Enable** check box to synchronize scrolling between the selected databases.
9. Click **Previous** or **Next** to navigate line-by-line in the **Editable Zone DB** area.
10. Click **Close**.

    To merge two zone databases, refer to "Merging two zone databases" on page 802.

Managing zone configuration comparison alerts

You can turn off the automatic zone configuration comparison function if you no longer want to see two of the alert messages that the comparison can produce. When a zone configuration is successfully activated, the comparison function can display an alert icon if either of two conditions exist.

The messages are "The active zone configuration does not exist in the zone database" and "The active zone configuration does not match <*zone configuration*> in the zone database." To turn off the icons and the messages, complete the following steps.

1. After successfully activating a zone configuration, click the **Active Zone Configuration** tab in the **Zoning** dialog box.
2. Select the **Turn off the comparison alerts between the active zone configuration and the zone database** check box.

   Any existing alert icons and messages are cleared and further comparisons are prevented.

Setting change limits on zoning activation

Use this procedure to set a limit on the number of changes a user can make to the zone database before activating a zone configuration. If the user exceeds the limit, zone configuration activation is not allowed. By default, all fabrics allow unlimited changes. Changes include adding, removing, or modifying zones, aliases, and zone configurations.

Use this procedure to set the following limits:

- Set a different limit for each fabric.
- Set limits on some fabrics while allowing other fabrics to have unlimited changes.
- Set a limit for fabrics that will be discovered later.

  NOTE
  You must have the Zoning Set Edit Limits privilege to perform this task.

1. Select **Configure > Zoning > Set Change Limits**.

   The **Set Change Limits for Zoning Activation** dialog box displays, as shown in Figure 400.

   FIGURE 400 Set Change Limits for Zoning Activation dialog box

2. Click **Change Count** for the fabric on which you want to set limits.

   The field changes to an editable field.

3. Enter the maximum number of zone database changes that can be made for that fabric before a zone configuration is activated.

   To set a limit, enter a positive integer.

   To allow unlimited changes, enter 0.

4. Repeat step 1 and step 3 for each fabric on which you want to set limits.
5. To set a limit for new, undiscovered fabrics, enter a value in the **Default Change Count for New Fabrics** field.

   This limit is enforced on all new fabrics as they are discovered. The default value is 0 (Unlimited).

6. Select the **Enforce change limits during zone activation** check box to enforce the change limits.

   If you want to set the limits now, but turn on enforcement of the limits at a later time, make sure the check box is clear.

7. Click **OK** to save your changes and close the dialog box.

Clearing the fabric zone database

ATTENTION
Clearing the zone database removes all zoning configuration information, including all aliases, zones, and zone configurations, in the fabric.
Clearing the fabric zone database is disruptive to the fabric.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning databases for the selected entity.

3. Select the Fabric Zone DB from the **Zone DB** list.
4. Select **Clear All** from the **Zone DB Operation** list.
5. Click **Yes** on the confirmation message.

   The message closes and the Fabric Zone DB is cleared of all zoning configurations.

6. Click **OK** to close the **Zoning** dialog box.

Removing all user names from a zone database

Use this procedure to remove all user names from the selected offline zone database.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning databases for the selected entity.

3. Select a zone database that you have checked out (your user name is in the **Current User** column) in the **Zone DB** list.
4. Select **Undo CheckOut** from the **Zone DB Operation** list.
5. Click **Yes** in the confirmation message.

   This removes the user names of users currently logged in to the client from the **Current User** column for this zone database.

6. Click **OK** to save your work and close the **Zoning** dialog box.

   Any zones or zone configurations you have changed are saved in the zone database.

Finding a member in one or more zones

Use this procedure to locate all instances of a member in the **Zones** list on the **Zone DB** tab.

1. Select **Configure > Zoning > Fabric**.

   For LSAN zones, select **Configure > Zoning > LSAN Zoning (Device Sharing)**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. If you want to show all fabrics discovered in the **Potential Members** list, right-click in the **Potential Members** list and select **Display All**.
5. Select the devices or ports you want to find in the **Potential Members** list.
6. Click **Find >** between the **Potential Members** list and the **Zones** list.

   If the member is found, all instances of the zone member found are highlighted in the **Zones** list.

Finding a zone member in the potential member list

Use this procedure to locate a zone member in the **Potential Members** list on the **Zone DB** tab.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select the zone member in the **Zones** list that you want to find in the **Potential Members** list. (Press **SHIFT** or **CTRL** and click each zone member to select more than one zone member.)

5. Click **Find <** between the **Potential Members** list and the **Zones** list.

   If the member is found, it is highlighted in the **Potential Members** list.

Finding zones in a zone configuration

Use this procedure to locate all instances of a zone in the **Zone Configurations** list on the **Zone DB** tab.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select the zone you want to find in the **Zones** list. (Press **SHIFT** or **CTRL** and click each zone to select more than one zone.)
5. Click **Find >** between the **Zones** list and the **Zone Configurations** list.

   If the zone is found, all instances of the zone are highlighted in the **Zone Configurations** list.

Finding a zone configuration member in the zones list

Use this procedure to locate a zone configuration member in the **Zones** list on the **Zone DB** tab.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Select the zone configuration member (for example, the zone) in the **Zone Configurations** list that you want to find in the **Zones** list. (Press **SHIFT** or **CTRL** and click each zone configuration member to select more than one zone configuration member.)
5. Click **Find <** between the **Zones** list and the **Zone Configurations** list.

   If the zone is found, it is highlighted in the **Zones** list.

Listing zone members

Use this procedure to identify the zone in the active zone configuration of the fabric to which an individual port belongs and the members of that zone.

For Peer zones, List Zone Members allows you to identify the zone in the active zone configuration of the fabric to which an individual port belongs and the members of that zone that are defined to communicate with it.

If the seed switch is running Fabric OS 7.0 or later, the **List Zone Members** dialog box also displays any active TI zones to which the port belongs.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Right-click in the **Potential Members** list and select **List Zone Members**.

   The **List Zone Members** dialog box displays. If the port is a member of a zone, the fabric name, the port name, and WWN zone members display.

   For Peer zones, the list zone members are based on principal and peer member connectivity as follows:

   - If the port is a principal member, only peer members (including offline members) display in the **List Zone Members** dialog box.
   - If the port is a peer member, only principal members (including offline members) display in the **List Zone Members** dialog box.
   - Note that the Property Member does not display in the **List Zone Members** dialog box.

5. Select **Historical Graph** or **Real-Time Graph** from the **Performance Graph** list.

   Historical Graph and **Real-Time Graph** will be launched for the attached switch ports, for the selected zone members.

6. Click **Close** to exit the **List Zone Members** dialog box.

Listing un-zoned members

Use this procedure to identify the device ports in the current fabric that are not part of the active zone configuration.

You can use this procedure for standard zones as well as LSAN zones.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Right-click in the **Potential Members** list and select **List Un-Zone Members**.

   The **Un-Zone Members** dialog box displays. The dialog box shows the fabric name and the connected device ports that are not part of the active zone configuration.

5. Click **Close** to exit the **Un-Zone Members** dialog box.

Removing an offline device

The Management application enables you to remove an offline device from all zones and zone aliases in the selected zone DB.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

3. Select **Offline Utility** from the **Zone DB Operation** list.

   The **Offline Device Management** dialog box displays.

4. Select the check box for the offline device you want to remove in the **Remove** column.

   Select the **Remove** check box to select all offline devices.

5. Click **OK** on the **Offline Device Management** dialog box.

   A warning message displays informing you that the selected zone members will be replaced from all zones and aliases in the selected zone DB.

6. Click **OK** on the message.

7. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Replacing zone members

You can replace one instance of a zone member in one zone, or all instances of the zone member in all the zones to which it belongs.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.

3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Right-click the zone member you want to replace in the **Zones** list and select one of the following options from the shortcut menu that displays:

   - **Replace** - To replace the zone member in a selected zone.
   - **Replace All** - To replace all instances of the selected zone member.

   The **Replace Zone Member** dialog box displays.

5. Select the option from the **Member Type** list that you want to use to identify the replacement zone member.

6. Enter the WWN, name, domain and port index numbers, or alias — whichever is appropriate for the method you chose in step 5.

   When you choose the WWN method, you may define a name for the replacement zone member.

7. Click **OK**.

   The new zone member replaces the old zone member in the **Zones** list and the **Replace Zone Member** dialog box closes.

8. Click **OK** or **Apply** to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Replacing an offline device by WWN

The Management application enables you to replace an offline device by WWN from all zones and zone aliases in the selected zone DB.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

3. Select **Offline Utility** from the **Zone DB Operation** list.

   The **Offline Device Management** dialog box displays.

4. Clear the **Remove** column check box for the offline device you want to replace.
5. Select **WWN** (default) in the corresponding **Replace Using** list.
6. Enter the WWN or select the name of the offline device in the corresponding **Replace Value** list.

   If the selected name has multiple device or device port WWNs assigned (names are set to non-unique in the Management application), the **Device or Device Port WWN of Non-unique Name** dialog box displays. The **WWN** list includes all device and device port WWNs assigned to the selected name.

7. Click **OK** on the **Offline Device Management** dialog box.

   A warning message displays informing you that the selected zone members will be removed from all zones and aliases in the selected zone DB.

8. Click **OK** on the message.
9. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

   Any zones or zone configurations you have changed are saved in the zone database.

Replacing an offline device by name

The Management application enables you to replace an offline device by name from all zones and zone aliases in the selected zone DB.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

3. Select **Offline Utility** from the **Zone DB Operation** list.

   The **Offline Device Management** dialog box displays.

4. Clear the **Remove** column check box for the offline device you want to replace.
5. Select **Name** in the corresponding **Replace Using** list.
6. Select the name of the offline device in the corresponding **Replace Value** list.

If the selected name has multiple device or device port WWNs assigned (names are set to non-unique in the Management application), the **Device or Device Port WWN of Non-unique Name** dialog box displays. The **WWN** list includes all device and device port WWNs assigned to the selected name.

7. Select the WWN you want to use from the **WWN** list and click **OK**.
8. Click **OK** on the **Offline Device Management** dialog box.

   A warning message displays informing you that the selected zone members will be removed from all zones and aliases in the selected zone DB.
9. Click **OK** on the message.
10. Click **OK** or **Apply** on the **Zoning** dialog box to save your changes.

    Any zones or zone configurations you have changed are saved in the zone database.

Peer zones

Peer zoning allows you to communicate between principal members and non-principal or peer members within a zone. You can enable communication between a principal member device and a peer member device in a Peer zone. You cannot enable communication for devices within principal members or peer members. Beginning with Fabric OS 7.4.0 or later, you can view, create, edit, rename, and delete the Peer zones and replace the peer members present in the zone through the **Zoning** dialog box.

Creating a Peer zone

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.
2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

   NOTE
   Peer zones are not supported for Network OS fabrics. Peer zones will be enabled only if the fabric selected has at least one switch running Fabric OS v7.4.0 or later.
4. Select **New Peer Zone** from the **New Zone** list.

   The **Add Peer Zone** dialog box displays, as shown in Figure 401. The **Add Peer Zone** dialog box will launch with the member **Type** selected on the **Type** list in the **Zoning** dialog box.

FIGURE 401 Add Peer Zone dialog box

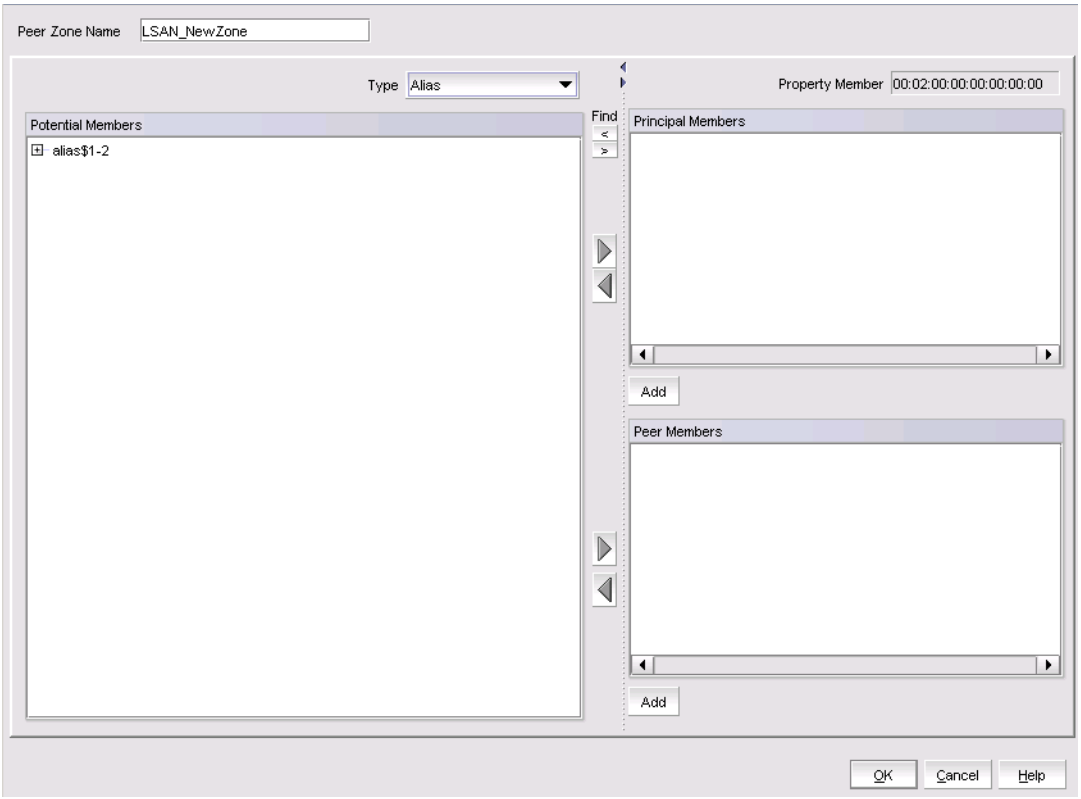

The following list shows the availabe member types:

- WWN
- WWN-Fabric Assigned
- Domain, Port Index
- Alias

Alias member type is supported for switches running Fabric OS 8.1.0 or later.You can switch between the available member types if the member included in the Peer zone belongs to same member type.

The **Property Member** value is auto-generated by the Management application whenever the Peer zone members are modified as follows:

- When you set or modify the member type as **WWN**, **Domain, Port Index**., or **Alias**
- When you create, edit, or remove the principal members.

You cannot edit or delete the property members.

5. Select the members from the **Potential Members** list and click the right arrow button to move the members to the **Principal Members** list.
6. Select the members from the **Potential Members** list and click the right arrow to move the members to the **Peer Members** list.

   NOTE
   You cannot add the same member to both the **Principal Members** and **Peer Members** lists.

7. Click **OK** or **Apply** to save your changes.

The Peer zone is saved to the Active Zone DB. To activate the Peer zone, you must move it to a zone configuration and activate the configuration. You can add an offline principal member and peer member using the **Add** button. For specific information, refer to "Adding offline members to a Peer zone" on page 828. The default logical display order of the Peer zone members is as follows:

- Property member - A WWN zone member which holds the definition of the peer zone like whether it is a peer zone or target driven zone, number of principal members, and peer zone member type. The **Alias** type property member indicates the total number of principal members present before and after expansion of the alias. The following lists the Property member generation logic for the Alias support in Peer zones.

  Property member: AA:BB:CC:DD:EE:FF:GG:HH

  - BB: 01 - Target created, 02 - User created (Peer Zone)
  - FF: 02 - Domain, Port Index Peer zone type, 03 - WWN Peer zone type
  - GG: Indicates the total number of principal members before expanding the alias.
  - HH: Indicates the total number of principal members after expanding the alias.
- Principal members - A Peer zone member that is allowed to communicate with any other Peer member in that zone. Principal and Principal members are synonyms. Principal members in a peer zone are not allowed to communicate with each other.
- Peer members - A peer zone member that is allowed to communicate with only Principal members of that peer zone. Peer, Peer member, Non-principal, Non-principal peer, or Non-principal members are synonyms.

Peer zone icons

The following table lists the Peer zone icons that display in the Zoning dialog boxes.

Icon	Description	Icon	Description
	Peer zone		Active Peer zone

Peer zone member icons

The following table lists the Peer zone member icons that display in the Zoning.

Icon	Description
	Property Member
	Initiator Principal Member
	Target Principal Member
	Unknown Principal Member
	E-port Principal Member (Director Type)
	E-port Principal Member (Pizza Box Type)
	AG port Principal Member
	Loop Port Principal Member
Peer Members display with normal port icons. Refer to "SAN port icons" on page 305.	

Adding offline members to a Peer zone

You can add an offline principal member or peer member to the Peer zone. For creating a new Peer zone, refer to "Creating a Peer zone" on page 825.

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning.

4. (*Optional*) If you want to show all discovered fabrics in the **Potential Members** list, right-click in the **Potential Members** list and select **Display All**.
5. Launch the **New Peer Zone** dialog box by performing one of the following options:
   - Select **New Peer Zone** from the **New Zone** list.
   - Right-click a zone in the **Zones** list and select **New Peer Zone**.

   The **Add Peer Zone** dialog box displays.

6. Click **Add** under the **Principal Members** list or the **Peer Members** list of the **Add Peer Zone** dialog box.

   The **Add Zone Member** dialog box displays. Perform the following actions:

   a. Select the **Member Type** from the list.

      NOTE
      Alias type is supported from Fabric OS 8.1.0 or later.

   b. Select the **Member ID** by performing one of the following options:
      - Select the **Existing End Device Node/Port Name.**

      OR

      - Enter the **End Device Node/Port WWN** address.

   c. Enter a name in the **Assign Name** field.

7. Click **OK**.

   The offline principal member or peer member is added to the **Principal Members** list or **Peer Members** list on the **Add Peer Zone** dialog box.

Viewing Peer zone properties

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.

4. Right-click the Peer zone you want to review in the **Zones** list and select **Properties**.

   The **Peer Zone Properties** dialog box displays.

5. Review the Peer zone properties.

   Note that when any modifications are made to an active zone, the **Peer Zone Properties** dialog box continues to show the status as Active until the changes are saved to the zone database.

   You can change the zone name by double-clicking the name and then modifying the name in the editable field.

   Number of Aliases and Zone Members Contained by Aliases are supported from Fabric OS 8.1.0 or later.

   The new property “Number of Principal Members” displays in the **Peer Zone Properties** dialog box. For other zones, the “Number of Principal Members” property value displays as **NA**.

6. Click **OK** to close the **Peer Zone Properties** dialog box.

Editing a Peer zone

1. Select **Configure > Zoning > Fabric**.

   The **Zoning** dialog box displays.

2. Click the **Zone DB** tab if that tab is not automatically displayed.
3. Right-click the Peer zone you want to edit in the **Zones** list and select **Edit**.

   The **Edit Peer Zone** dialog box displays. You can edit the following:

   a. Select the member from **Type** list.
   b. Select one or more members that you want to add from the **Potential Members** list and click the right arrow button to add to the **Principle Members** list or **Peer Members** list. (Press **Shift** or **Ctrl** and click each member to select more than one member.)
   c. Show all discovered fabrics in the **Potential Members** list by right-clicking in the **Potential Members** list and selecting **Expand All**.

4. Click **OK** on the **Edit Peer Zone** dialog box to save your changes.
5. Click **OK** on the **Zoning** dialog box to save your changes.

Merging Peer zone members

You can compare but not merge Peer zone members within the Peer zone or with other zone members. The **Compare/Merge Zone DBs** dialog box allows you to delete, assign, or unassign Peer zones to or from one or more zone configurations. Refer to “Merging fabrics” on page 801 for comparing and merging fabrics.

Renaming a Peer zone

Peer zones are renamed the same way that standard zones are renamed. For specific information, refer to “Renaming a zone” on page 788.

Listing Peer zone members

List Zone Members allows you to identify the zone in the active zone configuration of the fabric to which an individual port belongs and the members of that zone that are defined to communicate with it. For Peer zones, the list zone members are based on principal and peer member connectivity as follows:

- If the port is a principal member, only peer members (including offline members) display in the **List Zone Members** dialog box.
- If the port is a peer member, only principal members (including offline members) display in the **List Zone Members** dialog box.
- Note that the Property Member does not display in the **List Zone Members** dialog box.

For the specific procedure, refer to "Listing zone members" on page 821.

Replacing a Peer zone member

You can replace the principal members and the peer members. Peer zone members are replaced the same way that standard zones are replaced. For specific information, refer to "Replacing zone members" on page 823. The Management application enables you to replace an offline device by WWN from all zones and zone aliases in the selected zone DB. For specific information, refer to "Replacing an offline device by WWN" on page 824.

NOTE
You cannot replace the property member and **Alias** type member in a Peer zone.

Importing a Peer zone

You can import a Peer zone and save it as an offline zone database. For specific information, refer to "Importing an offline zone database" on page 804.

Exporting a Peer zone

You can export a Peer zone and save it as an offline zone database. Peer zone members are exported the same way that standard zones are exported. For specific information, refer to "Exporting an offline zone database" on page 804.

Deleting a Peer zone

Peer zones are deleted the same way that standard zones are deleted. For specific information refer to "Deleting a zone" on page 789.

LSAN Peer zones

Beginning with the 12.4.0 release, the LSAN Peer zoning is supported for MetaSAN running Fabric OS v7.4.0 or later. The LSAN Peer zone is supported only if the backbone or edge fabrics have at least one switch running on the supported Fabric OS v7.4.0 or later. It is recommended to create the LSAN Peer zone from the **LSAN Zoning** dialog box. For more information, refer to "Creating an LSAN zone" on page 806. You can view, create, edit, rename, and delete the Peer zones from the **Zoning** dialog box, with the LSAN scope. The LSAN peer zone combines the properties of both LSAN Zoning and Peer zone. The LSAN Peer zone is displayed similar to basic LSAN zones in the Zone trees of the **Zoning** dialog box and the **Activate LSAN Zones** dialog box.

The Property member will not be a part of the offline members of the **Zone DB** list. The display order of the LSAN Peer zone members is as follows:

- Property member

- Principal members
- Peer members

The LSAN Peer zone supports only WWN and VPWWN members. The offline zone members supports only WWN members.

The LSAN Peer zones from different fabrics of the MetaSAN scope will be displayed similarly to basic LSAN zones. If there are multiple LSAN Peer zones with the same name, they will not be merged and displayed as a single entity like regular LSAN zones. The LSAN Peer zones belonging to each fabric will be displayed with their fabric name appended to it within parentheses to indicate its source fabric. The fabric names displayed within the parentheses are only for display purposes and are not considered part of the LSAN Peer zone names. When the zone or its name are edited, only the actual zone name (the fabric name) can be edited and displayed in the **Edit** dialog box. If the name of the LSAN Peer zone is modified and is different from the other LSAN Peer zone with same name, the appended fabric name will not be displayed.

Creating an LSAN Peer zone

1. Select a backbone fabric from the Connectivity Map or Product List.
2. Select **Configure > Zoning > LSAN Zoning (Device Sharing)**.

   The **Zoning** dialog box displays, with the LSAN scope.
3. Select **New Peer Zone** in the **New Zone** list.

   The **Add Peer Zone** dialog box for LSAN zone displays, as shown in Figure 402. The prefix LSAN_ is automatically added in the text field.

FIGURE 402 Add Peer Zone dialog box for LSAN zone

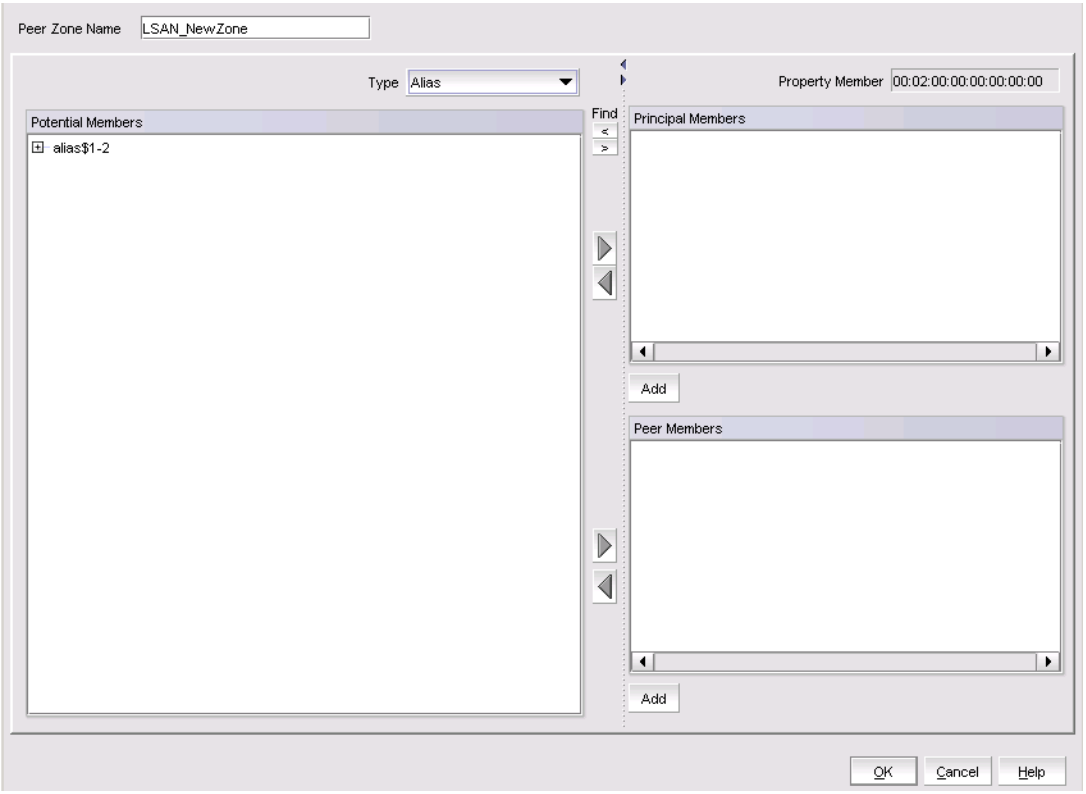

4. Add members to the Peer zone. Refer to "Creating a Peer zone" on page 825 to add, activate, and save the zone to its respective zone database.

Viewing LSAN Peer zone properties

1. Select a backbone fabric from the Connectivity Map or Product List.
2. Select **Configure > Zoning > LSAN Zoning (Device Sharing)**.

   The **Zoning** dialog box displays, with the LSAN scope.
3. Select a fabric from the **Zoning Scope** list.

   This identifies the target entity for all subsequent zoning actions and displays the zoning database for the selected entity.
4. Right-click the LSAN Peer zone you want to review in the **Zones** list and select **Properties**.

   The **LSAN Peer Zone Properties** dialog box displays.
5. Review the LSAN Peer zone properties.

   Note that when any modifications are made to an active zone, the **LSAN Peer Zone Properties** dialog box continues to show the status as Active until the changes are saved to the zone database.

   You can change the zone name by double-clicking the name and then modifying the name in the editable field.

   The new property "Number of Principal Members" displays in the **LSAN Peer Zone Properties** dialog box. For other zones, the "Number of Principal Members" property value displays as **NA**.
6. Click **OK** to close the **LSAN Peer Zone Properties** dialog box.

Editing an LSAN Peer zone

Use this procedure to edit an LSAN Peer zone. The **Edit** option is not available for a basic LSAN zone.

1. Select a backbone fabric from the Connectivity Map or Product List.
2. Select **Configure > Zoning > LSAN Zoning (Device Sharing)**.

   The **Zoning** dialog box displays, with the LSAN scope.
3. Right-click the LSAN Peer zone you want to edit in the **Zones** list and select **Edit**.

   The **Edit Peer Zone** dialog box displays. For specific information, refer to "Editing a Peer zone" on page 829.

Renaming an LSAN Peer zone

The LSAN Peer zones are renamed the same way that standard zones are renamed. For specific information, refer to "Renaming a zone" on page 788.

Replacing an LSAN Peer zone

You can replace the principal members and the peer members of the LSAN Peer zone. LSAN Peer zones members are replaced the same way that standard zones are replaced. For specific information, refer to "Replacing zone members" on page 823.

NOTE
You cannot replace the property member in an LSAN Peer zone.

Deleting an LSAN Peer zone

LSAN Peer zones are deleted the same way that standard zones are deleted. For specific information, refer to"Deleting a zone" on page 789.

Target Driven Peer zones

A Target Driven Peer Zone is a Peer Zone which is configured in a fabric through a target. You cannot create, rename, duplicate, replace, or edit the Target Driven Peer zone members. A Target Driven Peer zone can be created only through target members attached to the fabric. Target Driven Peer zone supports only WWN member type and only one principal member, but cannot be configured through the Management application. The target zones are displayed on the zone tree of the **Zoning** dialog box, as shown in Figure 403.

FIGURE 403 Target zones

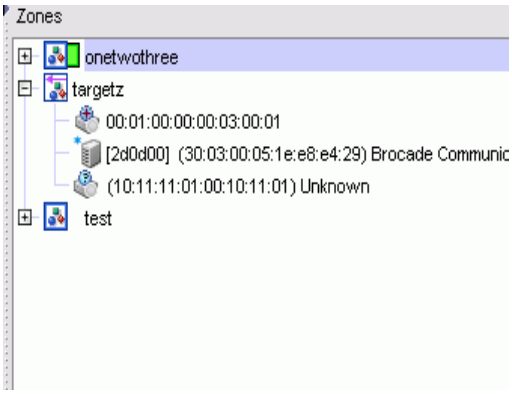

Viewing Target Driven Peer zone properties

Target Driven Peer zones properties can be viewed the same way that standard zones are viewed. For specific information, refer to "Viewing zone properties" on page 785.

Merging Target Driven Peer zone members

You can compare but not merge Target Driven Peer zone members within the Target Driven Peer zone or with other zone members. The **Compare/Merge Zone DBs** dialog box allows you to delete, assign, or unassign a Target Driven Peer zone from or to one or more zone configurations. Refer to "Merging fabrics" on page 801 for comparing and merging fabrics.

Importing a Target Driven Peer zone

You can import a Target Driven Peer zone and save it in an offline zone database, but you cannot save the zone database to a switch unless the switch has same Target Driven Peer Zone. For specific information, refer to "Importing an offline zone database" on page 804.

Exporting a Target Driven Peer zone

You can export a Target Driven Peer zone to an offline database, but you cannot save the zone database to a switch unless the switch has same Target Driven Peer Zone. The Target Driven Peer zone members are exported the same way that standard zones are exported. For specific information, refer to "Exporting an offline zone database" on page 804.

Deleting a Target Driven Peer zone

Target Driven Peer zones are deleted the same way that standard zones are deleted. For specific information, refer to "Deleting a zone" on page 789.

Fibre Channel over IP

FCIP services licensing

Most of the FCIP extension services described in this chapter require the High Performance Extension over FCIP/FC license. FICON emulation features require additional licenses.

The 16 Gbps Extension switch requires the WAN_Rate_Upgrade 1 and WAN_Rate_Upgrade 2 licenses. You must have the WAN_Rate_Upgrade 2 license to configure 40 Gbps ports.

The WAN_Rate_Upgrade 1, WAN_Rate_Upgrade 2, and Advanced FICON Acceleration licenses are not required to configure 32 Gbps, Router Extension blade.

The following features and licensing apply to the 8 Gbps Extension platforms.

- FCIP Adaptive Rate Limiting requires the FTR_AE (Advanced Extension) license.
- FCIP trunking requires FTR_AE license.
- IBM z/OS Global Mirror emulation (formerly eXtended Remote Copy or XRC) requires the FTR_AFA (Advanced FICON Acceleration) license.

The 10 Gigabit FCIP/Fibre Channel (FTR_10G) license is required for 10 GbE ports.

The Extension switch upgrade license enables full hardware, FCIP, and open systems tape pipelining on Fabric OS Extension Switches.

Use the **licenseShow** command to verify the needed licenses are present on the hardware used on both ends the FCIP tunnel. If required licenses are not installed, an error message will display while configuring the tunnel or circuit.

FCIP Concepts

Fibre Channel over IP (FCIP) is a tunneling protocol that enables you to connect Fibre Channel SANs over IP-based networks. Fabric OS Extension Switches and Extension Blades use FCIP to encapsulate Fibre Channel frames within IP frames that can be sent over an IP network to a partner Fabric OS Extension Switch or Extension Blade. When the IP packets are received, the Fibre Channel frames are reconstructed. FCIP uses a TCP transport that guarantees in-order delivery. The Fibre Channel fabric and all Fibre Channel targets and initiators are unaware of the presence of the IP network.

Because an FCIP tunnel uses an existing IP network, configuring and managing an FCIP tunnel requires knowledge of general IP networking concepts, and specific knowledge about the IP network that will be used for the tunnel. Because the IP network may be used to transport data over very long distances, and because the IP network is not designed exclusively for large data transfers, latency is an issue. Features such as data compression, trunking, FastWrite, Adaptive Rate Limiting (ARL), and Open Systems Tape Pipelining (OSTP) can reduce latency, and help manage tunnel bandwidth more effectively.

IP network considerations

Because FCIP uses TCP connections over an existing IP network, consult with the IP network administrator to be sure that the network hardware and software equipment operating in the data path can support those connections. Routers and firewalls that are in the data path need to be configured to pass layer 3 protocols 0800 (IP), 0806 (ARP), and 0001 (ICMP). Also, process layer ports for FTP (ports 20 and 21) Telnet (port 23), and SNMP (ports 161 and 162) should be configured on the management IP network to enable support personnel to access and transmit troubleshooting information.

FCIP platforms and supported features

The following Fabric OS platforms that support FCIP:

- 8 Gbps Extension switch
- 8 Gbps Extension blade (8-slot Backbone Chassis, 4-slot Backbone Chassis).

  NOTE
  The 8 Gbps Extension blade is supported in 16 Gbps Backbone and Director Chassis.

- SX6 Extension blade
- 16 Gbps 24-FC port, 18 GbE port switch

IPv6 addressing is not supported in conjunction with IPsec on all platforms in Fabric OS version v7.0, but will be supported in a later version.Table 63 summarizes FCIP capabilities per platform.

TABLE 63 FCIP capabilities

Capabilities	8 Gbps Extension switch	8 Gbps Extension blade
FCIP trunking	Yes	Yes
Adaptive Rate Limiting	Yes	Yes
10 GbE ports	No	Yes
FC ports up to 8 Gbps	Yes	Yes
Compression	Yes	Yes
Open Systems Tape Pipelining (OSTP) • FCIP Fastwrite • Tape Acceleration	Yes	Yes
FICON extension	Yes	Yes
IPSec for tunnel traffic	Yes	Yes
Diffserv priorities	Yes	Yes
VLAN tagging	Yes	Yes
VEX_Ports	Yes	Yes
Support for third party WAN optimization hardware	No[1]	No[1]
IPv6 addresses for FCIP tunnels[2]	Yes	Yes
Support for jumbo frames	No[1] MTU of 1500 is maximum	No[1] MTU of 1500 is maximum

1. Support is planned for a later release.
2. IPv6 addressing is not supported in conjunction with IPsec in Fabric OS version v7.0, but will be supported in a later version.

The way FCIP tunnels and virtual ports map to the physical GbE ports depends on the switch or blade model. The 8 Gbps Extension Switch and 8 Gbps Extension Blade tunnels are not tied to a specific GbE port, and may be assigned to any virtual port within the allowed range. The mapping of GbE ports to tunnels and virtual port numbers is summarized in Table 64.

TABLE 64 GbE port mapping

Switch or Blade Model	GbE ports	Tunnels	Virtual ports (VE_Ports, VEX_Ports)
8 Gbps Extension Switch	GbE ports 0-5	0-8	16-23
8 Gbps Extension blade	GbE ports 0-9 10GbE ports 10, 11	0-20	12-21 used by GbE ports (0-9) and by XGE1 22-31 used by XGE0

The 16 and 32 Gbps Extension Switch and 8 Gbps Extension Blade tunnels are not tied to a specific GbE port, and may be assigned to any virtual port within the allowed range. The mapping of GbE ports to virtual port numbers is summarized in Table 65.

TABLE 65 16 and 32 Gbps GbE port mapping

Switch or Blade Model	20 VE mode VE_Ports	10 VE mode VE_Ports
16 Gbps 24-FC port, 18 GbE port Switch	DPO 24-33 DPO 34-43	DPO 24-29 DPO 34-39
32 Gbps, Router Extension blade	DPO 16-25 DPO 26-35	DPO 16-20 DPO 26-30

FCIP trunking

FCIP Trunking is a method for managing the use of WAN bandwidth and providing redundant paths over the WAN to protect against transmission loss. This feature is available only on the 8 Gbps Extension Switches and 8 Gbps Extension Blades. Trunking is enabled by creating logical circuits within an FCIP tunnel. A tunnel may have multiple circuits. Each circuit is a connection between a pair of IP addresses that are associated with source and destination endpoints of an FCIP tunnel, as shown in Figure 404. Each circuit represents a portion of the available Ethernet bandwidth provided by the GbE ports that are connected to the WAN.

FIGURE 404 FCIP tunnel and FCIP circuits

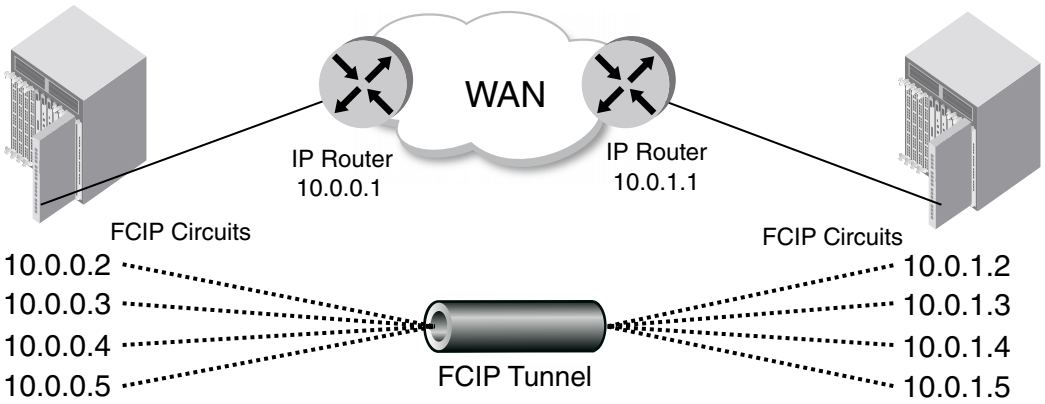

Design for redundancy and fault tolerance

Multiple FCIP tunnels can be defined between pairs of Extension Switches and Blades, but doing so defeats the concept of a multiple circuit FCIP tunnel. Defining two tunnels between a pair of switches or blades rather than one tunnel with two circuits is not as redundant or fault tolerant as having one multiple circuit tunnel.

FCIP tunnel restrictions for FCP and FICON emulation features

Multiple FCIP tunnels are not supported between pairs of Extension Switches and Blades when any of the FICON or FCP emulation features are enabled on the tunnel unless TI Zones or LS/LF configurations are used to provide deterministic flows between the switches. The emulation features require deterministic FC Frame routing between all initiators and devices over multiple tunnels. If there are non-controlled parallel (equal cost) tunnels between the same SID/DID pairs, emulation (Fast Write, Tape Pipelining, IBM z/OS Global Mirror (z Gm) or FICON Tape Pipelining) will fail when a command is routed via tunnel 1 and the responses are returned via tunnel 2. Therefore multiple equal cost tunnels are not supported between the switch pairs when emulation is enabled on any one or more tunnels without controlling the routing of SID/DID pairs to individual tunnels using TI Zones or LS/LF configurations.

FCIP Trunk configuration considerations

There are several points to consider when configuring an FCIP trunk:

- Each FCIP circuit is assigned a pair of IP addresses, one source IP address, and one destination IP address.
- The source IP address is used to determine which GbE interface to use. The GbE IP address must be on the same IP subnet as the source IP address. IP subnets cannot span across the GbE interfaces.
- The destination IP address is used to determine routing. If the destination IP address is also on the same subnet as the GbE interface, packets are routed over that subnet. If the destination IP address is on a different subnet, traffic must be routed to an IP gateway address.
- An FCIP circuit can have a maximum commit rate of 1,000,000 Kbps.
- In a scenario where a FCIP tunnel has multiple circuits of different metrics the data will flow over the lower metric circuits unless a failover condition occurs, as described in "FCIP circuit failover capabilities".
- The maximum bandwidth for a single circuit is 1 Gbps. However, a maximum of 10 Gbps per circuit is allowed between 10 GbE ports on 8 Gbps Extension Blades when both blades are running Fabric OS 7.0 or greater.

FCIP circuit failover capabilities

Each FCIP circuit is assigned a metric, which is used in managing failover for FC traffic. Typically, the metric will be either 0 or 1. If a circuit fails, FCIP Trunking tries first to retransmit any pending send traffic over another lowest metric circuit. In Figure 405, circuit 1 and circuit 2 are both lowest metric circuits. Circuit 1 has failed, and transmission fails over to circuit 2, which has the same metric. Traffic that was pending at the time of failure is retransmitted over circuit 2. In order delivery is ensured by the receiving Extension Switch or Blade.

FIGURE 405 Link loss and retransmission over peer lowest metric circuit

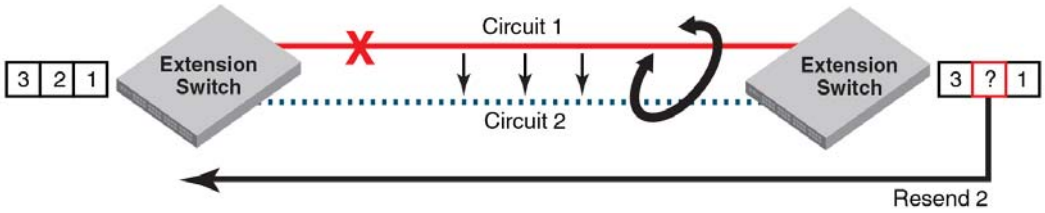

In Figure 406, circuit 1 is assigned a metric of 0, and circuit 2 is assigned a metric of 1. In this case, circuit 2 is a standby that is not used unless there are no lowest metric circuits available. If all lowest metric circuits fail, then the pending send traffic is retransmitted over any available circuits with the higher metric,

FIGURE 406 Failover to a higher metric standby circuit

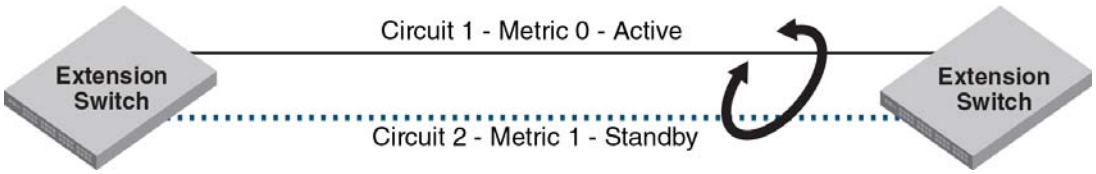

You can configure a tunnel with standby metric 1 circuits that operate when all circuits configured with metric 0 fail. However, these configurations can lead to condition where the tunnel is active, but operating in a degraded mode because all metric 1 circuits (and failed metric 0 circuits) will not be transferring tunnel data until all metric 0 circuits fail. Only at the point when all metric 0 circuits fail, do available metric 1 circuits over data transfer. Consider configuring "Circuit Failover Grouping" to avoid this problem.

Bandwidth calculation during failover

The bandwidth of higher-metric circuits is not calculated as available bandwidth on an FCIP tunnel until all lowest metric circuits have failed. Following is an example.

Assume the following configurations for circuits 0 through 3:

- Circuits 0 and 1 are created with a metric of 0. Circuit 0 is created with a maximum transmission rate of 1 Gbps, and circuit 1 is created with a maximum transmission rate of 500 Mbps. Together, circuits 0 and 1 provide an available bandwidth of 1.5 Gbps.
- Circuits 2 and 3 are created with a metric of 1. Both are created with a maximum transmission rate of 1 Gbps, for a total of 2 Gbps. This bandwidth is held in reserve.

The following actions occur during circuit failures:

- If either circuit 0 or circuit 1 fails, traffic flows over the remaining circuit while the failed circuit is being recovered. The available bandwidth is still considered to be 1.5 Gbps.
- If both circuit 0 and circuit 1 fail, there is a failover to circuits 2 and 3, and the available bandwidth is updated as 2 Gbps.
- If a low metric circuit becomes available again, the high metric circuits return to standby status, and the available bandwidth is updated again as each circuit comes online. For example, if circuit 0 is recovered, the available bandwidth is updated as 1 Gbps. If circuit 1 is also recovered, the available bandwidth is updated as 1.5 Gbps.

Circuit Failover Grouping

With Circuit Failover Grouping you can better control which metric 1 circuits will be activated if a metric 0 circuit fails. For this feature, you define a set of metric 0 and metric 1 circuits that are part of the same failover group. When all metric 0 circuits in the group fail, metric 1 circuits will take over data transfer, even if there are metric 0 circuits still active in other failover groups.

Typically, you would only define one metric 0 circuit in the group so that a specific metric 1 circuit will take over data transfer when the metric 0 fails. This will also avoid the problem of the tunnel operating in a degraded mode with less than the defined circuits before multiple metric 0 circuits fail.

Configure failover groups using the **Add FCIP Circuit** dialog box. Refer to "Adding an FCIP circuit" on page 859 for instructions.

Considerations and limitations

Failover groups operate under the following conditions;

- Each failover group is independent and operates autonomously.
- All metric 0 circuits in a group must fail before the metric 1 circuits are used.
- All metric 1 circuits in a group are used if all metric 0 circuits in the group fail or there is no metric 0 circuit in the group.
- Circuits can be part of only one failover group
- Failover circuit groups are only supported at Fabric OS v7.2.0 or greater.
- Both ends of the FCIP tunnel must have the same failover circuit groups defined for the feature to function properly.
- When a tunnel activates or circuits are modified, tunnel and circuit states will indicate a misconfiguration error if failover circuit group configurations are not valid.
- Modifying of the failover group ID is a disruptive operation, similar to modifying the metric.
- Circuit failover groups are not used to define load balancing over metric 0 circuits, (ONLY failover rules). Circuits of metric 0 will be load balanced over regardless of failover grouping.
- When no FCIP circuit failover groups are defined, failover reverts to default operation - all metric 0 circuits must fail before failing over to metric 1 circuit(s). In order to change default failover operation, a failover group should include at least one metric 0 and at least metric 1 circuit.
- A valid failover group requires at least one metric 0 circuit and at least one metric 1 circuit. If you do not configure these, a warning will display. If there is no metric 0 circuit and only a metric 1 circuit, the metric 1 circuit will be used, regardless of whether there are metric 0 circuits in another failover group.
- The number of valid failover groups defined per tunnel is limited by the number of circuits that you can create for the switch model. For an 8 Gbps Extension Blade, you can configure up to 5 valid groups on an 10-circuit tunnel. On an 8 Gbps Extension Switch, you could have up to 3 valid groups because you can only configure 6 circuits per tunnel.
- Consider available WAN bandwidth requirements when configuring failover circuit groups. Refer to "Bandwidth calculation during failover" on page 840.

Examples of circuit failover in groups

Tables Table 66 through Table 68 provide examples of how failover occurs on circuits with different bandwidths configured in failover groups.

Table 66 illustrates circuit failover in a tunnel with two failover groups, each with two circuits. All data through the tunnel is initially load balanced over Circuits 1 and 2. The following occurs during circuit failover:

- If circuit 1 fails, circuit 3 becomes active and data is load balanced over circuit 2 and 3.
- If circuit 2 fails, circuit 4 becomes active and data is load balanced over circuit 1 and 4.
- If both circuit 1 and 2 fail, circuit 3 and.4 become active and data is load balanced over both circuits.

TABLE 66 Tunnel with two failover groups with two circuits

Circuits in tunnel	Failover group ID	Circuit bandwidth	In use for tunnel data
Circuit 1 Metric 0	1	500Mb	If active, yes.
Circuit 2 Metric 0	2	1000Mb	If active, yes.

TABLE 66 Tunnel with two failover groups with two circuits

Circuits in tunnel	Failover group ID	Circuit bandwidth	In use for tunnel data
Circuit 3 Metric 1	1	500Mb	Only when circuit 1 fails.
Circuit 4 Metric 1	2	1000Mb	Only when circuit 2 fails.

Table 67 illustrates circuit failover in a tunnel with one failover group containing three circuits, In this case, failover occurs as if circuits are not part of a failover group. Circuit 2 and 3, both with metric 1, become active only after circuit 1 with metric 0 fails.

TABLE 67 Tunnel with one failover groups with three circuits

Circuits in tunnel	Failover group ID	Circuit bandwidth	In use for tunnel data
Circuit 1 Metric 0	1	1000Mb	If active, yes.
Circuit 2 Metric 1	1	500Mb	Only when circuit 1 fails.
Circuit 3 Metric 1	1	500Mb	Only when circuit 1 fails.

Table 68 illustrates circuit failover in a tunnel with circuits in failover groups and circuits that are not part of failover groups. In this configuration, all data is initially load balanced over circuit.1, circuit 2, and circuit 3 (when they are all active). The following occurs during circuit failover:

- If circuit1 fails, circuit 4 becomes active and data is load balanced over circuit 2, circuit 3, and circuit 4.

  Reason: Circuit 1 fails over to circuit 4 (both are in failover group 1) and circuit 3 is active with 500Mb bandwidth.
- If circuit 2 fails, data is load balanced over circuit 1 and circuit 3, and no other circuit becomes active.

  Reason: Circuit 1 and 3 are the only active circuits since circuit 4 and 5 only become active when circuits 3 or 1 fail.
- If circuit 2 and circuit 3 fail, circuit 5 becomes active and data is load balanced over circuit 1 and circuit 5.

  Reason: Ungrouped circuits 2 and 3 fail over to ungrouped circuit 5, which has a metric of 0.
- If circuit 1, circuit 2, and circuit 3 fail, circuit 4 and circuit 5 become active and data is load balanced over both.

  Reason: Circuit 1 fails over to circuit 4, which is the failover circuit for group 1 with a metric of 0. Ungrouped circuit 5 is the failover circuit for ungrouped, failed circuits 2 and 3.

TABLE 68 Tunnel with failover groups and non-grouped circuits

Circuits in tunnel	Failover group ID	Circuit bandwidth	In use for tunnel data
Circuit 1 Metric 0	1	500Mb	If active, yes.
Circuit 2 Metric 0	Not defined.	500Mb	If active, yes.
Circuit 3 Metric 0	Not defined.	500Mb	If active, yes.
Circuit 4 Metric 1	1	500Mb	Only when circuit 1 fails.
Circuit 5 Metric 1	Not defined	1000Mb	Only when circuit 2 and 3 fails.

Adaptive Rate Limiting

Adaptive Rate Limiting (ARL) is performed on FCIP tunnel connections to change the rate in which the FCIP tunnel transmits data through the TCP connections. This feature is available only on the 8 Gbps Extension Switches and 8 Gbps Extension Blades. ARL uses information from the TCP connections to determine and adjust the rate limit for the FCIP tunnel dynamically. This allows FCIP connections to utilize the maximum available bandwidth while providing a minimum bandwidth guarantee.

ARL applies a minimum and maximum traffic rate, and allows the traffic demand and WAN connection quality to dynamically determine the rate. As traffic increases, the rate grows towards the maximum rate, and if traffic subsides, the rate reduces towards the minimum. If traffic is flowing error-free over the WAN, the rate grows towards the maximum rate. If TCP reports an increase in retransmissions, the rate reduces towards the minimum.

FSPF link cost calculation when ARL is used

Fabric Shortest Path First (FSPF) is a link state path selection protocol that directs traffic along the shortest path between the source and destination based upon the link cost. When ARL is used, The link cost is equal to the sum of maximum traffic rates of all established, currently active low metric circuits in the tunnel. The following formulas are used:

- If the bandwidth is greater than or equal to 2 Gbps, the link cost is 500.
- If the bandwidth is less than 2 Gbps, but greater than or equal to 1 Gbps, the link cost is 1000000 divided by the bandwidth.
- If the bandwidth is less than 1 Gbps, the link cost is 2000 minus the bandwidth

QoS SID/DID priorities over an FCIP trunk

QoS SID/DID traffic prioritization is a capability of Fabric OS Adaptive Networking licensed feature. This feature allows you to prioritize FC traffic flows between hosts and targets.

Four internal TCP connections provide internal circuits for managing QoS SID/DID priorities over an FCIP tunnel, as illustrated Figure 407. The priorities are as follows:

- F class - F class is the highest priority, and is assigned bandwidth as needed at the expense of lower priorities, if necessary.
- QoS high - The QoS high priority gets at least 50% of the available bandwidth.
- QoS medium - The QoS medium priority gets at least 30% of the available bandwidth.
- QoS low - The QoS low priority gets at least 20% of the available bandwidth.

  NOTE
  The QoS high (50%), medium (30%), and low (20%) values are default values which you can change using procedures under "Configuring QoS Priorities" on page 844. These priorities are enforced only when there is congestion on the network. If there is no congestion, all traffic is handled at the same priority.

FIGURE 407TCP connections for handling QoS SID/DID-based FC traffic prioritization

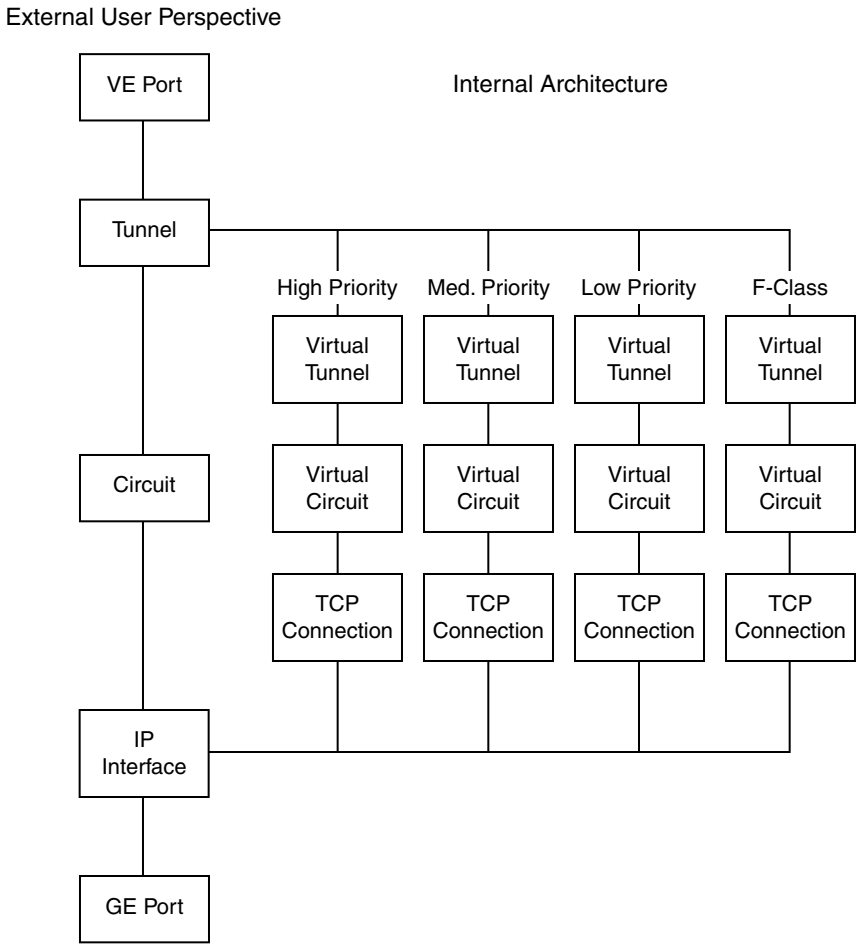

Configuring QoS Priorities

For 8 Gbps platforms only, you can change QoS priorities from the default settings using the following steps:

1. Select **Configure > FCIP Tunnels**.

   The **FCIP Tunnels** dialog box is displayed. All discovered fabrics with Extension Switches are listed under devices, and all existing FCIP tunnels are displayed.

2. Select the switch you want to configure under **Products**.
3. Click the **Add** button, or right-click on the switch and select **Add Tunnel.**

   The **Add FCIP Tunnel** dialog box is displayed.

4. Click **Advanced Settings.**

   The Advanced Settings dialog box is displayed. This dialog box has a Transmission tab, Security tab, and FICON Emulation tab. Configure QoS percentages on the Transmission tab (Figure 408**).**

FIGURE 408 Advanced Settings Transmission Tab

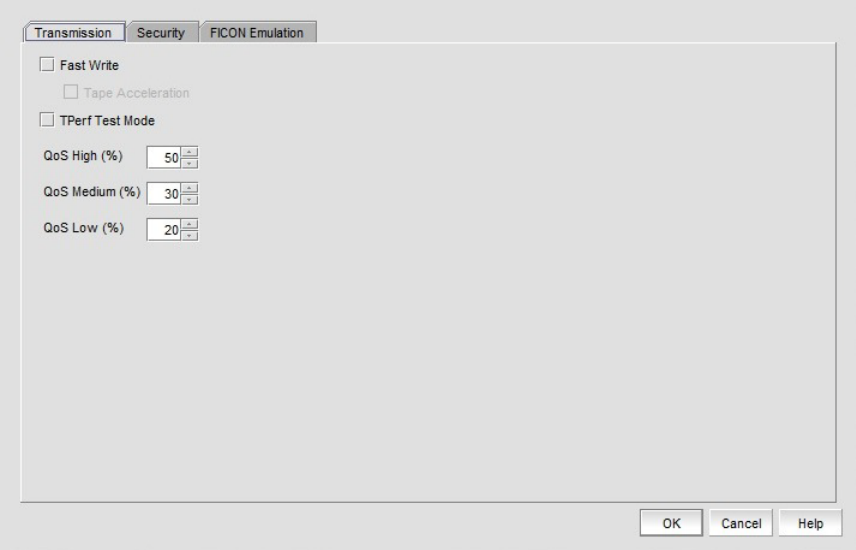

5. Click the up or down arrows by QoS High, QoS Medium, and QoS Low to increment values by 1% and override the default values of 50% (high), 30% (medium), and 20% (low). The three values must equal 100%. A minimum of 10% is required for each level.

NOTE
Editing QoS values is a disruptive operation, so a warning message displays when you make changes.

IPsec and IKE implementation over FCIP

Internet Protocol security (IPsec) uses cryptographic security to ensure private, secure communications over Internet Protocol networks. IPsec supports network-level data integrity, data confidentiality, data origin authentication, and replay protection. It helps secure your SAN against network-based attacks from untrusted computers, attacks that can result in the denial-of-service of applications, services, or the network, data corruption, and data and user credential theft. IPsec does not require you to configure separate security for each application that uses TCP/IP.

When configuring for IPsec, however, you must ensure that the same policies are defined in the switches or blades at each end of the FCIP tunnel. IPsec works on FCIP tunnels with or without compression, FCIP Fastwrite, and tape acceleration. IPsec can only be created on tunnels using IPv4 addressing.

IPsec for the 4 Gbps platforms

IPsec uses some terms that you should be familiar with before beginning your configuration. These are standard terms, but are included here for your convenience.

TABLE 69

Term	Definition
AES	Advanced Encryption Standard. FIPS 197 endorses the Rijndael encryption algorithm as the approved AES for use by US Government organizations and others to protect sensitive information. It replaces DES as the encryption standard.
AES-XCBC	Cipher Block Chaining. A key-dependent one-way hash function (MAC) used with AES in conjunction with the Cipher-Block-Chaining mode of operation, suitable for securing messages of varying lengths, such as IP datagrams.

TABLE 69

Term	Definition
AH	Authentication Header - like ESP, AH provides data integrity, data source authentication, and protection against replay attacks but does not provide confidentiality.
DES	Data Encryption Standard is the older encryption algorithm that uses a 56-bit key to encrypt blocks of 64-bit plain text. Because of the relatively shorter key length, it is not a secured algorithm and no longer approved for Federal use.
3DES	Triple DES is a more secure variant of DES. It uses three different 56-bit keys to encrypt blocks of 64-bit plain text. The algorithm is FIPS-approved for use by Federal agencies.
ESP	Encapsulating Security Payload is the IPsec protocol that provides confidentiality, data integrity and data source authentication of IP packets, and protection against replay attacks.
IKE	Internet Key Exchange is defined in RFC 2407, RFC 2408 and RFC 2409. IKEv2 is defined in RFC 4306. IKE uses a Diffie-Hellman key exchange to set up a shared session secret, from which cryptographic keys are derived and communicating parties are authenticated. The IKE protocol creates a security association (SA) for both parties.
MD5	Message Digest 5, like SHA-1, is a popular one-way hash function used for authentication and data integrity.
SHA	Secure Hash Algorithm, like MD5, is a popular one-way hash function used for authentication and data integrity.
MAC	Message Authentication Code is a key-dependent, one-way hash function used for generating and verifying authentication data.
HMAC	A stronger MAC because it is a keyed hash inside a keyed hash.
SA	Security Association is the collection of security parameters and authenticated keys that are negotiated between IPsec peers.

The following limitations apply to using IPsec:

- IPsec-specific statistics are not supported.
- To change the configuration of a secure tunnel, you must delete the tunnel and recreate it.
- There is no RAS message support for IPsec.
- IPsec can only be configured on IPv4 based tunnels.
- Secure Tunnels cannot be defined with VLAN Tagged connections.

IPSec for the 8 Gbps platforms

The 8 Gbps platforms use AES-GCM-ESP as a single, pre-defined mode of operation for protecting all TCP traffic over an FCIP tunnel. AES-GCM-ESP is described in RFC-4106. Key features are listed below:

- Encryption is provided by AES with 256 bit keys.
- The IKEv2 key exchange protocol is used by peer switches and blades for mutual authentication.
- IKEv2 uses UDP port 500 to communicate between the peer switches or blades.
- All IKE traffic is protected using AES-GCM-ESP encryption.
- Authentication requires the generation and configuration of 32 byte pre-shared secrets for each peer switch or blade.
- An SHA-512 hash message authentication code (HMAC) is used to check data integrity and detect third party tampering.
- PRF is used to strengthen security. The PRF algorithm generates output that appears to be random data, using the SHA-512 HMAC as the seed value.
- A 2048 bit Diffie-Hellman (DH) group is used for both IKEv2 and IPSec key generation.
- The SA lifetime limits the length of time a key is used. When the SA lifetime expires, a new key is generated, limiting the amount of time an attacker has to decipher a key. Depending on the length of time expired or the length of the data being transferred, parts of a message maybe protected by different keys generated as the SA lifetime expires. For the 8 Gbps Extension Switch and Blade, the SA lifetime is approximately eight hours, or two gigabytes of data, whichever occurs first.

- ESP is used as the transport mode. ESP uses a hash algorithm to calculate and verify an authentication value, and also encrypts the IP datagram.

QOS, DSCP, and VLANs

Quality of Service (QoS) refers to policies for handling differences in data traffic. These policies are based on data characteristics and delivery requirements. For example, ordinary data traffic is tolerant of delays and dropped packets, but voice and video data are not. QoS policies provide a framework for accommodating these differences in data as it passes through a network.

QoS for Fibre Channel traffic is provided through internal QoS priorities. Those priorities can be mapped to TCP/IP network priorities. There are two options for TCP/IP network-based QoS:

- Layer three DiffServ code Points (DSCP).
- VLAN tagging and Layer two class of service (L2CoS).

DSCP quality of service

Layer three class of service DiffServ Code Points (DSCP) refers to a specific implementation for establishing QoS policies as defined by RFC2475. DSCP uses six bits of the Type of Service (TOS) field in the IP header to establish up to 64 different values to associate with data traffic priority.

DSCP settings are useful only if IP routers are configured to enforce QoS policies uniformly within the network. IP routers use the DSCP value as an index into a Per Hop Behavior (PHB) table. Control connections and data connections may be configured with different DSCP values. Before configuring DSCP settings, determine if the IP network you are using implements PHB, and consult with your WAN administrator to determine the appropriate DSCP values.

VLANs and layer two quality of service

Devices in physical LANs are constrained by LAN boundaries. They are usually in close proximity to each other, and share the same broadcast and multicast domains. Physical LANs often contain devices and applications that have no logical relationship. Also, when logically related devices and applications reside in separate LAN domains, they must be routed from one domain to the other.

A VLAN is a virtual LAN network. A VLAN may reside within a single physical network, or it may span several physical networks. Related devices and applications that are separated by physical LAN boundaries can reside in the same VLAN. Also, a large physical network can be broken down into smaller VLANs. VLAN traffic is routed using 802.1Q-compliant tags within an Ethernet frame. The tag includes a unique VLAN ID, and Class of Service (CoS) priority bits. The CoS priority scheme (also called Layer two Class of Service or L2CoS), uses three Class of Service (CoS or 802.1P) priority bits, allowing eight priorities. Consult with your WAN administrator to determine usage.

When both DSCP and L2CoS are used

If an FCIP tunnel or circuit is VLAN tagged, both DSCP and L2CoS are relevant, unless the VLAN is end-to-end, with no intermediate hops in the IP network. The following table shows the default mapping of DSCP priorities to L2Cos priorities. This may be helpful when consulting with the WAN administrator. These values may be modified per FCIP tunnel.

TABLE 70 Default Mapping of DSCP priorities to L2Cos Priorities

DSCP priority/bits	L2CoS priority/bits	Assigned to:
46 / 101110	7 / 111	Class F
7 / 000111	1 / 001	Medium QoS
11 / 001011	3 / 011	Medium QoS

TABLE 70 Default Mapping of DSCP priorities to L2Cos Priorities (Continued)

DSCP priority/bits	L2CoS priority/bits	Assigned to:
15 / 001111	3 / 011	Medium QoS
19/ 010011	3 / 011	Medium QoS
23 / 010111	3 / 011	Medium QoS
27 / 011011	0 / 000	Class 3 Multicast
31 / 011111	0 / 000	Broadcast/Multicast
35 / 100011	0 / 000	Low Qos
39 / 100111	0 / 000	Low Qos
43 / 101011	4 / 100	High QoS
47 / 101111	4 / 100	High QoS
51 / 110011	4 / 100	High QoS
55 / 110111	4 / 100	High QoS
59 / 111011	4 / 100	High QoS
63 / 111111	0 / 000	-

Open systems tape pipelining

Open Systems Tape Pipelining (OSTP) can be used to enhance open systems SCSI tape write I/O performance. To implement OSTP over FCIP, you must enable the following two features:

- FCIP Fastwrite and Tape Acceleration.
- FC Fastwrite.

FCIP Fastwrite and Tape Acceleration

When the FCIP link is the slowest part of the network, consider using FCIP Fastwrite and Tape Read and Write Pipelining. FCIP Fastwrite and Tape Acceleration are two features that provide accelerated speeds for read and write I/O over FCIP tunnels in some configurations:

- FCIP Fastwrite accelerates the SCSI write I/Os over FCIP.
- Tape Acceleration accelerates SCSI read and write I/Os to sequential devices (such as tape drives) over FCIP, which reduces the number of round-trip times needed to complete the I/O over the IP network and speeds up the process. To use Tape Acceleration, you must also enable FCIP Fastwrite.

Both sides of an FCIP tunnel must have matching configurations for these features to work. FCIP Fastwrite and Tape Acceleration are enabled by turning them on during the tunnel configuration process. They are enabled on a per-FCIP tunnel basis.

Consider the constraints described in Table 71 when configuring tunnels to use OSTP.

TABLE 71 OSTP constraints

FCIP Fastwrite	Tape Acceleration
Each GbE port supports up to 2048 simultaneous accelerated exchanges, which means *a total of 2048 simultaneous exchanges combined* for Fastwrite and Tape Acceleration.	Each GbE port supports up to 2048 simultaneous accelerated exchanges, which means *a total of 2048 simultaneous exchanges combined* for Fastwrite and Tape Acceleration.
Does not natively support multiple equal-cost path configurations. Traffic isolation zoning can be used to support these configurations.	Does not natively support multiple equal-cost path configurations or multiple non-equal-cost path configurations. Traffic isolation zoning can be used to support these configurations.

TABLE 71 OSTP constraints

FCIP Fastwrite	Tape Acceleration
Class 3 traffic is accelerated with Fastwrite.	Class 3 traffic is accelerated between host and sequential device.
	With sequential devices (tape drives), there are 1024 initiator-tape (IT) pairs per GbE port, but 2048 initiator-tape-LUN (ITL) pairs per GbE port. The ITL pairs are shared among the IT pairs. For example: Two ITL pairs for each IT pair as long as the target has two LUNs. If a target has 32 LUNs, 32 ITL pairs for IT pairs. In this case, only 64 IT pairs are associated with ITL pairs. The rest of the IT pairs are not associated to any ITL pairs, so no Tape Acceleration is performed for those pairs. By default, only Fastwrite-based acceleration is performed on the unassociated pairs.
	Does not support multiple non-equal-cost path between host and sequential device

FICON emulation features

FICON emulation supports FICON traffic over IP WANs using FCIP as the underlying protocol. FICON emulation features support performance enhancements for specific applications. If you are using FCIP for distance extension in a FICON environment, evaluate the need for these features before you run the FCIP configuration wizard. FICON emulation may be configured by selecting **Advanced Settings** on the **Add Tunnel** or **Edit Tunnel** dialogs. The following features are available:

- IBM z/OS Global Mirror (z Gm) emulation.
- Tape write pipelining.
- Tape read pipelining.
- Teradata pipelining

IBM z/OS Global Mirror (z Gm) emulation

The IBM z/OS Global Mirror (z Gm) application, formerly known as eXtended Remote Copy (XRC), is a DASD application that implements disk mirroring, as supported by the disk hardware architecture and a host software component called System Data Mover (SDM). The primary volume and the secondary mirrored volume may be geographically distant across an IP WAN. The latency introduced by greater distance creates delays in anticipated responses to certain commands. The FICON pacing mechanism may interpret delays as an indication of a large data transfer that could monopolize a shared resource, and react by throttling the I/O. IBM z/OS Global Mirror (z Gm) emulation provides local responses to remote hosts, eliminating distance related delays. A FICON XRC Emulation License is required to enable IBM z/OS Global Mirror (z Gm) Emulation.

Tape write pipelining

FICON tape write pipelining improves performance for a variety of applications when writing to tape over extended distances. FICON tape write pipelining locally acknowledges write data records, enabling the host to generate more records while previous records are in transit across the IP WAN. If exception status is received from the device, the writing of data and emulation is terminated. The FICON Tape Emulation License is required to enable FICON Tape Write Pipelining.

Tape read pipelining

FICON tape read pipelining improves performance for certain applications when reading from FICON tape over extended distances. FICON tape read pipelining reads data from tape directly from the tape device. Reading of tape continues until a threshold is reached. The buffered data is forwarded to the host in response to requests from the host. When the host sends the status accept frame indicating that the data was delivered, the read processing on the device side credits the pipeline and requests more data from the tape. If exception status is received from the device, the reading of data and emulation is terminated. The FICON Tape Emulation License is required to enable FICON Tape Read Pipelining.

Teradata pipelining

Teradata emulation reduces latency on links to Teradata warehouse systems caused by WAN propagation delays and bandwidth restrictions. It accomplishes this by processing selected FICON commands for associated control, data, and status responses. FICON Teradata Emulation is supported between FICON Channels and FICON Teradata controllers. This feature is available only on 8 Gbps Extension Switch and Blade platforms operating with Fabric OS 7.0 and later.

Write pipelining

For write commands, control and status frames are generated for the host side of the WAN to pipeline write commands over the same or multiple exchanges.

Read pipelining

For read operations received by the device side of the WAN, a number of anticipatory read commands arAdde generated and transferred to the device. The data and status associated with these commands are sent to the host side of the WAN and queued in anticipation of host-generated read commands.

Connecting cascaded FICON fabrics over FCIP

This section provides a basic guide of IP best practices for connecting cascaded FICON fabrics over an IP network through FCIP and merging the fabrics. Included are planning considerations, steps for configuring an IP link between two Extension Switches and merging them into one fabric, and steps for configuring DWDM links to use R_RDYs.

IP best practice for connecting the fabrics is to perform the following steps in order:

1. Configure all IP tunnels and circuits between the fabrics.
2. Merge the FICON fabrics.

   NOTE
   Merging two cascaded FICON fabrics may be disruptive to current I/O operations in both fabrics, as it needs to disable and enable the switches in both fabrics. The merge process will not make any configuration changes on the primary (production) fabric that are disruptive.

3. Configure FICON Emulation features, if applicable.

   NOTE
   Consult with a qualified support specialist before implementing the FICON Acceleration feature.

The following procedures apply to configuring an IP connection between two Extension Switches or Blades, then merging the fabrics to which they belong. Before fabrics are linked and merged, the individual fabrics would appear as in Figure 409.

FIGURE 409 Switch display on Connectivity Map before fabric merge

After creating an IP link between sw31 and sw30, then merging the Test_FID_02 fabric is into the FID_02_Fabric, the merged fabric will look similar to that shown in the following figure.

FIGURE 410 Switch display on Connectivity Map when fabrics merged

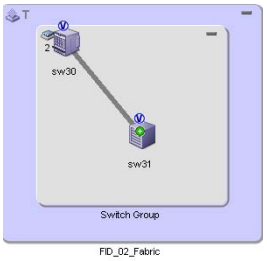

Procedures in this section may refer to additional sections in this chapter or additional chapters in this manual for more detailed information. This section assumes that the switches in the fabrics to be merged have been configured for FICON operation using procedures under "Configuring a switch for FICON operation" "FICON Environments".

Planning the configuration

Create a drawing to summarize the following elements of your planned configuration.

- IP network connections
  - Tunnels
  - Addresses
  - Bandwidth requirements for all circuits
  - Label all circuits and tunnels

  Determine how the IP network will be used by identifying redundant routes, network distance for each route, and minimum and maximum bandwidth requirements. The FICON acceleration feature is required for distances greater than 300 km. Before configuring this feature, Fabric OS professional services are highly recommended.
- Network distance

  Make sure network distance is measured in actual network delay. The FICON Acceleration license is required if distance exceeds 300 meters.
- Traffic Isolation (TI) zones.

  Determine the exact ports to use for TI zones.

  TI zones are used to segregate traffic such as tape backup and production DASD traffic in cascaded fabrics. If using TI zones, determine if zones should have failover disabled or enabled.

The FICON acceleration feature emulates the device it for which it is enabled. Although it effectively acts like the control unit cache, the control unit has a common processor that coordinates data written to it from different interfaces. Therefore, you must force traffic to a specific path using TI zones with failover disabled to ensure data is delivered in order. An alternative to TI zones is to use independent fabrics to ensure only one path is available. You can also use independent virtual fabrics.

- Buffer-to-buffer credit management for long-distance links.
- Use of dense wavelength division multiplexing (DWDM) or time division multiplexing (TDM) interfaces and buffer-to-buffer credit management for these interfaces.

  Typically the long distance BB credits are supplied if DWDM is used. Some older DWDM interfaces do not supply BB credits (R_RDY) so check with the DWDM vendor. You may need to calculate the correct number of BB credits required if using DWDM that does not provide BB credits. Note that BB credits depend not only on distance, but average frame size as well. Be sure and contact a Fabric OS support professional for assistance.

  Double check the type of optics required since long wave optics are commonly ordered for mainframe environments and occasionally DWDM interfaces use shortwave optics. Also find out if a TDM card is being used as you will need to follow procedures under "Configuring DWDM links to use R_RDYs" on page 854.

Configuring IP links and merging the fabrics

Use the following procedures to configure an IP connection between two Extension Switches or Blades, then merging the fabrics to which they belong.

1. Perform all tasks under "FCIP configuration guidelines" on page 855.
2. Configure tunnels circuits between the switches by following steps under "Configuring an FCIP tunnel" on page 856
3. Follow these guidelines when configuring tunnels using the **Add FCIP Tunnel** dialog box:
   - You can configure either switch as switch 1 or switch 2.
   - Specifications for FCIP circuits per tunnel, number of IP addresses per port, and other trunking capacities for the 8 Gbps Extension Switch and Blade are detailed in the *Fabric OS FCIP Administrator's Guide*.
   - For configuring **Port Type** on the **Add FCIP Tunnel** dialog box, VEX connections are for Fibre Channel Routing (FCR) and are not supported for FICON. Select **VE Port** as this refers to an E_Port connected to an IP instead of a Fibre Channel link.
4. On the **FCIP Tunnel Advanced Settings** dialog box **Transmission** tab (access by selecting **Advanced Settings** on the **Add FCIP Tunnel** dialog box), follow these guidelines:
   - Best practice is to **Enable Compression** and set **Compression Mode** to **Auto**.
   - **Fast Write** is not necessary for FICON. Keep in mind that disk-to disk mirroring is native FCP even if the front-side ports are FICON. If sharing FICON and FCP on the same tunnel, you can enable **Fast Write**. Enabling Fast Write depends on the application being extended over FCIP. Refer to "When to enable Fast Write" on page 854. As with any feature, if it is not needed. the best practice is to disable it.
   - The **Tape Acceleration** option is for open systems tape, not FICON tape emulation.
5. On the **FCIP Tunnel Advanced Settings** dialog box **Security** tab (access by selecting **Advanced Settings** on the **Add FCIP Tunnel** dialog box), follow these guidelines:
   - Recommended best practice is to enable IPSec. IPSEC on an 8 Gbps Extension Switch is Advanced Encryption Standard (AES) 256 only.
   - The **preShared Key** should be 32 alphanumeric characters and must match in tunnel configurations for both switches.
6. On the **FCIP Tunnel Advanced Settings** dialog box **FICON Emulation** tab (access by selecting **Advanced Settings** on the **Add FCIP Tunnel** dialog box), follow these guidelines:

- The recommended best practice is to complete all configuration for the IP connection and merge the fabrics before configuring FICON emulation. Configure settings in this tab after merging the fabrics.
- FICON Acceleration features require a license. These features include FICON Tape emulation, FICON XRC emulation, and FICON teradata pipelining.
- Select **Populate Default Values** unless recommended otherwise by a qualified Fabric OS support professional.
- Only select the features you require.
- Whenever selecting a FICON emulation feature, also select **Enable FICON Tin Tir Emulation** and **Enable FICON Device Level Ack Emulation**.
- Set only one type of acceleration feature per tunnel. Tape and XRC Emulation must not be enabled on the same tunnel.
- Except for tape and XRC emulation, it is not necessary to isolate traffic for emulation features.

7. Configure circuits for tunnels using steps under and "Adding an FCIP circuit" on page 859
8. Follow these guidelines when configuring circuits through the **Add FCIP Circuit** dialog box.
   - Start by configuring circuit 0, and then add additional circuits if desired.
   - Be sure to select **Verify IP Connectivity** to test the connection between both switches. IP connectivity is tested with the ping utility.
   - Make changes to IP settings by selecting **Advanced Settings** to display the **FCIP Circuit Advanced Settings** dialog box. Make changes to this dialog box only under direction of network administrators.
9. After you complete tunnel and circuit configuration between the fabrics, merge the fabrics using the Cascade FICON Fabrics Merge wizard by following procedures under "Cascaded FICON fabric merge" on page 928.
10. Consider the following when merging fabrics:
    - When merging fabrics, the primary fabric is the production fabric where disruption should not occur. The merge process will not make any disruptive configuration changes on the primary fabric. The secondary fabric is merged into the primary fabric.
    - Any CHPIDs with local connections in the secondary fabric should be configured offline.
    - The merged fabric will retain zone configurations from the primary fabric, so any zone configurations involving ports on the secondary fabric must be redone after fabric merge.
    - If the configuration wizard was used previously, the fabrics they will not merge. This is because the wizard sets the fabric security policies based on the fabric that is present at the time. Typically, this happens when a DR site is tested and validated before merging the fabrics. If this occurs you will get fabric security violation errors and the ports will automatically disable. To resolve this, keep working through the wizard regardless of any error messages. Do not bother to set any long distance modes. Re-enable the ports that were disabled due to a security violation as illustrated, then repeat the process.
    - There are no long-distance parameters to configure for IP links. Except for CWDM, most DWDM equipment provides the required buffer credits Typically, it is only necessary to set long distance mode when there are direct fibre runs.
    - For other considerations and a description of the merge process, refer to "Cascaded FICON fabric merge" on page 928.
11. After fabrics are successfully merged, configure FICON Emulation features as required. Refer to step 6.
12. Rezone the fabric as zoning was removed from the secondary fabric that you merged.
13. Configure traffic isolation (TI) zoning. Refer to the information on TI zones under "Planning the configuration" on page 851 and the "Traffic Isolation zones" section of "Zoning".
14. Clear error counters, which are common during switch configuration, by right-clicking the switch in the Connectivity Map or Product List and selecting **Performance > Clear Counters**. When merging fabrics in production environments, always check with the system administrator before clearing error counters.

Configuring DWDM links to use R_RDYs

TDM requires that you configure DWDM links to use R_RDYs and not VC_RDYs. The only way to turn off VC_RDYs is to start with QoS "OFF," and then turn on ISL R_RDY mode. Execute the following Fabric OS commands on E_Ports (ISL connections).

1. Enter the following command to disable credit recovery on a port.

   portcfgcreditrecovery --disable *slot/port speed*

2. Enter the following command to set the speed for the link. Only speeds supported by the installed SFP are supported. Use 0 to set back to automatic sensing mode.

   portcfgspeed *slot/port speed*

3. Enter the following command to disable QoS.

   portcfgqos --disable *slot/port*

4. Enter the following command to enable ISL R_RDY mode.

   portcfgislmode *slot/port,* **1**

5. Enter the following command to disable trunking on the port.

   portcfgtrunkport *slot/port,* **0**

6. Enter the following command to display port settings:

   portcfgshow

Extending RDR applications over FCIP

This section provides considerations for configuring tunnels and circuits when extending remote data replication (RDR) applications over FCIP.

When to enable Fast Write

Enabling Fast Write depends on the application that you are extending over FCIP. Use the following table to determine if Fast Write should be enabled for a tunnel configuration. Enable Fast Write through the **FCIP Tunnel Advanced Settings** dialog box. Access this dialog box by selecting **Advanced Settings** on the **Add FCIP Tunnel** dialog box.

TABLE 72 Using Fast Write for extended applications

Manufacturer	RDR Application	Platform	Type	Use Fast Write
IBM	Global Mirror (PPRC)	DS	Async	No
IBM	Metro Mirror	DS	Sync	No
IBM	XIV	XIV	Sync	Yes
IBM	Global Mirror	SVC	Async	No
IBM	Metro Mirror	SVC	Sync	No
EMC	SRDF/A	Symmetrix	Async	Yes
EMC	SRDF/S	Symmetrix	Sync	Yes (SiRT disabled)
EMC	SRDF Adaptive Copy	Symmetrix	Async	Yes
EMC	MirrorView	CLARiiON	Async	Yes

TABLE 72 Using Fast Write for extended applications

Manufacturer	RDR Application	Platform	Type	Use Fast Write
EMC	MirrorView	CLARiiON	Sync	Yes
EMC	SANcopy	CLARiiON	Async	Yes
HDS	Universal Replicator (HUR)	All	Async	No
HDS	TrueCopy	All	Async	No
HP	Continuous Access	EVA	Hybrid	No
*	OSTP	Tape	Tape	Yes (required for OSTP)

Compression mode

More aggressive compression modes can be used for asynchronous mirroring. For synchronous mirroring, only hardware or standard compression should be used. This is because more aggressive algorithms work by receiving additional frames to find compressible patterns on larger blocks of data. The time it takes to read in these additional frames add latency, which may not be tolerated by synchronous mirroring. Set compression modes on the **Add FCIP Tunnel** dialog box.

Circuit Keep Alive Time Out values

The circuit **Keep Alive Time Out** value, located on the **Transmission** tab of **FCIP Circuit Advanced Settings** dialog box, should be less than the protocol timeout for the application being extended. This allows circuit failover to be non-disruptive. By default, the circuit keep alive is 10 seconds (10000 ms) and 1 second (1000 ms) for FICON. Set this to 6 seconds (6000 ms) for IBM peer-to-peer remote copy (PPRC). All other applications should use the default.

SRDF considerations

Use SRDF/S SiRT (Single Roundtrip) or SRDF/A with FCIP-FW but not both.?Using SRDF/A and SRDF/S on the same remote adapter (RA) ports on the array is not recommended.?Use different VE_Ports for the tunnels, as if the tunnel destinations are different. If there is only one destination (SRDF/A and SRDF/S are going to the same place), isolate traffic from the SRDF/A and SRDF/S RA ports using TI zones and configure the tunnels accordingly. Note that there may be differences in bandwidth, Fast Write, and compression mode tunnel parameters.

FCIP configuration guidelines

FCIP configuration always involves two or more Extension Switches. The following should take place first before you configure a working FCIP connection from the Management application:

- The WAN link should be provisioned and tested for integrity.
- Cabling within the data center should be completed.
- Equipment should be physically installed and powered on.
- The Management application must have management port access to the Extension Switches.
- The Management application must be able to discover the fabrics the contain the Extension Switches.
- The Extension Switches should be physically connected to the IP network they will be using to pass data, and the connection should be active and working.
- Identify all the devices in the data path between the Extension Switches, including Ethernet switches, Ethernet routers, firewalls, and common carrier equipment. A network diagram is very helpful. Support engineers may ask you to provide a network diagram when troubleshooting problems.

- Routers and firewalls must be configured to pass ARP, ICMP, and IP layer 3 protocols.
- Persistently disable the VE_Ports before you configure them. Ports on a new Extension Switch or Extension Blade are persistently disabled by default.
- Determine which features you are implementing, and gather the information needed to implement those features. Table 63 summarizes feature support per FCIP platform.

Virtual Port Types

Virtual ports may be defined as VE_Ports or VEX_Ports.

VE_Ports

VE_Ports (virtual E_Ports) are used to create interswitch links (ISLs) through an FCIP tunnel. If VE_Ports are used on both ends of an FCIP tunnel, the fabrics connected by the tunnel are merged.

VEX_Port

A VEX_Port enables FC-FC Routing Service functionality over an FCIP tunnel. VEX_Ports enable interfabric links (IFLs). If a VEX_Port is on one end of an FCIP tunnel, the fabrics connected by the tunnel are not merged. The other end of the tunnel must be defined as a VE_Port.

Configuring an FCIP tunnel

When you configure an FCIP extension connection, you create FCIP tunnels and FCIP circuits between two Extension Switches.

NOTE
In Fabric OS 8.0.1 release, you can configure FCIP tunnel between two 32 Gbps, Router Extension blades and between 32 Gbps, Router Extension blade and 16 Gbps 24-FC port, 18 GbE port device.

NOTE
Beginning with Fabric OS 8.1.0 or later, you can add a maximum of 10 FCIP circuit per tunnel for Brocade 7840 and Brocade SX6 Extension Blade.

1. Select **Configure > FCIP Tunnels**.

   The **FCIP Tunnels** dialog box is displayed.

   The dialog box displays a tree structure of all FCIP-capable switches in the discovered fabrics, Extension Switches, and configured tunnels. Details such as number of circuits configured for tunnels, connected switches in tunnels, and tunnel status display in the right columns.

2. To add an FCIP tunnel and circuits between switches, follow these steps:

   a. Select the switch you want to configure under the **Products tree**.

   b. Click the **Add** button, or right-click on the switch and select **Add Tunnel**.

      The **Add FCIP Tunnel** dialog box is displayed. The name of the switch you selected is displayed in the **Switch** field under **Switch One**. This dialog box allows you to configure settings for both switches on either end of the tunnel.

A **Circuits** properties table displays at the bottom of the dialog box. For 8 Gbps platforms, this may contain columns for multiple circuits. Actual as well as cached circuits display. You can configure circuits using the **Add**, **Edit**, **Delete**, **Enable**, and **Disable** buttons to the right of the table. For 4 Gbps platforms, the **Delete**, **Enable**, and **Disable** buttons do not display. In addition, the **Edit** operation is only supported for cached circuits.

c. Click **Select Switch Two** under **Switch Two** on the **Add FCIP Tunnel** dialog box to display the **Select Switch** dialog box.

NOTE
In the devices running Fabric OS 8.0.1, an error message displays when a non-compatible device running Fabric OS 8.0.0 or earlier is selected in **Switch Two**.

The **Select Switch** dialog box displays discovered Extension Switches.

d. Select the switch you want to connect to switch one and click **OK**.

The switch and fabric names display in the **Switch Two area of** the **Add FCIP Tunnel** dialog box.

e. Select the desired tunnel ID from the **Tunnel** list.

NOTE
You cannot assign a **Tunnel** ID until at least one circuit is configured. The **Add Circuit** dialog box returns you to the **Add FCIP Tunnel** dialog box to allow you to select the **Tunnel ID**. In the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, you must select **Tunnel ID** before adding any circuit because you cannot change the **Tunnel ID** after adding circuits.

f. Enter a description of the tunnel in the **Description** field.

g. Skip to step 7 and continue configuration.

h. Configure the **Port Type** by choosing **VE Port** or **VEX Port**. If **VEX Port** is selected, enter the **Fabric ID** and select the **Interop Mode**.

NOTE
In the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, the tunnel creation or modification fails if the VE port is persistently disabled or the switch is disabled with an error message showing Port enable failed. Port is persistently disabled.

NOTE
VEx ports are not supported on the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch.

3. Select the desired FC compression mode from the **FC Compression Mode** list.

The default option is **Off**. A **Standard** option provides hardware compression.

For 8 Gbps Extension switches and the 8 Gbps Extension blades, the additional options for compression are as follows:

- Select the **Moderate** option to enable a combination of hardware and software compression that provides more compression than hardware compression alone. This option supports up to 8 Gbps of FC traffic.
- Select the **Aggressive** option to enable a software-only compression option that provides a more aggressive algorithm. This option supports up to 2.5 Gbps of FC traffic.
- Select the **Auto** option to enable the system to set the best compression mode based on the tunnel's configured bandwidth and the bandwidth of all tunnels in the system.

For the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, the default option is **Off**. The additional options for compression are as follows:

- Select the **Fast Deflate** option to enable compression, handled by the FPGA, on a frame-by-frame basis and uses an LZ algorithm to compress.
- Select the **Aggressive Deflate** option to enable software-only compression that provides a more aggressive algorithm.
- Select the **Deflate** option to enable multiple FC frames to group together and send for compression.

4. (Fabric OS 16 Gbps 24-FC port, 18 GbE port switch only) To enable IP compression, select the **IP Extension Mode** enable check box and select the desired compression mode from the **IP Compression Mode** list.

   The default option is **Off**. The additional options for compression are as follows:

   - Select the **Aggressive Deflate** option to enable software-only compression that provides a more aggressive algorithm.
   - Select the **Deflate** option to enable multiple frames to group together and send for compression.

5. Select **Advanced Settings from the Add FCIP Tunnel dialog box.**

   The **Transmission** tab of the **FCIP Tunnel Advanced Settings** dialog box displays (Figure 411).

FIGURE 411 FCIP Tunnel Advanced Settings dialog box

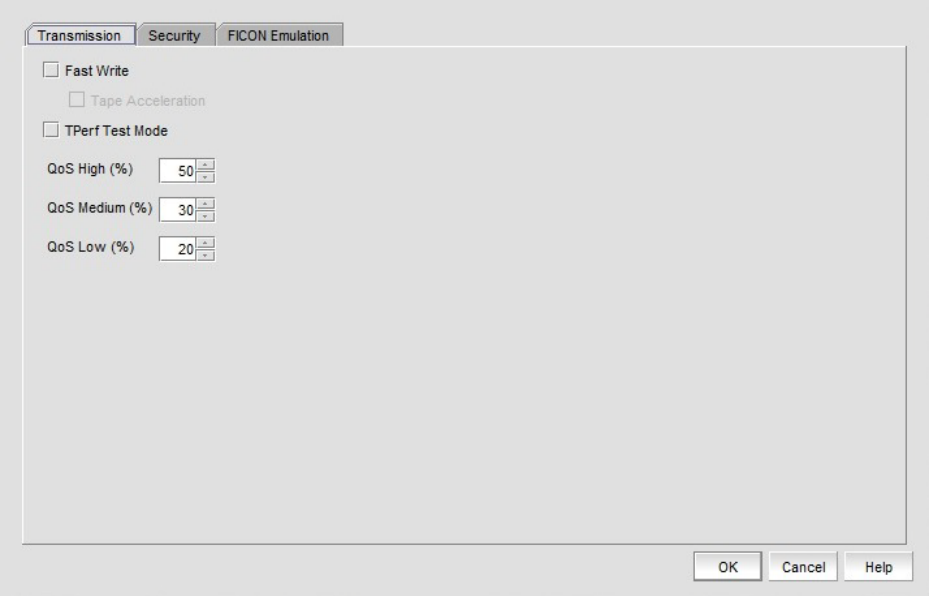

- Select the **Fast Write** check box to reduce delays caused by latency. Refer to "Enabling Open Systems Tape Pipelining" on page 864 for more information.
- Select the **Tperf Test Mode** check box for testing and troubleshooting tunnel. Refer to "Enabling Tperf test mode" on page 865 for more information.
- Select **L2CoS** and **DSCP** priorities. Refer to "QOS, DSCP, and VLANs" on page 847 for more information.
- Select **OK** to save the settings and close the dialog box.

6. To edit the configuration for an existing FCIP tunnel and circuits between two switches, follow these steps:

   a. From the **FCIP Tunnels** dialog box (refer to step 1), select the FCIP tunnel that you want to configure under the **Products** tree.

   b. Click **Edit**

The **Edit FCIP Tunnel** dialog box displays. This dialog box allows you to edit configurations on both switches on either end of the tunnel.

A **Circuits** properties table displays at the bottom of the dialog box. For 8 Gbps platforms, this may contain columns for multiple circuits. Actual as well as cached circuits display. You can configure circuits using the **Add**, **Edit**, **Delete**, **Enable**, and **Disable** buttons to the right of the table. For 4 Gbps platforms, the **Delete**, **Enable**, and **Disable** buttons do not display. In addition, the **Edit** operation is only supported for cached circuits.

c. Change configuration settings as required using the following steps.

7. To add a circuit, click **Add** to the right of the **Circuits** properties table at the bottom of the dialog box.

   The **Add FCIP Circuit** dialog box is displayed. Continue with "Adding an FCIP circuit".

Logical switch function on the FCIP Tunnels dialog box

The display and function of tunnels and circuits created on logical switches and with shared GigE ports varies on the **FCIP Tunnels** dialog box according to the discovery of the default switch and user-configured logical switches as follows:

- If the default and user-configured logical switches are discovered, all tunnels and circuits created on the logical switch display, including circuits with shared GigE ports.
- If the user-configured logical switch is discovered and the default logical switch is not discovered, the circuits and tunnels with shared GigE ports will be listed in the tunnel, but they cannot be edited or deleted.
- In a fabric with two logical switches that have a shared GigE port and only the default logical switch for one logical switch is discovered, the circuits and tunnels with shared GigE ports will be listed in the tunnel, but they cannot be edited or deleted.

For details on configuring FCIP with logical switches, use the following references:

- "Using FCIP with logical switches" section in the *Fabric OS FCIP Administrator's Guide*.
- "Virtual Fabrics".

Adding an FCIP circuit

When adding a new FCIP tunnel, you can add an FCIP circuit by selecting the **Add** button to the right of the **Circuits** properties table on the **Add FCIP Tunnel** dialog box. For 8 Gbps platforms, you can add multiple FCIP circuits to the tunnel with this button.

You can add circuits to existing FCIP tunnels through the **Edit FCIP Tunnel** dialog box. To display this dialog box, right-click a tunnel on the **FCIP Tunnels** dialog box and select **Edit Tunnel** or select a tunnel and click the **Edit** button. For details, refer to "Configuring an FCIP tunnel" on page 856.

FIGURE 412 Add FCIP Circuit dialog box

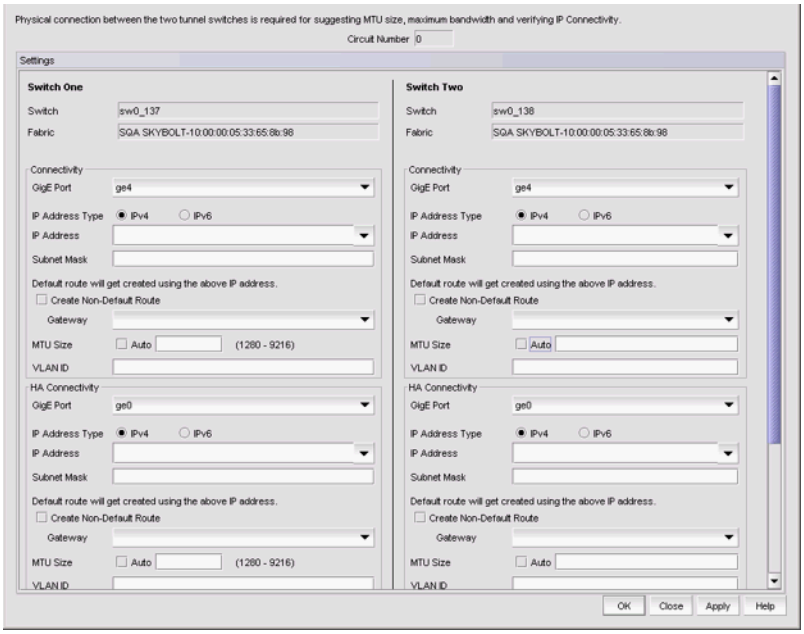

Use the following steps to add a circuit.

1. Select the **GigE Port** used for the Ethernet connection on each switch. The choices available depend on the Extension Switch or Blade model.

   For the 8 Gbps Extension blade, GbE ports display according to the operating mode set for the blade:

   - 1 Gbps mode - Ports ge0 through ge9
   - 10 Gbps mode - Ports xge0 and xge1
   - Dual mode - Ports ge0 through ge9 and xge0

   NOTE
   For the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, ge0 and ge1 are 40G ports and ge2 through ge17 are 1/10G ports. You cannot create circuits with 40G ports unless you have WAN rate upgrade 2 license.

2. Select **Cross port circuit** to configure the 10 GigE port on an 8 Gbps Blade platform as a 10 Gbps lossless failover circuit.
3. Select the **IP Address Type**. The implementation is a dual IP layer operation implementation as described in RFC 4213. IPv6 addresses can exist with IPv4 addresses on the same interface, but the FCIP circuits must be configured as IPv6 to IPv6 and IPv4 to IPv4 connections. IPv6-to-IPv4 connections are not supported. Likewise, encapsulation of IPv4 in IPv6 and IPv6 in IPv4 is not supported.
4. Select the **IP Address** for each port. This implementation of IPv6 uses unicast addresses for the interfaces with FCIP circuits. The unicast address must follow the RFC 4291 IPv6 standard and use the IANA-assigned IPv6 Global Unicast address space (2000::/3).

   NOTE
   For the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, you can configure two IP addresses, one as primary in the Connectivity group and the other as secondary in the HA Connectivity group. The primary address is mandatory.

5. For IPv4 addresses, specify the **Subnet Mask**. For IPv6 addresses, specify the prefix length.

The default is created from the IP address and Subnet Mask. If you want to create a route through a gateway router, click **Create Non-Default Rout**e, and select a **Gateway** address.

6. Enter the **MTU Size**.

   For SAN traffic, the largest possible Maximum Transmission Unit (MTU) size is generally the most efficient. MTU rates must match on both ends of the tunnel.

   For 4 Gbps platforms, enter a value from 1260 through 2348.

   For 8 Gbps platforms, enter a value from 1260 through 1500.

   For the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, enter a value from 1280 through 9216 or select the **Auto** check box to set auto mode (1). The **Auto** check box is only available for new circuits. If you created the circuit using the CLI, the **Auto** check box is not available. If you select auto mode, the MTU size value displays as 1 and is not editable.

   If you have an active connection between switch one and switch two, click **Verfiy IP Connectivity** under **Switch One Settings** to test the connection. To determine a suggested size, packets are sent across the FCIP tunnel, starting at the largest possible size packet that can be sent over IP. If a valid connection response is not received, a smaller packet is sent. This continues until a valid connection response is received, and that size becomes the suggested MTU size. MTU settings must match at both ends of the tunnel, and the setting specified under **Switch One Settings** is automatically applied to switch two.

   NOTE
   The function of the **Verify IP Connectivity** button function requires an active IP connection. The button is not available for the **Add FCIP Circuit** and **Edit FCIP Circuit** dialog boxes for 8 Gbps Extension platforms.

   NOTE
   Verify IP Connectivity button is not supported on the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch.

7. If a VLAN ID is used to route frames between the switches over the physical connection, enter the **VLAN ID** under **Switch One Settings**. You must assign the same VLAN ID to both the switches. If the VLAN ID is not same, an error message displays.

   The VLAN ID is an integer value from 1 through 4094 which sets the VLAN tag value in the header, assigning the traffic to that specific VLAN. Layer 2 class of service (L2CoS) values may be assigned to establish traffic priorities over a VLAN.

8. The **Metric** list is used to identify a failover circuit. By assigning a non-zero metric (1), you identify the circuit as a failover circuit. By default, a circuit is assigned a metric of 0. If a metric 0 circuit fails, FCIP trunking tries first to retransmit any pending send traffic over another circuit with a metric of 0. If no circuits with a metric of 0 are available, then the pending send traffic is retransmitted over any available circuit with a metric of 1.

   The default metric value for a crossport circuit configuration is 1. If a failover circuit is created with a metric of 0, it will be used for load balancing and not for failover.

9. Designate a **Failover Group ID** for the circuit from 0 through 9. A value of 0 designates the default failover group or no failover group.

   With circuit failover groups you can better control which metric 1 circuits will be activated if a metric 0 circuit fails. For Circuit Failover Grouping, you define a set of metric 0 and metric 1 circuits that are part of the same failover group. When all metric 0 circuits in the group fail, metric 1 circuits will take over data transfer, even if there are metric 0 circuits still active in other failover groups. For more information on Circuit Failover Grouping, refer to "Circuit Failover Grouping" on page 840.

   NOTE
   Failover Group ID will only be enabled when the switch or chassis is using Fabric OS v7.2 and later.

10. Select values for bandwidth settings. An uncommitted bandwidth is not allowed on an FCIP circuit. You must select **Committed** bandwidth. If you want to use ARL, set **Minimum** and **Maximum** bandwidth values. Bandwidth grows towards the maximum and reduces towards the minimum based on traffic conditions. If you do not want to use ARL, set **Minimum** and **Maximum** to the same value to set a single committed bandwidth. Refer to "Adaptive Rate Limiting" on page 843 for more information about ARL.

    NOTE
    The **Committed** value range in the **Add FCIP Circuit** dialog box depends on the Extension Switch or Blade platform.

    NOTE
    For the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, the **Minimum** bandwidth value is 20 Mbps and **Maximum** bandwidth value is based on the available license and port speed. For example, if user selects 1G port for connectivity and 10G port for HA connectivity, the maximum allowed bandwidth on the circuit is 1 Gbps but for 10G and 40G combination, the maximum allowed bandwidth on the circuit is 10 Gbps.

11. If the physical connection exists, click **Verify IP Connectivity** to test the connection between switch one and switch two. The IP connectivity of the connection is tested with the ping utility.
12. Select **Advanced Settings from the Add FCIP Circuit dialog box and continue if you want to do any of the following:**

    - Disable selective acknowledgment if your system cannot support selective acknowledgment.
    - Set the keep alive timeout to a value other than the default of 10 seconds.
    - Set the minimum retransmission time to a value other than the default of 100 ms.
    - Set the maximum retransmits to a value other than the default.
    - Use TCP/IP DSCP or L2CoS to prioritize FC traffic.
    - Set an adaptive rate limiting (ARL) algorithm mode to a value other than the default of Auto.

    If you select **Advanced Settings**, the **Transmission** tab of the **FCIP Circuit Advanced Settings** dialog box displays (Figure 413).

FIGURE 413 FCIP Circuit Advanced Settings dialog box

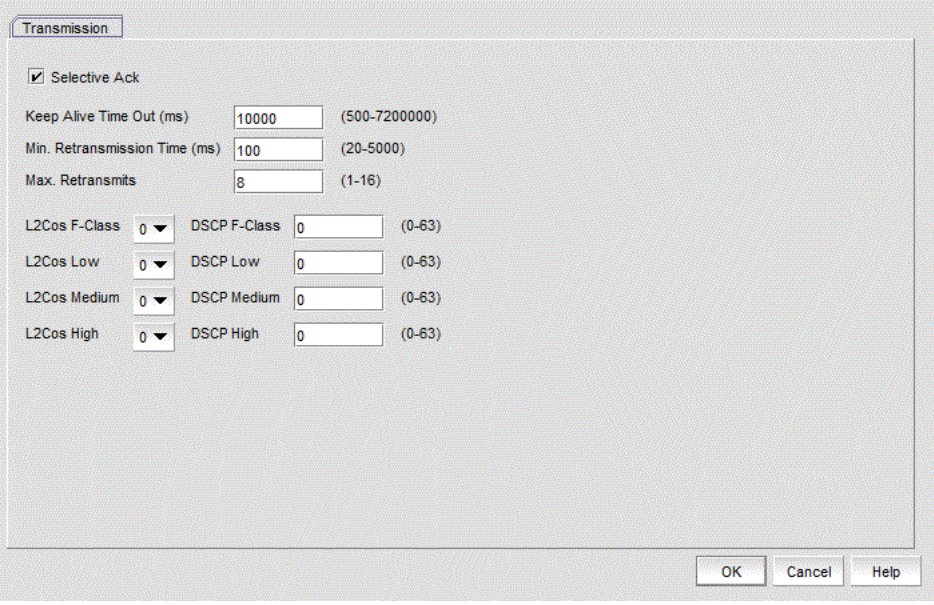

a. Clear the **Selective Ack** check box to disable selective acknowledgement. Do not clear the check box if your system does not support selective acknowledgement..

b. Enter a value in the **Keep Alive Time Out (ms)** field to override the default value of 10000 ms. As shown, the range is from 500 through 7200000.

c. Enter a value in the **Min. Retransmission Time (ms)** field to override the default value of 100 ms. As shown, the range is from 20 through 2000.

d. Enter a value in the **Max. Retransmits** field to override the default value of 8. As shown, the range is from 1 through 16.

e. Select **L2CoS** and **DSCP** priorities. Refer to "QOS, DSCP, and VLANs" on page 847 for more information.

NOTE
For the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, select the **L2CoS** value from the list because the default value is not predefined.

f. Select an adaptive rate limiting (ARL) algorithm mode from the **ARL Algorithm** list.
This field is only applicable to the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch. AARL algorithm options include:

- **Auto** (default) — Select to allow ARL to determine the best method.
- **Reset** — Select to reset all connections to the minimum rate. Best used with less error tolerant devices.
- **Step-down** — Select to have connections step down incrementally. Best used with shorter links on error tolerant devices.
- **Timed-step-down** — Select to have connections step down incrementally in specific time slices. Best used with long latency links.

g. Click **OK** to save the settings and close the dialog box.

13. Click **Apply** on the **Add FCIP Circuit** dialog box to add the circuit and leave the dialog box open to add additional circuits. Click **OK** to add the circuit and close the dialog box.

NOTE
If you add more than the allowed number of circuits to the **Add FCIP Circuit** dialog box, an information message "The circuits are saved successfully and the maximum number of circuits (x) has been reached for this tunnel" displays.

14. Click **OK** to close the **Add FCIP Tunnel** dialog box.

Logical switch function in the Add FCIP Circuit dialog box

The display and function of circuits created on logical switches and with shared GigE ports varies on the **Add FCIP Circuit** dialog box according to the discovery of the default switch and user-configured logical switches as follows.

- If both the default logical switch and user-configured logical switch are discovered:
  - The **GigE Port** list displays all GigE ports in the logical switches, including ports from the default logical switches and user-configured logical switches.
  - A circuit created with a shared GigE port will create an interface on the default logical switch, but the circuit will be created on the selected logical switch.
  - Selecting **Verify IP Connectivity** verifies the connectivity using the default logical switch because the interface is on this switch.
- If the user-configured logical switch is discovered and the default logical switch is not discovered:
  - On adding a circuit, only the GigE ports present in the logical switch will display.
  - You cannot display or edit shared circuits of the default logical switch.

- In a fabric with two logical switches that have a shared GigE port and only the default logical switch for one logical switch is discovered on:
  - On adding a circuit, only the GigE ports present in the logical switch with the discovered default logical switch will be listed.

For details on configuring FCIP with logical switches, use the following references:

- "Using FCIP with logical switches" section in the *Fabric OS FCIP Administrator's Guide*.
- "Virtual Fabrics".

Circuit configuration failure

When a tunnel cannot be created because the process for adding a new circuit configuration fails, the **FCIP Tunnel/Circuit Configurations** dialog box displays. Using this dialog box, you can perform the following tasks:

- Roll back the current changes to the circuit configuration.
- Elect to not roll back current circuit configuration changes.
- Continue configuring additional circuits at this point.
- Stop configuring additional circuits.

Configuring FCIP tunnel advanced settings

Compression, FCIP fast write and tape pipelining, IPsec and IKE policies, and FICON emulation features are configured as advanced settings.

1. Click **Advanced Settings on the Add FCIP Tunnel or Edit FCIP Tunnel dialog box.**

   The Advanced Settings dialog box is displayed. This dialog box has a Transmission tab, a Security tab, and a FICON Emulation tab.

2. Click **OK** to close **Advanced Settings** when you have configured the features that you want to implement.
3. Click **OK** to close the **Add FCIP Tunnel** dialog box.

Enabling Open Systems Tape Pipelining

Latency introduced by a long distance IP connection can negatively impact tape I/O performance. Open Systems Tape Pipelining (OSTP) may be used to improve performance on SCSI write I/Os to sequential devices (such as tape drives). When OSTP is used, the Extension Blades or Switches emulate write commands and responses locally to reduce delays caused by latency. Both sides of an FCIP tunnel must have matching configurations for these features to work. OSTP may be configured by selecting **Advanced Settings** on the **Add FCIP Tunnel** dialog box. OSTP options are available on the **Transmission** tab.

To enable OSTP, complete the following steps:

1. Select **Advanced Settings** on the **Add FCIP Tunnel or Edit FCIP Tunnel dialog box to display the Advanced Settings dialog box.**
2. From the **Transmission** tab, select the **Fast Write** check box.

   This enables the **Tape Acceleration** check box.

NOTE
For the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, the **Tape Accelartion** list contains Disabled, Write-only, and Read/Write options.

3. Select the **Tape Acceleration** check box.
4. Click **OK**.

Enabling Tperf test mode

To enable Tperf test mode, complete the following steps:

1. Select **Advanced Settings** on the **Add FCIP Tunnel or Edit FCIP Tunnel dialog box to display the Advanced Settings dialog box.**
2. From the **Transmission** tab, select the **TPerf Test Mode** check box.
3. Select the **Tape Acceleration** check box.
4. Click **OK**.

Tperf test mode should not be enabled during normal operations. It is only used for testing and troubleshooting tunnels. Refer to the *Fabric OS FCIP Administrator's Guide* for information about Tperf.

Configuring QoS percentages

For 8 Gbps platforms, you can adjust Quality of Service (QoS) priority percentages from the preset default values of 50 percent (High), 30 percent (Medium), and 20 percent (Low). Values for the three priority levels must equal 100 percent. A minimum of 10 percent is required for each level. You can adjust percentages in increments of 1 percent. To configure QoS percentages, complete the following steps:

1. Select **Advanced Settings** on the **Add FCIP Tunnel or Edit FCIP Tunnel dialog box to display the Advanced Settings dialog box.**
2. From the **Transmission** tab, click the up and down arrows by the **QoS (High)**, **QoS (Medium)**, and **QoS (Low)** percentage values to increase and decrease values.

Configuring the ARL mode

To set an adaptive rate limiting (ARL) algorithm mode to a value other than the default of Auto, complete the following steps.

1. Select **Advanced Settings** on the **Add FCIP Tunnel or Edit FCIP Tunnel dialog box.**

   The Advanced Settings dialog box displays.
2. Click the **Transmission** tab.
3. Select an adaptive rate limiting (ARL) algorithm mode from the **ARL Algorithm** list.
   This field is only applicable to the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch. ARL algorithm options include:
   - **Auto** (default) — Select to allow ARL to determine the best method.
   - **Reset** — Select to reset all connections to the minimum rate. Best used with less error tolerant devices.
   - **Step-down** — Select to have connections step down incrementally. Best used with shorter links on error tolerant devices.
   - **Timed-step-down** — Select to have connections step down incrementally in specific time slices. Best used with long latency links.

4. Click **OK** to save the settings and close the dialog box.

Configuring IPsec and IKE policies

IPsec and IKE policies are configured from the **Security** tab. The options and procedures are platform-dependent. Figure 414 on page 866 shows the **Security** tab for the 8 Gbps Extension Switch and 8 Gbps Extension Blade.

1. Select **Advanced Settings** on the **Add FCIP Tunnel or Edit FCIP Tunnel dialog box to display the Advanced Settings dialog box.**
2. Select the **Security** tab.

FIGURE 414 Advanced Settings Security Tab for the 8 Gbps Extension Switch and Blade

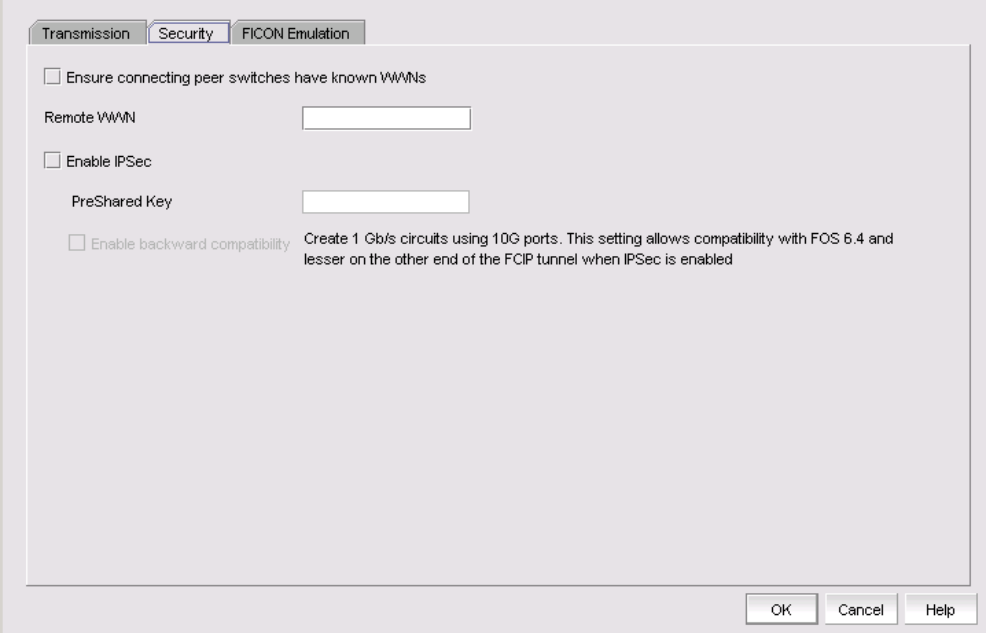

3. As an option, click **Ensure connecting peer switches have known WWNs**. This provides an added measure of security.
4. Enter the WWN for the remote switch.
5. Assign IKE and IPsec policies. For the 4 Gbps Extension Switch and Blade, you must choose from a list of policies. The 8 Gbps Extension Switch and Blade have predefined IKE and IPsec policies. These policies are enabled by selecting the **Enable IPSec** check box. Matching policies are applied to the remote switch. Note that the **Enable IPSec** check box is unavailable while editing the tunnels because the IPsec settings cannot be edited for the secured tunnels.

   NOTE
   IPsec settings cannot be edited. If you want to change settings, you must delete the tunnel and then create a new tunnel with the new settings.

6. In the **PreShared Key** field, specify the key for IKE authentication. Use the following specifications, depending on your extension platform:
   - For the 4 Gbps Extension Switch and Blade and the 8 Gbps Extension Blade, the key value must be from 12 through 32 alphanumeric characters. The length depends on the chosen IKE policy.
   - For the 8 Gbps Extension switch, the key value must be a minimum of 32 alphanumeric characters.

These policies are used to make the connection more secure through authentication and encryption. When you select a policy for the local switch, a matching policy is automatically selected on the remote switch. If no matching policy is found, you must manually configure the policy on the remote switch.

NOTE
For the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch, select a name from the **Policy Name** list. The list displays all the predefined policy names. Select the ellipses button to the right of the **Policy Name** field to configure an IPsec policy name using the **Configure IPsec Policy** dialog box. Minimum 16 and maximum 64 characters, special characters ~ @ % - _ + [ ] : are allowed for the policy name.

7. You can activate the **Enable backward compatibility feature** on 8 Gbps platforms if IPsec is enabled. This allows multiple 1 Gbps circuits to be created using 10 Gbps ports even if the switch at one end of the tunnel is using Fabric OS 7.0 and the switch at the other end is using versions earlier than Fabric OS v7.0. Note that this feature can only be enabled when IPsec is enabled and when circuits are configured without any advanced 10 Gbps features, such as lossless failover, multi-gigabit circuits, or 10 Gbps Adaptive Rate Limiting (ARL).

Configuring FICON emulation

FICON emulation and acceleration features and operating parameters are configured from the **FICON Emulation** tab (Figure 415). Before you configure these features, you must decide which features you want to implement, and you must look closely at the operational parameters to determine if values other than the default values are better for your installation.

1. Select **Advanced Settings** on the **Add FCIP Tunnel or Edit FCIP Tunnel dialog box to display the Advanced Settings dialog box.**
2. Select the **FICON Emulation** tab.

FIGURE 415 Advanced Settings FICON Emulation Tab

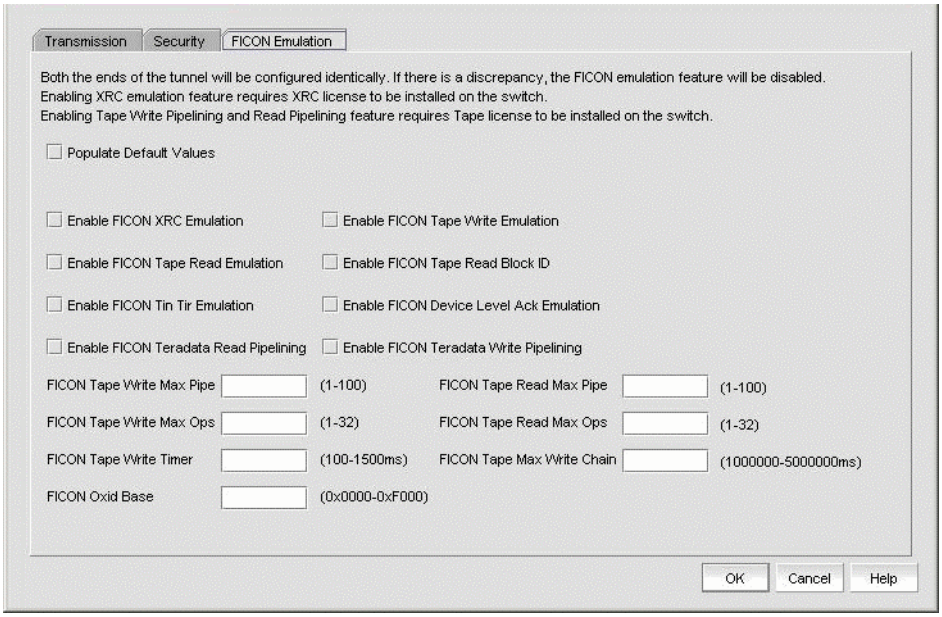

3. Select the check boxes for the FICON emulation features you want to implement.

The primary FICON emulation features are FICON XRC Emulation (IBM z/OS Global Mirror emulation), tape write pipelining, tape read pipelining, TIN/TUR emulation and device level ACK emulation provide support for the primary features. If you select any of the primary features, you must also select TIN/TUR emulation and device level ACK emulation.

For 8 Gbps platforms operating with Fabric OS 7.0 and later, you can also enable FICON Teradata read pipelining and FICON Teradata write pipelining.

4. Select **Populate Default Values** at the top of the dialog box to set all operational parameters for FICON emulation to default values. This option is not be enabled if existing values are configured for the tunnel.
5. Select individual operational parameters for FICON emulation.
   - **FICON Tape Write Max Pipe** defines a maximum number of channel commands that may be outstanding at a given time during write pipelining. Too small of a value will result in poor performance. The value should be chosen carefully based upon the typical tape channel program that requires optimum performance. The range is 1-100.
   - **FICON Tape Read Max Pipe** defines a maximum number of channel commands that may be outstanding at a given time during read pipelining. Too small of a value will result in poor performance. The value should be chosen carefully based upon the typical tape channel program that requires optimum performance. The range is 1-100.
   - **FICON Tape Write Max Ops** defines a maximum number of concurrent emulated tape write operations. The range is 1-32.
   - **FICON Tape Read Max Ops** defines a maximum number of concurrent emulated tape read operations. The range is 1-32.
   - **FICON Tape Write Timer** defines a time limit for pipelined write chains. This value is be specified in milliseconds (ms). If a pipelined write chain takes longer than this value to complete, the ending status for the next write chain will be withheld from the channel. This limits processing to what the network and device can support. Too small a value limits pipelining performance. Too large a value results in too much data being accepted for one device on a path. The range is 100-1500.
   - **FICON Tape Max Write Chain** defines the maximum amount of data that can be contained in a single CCW chain. If this value is exceeded, emulation is suspended. The range is 1,000,000 to 5,000,000 ms.
   - **FICON Oxid Base** defines the base value of an entry pool of 256 OXIDs supplied to emulation generated exchanges. It should fall outside the range used by FICON channels and devices to avoid conflicts. The range is 0x0000 to 0xF000.

Configuring Load Balance

Load Balancing features for a tunnel are configured from the **Load Balance** tab (Figure 416).

NOTE
You can configure **Spillover** only for the Fabric OS devices running 8.0.1 or later.

1. Select **Advanced Settings** on the **Add FCIP Tunnel or Edit FCIP Tunnel dialog box to display the Advanced Settings dialog box.**
2. Select the **Load Balance** tab.

FIGURE 416 Advanced Settings Load Balance Tab

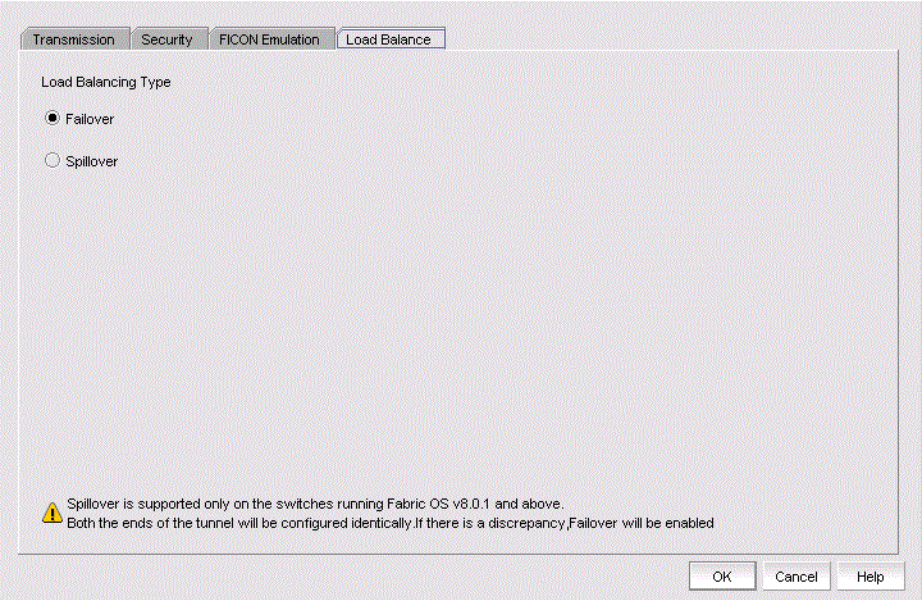

3. Select **Failover** or **Spillover** option in load balancing type. **Failover** is the default value.

Viewing FCIP connection properties

The FCIP connection properties show properties of the blades or switches on both sides of a connection. To view FCIP connection properties, right-click the connection between two Extension Blades or Switches and select **Properties** (Figure 417).

FIGURE 417FCIP connection properties

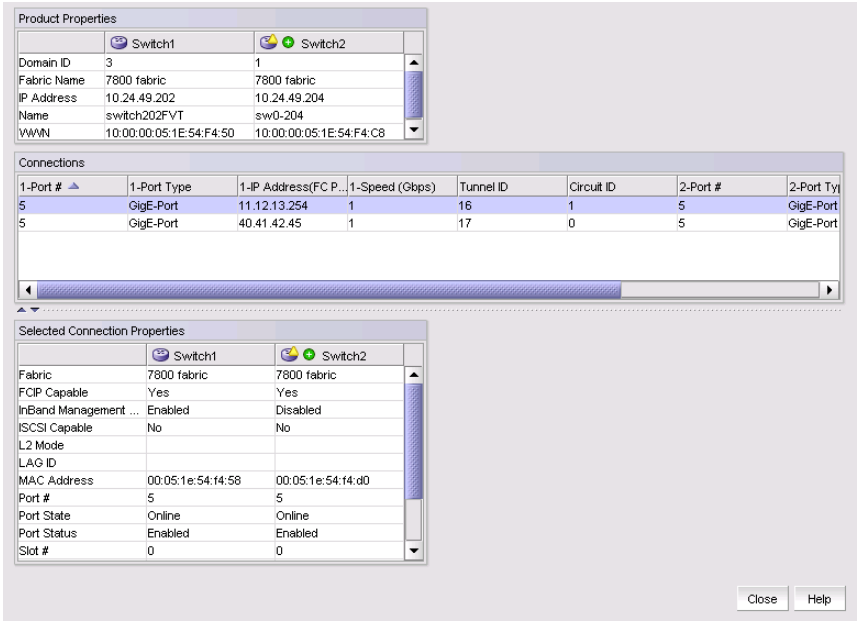

If the default logical switch is not discovered the dialog box for shared GbE links will display VE_Port information instead of GbE port information. Refer to Figure 417.

FIGURE 418FCIP connection properties (default switch not discovered)

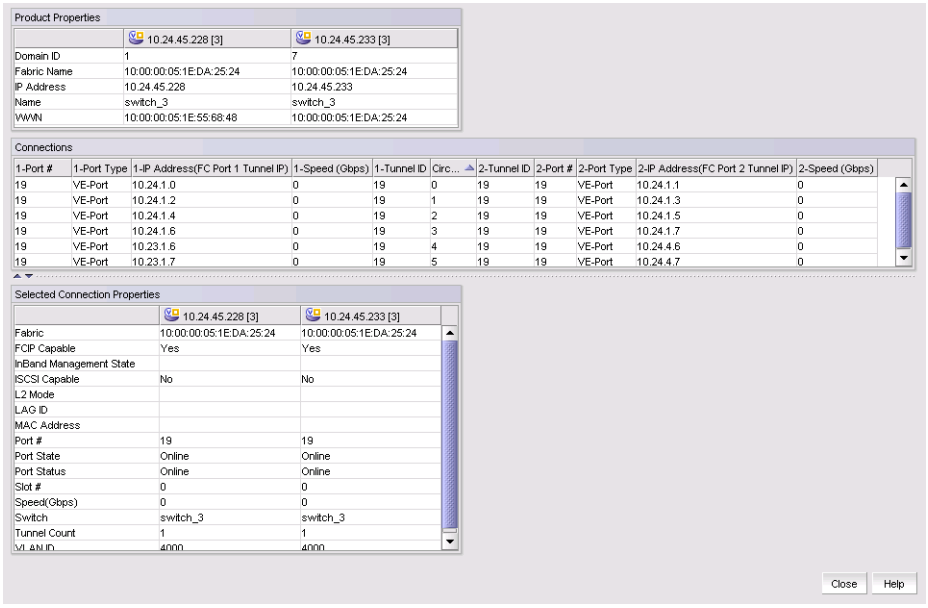

Viewing General FCIP properties

Use the following steps to view general FCIP properties for a switch or blade.

1. Right click an Extension Blade or Switch from the Fabric Tree structure or on the Connectivity Map, and select **Properties**.
2. Select the **Properties** tab.

FIGURE 419 General FCIP properties tab (Extension Switch or Blade)

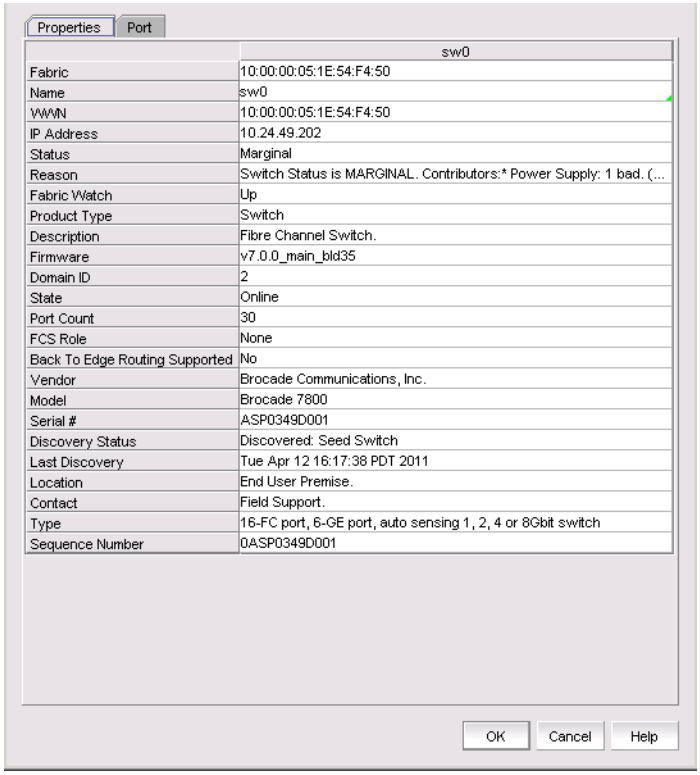

Use the following steps to view the properties of a chassis where an Extension Blade is installed.

1. Right click the chassis in the Switch group in Fabric Tree structure or on the Connectivity Map where the Extension Blade is installed, and select **Properties**.
2. Select the **Properties** tab.

FIGURE 420 General FCIP properties tab (blade chassis)

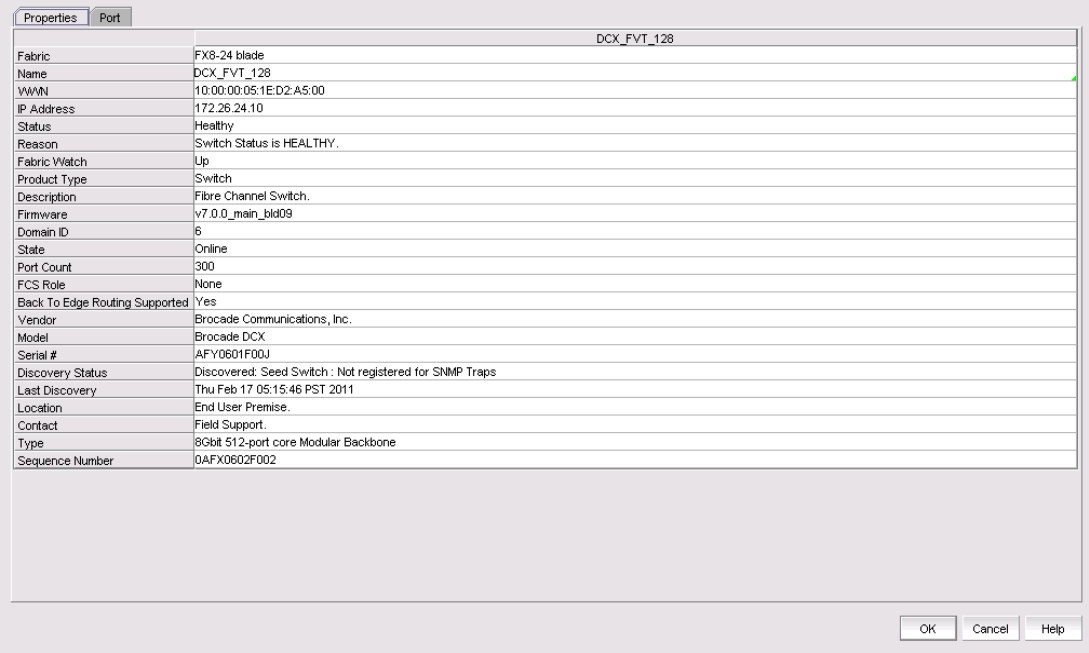

Viewing FCIP port properties

Take the following steps to view FCIP FC. VE/VEX, and GbE port properties.

1. Right click an Extension Blade or Switch from the Fabric Tree structure or on the Connectivity Map, and select **Properties**.
2. Select the **Port** tab.
3. To view FC port information, select the **FC** from the **Type** drop-down list (Figure 421).

FIGURE 421 FC ports properties

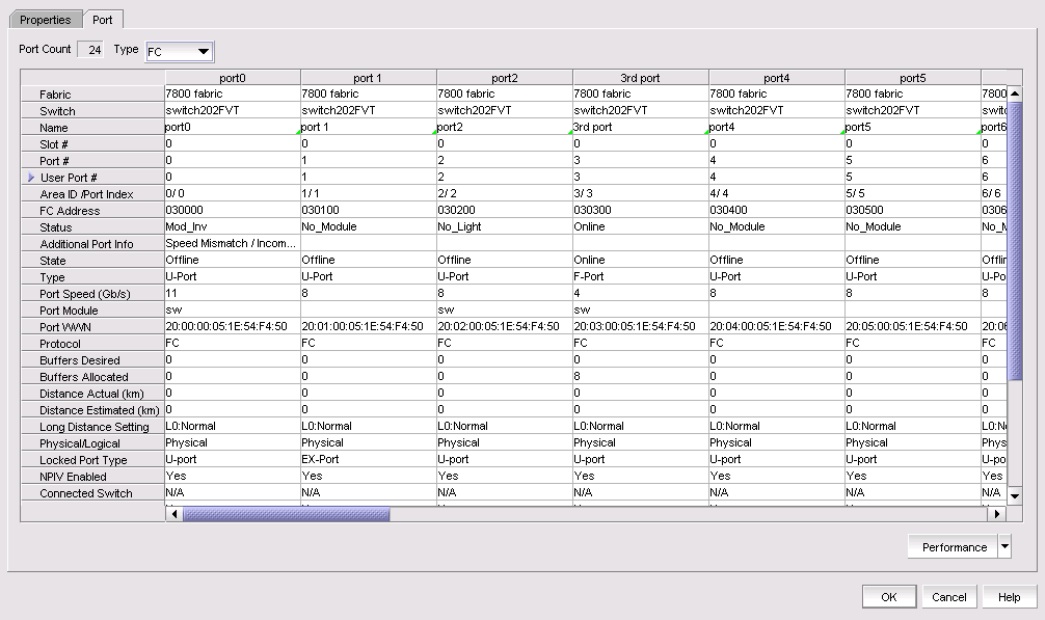

4. To view VE and VEX port information, select the **VE/VEx** from the **Type** drop-down list (Figure 422).

FIGURE 422 VE/VEx port properties

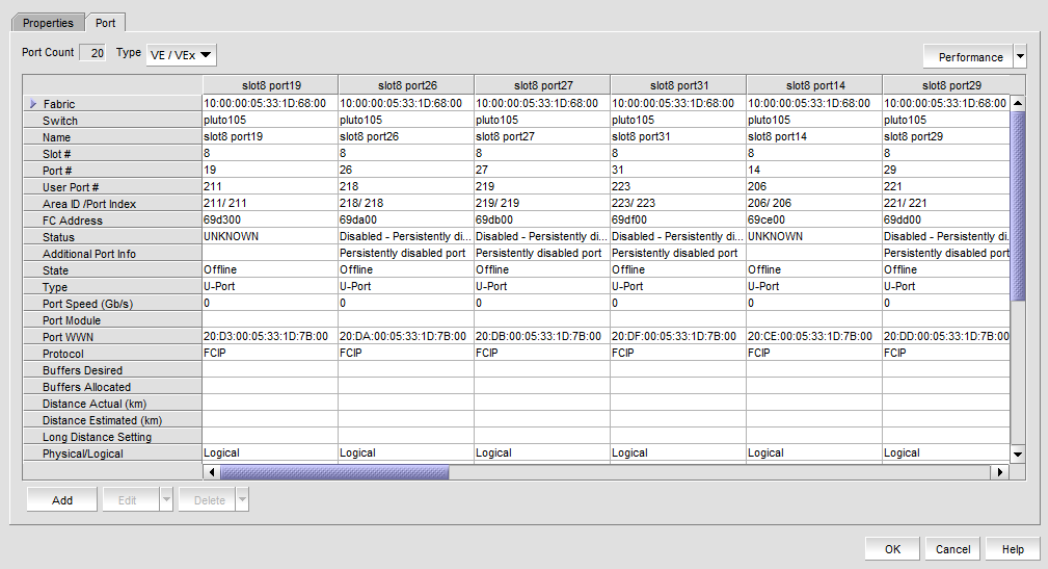

5. To view GbE (Ethernet) port information, select the **GigE** from the **Type** drop-down list (Figure 423).

FIGURE 423 GbE port properties

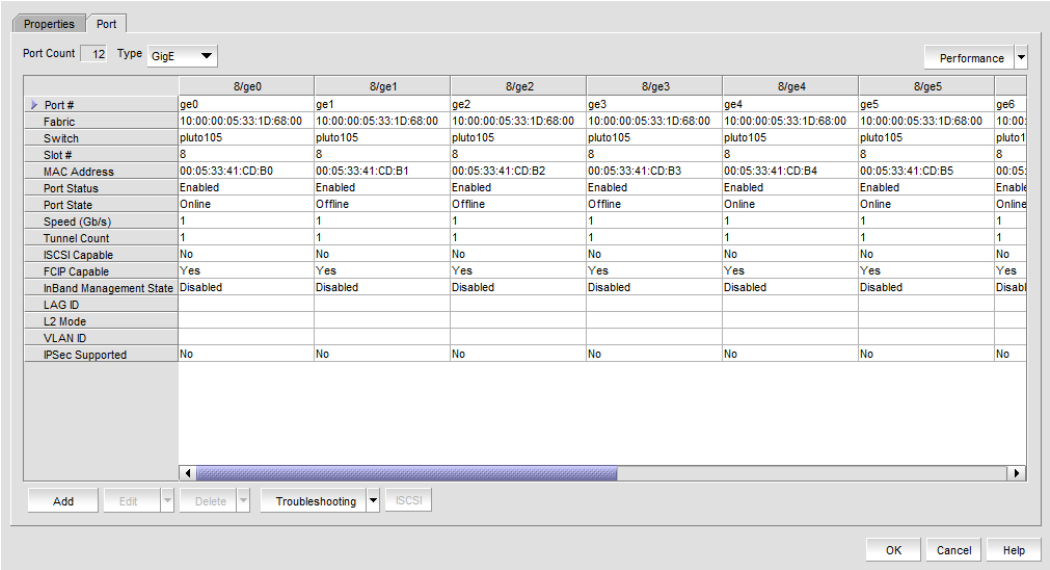

Editing FCIP circuits

FCIP circuit settings may be edited from the **Edit FCIP Circuit** dialog box. The procedure for launching this dialog box for the 4 Gbps Extension Switch and Blade is different than the procedure for the 8 Gbps Extension Switch and the 8 Gbps Extension Blade. Also note the following differences for these platforms:

- The 4 Gbps Extension Switch and Blade have only one circuit per tunnel, and the circuit is edited as part of the tunnel. For 4 Gbps platforms, the **Delete**, **Enable**, and **Disable** buttons do not display. In addition, the **Edit** operation is only supported for cached circuits.
- The 8 Gbps Extension Switch and 8 Gbps Extension Blade may have multiple circuits per tunnel, and circuits may be selected individually.

For the 4 Gbps Extension Switch and Blade:

1. From the **FCIP Tunnels** dialog box, select the tunnel you want to edit.
2. Select **Edit**.

   The **Edit FCIP Tunnel** dialog box displays.
3. Select **Edit** to the right of the **Circuits** properties table at the bottom of the dialog box.

   The **Edit FCIP Circuit** dialog box displays.

For the 8 Gbps Extension Switch and the 8 Gbps Extension Blade:

1. Select **Edit**.

   The **Edit FCIP Tunnel** dialog box displays.
2. Select a circuit that you want to edit from the **Circuits** properties table at the bottom of the dialog box and select **Edit**.

   The **Edit FCIP Circuit** dialog box displays (Figure 424).

FIGURE 424 Edit FCIP Circuit dialog box

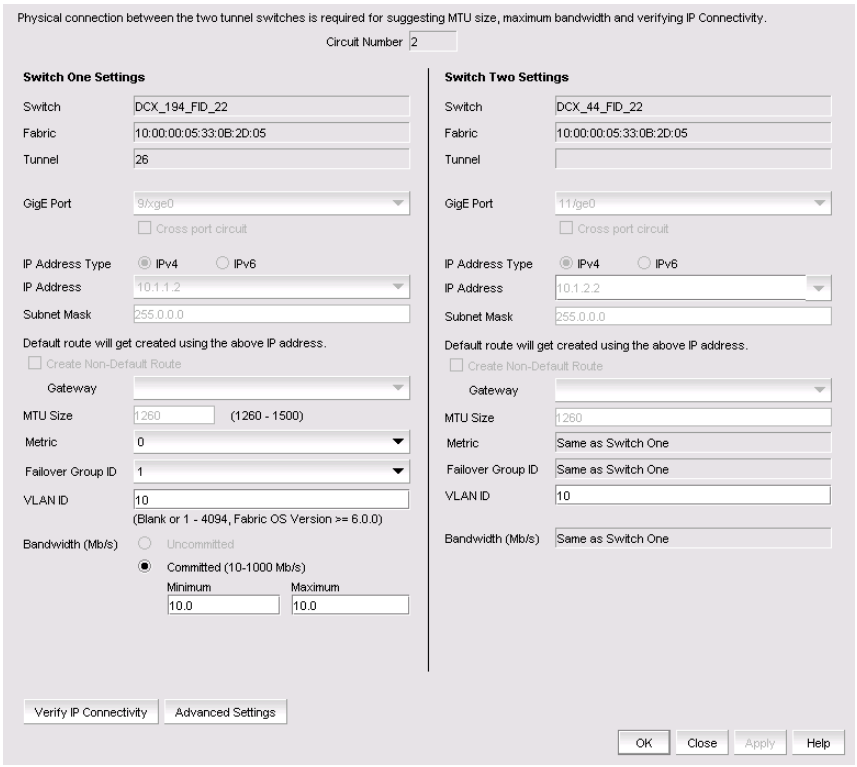

3. Fields and parameters are as described in "Adding an FCIP circuit". You can edit all editable fields and parameters.

Disabling FCIP tunnels

1. From the **FCIP Tunnels** dialog box, select the tunnel you want to disable.
2. Select **Disable**.

   A confirmation dialog box displays showing the switches on both ends of the tunnel and tunnel number.
3. Click **Yes** to disable the tunnel.

Enabling FCIP tunnels

1. From the **FCIP Tunnels** dialog box, select the tunnel you want to enable.
2. Select **Enable**.
3. Click **OK** to enable the tunnel.

Deleting FCIP tunnels

1. From the **FCIP Tunnels** dialog box, select the tunnel you want to delete.
2. Select the **Delete**.

A confirmation dialog box displays, warning you of the consequences of deleting a tunnel.

3. Click **OK** to delete the tunnel.

Disabling FCIP circuits

1. From the **FCIP Tunnels** dialog box, select the tunnel that contains the circuit.
2. Select **Edit**.

   The **Edit FCIP Tunnel** dialog box displays.
3. Select **the circuit that you want to disable from the Circuit properties table at the bottom of the dialog box**.
4. Select **Disable**.
5. For tunnels with multiple circuits, select additional circuits from the table to disable and select **Disable** after each selection.
6. Click **OK** to disable the circuit(s).

Enabling FCIP circuits

1. From the **FCIP Tunnels** dialog box, select the tunnel that contains the circuit.
2. Select **Edit**.

   The **Edit FCIP Tunnel** dialog box displays.
3. Select **the circuit that you want to disable from the Circuit properties table at the bottom of the dialog box**.
4. Select **Enable**.
5. For tunnels with multiple circuits, select additional circuits from the table to enable and select **Enable** after each selection.
6. Click **OK** to enable the circuit(s).

Deleting FCIP Circuits

1. From the **FCIP Tunnels** dialog box, select the tunnel that contains the circuit.
2. Select **Edit**.

   The **Edit FCIP Tunnel** dialog box displays.
3. Select **the circuit that you want to delete from the Circuit properties table at the bottom of the dialog box**.
4. Select **Delete**.
5. For tunnels with multiple circuits, select additional circuits from the table to delete and select **Delete** after each selection.
6. Click **OK** to delete the circuit(s).

Displaying FCIP performance graphs

You can display performance graphs by clicking the **Performance** button on the FCIP Tunnels dialog box. You can also display performance graphs from Properties, as described in the following sections.

Displaying performance graphs for FC ports

1. Right-click a blade an Extension Blade or Switch from the Fabric Tree structure or Connectivity Map, and select **Properties**.
2. Select the **Port** tab.
3. Select **FC** from the **Type** drop-down list.
4. Click **Performance > Real Time Graph**.

Displaying FCIP performance graphs for Ethernet ports

1. Right-click a blade an Extension Blade or Switch from the Fabric Tree structure or Connectivity Map, and select **Properties**.
2. Select the **Port** tab.
3. Select **GigE** from the **Type** drop-down list.
4. Click **Performance > Real Time Graph**.

Displaying tunnel properties from the FCIP tunnels dialog box

Tunnel properties can be displayed from the **FCIP Tunnels** dialog box.

1. Select a tunnel from the **FCIP tunnels** dialog box.
2. Select the **Tunnel** tab.

   Tunnel properties are displayed.

FIGURE 425 Tunnel properties on the FCIP Tunnels dialog box

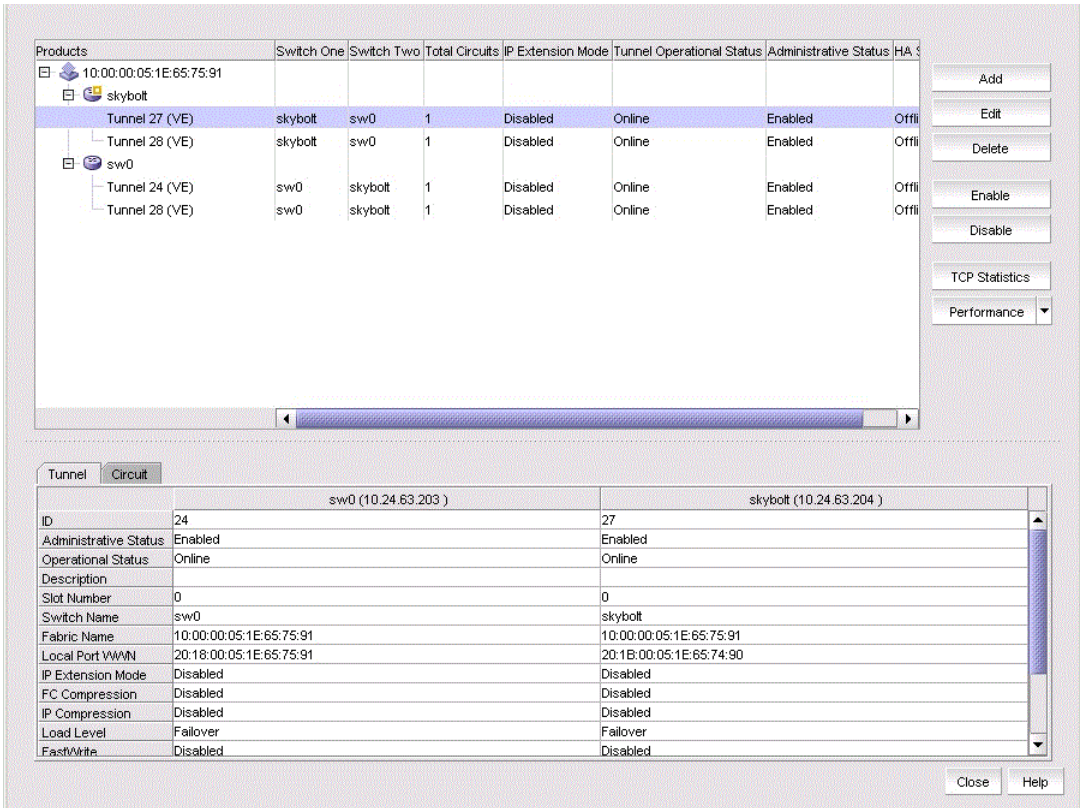

Displaying FCIP circuit properties from the FCIP tunnels dialog box

Tunnel properties can be displayed from the **FCIP Tunnels** dialog box using the following procedure.

1. Select a tunnel from the **FCIP tunnels** dialog box.
2. Select the **Circuit** tab.

   Circuit properties are displayed (Figure 426).

FIGURE 426 Circuit properties on the FCIP Tunnels dialog box

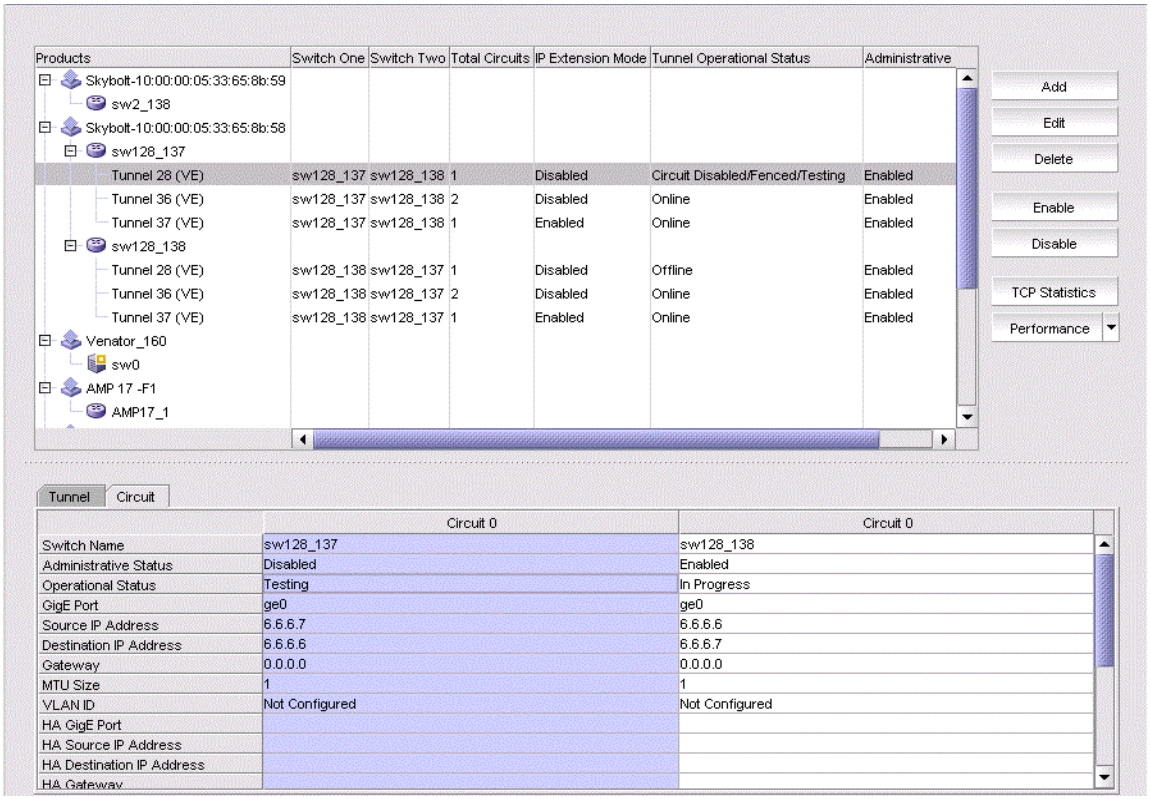

Displaying switch properties from the FCIP Tunnels dialog box

Switch properties are displayed on the **FCIP Tunnels** dialog box when you select a switch (Figure 427).

FIGURE 427 Switch properties on the FCIP Tunnels dialog box

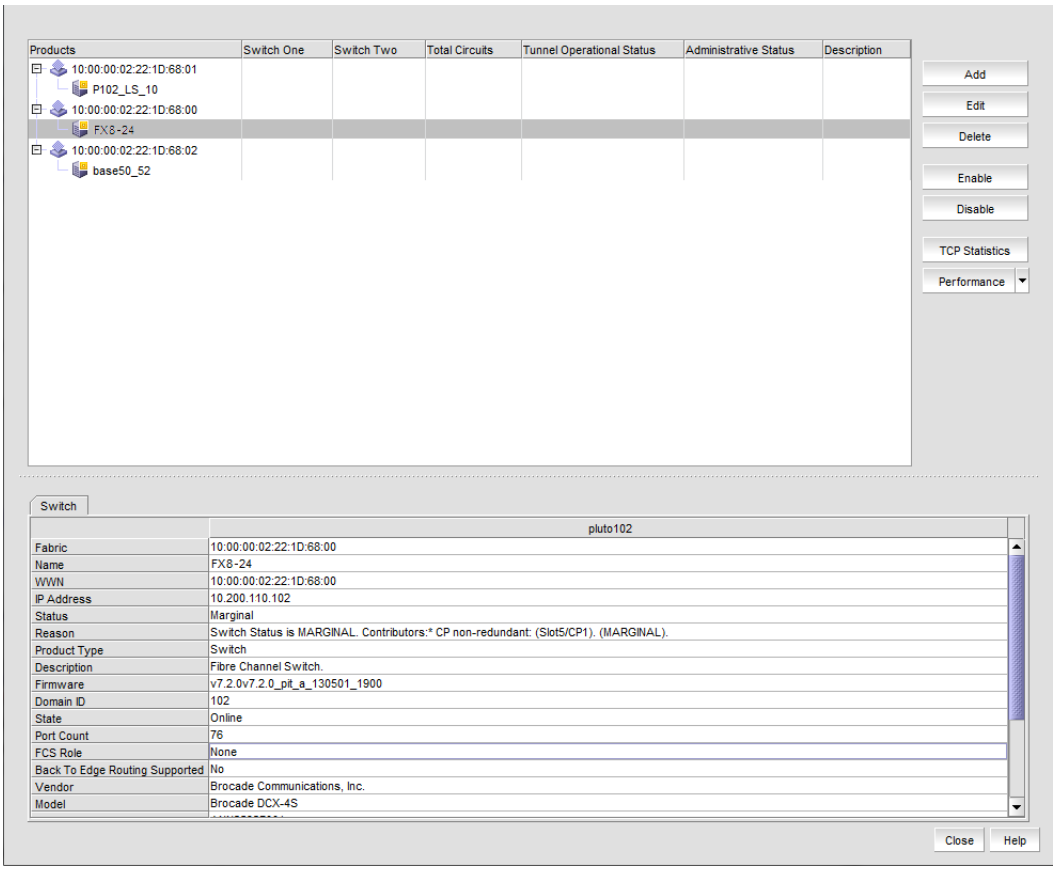

Displaying fabric properties from the FCIP Tunnels dialog box

Fabric properties are displayed on the **FCIP Tunnels** dialog box when you select a fabric. (Figure 428).

FIGURE 428 Fabric properties on the FCIP Tunnels dialog box

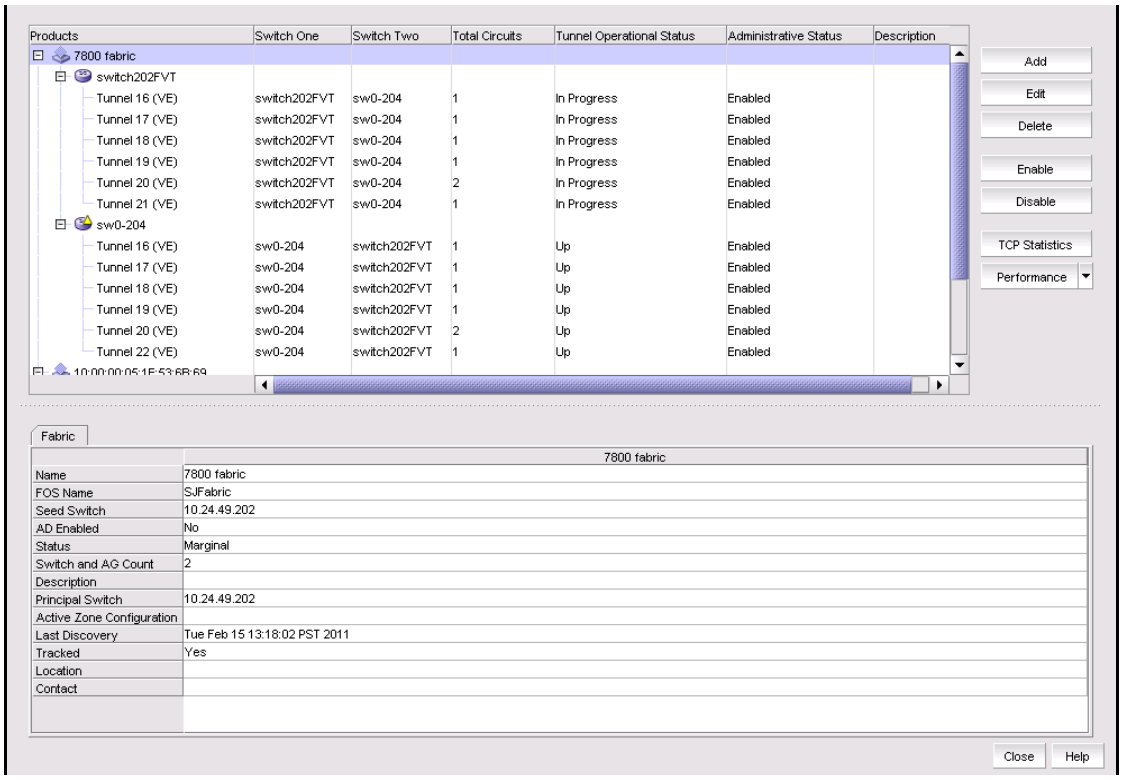

Troubleshooting FCIP Ethernet connections

1. Right-click a blade an Extension Blade or Switch from the Fabric Tree structure or Connectivity Map, and select **Properties**.
2. Select the **Port** tab.
3. Select **GigE** from the **Type** drop-down list.
4. Select an Ethernet port.
5. Click **Troubleshooting**.

   The following options are presented:

   - **IP Ping** — Tests connections between a local Ethernet port (ge0 or ge1) and a destination IP address.
   - **IP Traceroute** — Traces routes from a local Ethernet port (ge0 or ge1) to a destination IP address.

Fabric Binding

Fabric Binding overview

The fabric binding feature enables you to configure whether switches can merge with a selected fabric. This provides security from accidental fabric merges and potential fabric disruption when fabrics become segmented because they cannot merge.

Enabling Fabric Binding activates Switch Connection Control (SCC) policy and sets Fabric Wide Consistency Policy (FWCP) and insistent domain ID. Disabling Fabric Binding on Fabric OS devices deletes SCC policy and sets FWCP to absent.

NOTE
In a pure Fabric OS fabric, enabling insistent domain ID is mandatory.

Viewing fabric binding membership

1. Select **Configure > Fabric Binding**.

   The **Fabric Binding** dialog box displays (Figure 429).

FIGURE 429 Fabric Binding dialog box

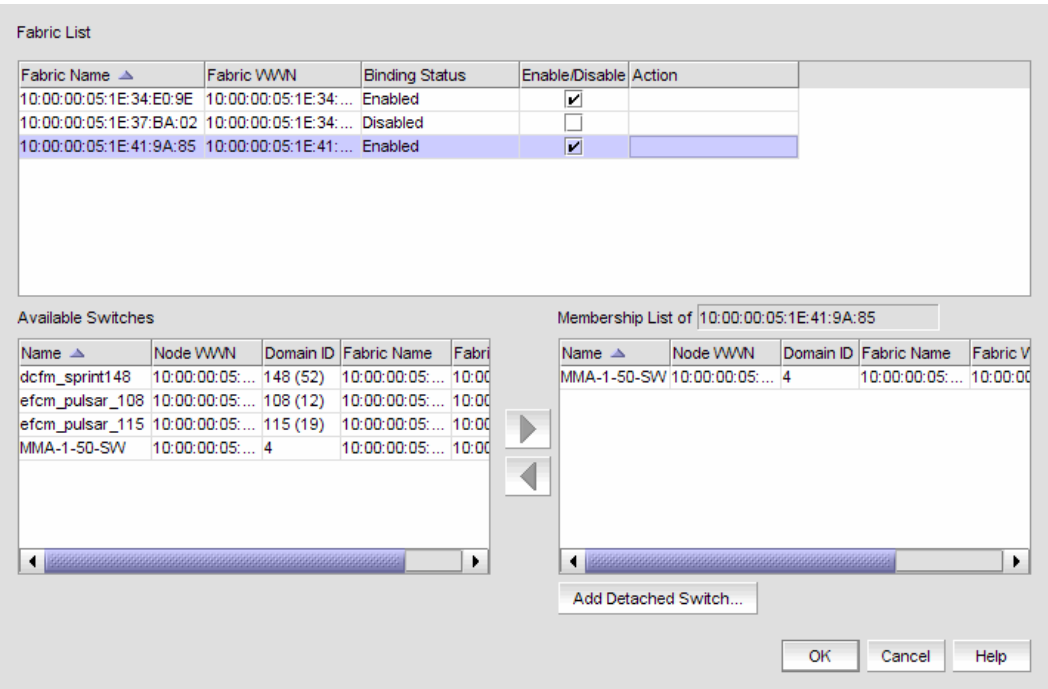

2. Review the fabric binding membership details.

- **Fabric List** table — Lists the fabrics in your network.
  - **Fabric Name** — The name of the fabric.
  - **Fabric WWN** — The world wide name of the fabric.
  - **Binding Status** — The binding status (enabled/disabled) of the fabric.
  - **Enable/Disable** check box — Indicates whether fabric binding is enabled. Select to enable a fabric binding for the fabric. For step-by-step instructions, refer to "Enabling fabric binding" on page 885 and "Disabling fabric binding" on page 885.
  - **Action** — Displays any actions on the fabric.
- **Available Switches** table — Lists the switches available to add to the fabric binding membership list. For step-by-step instructions, refer to "Adding switches to the fabric binding membership list" on page 886.
  - **Name** — The name of the switch fabric.
  - **Node WWN** — The node WWN of an available or member switch.
  - **Domain ID** — The domain ID of an available or member switch.

    NOTE
    You can copy (Ctrl+C) and paste (Ctrl+V) the node WWN into the **Node WWN** field. It does not matter if the copy source contains colons (11:22:33:44:55:66:77), only the numbers are pasted (11223344556677) in the **Node WWN** field.

  - **Fabric Name** — The name of the fabric.
  - **Fabric WWN** — The world wide name of the fabric.
- **Membership List of *Fabric_Name*** table — The current Fabric Membership List (FML) of the highlighted fabric, including the following details:
  - **Name** — The name of the switch fabric.
  - **Node WWN** — The node WWN of an available or member switch.
  - **Domain ID —** The domain ID of an available or member switch.
  - **Fabric Name** — The name of the fabric.
  - **Fabric WWN** — The world wide name of the fabric.
  - **Attached** — Whether or not the switch is attached.

  If you have never configured the FML, a default list with all the member switches displays. To remove a switch from the membership list, refer to "Removing switches from fabric binding membership" on page 886.
- **Add Detached Switch** button — Click to enter the domain ID and WWN of the detached switch. For step-by-step instructions, refer to "Adding detached devices to the fabric binding membership list" on page 886.

  Domain IDs are between the values of 1 to 239. In HEX mode, Domain IDs are between the values of 01 to EF.

  Fibre Channel networks use world wide names to uniquely identify nodes and ports within nodes. For many devices, the 64-bit WWNs are fixed, and their assignment follows conventions established by the IEEE. For other devices, the WWNs may be set or modified by the user. World wide names are a special concern for the Management application because:

  - WWNs are used as the primary keys to identify network elements.
  - Experience has been that an ill-formed WWN is evidence of a malfunctioning device.

  Proper operation with the management application requires that WWNs be unique within the network and well-formed. This means they must be 64 bits in length and the first byte cannot be zero.

3. Click **OK** on the **Fabric Binding** dialog box.

Enabling fabric binding

Fabric binding is a security method for restricting switches within a multiple-switch fabric. Fabric Binding is required for FICON in mixed fabrics.

Fabric Binding is enabled through the **Fabric Binding** dialog box. After you have enabled Fabric Binding, use the **Fabric Membership List/Add Detached Switch** to add switches that you want to allow into the fabric.

NOTE
If FMS mode with HIF is enabled or disabled on switches running Fabric OS 7.3.0, an error message is displayed.

1. Select **Configure > Fabric Binding**.

   The **Fabric Binding** dialog box displays (Figure 429).

2. In the **Fabric List** table, select the **Enable/Disable** check box for fabrics for which you want to configure fabric binding.

   For instructions on adding and removing switches from the membership list, refer to "Adding switches to the fabric binding membership list" on page 886 and "Removing switches from fabric binding membership" on page 886.

3. Click **OK** on the **Fabric Binding** dialog box.

   The **Fabric Binding Status** dialog box displays with the following information:

   - **Fabric Name** — Displays the enabled fabric name selected for fabric binding.
   - **Applying Fabric Binding Changes for selected fabric** message — Displays the status of the fabric binding changes.
   - **Setting SCC Policy** message — The Switch Connection Control (SCC) policy prevents unauthorized devices from joining a fabric.
   - **Setting FWCP Policy** message — Fabric-wide consistency is necessary for FICON switch binding.
   - **List of possible reasons that could cause Fabric Binding failure** message — Refer to the *Fabric OS Administrator's Guide* for detailed information.

Disabling fabric binding

Fabric binding cannot be disabled while High Integrity Fabric (HIF) is active if the switch is offline. This disables fabric binding and High Integrity Fabric on the switch, but not the rest of the fabric. Disabled switches segment from the fabric.

1. Select **Configure > Fabric Binding**.

   The **Fabric Binding** dialog box displays (Figure 429).

2. In the **Fabric List** table, clear the **Enable/Disable** check box for fabrics for which you want to disable fabric binding.

   NOTE
   If FMS mode with HIF is enabled or disabled on switches running Fabric OS v7.3.0, an error message is displayed.

3. Click **OK** on the **Fabric Binding** dialog box.

   The **Fabric Binding Status** dialog box displays with the following information:

   - **Fabric Name** — Displays the enabled fabric name selected for fabric binding.
   - **Applying Fabric Binding Changes for selected fabric** message — Displays the status of the fabric binding changes.
   - **Setting SCC Policy** message — The Switch Connection Control (SCC) policy prevents unauthorized devices from joining a fabric.

- **Setting FWCP Policy** message — Fabric-wide consistency is necessary for FICON switch binding.
- **List of possible reasons that could cause Fabric Binding failure** message — Refer to the *Fabric OS Administrator's Guide* for detailed information.

Adding switches to the fabric binding membership list

Once you have enabled fabric binding (refer to "Enabling fabric binding" on page 885), you can add switches to the fabric binding membership list.

To add a switch to the fabric, complete the following steps.

1. Select **Configure > Fabric Binding**.

   The **Fabric Binding** dialog box displays (Figure 429).

2. Select the switches you want to add to the selected fabrics' Fabric Membership List (FML) in the **Available Switches** table.
3. Click the right arrow to move the switches to the **Membership List** table.
4. Click **OK** on the **Fabric Binding** dialog box.

Adding detached devices to the fabric binding membership list

To add a switch that does not have a physical connection and is not discovered to the fabric, complete the following steps.

1. Select **Configure > Fabric Binding**.

   The **Fabric Binding** dialog box displays (Figure 429).

2. Click **Add Detached Switch**.

   The **Add Detached Switch** dialog box displays.

3. Enter the domain ID of the switch in the **Domain ID** field.
4. Enter the node World Wide Name (WWN) of the switch in the **Node WWN** field.

   You can copy (Ctrl+C) and paste (Ctrl+V) the node WWN in the **Node WWN** field. It does not matter if the copy source contains colons (11:22:33:44:55:66:77); only the numbers are pasted (11223344556677) in the **Node WWN** field.

5. Click **OK** on the **Add Detached Switch** dialog box.

   The added switch displays in the **Membership List of** *Fabric_Name* table on the **Fabric Binding** dialog box.

6. Click **OK** on the **Fabric Binding** dialog box.

Removing switches from fabric binding membership

Once you have enabled fabric binding (refer to "Enabling fabric binding" on page 885), you can remove switches that are not part of the fabric from the membership list.

1. Select **Configure > Fabric Binding**.

   The **Fabric Binding** dialog box displays (Figure 429).

2. Select the switches you want to remove from the selected fabrics' Fabric Membership List (FML) in the **Membership List** table.

NOTE
The selected switch cannot be part of the fabric.

3. Click the left arrow to move the switches to the **Available Switches** table.
4. Click **OK** on the **Fabric Binding** dialog box.

High Integrity Fabric overview

The High Integrity Fabric (HIF) mode option automatically enables features and operating parameters that are necessary in multiswitch Enterprise Fabric environments. When HIF is enabled, each switch in the fabric automatically enforces a number of security-related features including Fabric Binding, Switch Binding, Insistent Domain IDs, and Domain Register for State Change Notifications (RSCNs).

HIF activates the Switch Connection Control (SCC) policy, sets the Insistent Domain ID, and sets the Fabric-Wide Consistency Policy (FWCP) for SCC in strict mode.

Activating HIF mode enables the following features:

- **Switch Connection Control —** This feature, enabled through a device's Element Manager, prevents unauthorized switches from joining a fabric.
- **Fabric-Wide Consistency Policy —** This feature makes sure that switches in the fabric enforce the same policies.
- **Insistent Domain ID —** This feature, enabled through a device's Element Manager, sets the domain ID as the active domain identification when the fabric initializes. When Insistent Domain ID is enabled, the switch isolates itself from the fabric if the preferred domain ID is not assigned as the switch's domain ID.

High Integrity Fabric requirements

High Integrity Fabric (HIF) refers to a set of strict, consistent, fabric-wide policies. There are several specific configuration requirements for high integrity fabrics:

- Insistent Domain ID (IDID) must be enabled in the participating switches.
- Port-based routing must be used on the participating switches.
- A policy must be set that limits connectivity to only the switches within the same fabric. Fabric binding is a security method for restricting switches that may join a fabric. For Fabric OS switches, fabric binding is implemented by defining a Switch Connection Control (SCC) policy that prevents unauthorized switches from joining a fabric.
- Dynamic Load Sharing (DLS) should be disabled. If DLS is not disabled, DLS automatically adjusts routes when a new ISL is added, and when an ISL is taken offline and brought online again. This process may result in dropped frames.

  NOTE
  Port binding is a security method for restricting devices that connect to particular switch ports. Port binding should never be used in FICON environments. The FICON channel cannot be added to the port binding list.

Activating high integrity fabrics

To activate a HIF, complete the following steps.

1. Select **Configure > High Integrity Fabric**.

   The **High Integrity Fabric** dialog box displays (Figure 430).

FIGURE 430 High Integrity Fabric dialog box

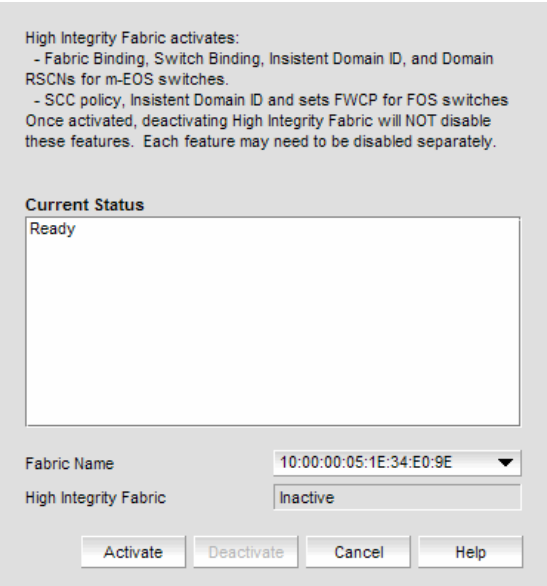

2. Select the fabric on which you want to activate HIF from the **Fabric Name** list.

   The HIF status displays in the **High Integrity Fabric** field.

3. Click **Activate**.

   HIF activates the Switch Connection Control (SCC) policy, sets Insistent Domain ID, and sets the Fabric-Wide Consistency Policy (FWCP) for SCC in strict mode.

Deactivating high integrity fabrics

NOTE
Deactivating high integrity fabrics is not supported in a pure Fabric OS environment.

NOTE
A warning message is displayed when you disable HIF mode.

To deactivate a HIF, complete the following steps.

1. Select **Configure > High Integrity Fabric**.

   The **High Integrity Fabric** dialog box displays (Figure 430).

2. Select the fabric on which you want to deactivate HIF from the **Fabric Name** list.

   The HIF status displays in the **High Integrity Fabric** field.

3. Click **Deactivate**.

   Deactivating HIF on a fabric does not deactivate the features on the individual switches, you must disable the SCC policy, Insistent Domain ID, and the Fabric Wide Consistency Policy for SCC in tolerant mode individually:

Port Fencing

About port fencing

NOTE
This feature is only available for Fabric OS devices running 7.3.X and earlier. It not supported on devices running 7.4.0 or later.

NOTE
All Fabric OS devices must have Fabric Watch and must be running firmware Fabric OS 7.0 or later.

NOTE
This feature requires a Trial or Licensed version.

Port Fencing allows you to protect your SAN from repeated operational problems experienced by ports. Use Port Fencing to set threshold limits for the number of specific port events permitted during a given time period on the selected object. For default threshold values for Fabric OS devices, refer to Chapter 7 of the ***Fabric Watch Administrator's Guide***.

Port Fencing objects include the SAN, Fabrics, Directors, Switches (physical), Virtual Switches, Ports, as well as Port Types (E_port, F_port, and FX_port). Use Port Fencing to directly assign a threshold to these objects. When a switch does not support Port Fencing, a "No Fencing Changes" message displays in the **Threshold** field in the **Ports** table.

If the port detects more events during the specified time period, the device firmware blocks the port, disabling transmit and receive traffic until you investigate, solve the problem, and manually unblock the port.

Physical fabrics, directors, switches, port types, and ports display when you have the privileges to manage that object and are indicated by the standard product icons.

NOTE
Port Fencing displays any existing thresholds discovered on manageable fabrics, directors, and switches running firmware version Fabric OS 7.0 or later.

Viewing port fencing configurations

NOTE
This feature is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

Port Fencing allows you to protect your SAN from repeated operational problems experienced by ports. Use Port Fencing to set threshold limits for the number of specific port events permitted during a given time period on the selected object. For default threshold values for Fabric OS devices, refer to Chapter 7 of the *Fabric Watch Administrator's Guide*.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays (Figure 432).

FIGURE 431 Port Fencing dialog box

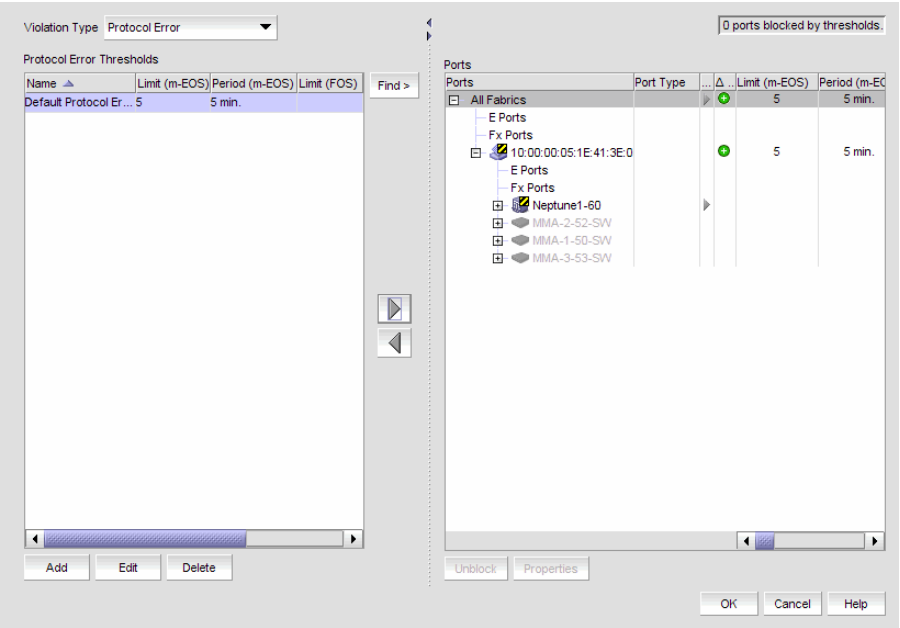

The Port Fencing dialog box contains the following field and components:

- **Violation Type list** — The name of the ISL, Link, or Security threshold currently active on this port. If the object does not support Port Fencing, this field displays "# No Fencing Support #". This field displays as inactive (grayed-out) when the object is only partially managed by the management application.
  (ISL Threshold only) If the port type is E_port, the ISL Threshold name displays in a bold font to indicate when the threshold is currently active on the port type.
- Thresholds table — List of configured thresholds based on the threshold type selected in the **Violation Type list.**
  - **Limit (Fabric OS) —** The number of events allowed for the assigned threshold.
    If the object has no fencing support or no fencing changes, this field displays two hyphens separated by a space (- -).
    When the object is only partially managed by the management application, this field displays as inactive (grayed-out).
  - **Period (Fabric OS) —** The time limit (in seconds or minutes) for the assigned threshold. This field displays as inactive (grayed-out) when the object is only partially managed by the management application.

  - **Ports Affected** — The total number of ports on all objects that could be affected by the threshold setting. It does not show the current number of ports affected. This value updates in real time as you add and subtract each threshold from each object.
- Find button — Select a threshold in the thresholds table and click the find button (>) to highlight each instance of the selected threshold in the **Ports** table.
- Right arrow button — Select a threshold in the thresholds table and click the right arrow button to add the selected threshold to the selected fabrics, switches, or switch ports (refer to "Assigning thresholds" on page 901).
- Left arrow button — Select a threshold in the Ports table and click the left arrow button to remove the selected threshold from the associated fabrics, switches, or switch ports (refer to "Removing thresholds from individual objects" on page 908).
- **Add** button — Click to add an ISL protocol threshold (refer to "Adding thresholds" on page 894).
- **Edit** button — Click to edit an ISL protocol threshold (refer to "Editing thresholds" on page 903).
- **Delete** button — Click to delete an ISL protocol threshold (refer to "Removing thresholds from the thresholds table" on page 908).
- **Ports** table — All managed fabric, director, switch, port type, and port objects (label and icon) in its hierarchical relationship to the other objects in the tree.
  - **Ports** — Displays all discovered fabrics, devices, and ports as both text and icons.
  - **Port Type** — The operational port type of the port. This field displays as inactive (grayed-out) when either the object's firmware does not support Port Fencing or the object is only partially managed by the management application.
  - **Directly Assigned** — A right arrow icon to indicate that the threshold is directly assigned to this object and is inherited by all objects below it in the tree. This field displays as inactive (grayed-out) when either the object's firmware does not support Port Fencing or the object is only partially managed by the management application.
  - **Changed** indicator — The change icons in real time when you change information in the dialog box. One change icon indicates a new threshold was applied (either directly or inherited) to the port, and another indicates that a threshold was removed from this object (during this session) and no threshold applies to the port.
  - *Threshold_type* **Threshold** — The name of the ISL threshold policy.
  - **Limit (Fabric OS) —** The number of events allowed for the assigned threshold.
    If the object has no fencing support or no fencing changes, this field displays two hyphens separated by a space (- -). When the object is only partially managed by the management application, this field displays as inactive (grayed-out).
  - **Period (Fabric OS) —** The time limit (in seconds or minutes) for the assigned threshold. This field displays as inactive (grayed-out) when the object is only partially managed by the management application.
  - **Operational State** — The operational state of the port.
  - **Blocked Configuration** — The current configuration of the port (Blocked or Unblocked).
  - **Port WWN** — The port world wide name of the port.
  - **Connected Product** — The device label of the connected object.
  - **Connected Port # —** The port number of the connected port.
  - **Connected Port WWN** — The port world wide name of the connected port.
  - **Connected Port Name** — The name of the connected port configured in the Element Manager.
  - **FC Address** — The FC address of the port.
- **Properties** button — Click to display the **Properties** dialog box for the fabric, switch, or port selected in the **Ports** table. The All Fabrics and Port Type objects do not have properties. For more information, refer to "Viewing SAN device properties" on page 1354.
- **Unblock** button — Click to unblock a blocked port after a warning message displays (refer to "Unblocking a port" on page 901). This button becomes active after you select a blocked port in the **Ports** table.

2. Click **OK** on the **Port Fencing** dialog box.

Thresholds

You can create thresholds, which you can then assign to available objects in the tree. Port Fencing threshold types include the following:

- C3 Discard Frames (Fabric OS only)
- Invalid CRCs (Fabric OS only)
- Invalid Words (Fabric OS only)
- Link Reset (Fabric OS only)
- Protocol Errors (Fabric OS)
- State Change (Fabric OS only)

NOTE
Fabric OS devices are allowed only 2 defined thresholds (one default and one custom) for each threshold type and only one of these thresholds can be active on the device.

During the dynamic operation of a Fabric, any port could be any type. For example, a technician could disconnect a port from a switch and reconnect that port to a storage port, or the port could change from an E_port to an F_port. Therefore, when calculating the **Affected Ports** value the Management application does not look for the current port type, but looks at the policy priority level in relation to the other policies currently assigned to this switch.

When there are two or more policies on a switch, the total number of **Affected Ports** may be more than the total number of ports on the switch (the same port may adopt different policies depending on changes in the port's port type).

For default threshold values for Fabric OS devices, refer to Chapter 7 of the *Fabric Watch Administrator's Guide*.

C3 Discard Frames threshold

NOTE
This threshold is only available for Fabric OS devices running 7.0 or later.

Use this type of threshold to block a port when a C3 Discard Frames violation meets the Fabric OS switch threshold. This threshold is only supported on Fabric OS directors, switches, and blades with a 4 Gbps, 8 Gbps, or 16 Gbps ASIC.

- 40-port, 8 Gbps FC Switch
- 8 Gbps 8-FC port, 10 GbE 24-DCB port Switch
- 8 Gbps 24-port edge switch
- 8 Gbps 40-port edge switch
- 8 Gbps 80-port edge switch
- 16 Gbps 24-port Edge switch
- 16 Gbps 48-port Edge switch
- 16 Gbps 96-port Edge switch
- Director Chassis
- 8-slot Backbone Chassis
- 4-slot Backbone Chassis
- 16 Gbps 8-slot Backbone Chassis
- 16 Gbps 4-slot Backbone Chassis

- 32 Gbps 8-slot Backbone Chassis
- 32 Gbps 4-slot Backbone Chassis
- 8 Gbps Encryption Switch
- Encryption Blade
- 10 Gbps FCoE Port Router Blade
- 8 Gbps Extension Blade
- FC 8 Gbps 16-port Blade
- FC 8 Gbps 32-port Blade
- FC 8 Gbps 48-port Blade
- FC 8 Gbps 64-port Blade
- FC 16 Gbps 32-port Blade
- FC 16 Gbps 48-port Blade
- 16 Gbps 64-port Blade
- 32 Gbps 48-port Blade
- 32 Gbps Router Extension blade

Invalid CRCs threshold

NOTE
This threshold is only available for Fabric OS devices.

Use this type of threshold to block a port when an Invalid CRCs violation meets the Fabric OS switch threshold.

Invalid words threshold

NOTE
This threshold is only available for Fabric OS devices.

Use this type of threshold to block a port when an Invalid Words violation meets the Fabric OS switch threshold.

Link Reset threshold

NOTE
This threshold is only available for Fabric OS devices.

Use this type of threshold to block a port when the link timeout errors meet the threshold.

Protocol error threshold

Use Protocol Error thresholds to block a port when one of the following protocol errors meet the threshold:

- ISL Bouncing–ISL has repeatedly become unavailable due to link down events.
- ISL Protocol Mismatch–ISL has been repeatedly put into the Invalid Attachment state due to a protocol error.

State Change threshold

NOTE
This threshold is only available for Fabric OS devices running 7.0 or later.

Use this type of threshold to block a port when a state change violation type meets the Fabric OS switch threshold.

For 4 Gbps Router, Extension Switches and Blades only, when you apply this threshold on an E Port, the threshold is also applied to the VE Ports (internally by Fabric OS).

Adding thresholds

NOTE
This feature requires a Trial or Licensed version.

The Management application allows you to create Invalid CRCs, Invalid words, Link, Link Reset, Protocol Error, Security, and Sync Loss thresholds.

Adding a C3 Discard Frames threshold

NOTE
This threshold is only available for Fabric OS devices running 7.0 or later.

Use to block a port when a **C3 Discard Frames** violation type meets the Fabric OS switch threshold. For default threshold values for Fabric OS devices, refer to Chapter 7 of the ***Fabric Watch Administrator's Guide***.

To add an C3 Discard Frames threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays (Figure 432).

FIGURE 432 Port Fencing dialog box

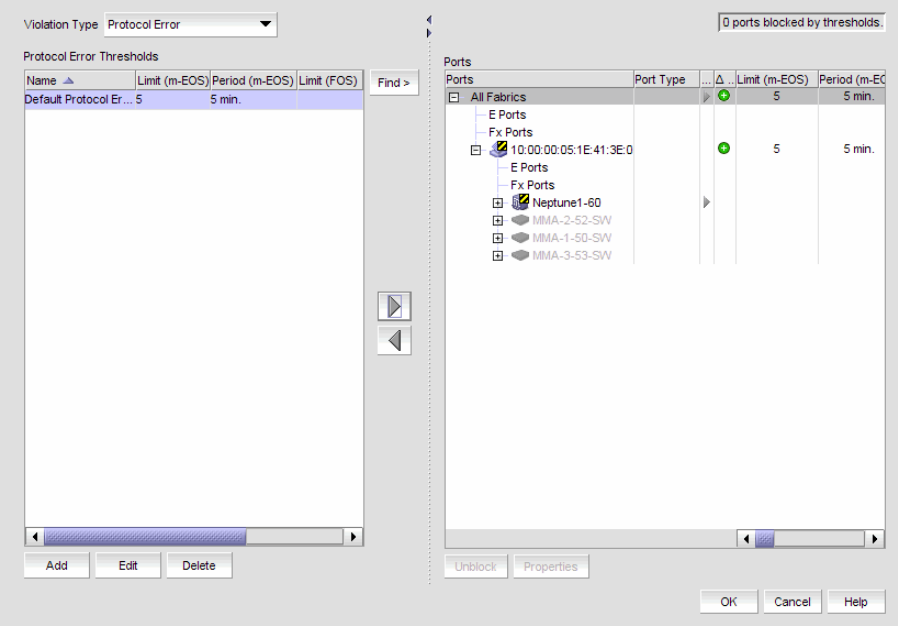

2. Select **C3 Discard Frames (Fabric OS only)** from the **Violation Type** list.
3. Click **Add**.

   The **Add C3 Discard Frames Threshold** dialog box displays.

FIGURE 433 Add C3 Discard Frames Threshold dialog box

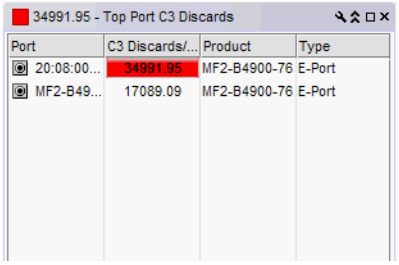

4. Enter a name for the threshold in the **Name** field.
5. Select one of the following options:
   - Default — Uses device defaults. Go to step 8.
   - Custom — Uses your selections. Continue with step 6.
6. Enter the number of C3 discarded frames allowed for the threshold in the **Threshold** errors field.
7. Select the time period for the threshold from the **errors per** list. The following choices are available:
   - None — the port is blocked as soon as the specified number of C3 discarded frames allowed is met.
   - Second — the port is blocked as soon as the specified number of C3 discarded frames allowed is reached within a second.
   - Minute — the port is blocked as soon as the specified number of C3 discarded frames allowed is reached within a minute.
   - Hour — the port is blocked as soon as the specified number of C3 discarded frames allowed is reached within a hour.

- Day — the port is blocked as soon as the specified number of C3 discarded frames allowed is reached within a day.

8. Click **OK** to add the C3 discarded frames threshold to the table and close the **Add C3 Discard Frames Threshold** dialog box.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.

9. Click **OK** on the **Port Fencing** dialog box.

Adding an Invalid CRCs threshold

NOTE
This threshold is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when an **Invalid CRC** violation type meets the Fabric OS switch threshold. For default threshold values for Fabric OS devices, refer to Chapter 7 of the *Fabric Watch Administrator's Guide*.

To add an Invalid CRCs threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **Invalid CRCs (Fabric OS only)** from the **Violation Type** list.
3. Click **Add**.

   The **Add Invalid CRCs Threshold** dialog box displays.

   FIGURE 434 Add Invalid CRCs Threshold dialog box

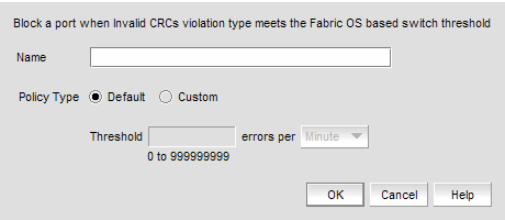

4. Enter a name for the threshold in the **Name** field.
5. Select one of the following options:
   - Default — Uses device defaults. Go to step 8.
   - Custom — Uses your selections. Continue with step 6.
6. Enter the number of invalid CRCs allowed for the threshold in the **Threshold** errors field.
7. Select the time period for the threshold from the **errors per** list. The following choices are available:
   - None — the port is blocked as soon as the specified number of invalid CRCs allowed is met.
   - Second — the port is blocked as soon as the specified number of invalid CRCs allowed is reached within a second.
   - Minute — the port is blocked as soon as the specified number of invalid CRCs allowed is reached within a minute.
   - Hour — the port is blocked as soon as the specified number of invalid CRCs allowed is reached within a hour.

- Day — the port is blocked as soon as the specified number of invalid CRCs allowed is reached within a day.

8. Click **OK** to add the Invalid CRCs threshold to the table and close the **Add Invalid CRCs Threshold** dialog box.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.

9. Click **OK** on the **Port Fencing** dialog box.

Adding an Invalid Words threshold

NOTE
This threshold is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when the **Invalid Words** violation type meets the Fabric OS switch threshold. For default threshold values for Fabric OS devices, refer to Chapter 7 of the ***Fabric Watch Administrator's Guide***.

To add an Invalid Words threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **Invalid Words (Fabric OS only)** from the **Violation Type** list.
3. Click **Add**.

   The **Add Invalid Words Threshold** dialog box displays.

   FIGURE 435 Add Invalid Words Threshold dialog box

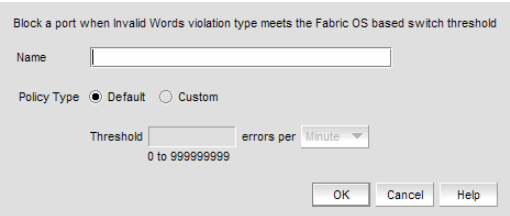

4. Enter a name for the threshold in the **Name** field.
5. Select one of the following options:
   - Default — Uses device defaults. Go to step 8.
   - Custom — Uses your selections. Continue with step 6.
6. Enter the number of invalid words allowed for the threshold in the **Threshold** errors field.
7. Select the time period for the threshold from the **errors per** list. The following choices are available:
   - None — the port is blocked as soon as the specified number of invalid words allowed is met.
   - Second — the port is blocked as soon as the specified number of invalid words allowed is reached within a second.
   - Minute — the port is blocked as soon as the specified number of invalid words allowed is reached within a minute.
   - Hour — the port is blocked as soon as the specified number of invalid words allowed is reached within a hour.

- Day — the port is blocked as soon as the specified number of invalid words allowed is reached within a day.

8. Click **OK** to add the Invalid Words threshold to the table and close the **Add Invalid Words Threshold** dialog box.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.

9. Click **OK** on the **Port Fencing** dialog box.

Adding a Link Reset threshold

NOTE
This threshold is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when the **Link Reset** violation type meets the Fabric OS switch threshold. For default threshold values for Fabric OS devices, refer to Chapter 7 of the ***Fabric Watch Administrator's Guide***.

To add a Link Reset threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **Link Reset (Fabric OS only)** from the **Violation Type** list.
3. Click **Add**.

   The **Add Link Reset Threshold** dialog box displays.

   FIGURE 436 Add Link Reset Threshold dialog box

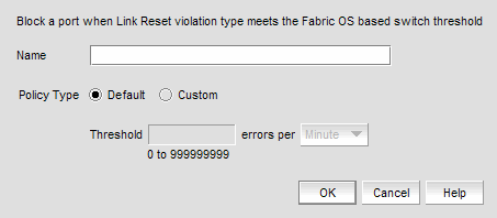

4. Enter a name for the threshold in the **Name** field.
5. Select one of the following options:
   - Default — Uses device defaults. Go to step 8.
   - Custom — Uses your selections. Continue with step 6.
6. Enter the number of link resets allowed for the threshold in the **Threshold** errors field.
7. Select the time period for the threshold from the **errors per** list. The following choices are available:
   - None — the port is blocked as soon as the specified number of link resets allowed is met.
   - Second — the port is blocked as soon as the specified number of link resets allowed is reached within a second.
   - Minute — the port is blocked as soon as the specified number of link resets allowed is reached within a minute.
   - Hour — the port is blocked as soon as the specified number of link resets allowed is reached within a hour.

- Day — the port is blocked as soon as the specified number of link resets allowed is reached within a day.

8. Click **OK** to add the Link Resets threshold to the table and close the **Add Link Reset Threshold** dialog box.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.

9. Click **OK** on the **Port Fencing** dialog box.

Adding a Protocol Error threshold

NOTE
This feature requires a Trial or Licensed version.

Use this type of threshold to block a port when one of the following ISL protocol errors meet the threshold:

- ISL Bouncing–ISL has repeatedly become unavailable due to link down events.
- ISL Segmentation–ISL has repeatedly become segmented.
- ISL Protocol Mismatch–ISL has been repeatedly put into the Invalid Attachment state due to a protocol error.

To add a Protocol Error threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **Protocol Error** from the **Violation Type** list.
3. Click **Add**.

   The **Add Protocol Error Threshold** dialog box displays.

   FIGURE 437 Add Protocol Error Threshold dialog box

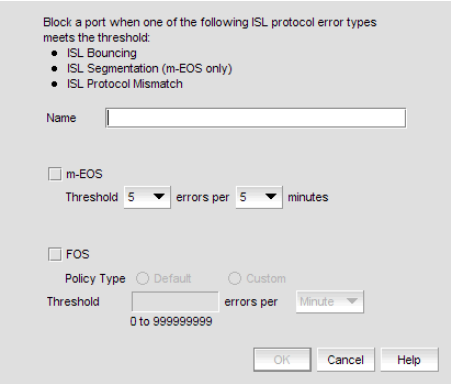

4. Enter a name for the threshold in the **Name** field.
5. Select the **Fabric OS** check box.
   a. Select one of the following options:
      - Default — Uses device defaults. Go to step 6.
      - Custom — Uses your selections. Continue with step b.
   b. Enter the number of protocol errors allowed for the threshold from the **Threshold** errors field.
   c. Select the time period for the threshold from the **errors per** list. The following choices are available:

- None — the port is blocked as soon as the specified number of protocol errors allowed is met.
- Second — the port is blocked as soon as the specified number of protocol errors allowed is reached within a second.
- Minute — the port is blocked as soon as the specified number of protocol errors allowed is reached within a minute.
- Hour — the port is blocked as soon as the specified number of protocol errors allowed is reached within a hour.
- Day — the port is blocked as soon as the specified number of protocol errors allowed is reached within a day.

6. Click **OK** to add the protocol errors threshold to the table and close the **Add Protocol Error Threshold** dialog box.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.

7. Click **OK** on the **Port Fencing** dialog box.

Adding a State Change threshold

NOTE
This threshold is only available for Fabric OS devices running 7.0 or later.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when a state change violation type meets the Fabric OS switch threshold. For 4 Gbps Router, Extension Switches and Blades only, when you apply this threshold on an E Port, the threshold is also applied to the VE Ports (internally by Fabric OS). For default threshold values for Fabric OS devices, refer to Chapter 7 of the *Fabric Watch Administrator's Guide*.

To add an State Change threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **State Change (Fabric OS only)** from the **Violation Type** list.
3. Click **Add**.

   The **Add State Change Threshold** dialog box displays.

   FIGURE 438 Add State Change Threshold dialog box

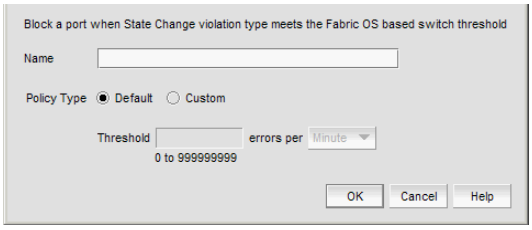

4. Enter a name for the threshold in the **Name** field.
5. Select one of the following options from the Policy Type field:
   - Default — Uses device defaults. Go to step 8.
   - Custom — Uses your selections. Continue with step 6.
6. Enter the number of state changes allowed for the threshold in the **Threshold** errors field.

7. Select the time period for the threshold from the **errors per** list. The following choices are available:
   - None — the port is blocked as soon as the specified number of state changes allowed is met.
   - Second — the port is blocked as soon as the specified number of state changes allowed is reached within a second.
   - Minute — the port is blocked as soon as the specified number of state changes allowed is reached within a minute.
   - Hour — the port is blocked as soon as the specified number of state changes allowed is reached within a hour.
   - Day — the port is blocked as soon as the specified number of state changes allowed is reached within a day.
8. Click **OK** to add the state changes threshold to the table and close the **Add State Change Threshold** dialog box.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.
9. Click **OK** on the **Port Fencing** dialog box.

Assigning thresholds

You can assign thresholds to any active object in the **Ports** table. You can only assign one threshold to an object at a time. If you assign a threshold to a switch, director, or fabric object, or to the All Fabrics object, the threshold is assigned to all subordinate objects (which do not have a directly assigned threshold) in the tree.

However, if an object inherits a threshold from another object above it in the hierarchy, you cannot remove that inherited threshold directly from the subordinate object. You must either remove the threshold from the higher object to which it was directly assigned or directly assign a different threshold to the subordinate object.

To assign an existing threshold to fabric, director, switch, port type, and port objects, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.
2. Select a threshold type from the **Violation Type** list.
3. Select the threshold you want to assign from the **Thresholds** table.
4. Select the objects (All Fabrics, Fabric, Director, Switch, Port Type, and/or Port) to which you want to assign the threshold from the **Ports** table.
5. Click the right arrow.

   A directly assigned icon ( ) displays next to the objects you selected in the **Ports** table to show that the threshold was applied at this level and was inherited by every subordinate object below it in the tree (if not affected by lower level direct assignments).

   An added icon ( ) appears next to every object in the tree to which the new threshold is applied.
6. Click **OK** on the **Port Fencing** dialog box.

Unblocking a port

The Management application allows you to unblock a port (only if it was blocked by Port Fencing) once the problem that triggered the threshold is fixed. When a port is blocked an Attention icon ( ) displays next to the port node.

To unblock a port, complete the following steps.

1. Select Monitor > Fabric Watch > Port Fencing.

   The Port Fencing dialog box displays.

2. Right-click anywhere in the Ports table and select Expand.
3. Select a blocked port from the Ports table.
4. Click Unblock.
5. Click OK on the message.

   If you did not solve the root problem, the threshold will trigger again.
6. Click OK on the Port Fencing dialog box.

Avoiding port fencing inheritance

When you directly assign a threshold to an object, the threshold is inherited by all subordinate objects in the tree (unless they already have directly assigned thresholds). You cannot remove an inherited threshold from a subordinate object. However, the Management application allows you to effectively avoid inheritance for individual subordinate objects while maintaining inheritance for other subordinate objects. To avoid inheritance for an individual subordinate object, you must create a new threshold with a maximum limit of events allowed and a minimum time period, then assign the new threshold to the subordinate object.

To turn off port fencing inheritance, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.
2. Select a threshold type from the **Violation Type** list.
3. Click **Add**.

   The **Add *Type* Threshold** dialog box displays.
4. Type a name for the new threshold (for example, AvoidProtocolError) in the **Name** field.
5. Select or enter the maximum number of errors or violations allowed in the **Threshold errors/violations** field.
6. Select the minimum time period available from the **Threshold minutes/seconds** list.
7. Click **OK** on the **Add *Type* Threshold** dialog box.
8. Click **OK** on the **Port Fencing** dialog box.

Editing thresholds

The Management application allows you to edit the name, number of events needed, and time period of ISL Protocol, Link, and Security thresholds.

Editing a C3 Discard Frames threshold

NOTE
This threshold is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when a **C3 Discard Frames** violation type meets the Fabric OS switch threshold.

To edit a C3 Discard Frames threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **C3 Discard Frames (Fabric OS only)** from the **Violation Type** list.
3. Select the threshold you want to change and click **Edit**.

   The **Edit C3 Discard Frames** dialog box displays.

4. Complete step 4 through step 7 in "Adding a C3 Discard Frames threshold" on page 894.
5. Click **OK** on the **Edit C3 Discard Frames Threshold** dialog box.

   If the threshold has already been assigned to ports, an "Are you sure you want to make the requested changes to this threshold on "X" ports?" message displays. Click **OK** to close.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.

6. Click **OK** on the **Port Fencing** dialog box.

Editing an Invalid CRCs threshold

NOTE
This threshold is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when the **Invalid CRCs Threshold** violation type meets the Fabric OS switch threshold.

To edit an Invalid CRCs threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **Invalid CRCs (Fabric OS only)** from the **Violation Type** list.

3. Select the threshold you want to change and click **Edit**.

   The **Edit Invalid CRCs Threshold** dialog box displays.

4. Complete step 4 through step 7 in "Adding an Invalid CRCs threshold" on page 896.
5. Click **OK** on the **Edit Invalid CRCs Threshold** dialog box.

   If the threshold has already been assigned to ports, an "Are you sure you want to make the requested changes to this threshold on "X" ports?" message displays. Click **OK** to close.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.

6. Click **OK** on the **Port Fencing** dialog box.

Editing an Invalid Words threshold

NOTE
This threshold is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when the **Invalid Word Threshold** violation type meets the Fabric OS switch threshold.

To edit an Invalid Words threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **Invalid Words (Fabric OS only)** from the **Violation Type** list.
3. Select the threshold you want to change and click **Edit**.

   The **Edit Invalid Words Threshold** dialog box displays.

4. Complete step 4 through step 7 in "Adding an Invalid Words threshold" on page 897.
5. Click **OK** on the **Edit Invalid Words Threshold** dialog box.

   If the threshold has already been assigned to ports, an "Are you sure you want to make the requested changes to this threshold on "X" ports?" message displays. Click **OK** to close.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.

6. Click **OK** on the **Port Fencing** dialog box.

Editing a Link Reset threshold

NOTE
This threshold is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when the **Link Reset** violation type meets the Fabric OS switch threshold.

To edit a Link Reset threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **Link Reset (Fabric OS only)** from the **Violation Type** list.
3. Select the threshold you want to change and click **Edit**.

   The **Edit Link Reset Threshold** dialog box displays.

4. Complete step 4 through step 7 in "Adding a Link Reset threshold" on page 898.
5. Click **OK** on the **Edit Link Reset Threshold** dialog box.

   If the threshold has already been assigned to ports, an "Are you sure you want to make the requested changes to this threshold on "X" ports?" message displays. Click **OK** to close.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.

6. Click **OK** on the **Port Fencing** dialog box.

Editing a Protocol Error threshold

NOTE
This threshold is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when one of the following ISL protocol errors meet the threshold:

- ISL Bouncing–ISL has repeatedly become unavailable due to link down events.
- ISL Segmentation–ISL has repeatedly become segmented.
- ISL Protocol Mismatch–ISL has been repeatedly put into the Invalid Attachment state due to a protocol error.

To edit a Protocol Error threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select **Protocol Error** from the **Violation Type** list.
3. Select the threshold you want to change and click **Edit**.

The **Edit Protocol Error Threshold** dialog box displays.

4. Complete step 4 through step 5 in "Adding a Protocol Error threshold" on page 899.
5. Click **OK** on the **Edit Protocol Error Threshold** dialog box.

   If the threshold has already been assigned to ports, an "Are you sure you want to make the requested changes to this threshold on "X" ports?" message displays. Click **OK** to close.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.
6. Click **OK** on the **Port Fencing** dialog box.

Editing a State Change threshold

NOTE
This threshold is only available for Fabric OS devices running 7.0 or later.

NOTE
This feature requires a Trial or Licensed version.

Use to block a port when a state change violation type meets the Fabric OS switch threshold. For 4 Gbps Router, Extension Switches and Blades only, when you apply this threshold on an E Port, the threshold is also applied to the VE Ports (internally by Fabric OS).

To edit an State Change threshold, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.
2. Select **State Change (Fabric OS only)** from the **Violation Type** list.
3. Select the threshold you want to change and click **Edit**.

   The **Edit State Change Threshold** dialog box displays.
4. Complete step 4 through step 7 in "Adding a State Change threshold" on page 900.
5. Click **OK** to add the state change threshold to the table and close the **Edit State Change Threshold** dialog box.

   To assign this threshold to fabrics, switches, or switch ports, refer to "Assigning thresholds" on page 901.
6. Click **OK** on the **Port Fencing** dialog box.

Finding assigned thresholds

The Management application allows you to find all ports with a specific threshold applied.

NOTE
This search is performed on the threshold name. Since Fabric OS devices do not retain the threshold name, the ability to search for a threshold on a Fabric OS device is not available in most cases.

To find assigned thresholds, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select a threshold type from the **Violation Type** list.
3. Select a threshold from the **Threshold** table.
4. Click **Find**.
5. Every port which uses the selected threshold is highlighted in the **Ports** table.
6. Click **OK** on the **Port Fencing** dialog box.

Viewing thresholds

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select a threshold type from the **Violation Type** list.
3. Review the **Thresholds** and **Ports** tables.
4. Repeat step 2 and step 3, as necessary.
5. Click **OK** on the **Port Fencing** dialog box.

Viewing all thresholds on a specific Fabric OS device

NOTE
This threshold is only available for Fabric OS devices.

NOTE
This feature requires a Trial or Licensed version.

To view all thresholds assigned to a specific switch, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Right-click anywhere in the **Ports** table and select **Expand**.
3. Right-click the device for which you want to view threshold information and select **Switch Thresholds**.

   The **Switch Thresholds** dialog box displays with a list of all thresholds assigned to the selected switch.

4. Review the **Thresholds** table.
   - **#** (Number) — The line number for each threshold in the table.
   - **Status** — The threshold status.
   - **Directly Assigned Indicator** — Whether or not the threshold was directly assigned.
   - **Name** — The threshold name.
   - **Limit** — The number of events required to trigger the threshold.
   - **Period** — The time limit required (for the number of events) to trigger a port blocking action.
   - **Area** — The threshold type.
   - **Class** — The port type.

- **Disabled on Ports** — The port numbers on which the threshold is disabled.

5. Click **Close** on the **Switch Thresholds** dialog box.
6. Click **OK** on the **Port Fencing** dialog box.

Removing thresholds

When you assign a new threshold to an object, the threshold that was active on that object is automatically removed. The Management application also allows you to remove thresholds from an individual Fabric, Switch, or Switch Port, from all Fabrics, Switches, and Switch Ports at once, as well as from the **Threshold** table.

Removing thresholds from individual objects

To remove thresholds from the All Fabrics object, an individual Fabric, Chassis group, Switch, or Switch Port, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select a threshold type from the **Violation Type** list.
3. Select the object with the threshold you want to remove in the **Ports** table.
4. Click the left arrow.

   NOTE
   If the selected object inherits a threshold assignment from an object higher in the tree, you cannot remove the threshold. However, you may assign a different threshold directly to the selected subordinate objects or change the assignment on the higher object.

   A removed icon ( ) displays next to every instance where the threshold was removed from a selected object and it does not inherits a threshold from higher in the tree.

   If an inherited threshold replaces the removed threshold, an added icon ( ) displays next to every instance where the threshold was replaced.

   A directly assigned icon ( ) displays next to each object with an assigned threshold which does not inherit a threshold from higher in the tree.

   NOTE
   If you remove a threshold from All Fabrics, it removes the threshold from individual Fabrics, switches, and switch ports in all Fabrics except for a Chassis group. You must remove repeat the procedure for the Chassis group.

5. Click **OK** on the **Port Fencing** dialog box.

Removing thresholds from the thresholds table

To remove thresholds from all Fabrics, Switches, and Switch Ports as well as the **Threshold** table, complete the following steps.

1. Select **Monitor > Fabric Watch > Port Fencing**.

   The **Port Fencing** dialog box displays.

2. Select a threshold type from the **Violation Type** list.
3. Select the threshold you want to remove in the **Thresholds** table.
4. Click **Delete**.

   A removed icon (⛔) displays next to the selected threshold in the **Thresholds** table when you click **Delete**.
5. Click **OK** on the **Port Fencing** dialog box.

FICON Environments

FICON configurations

IBM Fibre Connection (FICON) is a protocol used between IBM (and compatible) mainframes and storage. FICON configurations can be categorized into three types, based on complexity:

- Point-to-point configurations that do not use a switch.
- Switched point-to-point configurations, also called single switch configurations, connect a host channel to a storage control unit using a single switch. In this type of configuration, the channel is configured to use single-byte addressing.
- Cascaded configurations, also called high integrity fabrics, connect host channels and storage control units that reside in different domains. Cascaded FICON fabrics must be configured as high integrity fabrics. In this type of configuration, the channel is configured to use two-byte link addressing. Figure 439 and Figure 440 are examples of cascaded FICON configurations. IBM does not support configurations that have more than two domains in a path from a FICON Channel interface to a FICON Control Unit interface to Channel-to-Channel (CTC) except under special circumstances.

FIGURE 439 Cascaded configuration, two domains

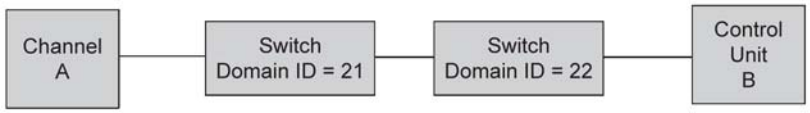

FIGURE 440 Cascaded configuration, three domains, but only two in a path

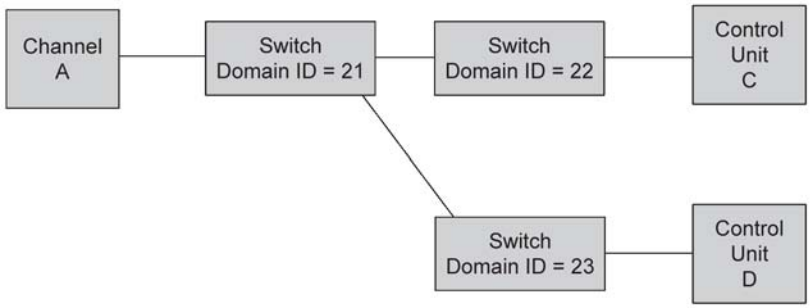

Configuring a switch for FICON operation

This section provides a basic guide for configuring a switch for FICON operation. Procedures assume that the switch is installed and IP addresses are assigned to the switch for discovery and access by the Management application. These procedures may refer to additional sections in this chapter or chapters in this manual for more detailed information.

Planning the configuration

Perform the following tasks to plan your configuration:

1. Obtain a high-level drawing of the intended fabric configuration.
2. Obtain all required license keys for the switch and Management application features.

   Licenses must be converted from transaction codes delivered with the switch. Access to a public internet connection is required. It is highly recommended that you obtain license keys before the scheduled configuration.

3. Obtain all versions of firmware for switches that the will be managed by the Management application so that you can add them to the firmware **Repository** in step 13.

   Although switches are loaded with the latest firmware at the time of manufacture, firmware may be out of date due to switch storage and transit times. If adding a switch to an existing fabric, you may need to upgrade the existing fabric, downgrade the new switch, or use a mixture of firmware in the fabric. Note that using firmware versions for switches in the same fabric that vary by one release is not recommended.

   Observe the following best practices:

   - Always check the version of firmware on a switch
   - Unless otherwise advised by a certified Fabric OS support professional, always load the most recently qualified firmware.
   - Before upgrading or downgrading firmware read the upgrade and downgrade considerations in the firmware release notes.

4. If incorporating more than one switch into a fabric, refer to planning steps in "Cascaded FICON fabric" on page 924.
5. Make a record of the following information for the switch:

   - Fabric name.
   - Switch name.
   - Domain ID (DID).

     Domain IDs are entered in either decimal or hexadecimal. If you enter the domain ID in decimal, ensure you use the correct hexadecimal equivalent. For example, if the first byte of the link address is 33, then the domain ID in decimal is 51. Also, use a domain ID that is the hexadecimal equivalent of the Switch ID in the IOCP. For example, for Switch ID 1F, set Domain ID to 31 in decimal or 1F in hexadecimal.

     The recommended best practice is the make the hexadecimal equivalent of the domain ID match the switch ID in HCD or IOCP. Also, use a unique domain ID for every switch, although this is obviously not possible in very large data centers.

   - Fabric ID (FID).

     Configure a FID if you are enabling a virtual fabric. A FID can be any number between 1-128, and all switches in the same fabric must have the same FID. Note that FMS cannot be enabled in the default switch on the 8-slot Backbone Chassis or 16 Gbps 8-slot Backbone Chassis. Therefore, the recommended best practice is leave the default switch FID at 128 and create a new logical switch for all FICON ports. A simple FID numbering scheme starting from 1 is recommended. There is no correlation between the FID and the DID.

- Management IP address.
- Administrator password.

  Although the Management application is typically configured for managing the switch as an admin user, root will also work. The default admin password is "password." You do not need to change the password during installation; however if the password is changed, the password for device discovery must be changed also. Although launched from the Management application, Element Manager (Web Tools) passwords do not propagate the Management application.

  The recommended best practice is to create identical passwords for all switches in the same fabric. This not only simplifies discovery, but in most cases since users are given access to a fabric, not an individual switch, there are fewer passwords to remember and maintain.
- Call home number.

  This may not apply. If using a call home service you will need the phone number for the service and an understanding of what is being covered in the service agreement.
- Required firmware for the switch. Refer to step 3.
- Port addressing.

  The port address is important because it is implemented in HCD or IOCP. The easiest port addressing scheme is to start from 0x00 at the bottom left of the port card, increment on ports going up the card, then continue starting numbering from the bottom right of the next column of ports. Any port addressing scheme is possible however.

6. If you are considering creating a cascaded switch configuration, consider connecting all ISLs between switches first. This will help simplify cascaded configuration. If this is not possible, you can merge cascaded fabrics later using steps in "Cascaded FICON fabric merge" on page 928.
7. If you are considering connecting cascaded switches over IP networks, refer to the planning considerations in the "Connecting cascaded FICON fabrics over FCIP" in "Fibre Channel over IP".

Configuring the switch

Perform the following steps to configure a switch for FICON operation.

1. Launch the Management application and select the **SAN** tab.

   NOTE
   The recommended best practice is to run the application client from a server other than the Management application server itself. Sometimes during installation this is not practical. To start a client on the Management application server, double click on the application icon. To open a client from a system other than the Management application server, open a browser and enter the IP address of the Management application server.
2. Configure the Management application display for FICON. Refer to the "Setting your FICON display" section of "Application Configuration".
3. Select the Decimal-Hex drop down selector on the tool bar at the top of the **SAN** tab to display domain IDs and port numbers in hex format.
4. Select **Discover > Fabrics**.

   The **Discover Fabrics** dialog box displays. If the switch is already in a fabric, it is automatically added and should display under the discovered fabric. If the switch does not display, perform step 5 and step 6.

5. Select **Add** on the **Discover Fabrics** dialog box.

   The **Add Fabric Discovery** dialog box displays.

   An error message “Discovery Failed. Fabric is busy, try again after sometime.” displays when the switch is busy. It is not recommended to continue with the other operations as the Management Application will not receive any updates from the fabric unless it is discovered. Refresh the Management Application and try again to discover the Fabric.

6. Perform one of the following tasks to configure a switch for discovery:

   - Add information for the switch in the **IP Address** tab and click **OK**.

     FIGURE 441 Add Fabric Discovery dialog box (IP Address tab)

     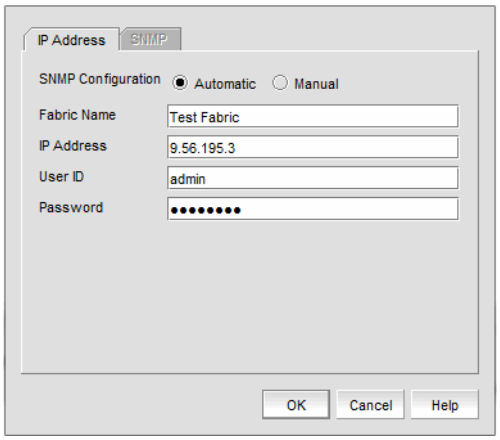

     NOTE
     Selecting **Automatic** to use the SNMPv3 profile is recommended.

   - To manually configure SNMP for discovery, select **Manual** to activate the **SNMP** tab, then select the **SNMP** tab. Fill out the fields as required.

     FIGURE 442 Add Fabric Discovery dialog box (SNMP tab)

     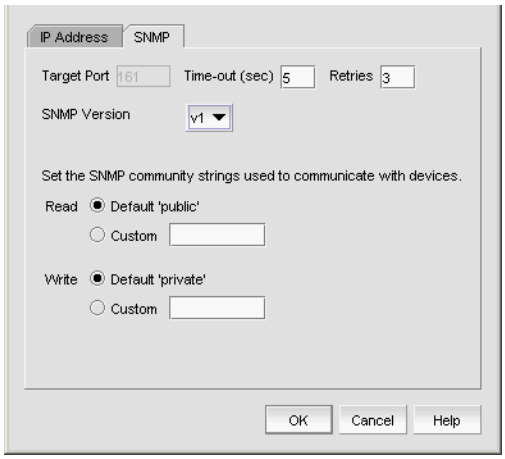

   Refer to the “SAN discovery overview” section in “Discovery” for more information on using these dialog boxes.

7. Add all required licenses to the switch using the following steps:

a. Select a discovered switch from the Product List panel, and then select **Element Manager > Admin**.

   The Web Works **Switch Administration** window displays.

b. Select the **License** tab and click **Add**.

   The **Add License** dialog box displays.

c. Past or enter the license key in the **License Key** field.

d. Click **Add License**.

e. Repeat steps b through d for additional licenses.

f. Click **Refresh** to display new licenses in the **License** tab.

8. As an optional step, manage switch users by selecting the **User** tab on the Web Works **Switch Administration** window. Use this tab to add users, change passwords, or perform other steps to manage switch users.

   NOTE
   If you change the password for a user that was used for Management application discovery, you must delete the switch from the **Discover Fabrics** dialog box, and then discover the switch again with the new login credentials.

9. To download firmware to the switch, select **Configure > Firmware Management** from the **SAN** tab on the Management application window as shown in Figure 443 on page 915.

   FIGURE 443 Selecting Firmware Management from Configure menu

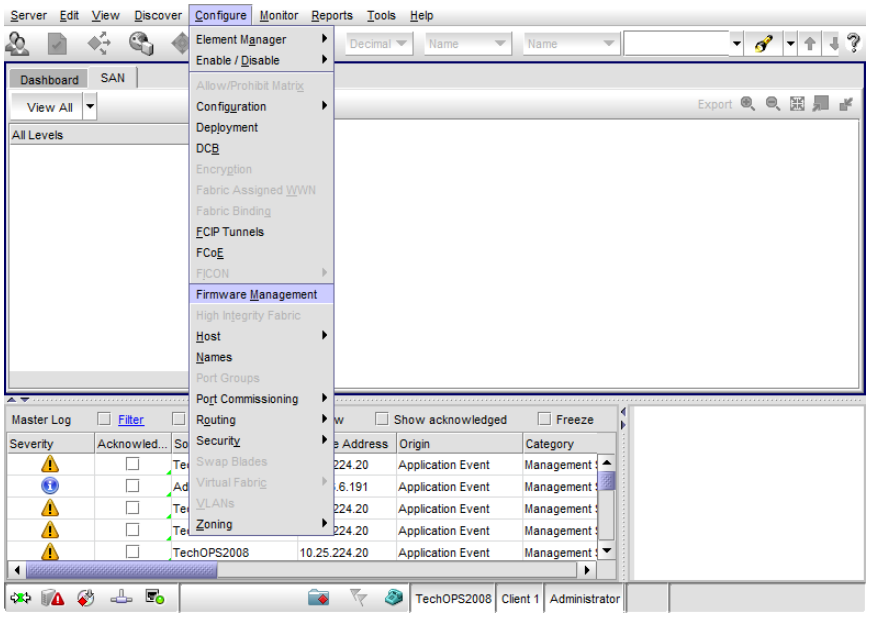

   The **Firmware Management** dialog box displays.

10. Select the **Download** tab (Figure 444).

FIGURE 444 Firmware download

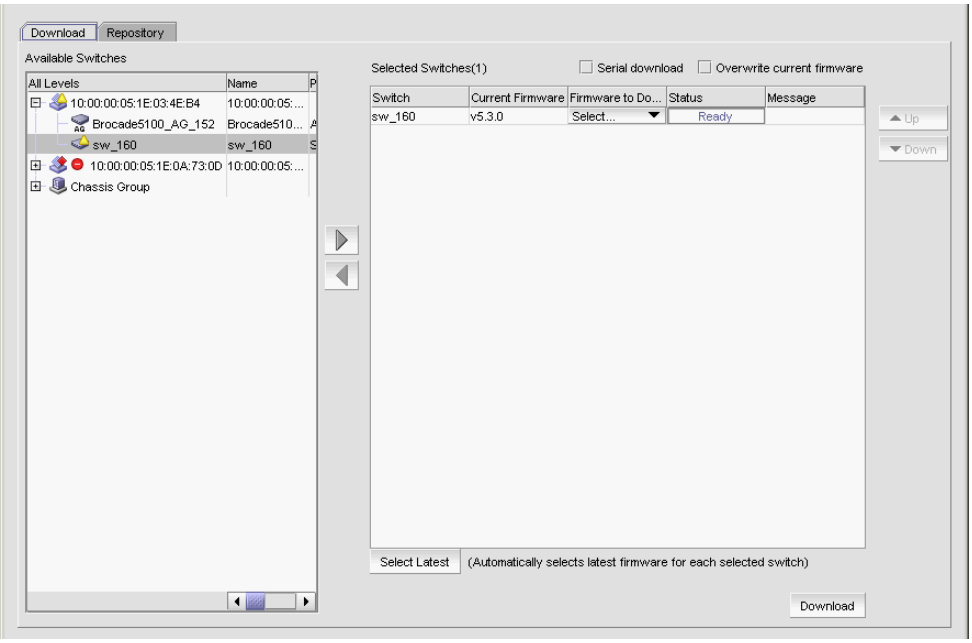

11. Select the switches in the **Available Switches** panel where you want to download firmware, and then click the right arrow to move them under **Selected Switches**.
12. Click **Download**.
13. Select the **Repository** tab to import new firmware files for downloads. Refer to the "Firmware management" section in "SAN Device Configuration" for more information on importing firmware.
14. If you are not using virtual fabrics or you do not plan to enable virtual fabrics and only use the default switch, skip to step 15. As an option at this point, you can configure virtual fabrics by referring to procedures in the following sections under "Configuring Virtual Fabrics" in the "Virtual Fabrics" chapter, then return to step 15.
    - "Enabling Virtual Fabrics"
    - "Creating a logical switch or base switch"
    - "Assigning ports to a logical switch"

    For best practices for configuring virtual fabrics, refer to **"FICON best practices for Virtual Fabrics"** on page 606.
15. To configure the switch as part of a fabric, follow procedures under "Configuring a cascaded FICON fabric" on page 925, then return to step 16.
16. If a name does not display for the switch after configuring the fabric, right click the switch icon in topology of the SAN tab and select **Properties**.

    The switch **Properties** dialog box displays.
17. Edit the switch name.
18. Define port fencing parameters for the switch using the following steps (optional):

    NOTE
    Although this is an optional step, best practice is to configure port fencing.

a. Configure thresholds that you require for the switch using steps under the "Adding thresholds" in "Port Fencing".

   Following are recommend parameters for the various thresholds:

   - C3 Discard Frames = 2 per minute.
   - Invalid Words = 25 per minute.
   - Invalid CRCs = 3 per minute. Note that it is not uncommon for an ISL to travel through a path that is more prone to noise than internal data center connections to control units and channels. Therefore, a slightly higher CRC threshold may be better for E-Port connections. In most cases the CRC is set to 3.
   - Link Reset = 2 per minute.
   - Protocol Error = 2 per minute.
   - State Change = 7 per minute.

b. Assign a threshold to the switch using steps under "Assigning thresholds" in "Port Fencing".

19. Set the zoning policy for the switch by referring to steps under "Enabling or disabling the default zone for fabrics" in "Zoning".

    The recommended policy is to disable the default zone (No Access). Although enabling the default zone (All Access) can be used for FICON environments, prohibiting connection between ports using the **Configure Allow/Prohibit Matrix** dialog box requires activating at least one zone. Even if you do not want to prohibit connections using the matrix, configuring a single zone containing all ports provides the same benefits as All Access, while providing flexibility to configure "prohibits" or more restrictive zoning in the future. In addition, when moving an ISL in the future, there will not need to modify zoning.

20. Configure zoning using steps under "Configuring zoning" in "Zoning".

    Be sure to reference the "Zoning and FICON" section of "Zoning" for more information on FICON environments.

21. Configure the Allow/Prohibit Matrix for the switch using steps under "Configuring an Allow/Prohibit Matrix" on page 918.

22. Configure Call Home by referring to procedures in "Call Home".

    NOTE
    The call home number and the events to trigger a call home depend on your service contract and service provider. Contact your service provider for additional information.

23. Enable bottleneck detection using the following Fabric OS **bottlneckmon** commands:

    - **bottleneckmon --cfgcredittools -intport -recover onLrOnly** - This command monitors for lost credits on links. This is necessary because occasional errors on links can cause lost credits that can result in IFCCs and poor performance over time.
    - **bottleneckmon --enable -alert** - This command causes AN-1004 RAS log messages to generate whenever congestion occurs and AN-1010 RAS log messages to generate whenever severe congestion occurs. The recommended best practice is to enable alerts now so that you don't forget when you merging the fabrics.

    The **bottlneckmon** command operates the entire chassis, regardless of the FID where it is executed.

24. Clear error counters (common during switch configuration) by right-clicking the switch in the Connectivity Map or Product List and selecting **Performance > Clear Counters**.

Configuring FICON display

You can set display settings for FICON display so that the columns of any table that contains end device descriptions to move the following eight columns to be the first columns: FC Address, Serial #, Tag, Device Type, Model, Vendor, Port Type, and WWN. For instructions, refer to "Setting your FICON display" on page 84.

Configuring an Allow/Prohibit Matrix

The Allow/Prohibit Matrix is a FICON port attribute that can be used to prohibit communication between specific ports. Allow/Prohibit Matrix are not recommended on E_Ports (inter-switch links).

The Allow/Prohibit Matrix can be manipulated by host-based management programs using FICON Control Unit Port (CUP), or from a Management application program to create policies and determine paths for data and command flows. Up to eight Allow/Prohibit matrices can be modified at the same time. Allow/Prohibit Matrix settings apply per switch rather than per fabric, and only work when an active zone configuration is present in the fabric.

Multiple configurations can be defined, edited, copied, or removed. Only one configuration can be active per switch.

Configuring the Allow/Prohibit matrix requires that a zone configuration be activated on the fabric. Prohibits can be configured without an active zone configuration, but they cannot be enforced until an effective zone is configured.

1. Select **Configure > Allow/Prohibit Matrix**.

   The **Configure Allow/Prohibit Matrix** dialog box displays.

2. Select a switch from **Available Switches**.

   Two default configurations (Active and IPL) are displayed in a tree structure under the switch. Existing configurations are also displayed.

3. Choose one of the following options:

   - Double-click a configuration file.
   - Select a configuration file and click the right arrow.

   A matrix displays in the **Allow/Prohibit Matrix** panel. The switch ports are displayed on both the vertical axis and horizontal axis. An Allow icon ( ) indicates communication is allowed between the ports, as shown in Figure 445 on page 918.

FIGURE 445 Active Configuration in Allow/Prohibit Matrix panel

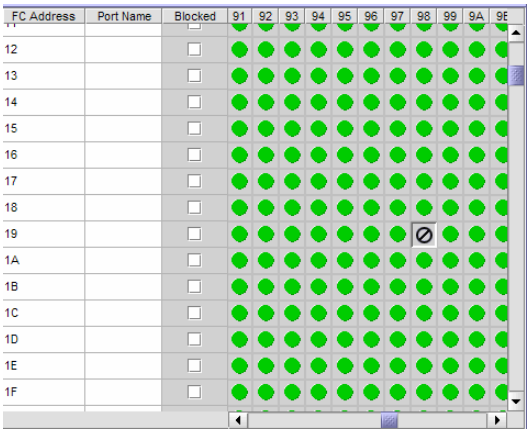

4. Prohibit a connection between two ports by clicking the intersection point between the ports.

   A prohibit icon (⊘) displays at the intersection point. If you know the port addresses of the ports for which you want to prohibit or allow communication and do not want to search the matrix for the exact port intersection point, use the procedure in "Configuring an Allow/Prohibit Matrix manually" on page 919.

5. Repeat step 4 as needed to create the matrix you want to apply. If you want to change a selection from prohibit to allow, click the intersection point to clear the prohibit icon.

6. When you have completed the matrix, click **Save** if you started with a new matrix, or **Save As** to save a copy of an existing matrix.

7. Click **Analyze Zone Conflicts**.

   This operation can be done before or after a configuration is saved. This operation checks the current zoning settings for conflicts with settings in the Allow/Prohibit Matrix. Zone conflict is analyzed against the switch for port zoning only. The table cells display in the red background if the two ports are not in the same zone in an active zone configuration.

8. Click **Close** on the **Configure Allow/Prohibit Matrix** dialog box.

Configuring an Allow/Prohibit Matrix manually

To configure to allow or prohibit communication between specific ports manually, complete the following steps.

1. Select **Configure > Allow/Prohibit Matrix**.

   The **Configure Allow/Prohibit Matrix** dialog box displays.

2. Select a switch from **Available Switches**.

   Two default configurations (Active and IPL) are displayed in a tree structure under the switch. Existing configurations are also displayed.

3. Choose one of the following options:

   - Double-click a configuration file.
   - Select a configuration file and click the right arrow.

   A matrix displays. The switch ports are displayed on both the vertical axis and horizontal axis. An Allow icon (●) indicates communication is allowed between the ports.

4. Click **Manual Allow/Prohibit**.

   The **Manual Allow/Prohibit** dialog box displays, as shown in Figure 446 on page 920.

FIGURE 446 Manual Allow/Prohibit dialog box

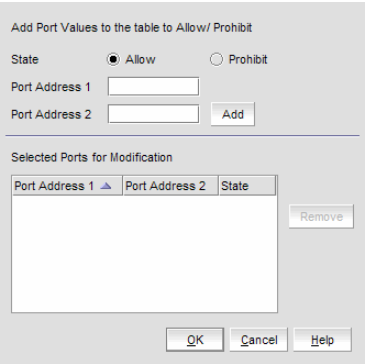

NOTE
The **Manual Allow/Prohibit** dialog box is only available for Fabric OS products.

5. Select one of the following options:
   - Select **Allow** to allow communication between two specific ports.
   - Select **Prohibit** to prohibit communication between two specific ports.
6. Enter the port number of the first port for which you want to allow or prohibit communication in the **Port Address 1** field.
7. Enter the port number of the second port for which you want to allow or prohibit communication in the **Port Address 2** field.
8. Click **Add**.

   The information displays in the **Selected Ports for Modification** list.

   To delete any of these manual configurations, select the configuration you want to delete in the **Selected Ports for Modification** list and click **Remove**.

   The **Selected Ports for Modification** list displays the following information:
   - **Port Address 1** column — The port number of the first port for which you want to allow or prohibit communication.
   - **Port Address 2** column — The port number of the second port for which you want to allow or prohibit communication.
   - **State** column — Whether you want to allow or prohibit communication.
9. Repeat step 5 through step 8 for each allow or prohibit configuration.
10. Click **OK** on the **Manual Allow/Prohibit** dialog box.
11. When you have completed the matrix, click **Save** if you started with a new matrix, or **Save As** if you edited a copy of an existing matrix.
12. Click **Analyze Zone Conflicts**.

    This operation can be done before or after a configuration is saved. This operation checks the current zoning settings for conflicts with settings in the Allow/Prohibit Matrix. Zone conflict is analyzed against the switch for port zoning only. The table cells display in the red background if the two ports are not in the same zone in an active zone configuration.
13. Click **Close** on the **Configure Allow/Prohibit Matrix** dialog box.

Saving or copying Allow/Prohibit Matrix configurations to another device

When copying or saving a configuration from a small switch (source switch with fewer ports; for example, 64 ports) to a larger switch (destination switch with a larger number of ports; for example, 256 ports) only the port address range of the smaller switch will be affected on the larger switch. All additional port addresses will display the default settings (port state defaults to “Allow” and the **Blocked** check box defaults to cleared).

Copying or saving a configuration from a larger switch to a smaller device only copies or saves the port address range that matches the smaller switch. Additionally a message displays that the additional port addresses from the larger switch are discarded.

When copying or saving a configuration from or to logical switches, the only ports affected are the port addresses defined in the logical switch. The FICONd CUP Daemon retains the full compliment of records regardless of the size of the logical switch. Therefore, copying or saving a configuration from or to logical switches should work the same as copying or saving between standard switches.

Copying an Allow/Prohibit Matrix configuration

To duplicate an existing Allow/Prohibit Matrix configuration, complete the following steps.

1. Select **Configure > Allow/Prohibit Matrix**.

   The **Configure Allow/Prohibit Matrix** dialog box displays.

2. Select the Allow/Prohibit Matrix configuration you want to copy.

   You can do this by expanding the view for the switch under **Available Switches** and selecting a configuration, or you can select the matrix under **Allow/Prohibit Matrix**.

3. Click **Duplicate**.

   The **Save As/Duplicate** dialog box displays, as shown in Figure 447 on page 921.

   FIGURE 447 Save As/Duplicate dialog box

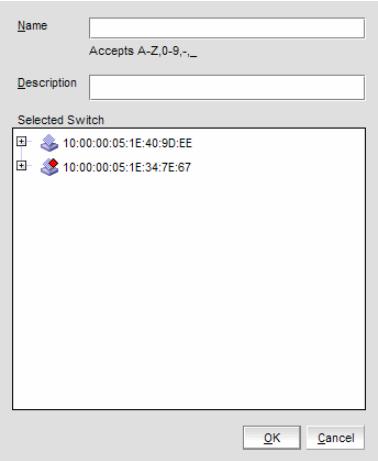

4. Enter a name for the configuration.
5. Enter a description for the configuration.
6. Select the check box for the switch to which you want to save the configuration in the **Selected Switch** list.

7. Click **OK**.

   A message displays stating that the outstanding port configuration is discarded when copying a configuration from the switch with more ports to a switch with fewer ports and vice versa. Click **OK** to close the message.

   The copied configuration displays in the **Available Switches** list under the selected switch. To edit this configuration, refer to "Configuring an Allow/Prohibit Matrix" on page 918 or "Configuring an Allow/Prohibit Matrix manually" on page 919.

Saving an Allow/Prohibit Matrix configuration to another device

To save an existing Allow/Prohibit Matrix configuration to another device, complete the following steps.

1. Select **Configure > Allow/Prohibit Matrix**.

   The **Configure Allow/Prohibit Matrix** dialog box displays.

2. Select the Allow/Prohibit Matrix configuration you want to save.

   You can do this by expanding the view for the switch under **Available Switches** and selecting a configuration, or you can select the matrix under **Allow/Prohibit Matrix**.

3. Click **Save As**.

   The **Save As/Duplicate** dialog box displays, as shown in Figure 448 on page 922.

   FIGURE 448 Save As/Duplicate dialog box

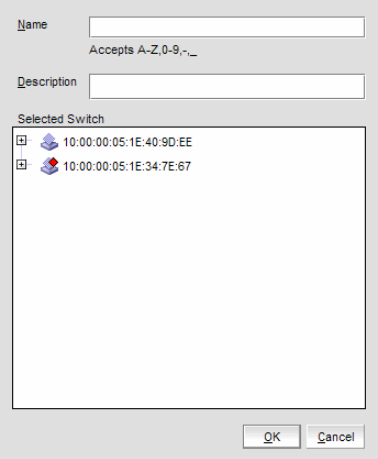

4. Enter a name for the configuration.
5. Enter a description for the configuration.
6. Select the check box for the device to which you want to save the configuration in the **Selected Switch** list.
7. Click **OK**.

   A message displays stating that the outstanding port configuration is discarded when copying a configuration from the switch with more ports to a switch with fewer ports and vice versa. Click **OK** to close the message.

   The saved configuration displays in the **Available Switches** table under the selected switch. To edit this configuration, refer to "Configuring an Allow/Prohibit Matrix" on page 918 or "Configuring an Allow/Prohibit Matrix manually" on page 919.

Activating an Allow/Prohibit Matrix configuration

You must have an active zone configuration before you can activate an Allow/Prohibit Matrix configuration.

1. Select **Configure > Allow/Prohibit Matrix**.

   The **Configure Allow/Prohibit Matrix** dialog box displays.

2. Select the Allow/Prohibit Matrix configuration you want to activate.
   You can do this by expanding the view for the switch under **Available Switches** and selecting a configuration, or you may select the matrix under **Allow/Prohibit Matrix**.

3. Click **Activate**.

   A confirmation message displays, as shown in Figure 449 on page 923.

   FIGURE 449 Activate Matrix Confirmation message

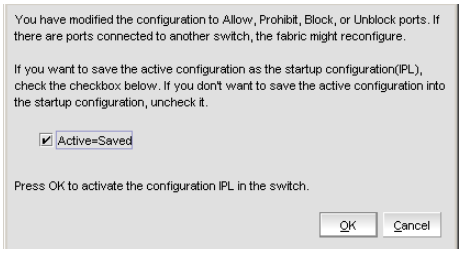

4. Select the **Active=Saved** check box to save the active configuration as the startup configuration (IPL).

5. Click **OK** to confirm.

   If you select the **Active=Saved** check box, the text [=Active] is appended to the IPL file in the **Configure Allow/Prohibit Matrix** dialog box.

   The **Active=Saved** check box and the IPL filename represent the current state of the Active=Saved Mode (ASM) bit on the switch. However, this is limited to changes done to the ASM configuration through the Management application. If changes occur through external means (such as, Webtools or the CLI) the changes are not reflected in the Management application until the **Configure Allow/Prohibit Matrix** dialog box is re-launched.

   NOTE
   Active=Saved means the matrix configuration will survive a power failure. If not selected, all ports can access each other after power is restored.

Deleting an Allow/Prohibit Matrix configuration

You cannot delete the active configuration, the IPL configuration, or a configuration that is marked as having uncommitted changes.

1. Select **Configure > Allow/Prohibit Matrix**.

   The **Configure Allow/Prohibit Matrix** dialog box displays.

2. Select the Allow/Prohibit Matrix configuration you want to delete.
   You can do this by expanding the view for the switch under **Available Switches** and selecting a configuration, or you can select the matrix under **Allow/Prohibit Matrix**.

3. Click **Delete**.

   A confirmation message displays.

4. Click **Yes** to confirm.

Changing the Allow/Prohibit Matrix display

You can modify the matrix display on the **Configure Allow/Prohibit Matrix** dialog box using the **Window Arrangement** list above the matrix display or the **Clear all port names** option below the display.

Changing window arrangement

There are three options for the **Allow/Prohibit Matrix** display on the **Configure Allow/Prohibit Matrix** dialog box located in the **Window Arrangement** list above the display.

- The matrix definitions may be cascaded (this is the default view).
- The matrix definitions may be tiled horizontally.
- The matrix definitions may be tiled vertically.

Perform the following steps to change the display to the desired format.

1. Select **Configure > Allow/Prohibit Matrix**.

   The **Configure Allow/Prohibit Matrix** dialog box displays.

2. Select **Cascade**, **Tile Horizontally**, or **Tile Vertically** from the **Window Arrangement** list.

Clearing port names

Use the following steps to clear all port names from the selected matrix.

1. Select **Clear Port Names** below the matrix display.

   A warning displays asking you to confirm the operation.

2. Select **Yes** to clear all port names from the matrix or select **No** to cancel the operation.

Cascaded FICON fabric

NOTE
You must have FICON Management privileges to configure a fabric for cascaded FICON.

NOTE
If HIF mode is not enabled and the FMS mode is deployed, then the fabric will set the HIF key.

The Management application enables you to easily configure a fabric for cascaded FICON. Note that configuring a fabric for cascaded FICON may be disruptive to current I/O operations in the fabric, as this involves disabling and enabling the switches in the fabric.

FICON configuration performs the following operations on the selected fabric:

- Turns on the insistent domain ID flag (IDID) on all switches.
- Sets High Integrity Fabric Configuration (HIFC) on the seed switch.
  - Fabric-wide consistency policy (FWCP) is configured to include SCC in strict mode.
  - SCC policy is created or modified to limit connectivity to only the switches in the selected fabric.

- Enables device-based routing on all switches.
- Enables In-Order Delivery (IOD) on all switches.
- Enables Dynamic Load Sharing (DLS) based on user selection and the firmware level.

  NOTE
  To enable DLS, all switches in the fabric must be 8 Gbps or faster and running Fabric OS 7.0 or later.

  NOTE
  You must enable FMS mode for logical fabrics running Fabric OS 7.3.0 or later.

- (Optional) Turns on FICON Management Server (FMS) mode on all switches.

  Consider the following information when enabling FMS mode:

  - If switches are running Fabric OS 7.0 and later, FMS mode will not be enabled unless the switches have an active CUP license.
  - If switches are running Fabric OS earlier than version 7.0 and do not have a CUP license, after successful configuration, you can access the Port Connectivity (Allow/Prohibit) matrix, but the host system cannot communicate with the FICON Management Server unless you install a CUP license. If a CUP license is later installed on these switches, then FMS mode must be re-enabled on these switches.
  - For logical fabrics running Fabric OS 7.1 or later, you can enable FMS mode when logical switches are configured to allow XISL use.

Configuring a cascaded FICON fabric

The FICON wizard automatically creates HIFC settings that support a cascaded FICON fabric.

1. Select **Configure > FICON > Configure Fabric** or right-click a fabric in the product tree and select **FICON > Configure Fabric**.

   The **Configure Cascaded FICON Fabric /Switch** dialog box displays, as shown in Figure 450.

   FIGURE 450 Configure Cascaded FICON Fabric /Switch dialog box

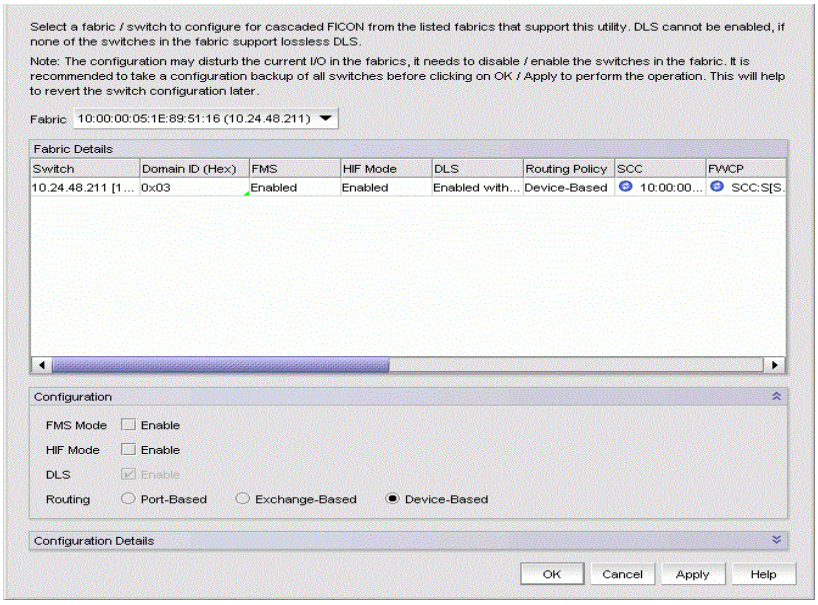

2. Use the **Fabric** list to select the fabric you want to configure. The **Fabric Details** table displays the current FICON active setting information. If you make any changes to the existing settings, the expected parameter changes are indicated with a modified icon( ). The modified icon always displays for **SCC** and **FWCP** policies. The **Fabric Details** table displays the following details for switches in the fabric:

   - **Switch** - The switch IP address.
   - **Domain ID (Hex)** - The Domain ID of the switch present in the selected fabric. You can edit the Domain ID. The allowed range is 1 through EF (1 through 239).
   - **FMS** - Whether FICON Management Server (FMS) mode is enabled.
   - **HIF Mode** - Whether High Integrity Fabric (HIF) mode is enabled.
   - **DLS** - Whether switch dynamic load sharing is enabled.
   - **Routing Policy** - The routing policy configured on the switch (Port-Based, Exchange-Based, or Device-Based). The default routing policy is Device-Based.
   - **SCC** - Switch Connection Control (SCC) policy.
   - **FWCP**- Fabric Wide Consistency Policy (FWCP).
   - **IDID** - Whether the switch insistent domain ID (IDID) is enabled or disabled.
   - **IOD** - Whether switch in-order delivery is enabled or disabled.
   - **Switch WWN** - The world wide name of the switch.
   - **Type** - The type of switch.
   - **Model number** - The model number of the switch.
   - **Manufacturer** - The manufacturer of the switch.
   - **Plant of Manufacture** - The switch manufacturing plant.
   - **RNID Sequence** - RNID sequence number for the switch.
   - **Tag** - The tag number of the switch.

   You can view the **Fabric Details** table, collapse or expand the **Configuration** panel, and the **Configuration Details** panel using the Collapse/Expand button. By default, the **Configuration Details** panel is collapsed. You can expand only two panels at a time.

   NOTE
   (Fabric OS switches only) All switches in a fabric must be running Fabric OS version 7.0 or later.

3. Select the **FMS Mode** check box to manage the fabric by a host-based management program using FICON CUP protocol.

   If you select **FMS Mode**, each switch is checked for a CUP license. Any switches that do not have a CUP license are listed, with a reminder that a CUP license is necessary to communicate with the FICON Management Server.

4. Select the **HIF Mode** check box to enable HIF mode on all switches running Fabric OS 7.3 and later. The status message will indicate if HIF mode is enabled.

   NOTE
   A warning message will be displayed when HIF is deactivated and FMS mode is enabled for one switch in the selected fabric.

5. Select the **DLS** check box to enable Dynamic Load Sharing (DLS) or Lossless DLS only on switches that support lossless DLS. For more information, refer to "Enabling DLS" on page 927. You must enable DLS to select routing policies.

6. Select one of the following options to enable port-based, exchange-based, or device-based routing on switches:

   - **Port-Based** enables port-based routing on 4 Gbps platform switches.

- **Exchange-Based** enables exchange-based routing for the fabric if all switches are 8 Gbps or greater platforms running Fabric OS 7.0 or later. If these requirements are not met, an error message displays.
- **Device-Based** enables device-based routing for the fabric if all switches in the fabric are 8 Gbps or greater platforms running Fabric OS 7.1 or later. If these requirements are not met, an error message displays.

  NOTE
  Exchange-based routing, port-based routing, or device-based routing is enabled on all switches of the selected fabric. You cannot enable a mixed routing policy.

7. Click **Apply** to save the configuration changes without closing the currently active **Configure Cascaded FICON Fabric /Switch** dialog box.
8. Click **OK** if you want to proceed.

   A warning message displays listing the switches of the selected fabric that are to be disabled and re-enabled in order to enable the desired routing policy and IDID. When you select another fabric from the **Fabric Details** table and change the configuration setting, but do not save the settings, a confirmation message displays. Click **Yes** to abort the changes made to the currently active fabric or click **No** to abort the fabric scope change and continue with the currently active fabric.
9. Click **Yes** to continue.

   If configuration is successful, a confirmation message displays.

   If **FMS Mode** was selected, each switch is checked for a CUP license. Any switches that do not have a CUP license are listed, with a reminder that a CUP license is necessary to communicate with the S/B FICON Management Server.

   NOTE
   FMS mode cannot be enabled on switches running Fabric OS 7.0 and later unless the switches have an active CUP license.

Enabling DLS

Consider the following when enabling Dynamic Load Sharing (DLS) in step 5:

- DLS requires DLS support on the switch. Lossless DLS requires Lossless DLS support on the switch.
- Enabling DLS will enable IOD without Lossless DLS on all other switches, enable DLS on switches that support DLS, and disable DLS on all other switches.
- DLS is only supported on the 40-port, 8 Gbps FC Switch, 80-port, 8 Gbps FC Switch, 512-port Backbone Chassis, and 4-slot Backbone Chassis.
- Enabling DLS may result in dropped frames when paths fail over. It is recommended that you set the preferred IOD delay time to minimize frame drops.
- To enable DLS, all switches in the fabric must be 8 Gbps or faster and running Fabric OS 7.0 or later.

Cascaded FICON fabric merge

The Management application provides a wizard to help you merge two fabrics for cascaded FICON. Note that merging two cascaded FICON fabrics may be disruptive to current I/O operations in both fabrics as this involves disabling and enabling the switches in both fabrics. The merge process will not make any configuration changes on the primary (production) fabric that are disruptive.

NOTE
It is recommended that you run a configuration backup on all switches before performing the fabric merge. This helps you to revert back the switch configurations later.

The cascaded FICON fabrics merge wizard performs the following operations:

- Checks the primary and secondary fabrics for any merge issues.
- Configures High Integrity Fabric Configuration (HIFC) on the seed switch of the primary and secondary fabric.
  - SCC policy will be created or modified to limit connectivity to switches from both fabrics.
  - Configures Fabric-Wide Consistency Policy (FWCP) on both fabrics.
  - FWCP is configured in tolerant mode for SCC for a Fibre Channel Routing (FCR) fabric.
- Enables Port-Based Routing (PBR) on all switches in the secondary fabric if all the switches in the primary fabric are found to be enabled for PBR. Note that a mixed policy of Exchanged-Based Routing (EBR), Device-Based Routing (DBR) and PBR cannot be enabled on a fabric.
- Enables Exchange-Based Routing (EBR) on all switches in the secondary fabric if all switches in the primary fabric are enabled for EBR. Note that EBR requires that switches operate at 8 Gbps or greater with Fabric OS 7.0 or later. If all the EBR-enabled switches in the primary fabric are found to meet these requirements and a switch in the secondary fabric does not meet these requirements, an error message displays. Note that a mixed policy of EBR and PBR cannot be enabled on a fabric.
- Enables Device-Based Routing (DBR) on all switches in the secondary fabric if all switches in the primary fabric are enabled for DBR. Note that DBR requires that switches operate at 8 Gbps or greater with Fabric OS 7.1 or later. If all the DBR-enabled switches in the primary fabric are found to meet these requirements and a switch in the secondary fabric does not meet these requirements, an error message displays. Note that a mixed policy of PBR, EBR, and DBR cannot be enabled on a fabric.
- (Optional) Turns on FICON Management Server (FMS) mode on all switches. If some switches already have FMS mode enabled, it is re-enabled.

  Consider the following information when enabling FMS mode.

  - If switches are running Fabric OS 7.0 and later, FMS will not be enabled unless the switches have an active CUP license.
  - If switches are running Fabric OS earlier than version 7.0 and do not have a CUP license, after successful configuration, you can access the Port Connectivity (Allow/Prohibit) matrix, but the host system cannot communicate with the FICON Management Server unless you install a CUP license. If a CUP license is later installed on these switches, then FMS mode must be re-enabled on these switches.
  - For logical fabrics running Fabric OS v7.1 or later, you can enable FMS mode when logical switches are configured to allow XISL use.
- (Optional) Configures long distance settings on selected ports of primary and secondary fabrics (requires an Extended Fabric license).

  NOTE
  If the distance between the merged fabrics is 10 km or greater, you must configure the connection as a long distance connection.

Note that the merge wizard does not enable primary fabric switches for DLS, In-Order Delivery (IOD), insistent domain ID flag (IDID), and Advanced Performance Tuning (APT).

- In-Order Delivery (IOD) will be enabled on all switches in the secondary fabric.
- Dynamic Load Sharing (DLS) will be enabled on switches in the secondary fabric that are operating at 8 Gbps or greater and are running Fabric OS 7.0 or later.

  NOTE
  To enable DLS, all switches in the fabric must be 8 Gbps or faster and running Fabric OS 6.3 or later.

- Primary fabric switches will not be disturbed for disruptive operations, such as IDID and APT. Instead, all primary fabric switches will be validated for current routing policies and the same policies will be enabled on all the secondary fabric switches.

The cascaded FICON fabrics merge wizard performs the following operations to avoid Active Directory (AD), Access Control List (ACL), and zone database merge conflicts between the two fabrics:

- Clears Admin Domain, Access Control Lists (ACLs), and zone databases, if they exist, from the secondary fabric that you select within the wizard.

  NOTE
  Clearing the ACL database in a large fabric can take a long time; for example, in a 50-switch fabric, this operation can take from 30 minutes to 1 hour.

- Sets the default zoning configuration on the secondary fabric to match the default zoning status of the primary fabric.
- Modifies ACL policy on the secondary fabric to match the primary fabric parameters, including Accept Distribution and FWCP.
- Sets FWCP in strict mode for SCC for the primary fabric.
- Sets FWCP in tolerant mode for the Fibre Channel Routing (FCR) fabric.

Merging two cascaded FICON fabrics

If you want to join two cascaded FICON fabrics, they must be merged. If the distance between fabrics is 10 km or more, an Extended Fabrics license is required, and an extra step is required to configure the connection as a long distance connection. To successfully configure a long distance connection, use the same E_Ports and cable distance values used when configuring Extended Fabrics. For long distance connections, it is recommended that you create the Extended Fabrics configuration first, have an active connection, and have the E_Port and cable distance values ready before you merge the fabrics.

1. Select **Configure > FICON > Merge Fabrics** or right-click a fabric in the product tree and select **FICON > Merge Fabrics**.

   The **Overview** screen of the cascade FICON fabrics merge wizard displays.

   NOTE
   The cascade FICON fabrics merge wizard is only available for Fabric OS products.

2. Click **Next**.

   The **Select fabrics** screen displays.

3. Select the two fabrics you want to merge under **Available Fabrics**, and click the right arrow to move them to **Selected Fabrics**. You may do this one fabric at a time, or select both by pressing **CTRL** and then clicking each fabric.

   NOTE
   All switches in a fabric must be running Fabric OS version 7.0 or later and must be reachable.

NOTE
Switches running Fabric OS 6.3 or earlier cannot be merged with switches running Fabric OS 6.4 or later.

NOTE
For 8 Gbps switches, all switches in the fabric must be 8 Gbps or faster. 8 Gbps switches cannot be merged with switches that have SFP transceivers with a speed less than 8 Gbps.

4. Click **Next**.

   The **Set up merge options** screen displays, as shown in Figure 451.

FIGURE 451 Set up merge options screen

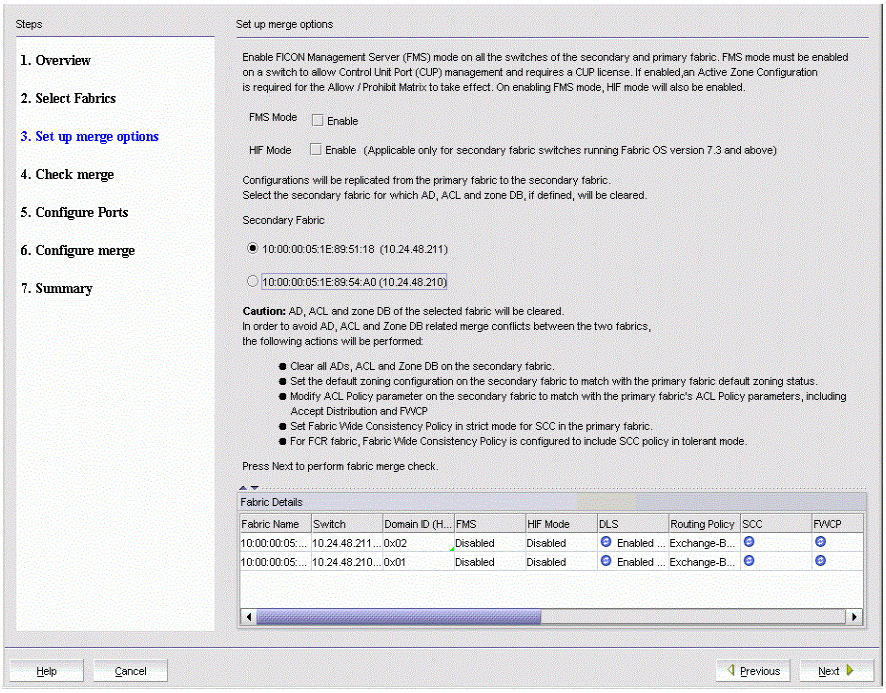

5. Select **FMS Mode** to manage the fabric by a host-based management program using FICON CUP protocol. Note that you cannot enable FMS mode on switches running Fabric OS 7.0 or later unless they have an active CUP license.
6. Select **HIF Mode** to enable HIF mode on all secondary fabric switches running Fabric OS version 7.3 and later.

   NOTE
   HIF mode is auto-enabled when FMS mode is selected.

   NOTE
   HIF mode is enabled for secondary fabric switches.

7. Select a secondary fabric where AD, ACL, and zone databases, if defined, will be cleared.
8. Read the bulleted list of actions so you understand the actions that are taken to avoid conflicts when the fabrics are merged.

   The **Fabric Details** table displays the details of the switches in the fabric. For **FMS**, **SCC**, and **FWCP**, the modified icon will display based on the primary fabric and secondary fabric selection. For the secondary fabric, the Domain ID field is editable.

Property fields containing a green triangle ( ) in the lower right corner are editable. To edit a field with a green triangle ( ), click in the field and make your changes.

9. Click **Next**.

   The **Check merge** screen displays.

   A **Status details** table shows progress through merge check points. A rotating arrow under **Status** indicates a merge check step is in progress. A green check mark indicates successful completion of that merge check. A blue information sign indicates a skipped check. A red stop sign indicates a failed step. If the configuration is successful, all configuration items have green check marks.

   NOTE
   Beginning with the 12.4.0 release, the HIF Mode Test can be performed for check merge.

10. If the merge fails, but is recoverable, click **Resolve**.
11. If desired, click **Check Merge Again** to run the merge check test again.
12. Click **Next** to continue.

    The **Configure Ports** screen displays, as shown in Figure 452. You can configure Encryption, Compression, and Long Distance for E_Ports, EX_Ports, and U_Ports.

FIGURE 452 Configure Ports screen

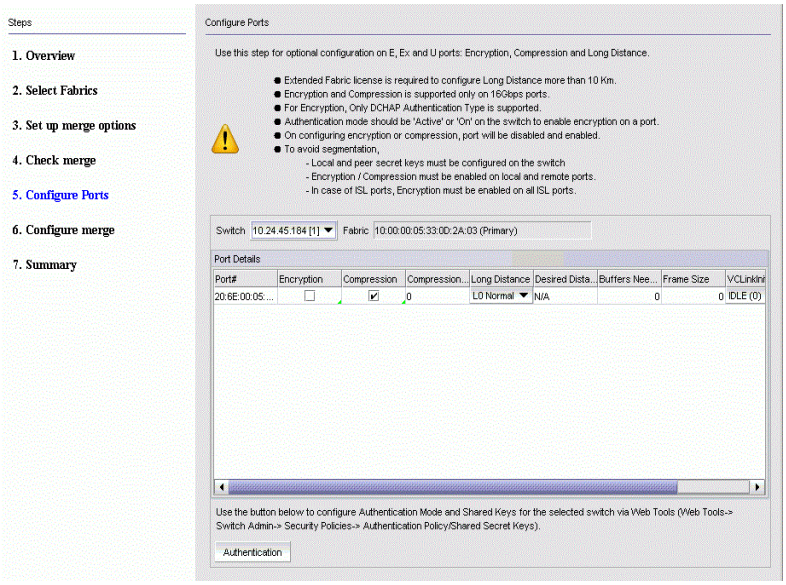

13. Read the bulleted list of actions so you understand the actions that are taken for Encryption, Compression, and Long Distance configuration.
14. Use the **Switch** list to select the primary or secondary switch you want to configure. The **Port Details** table lists the following information based on the switch selected:

    - **Port#** - Displays the port number based on the port label selected in the **SAN** tab.
    - **Encryption** - Displays whether encryption is enabled or disabled. Only supported on 16 Gbps ports.
    - **Compression** - Displays whether compression is enabled or disabled. Only supported on 16 Gbps ports.

- **Long Distance** - Displays whether the connection is considered to be "a normal or longer distance". Supported values include L0 - Normal, LE <= 10 KM, LD - Auto, or LS - Static.
- **Desired Distance** - Displays the desired link distance. This field is editable.
- **Buffer Needed** - The number of buffers needed. This field is editable.
- **Frame Size** - The size of the frame. This field is auto-populated.
- **VCLinkInit** - The fill words used on long distance links. Possible values include: IDLE (0) or Arbitary (1).
- **FillWord** - Select ARBs or IDLEs to configure the Fibre Channel Primitive Signal Fill Words. Possible values include IDLE-IDLE (0), ARBFF-ARBFF (1), IDLE-ARBFF (2), or AA-THEN-IA (3).

  NOTE
  For port encryption, the switch authentication policy should be **ON** or **Active** for the port. Click the **Authentication** button to activate the Authentication Mode and Shared Keys on the switch via Web Tools.

  NOTE
  To configure long distance, you can select only **Buffer Needed** or **Desired Distance** and **Frame Size**.

15. Click **Next**.
    The **Configure merge** screen displays.
16. Read and review the information on the **Configure merge** screen. If you understand and agree, click **Next** to confirm the information.

    A **Summary** screen displays.
17. Read the information, and click **Finish** to close the wizard.

Resolving merge conflicts

You can resolve the following types of switch configuration conflicts:

- Domain ID
- TOV
- Buffer To Buffer Credit
- Disable Device Probe

  NOTE
  This test will be skipped if all primary and secondary fabric switches are found to be running Fabric OS 7.0 and later.

- Route Priority Per Frame
- Sequence Level Switching
- Suppress Class F
- Long Distance Setting
- Data Field Size
- VC Priority

Note that not all tests support resolution. If a test supports resolution, the **Description** column contains the text "Resolvable".

To resolve merge conflicts, complete the following steps.

1. Select the failed test where the **Description** column contains the text "Resolvable".
2. Click **Resolve**.

   A "The switches in fabric *Name* will be disabled prior to making the configuration change. The switches will be reenabled after the configuration changes are applied. Please confirm to proceed." warning message displays.
3. Click **OK** on the warning message.

   The values of the fabric chosen on the **Set up merge options** screen are applied to all devices in the second fabric. Once the settings are applied, the test is run again and the merge results are updated.

   If the test passes, go to step 4.

   If an error occurs, an error message displays. You must use Web Tools or the CLI to resolve this conflict. Click **OK** on the error message and go to step 4.

   If you are resolving a domain ID error, there may be multiple switches involved. If multiple switches have the domain ID error, the **Configure Domain IDs** dialog box displays listing all devices that have the domain ID conflict.

   a. Select the device for which you want to resolve the domain ID in the **Available Switches** list and click the right arrow button.
   b. Select a new domain ID for the device from the **Domain ID** list.
   c. Repeat steps a and step b for each device in the **Available Switches** list.
   d. Click **OK** on the **Configure Domain IDs** dialog box.
4. Repeat step 1 through step 3 until all resolvable tests pass.
5. Perform step 12 through step 17 of the procedure "Merging two cascaded FICON fabrics" on page 929 to finish resolving a merge conflict.

Port groups

A port group is a group of FC ports from one or more switches within the same fabric. Port groups are user-specific; you can only view and manage port groups that you create.

The ports display in the order in which you add them to the port group. The order in which you add ports to a port group is persisted in both the port group and the Allow/Prohibit Matrix. While port groups can be at the fabric level (ports from multiple switches within the same fabric), the Allow/Prohibit Matrix is at the switch level. Therefore, when you view the Allow/Prohibit Matrix for a port group with ports from multiple switches, the matrix only shows the ports for the selected switch.

To reorder the ports, you must remove the ports, save your changes, then open the **Port Groups** dialog box and add the ports back to the port group in the new order.

Once you create a port group, you can view and edit the Allow/Prohibit Matrix for the port group. Allow/Prohibit Matrix is a FICON port attribute that can be used to prohibit communication between specific ports. For more information about the Allow/Prohibit Matrix, refer to "Configuring an Allow/Prohibit Matrix" on page 918.

Creating a port group

NOTE
At least one switch must be reachable to create a port group.

To create a port group, complete the following steps.

1. Select **Configure > Port Groups**.

   The **Port Groups** dialog box displays, as shown in Figure 453 on page 934.

   FIGURE 453 Port Groups dialog box

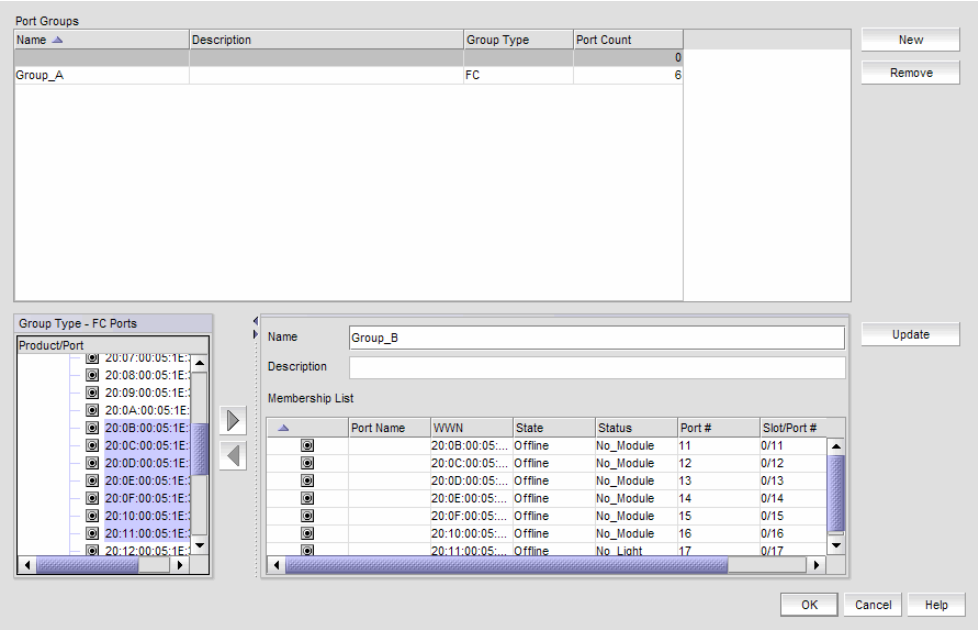

2. Click **New**.
3. Enter a name for the port group in the **Name** field.
4. Enter a description for the port group in the **Description** field.
5. Select one or more ports to add to the group in the **Group Type - FC Ports** list.

   A port group must have at least one port in the **Membership List**. All ports must be from switches in the same fabric.

6. Click the right arrow button.

   The selected ports display in the **Membership List**.

7. Click **Update**.

   The new port group displays in the **Port Groups** list.

8. Click **OK** to close the **Port Groups** dialog box.

Viewing port groups

To view port groups, complete the following steps.

1. Select **Configure > Port Groups**.

   The **Port Groups** dialog box displays only port groups defined by you.

   If a fabric becomes un-monitored, any port groups associated with that fabric do not display in the **Port Groups** list. Once the fabric becomes monitored again, the associated port groups display in the **Port Groups** list. For more information about monitoring and un-monitoring fabrics, refer to "SAN Fabric monitoring" on page 46.

If a fabric is removed from discovery, any port groups associated with that fabric are removed permanently from the **Port Groups** dialog box.

If a device is removed from a fabric, then all ports associated with that device are automatically removed permanently from the port group. If the port group only contains ports from the removed device, then the port group is removed permanently from the **Port Groups** dialog box.

If a fabric or device is added to the topology while the **Port Groups** dialog box is open, it does not display in the **Group Type - FC Ports** tree until you close and reopen the **Port Groups** dialog box.

2. Edit the port group, as needed.

   To edit a port group, refer to "Editing a port group" on page 935.

3. Delete the port group, as needed.

   To delete a port group, refer to "Deleting a port group" on page 935.

4. Click **OK**.

Editing a port group

To edit a port group, complete the following steps.

1. Select **Configure > Port Groups**.

   The **Port Groups** dialog box displays.

2. Select the port group you want to edit in the **Port Groups** list.

   The information for the selected port group displays in the update information area.

3. Change the name for the port group in the **Name** field, if necessary.

   NOTE
   If you change the port group name, it is the same as copying the existing port group with a new name.

4. Change the description for the port group in the **Description** field, if necessary.
5. Select one or more ports to add to the group in the **Group Type - FC Ports** list.
6. Click the right arrow button.

   The selected ports display in the **Membership List**.

7. Select one or more ports to remove from the group in the **Membership List**.
8. Click the left arrow button.

   The selected ports are removed from the **Membership List**.

9. Click **Update**.
10. Click **OK**.

Deleting a port group

To delete a port group, complete the following steps.

1. Select **Configure > Port Groups**.

   The **Port Groups** dialog box displays.

2. Select the port group you want to delete in the **Port Groups** list.
3. Click **Remove**.

   The selected ports are removed from the **Port Groups** list.

4. Click **OK**.

Swapping blades

> NOTE
> Blade-based port swap is mainly used for FICON and is only applicable for port blades. However, the Management application does not block blade-based port swap for other application blades, including the 8 Gbps 24-port blade.

You can swap all of the ports from one blade to another blade. During this operation, all ports in the selected blades are swapped. This operation disrupts the traffic on all ports for the selected blades. If GE_Ports are present on the blade, only the non-GE_Ports are swapped.

To swap blades, you must meet the following requirements:

- The chassis must be running Fabric OS 7.0 or later.
- You must have read and write access for the Product Administration privilege.
- The chassis must have at least two blades of the same type present.

Example

The source blade has ports sp1 and sp2, and the destination blade has ports dp1 and dp2. During the swap operation, the address sp1 is swapped with dp1 and address sp2 is swapped with dp2.

> NOTE
> To perform the swap blades function, you must have read and write access for the Product Administration privilege.

To swap blades, complete the following steps.

1. Select a chassis that contains at least two of the same type of blades.
2. Select **Configure > Switch > Swap Blades**.

   The **Swap Blades** dialog box displays.

3. Select the blade you want to replace from the first **Swap Blades** list.

   Once you select a blade, the second list automatically filters out the selected blade and any blade types that do not match the selected blade.

4. Select the blade with which you want to replace the first blade from the second **Swap Blades** list.
5. Select the **Enable ports after swap is complete** check box to enable ports on the destination blade after the swap is complete.
6. Click **OK**.

NOTE
This operation disrupts the traffic on all ports for the selected blades.

7. Click **Yes** on the confirmation message.

   Once the swap blade operation is complete, a “success” or “failure” message displays.

Deployment Manager

Introduction to the Deployment Manager

The Deployment Manager allows you to view, edit, duplicate, delete, deploy, and generate reports for the following types of deployment configurations:

- DCB
- VLAN
- STP
- Security

You cannot create configurations using the Deployment Manager. The deployment configurations must have been previously created and saved. Refer to the following sections for information about creating these types of configurations:

- "Fibre Channel over Ethernet" on page 507 (for DCB configurations)
- "VLAN Management" on page 1497 (for VLAN and STP configurations)
- "Security Management" on page 583

Deployments that were created through the legacy Configuration Wizard are listed in the **Deployment Manager** dialog box, but cannot be modified, deployed, or deleted. You can only launch reports for these deployments.

Editing a deployment configuration

1. Select **Configure > Task Scheduler**.

   The **Task Scheduler** dialog box displays, as shown in Figure 454.

FIGURE 454 Task Scheduler dialog box

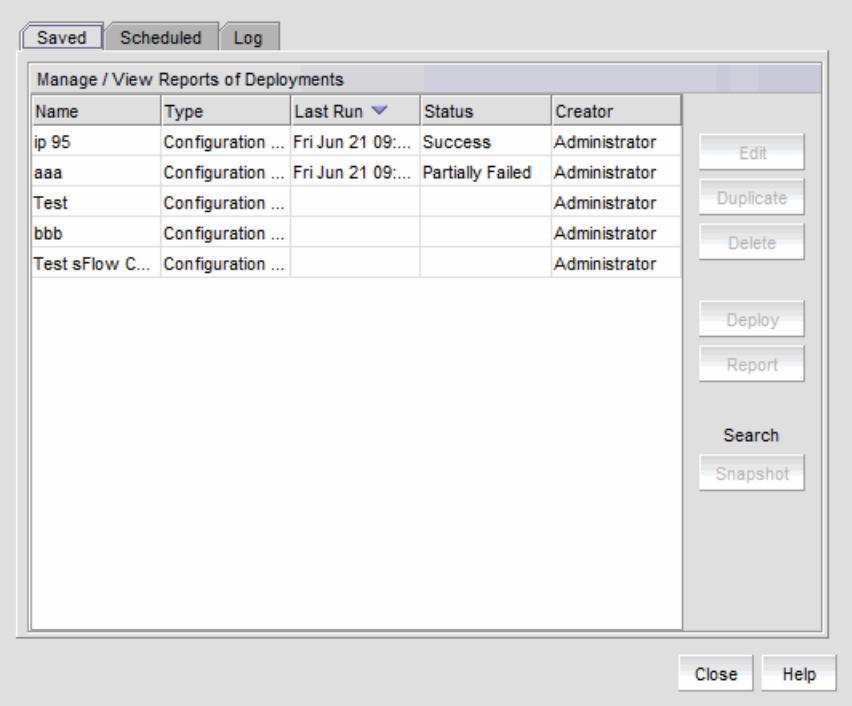

2. Select a deployment configuration in the **Saved** or **Scheduled** tab.

   Policy-based routing configurations cannot be edited.

3. Click **Edit**.

   A dialog box specific to the type of deployment displays. This is the same dialog box that was used when the deployment was created.

4. Update the dialog box with the information you want to change.

Duplicating a deployment configuration

1. Select **Configure > Task Scheduler**.

   The **Task Scheduler** dialog box displays.

2. Select a deployment configuration in the **Saved** or **Scheduled** tab.

   NOTE
   VLAN configurations and policy-based routing configurations cannot be duplicated.

3. Click **Duplicate**.

   A dialog box specific to the type of deployment displays. This is the same dialog box that was used when the original deployment was created.

4. Update the dialog box with any information you want to change.

   A copy of the deployment configuration is created with the name "***originalName* copy*n***". For example, if the original name is "test", the new name is "test copy1". If you duplicate "test" again, the name of the second duplicate is "test copy2".

Deleting a deployment configuration

1. Select **Configure > Task Scheduler**.

   The **Task Scheduler** dialog box displays.

2. Select a deployment configuration in the **Saved** or **Scheduled** tab.
3. Click **Delete**.
4. Click **Yes** in the confirmation dialog.

   The deployment configuration is deleted and removed from the **Task Scheduler** dialog box.

   If the deployment configuration is already in progress, it is not deleted.

Deploying a configuration

1. Select **Configure > Task Scheduler**.

   The **Task Scheduler** dialog box displays.

2. Select a deployment configuration in the **Saved** or **Scheduled** tab.
3. Click **Deploy**.

   The **Deployment Status** dialog box displays.

4. Click **Start**.

   The selected configuration is deployed.

   You cannot deploy configurations that are already in progress.

Viewing deployment logs

1. Select **Configure > Task Scheduler**.

   The **Task Scheduler** dialog box displays.

2. Click the **Log** tab.

   A list of deployment configurations that are executed and the status of each displays.

Generating a deployment report

1. Select **Configure > Task Scheduler**.

   The **Task Scheduler** dialog box displays.

2. Select a deployment in the **Saved**, **Scheduled**, or **Log** tab.
3. Click **Report**.

   An HTML report displays. You can click the Configuration Name or Deployment Time to see additional details.

Generating a deployment configuration snapshot report

1. Select **Configure > Task Scheduler**.

   The **Task Scheduler** dialog box displays.
2. Select a deployment in the **Saved** or **Scheduled** tab.
3. Click **Deploy**.

   The **Deployment Status** dialog box displays.
4. Click **Snapshot Report**.

   The **Configuration Snapshot Report** dialog box displays.
5. (*Optional*) If the configuration snapshot list is too long, you can filter the list.
   a. Select the start date and end date of the configuration snapshots you wish to view.
   b. Click **Find**.

      The Management application displays the list of snapshots that match the start date and end date you specified.
6. Select a product from the **Device Configuration** column to display the configuration snapshots that are available for that product.
7. Click **View** to display information for that deployment.

   The **View Pre/Post Configuration Snapshot** dialog box displays details of the selected configuration.

Searching the configuration snapshots

1. Select **Configure > Task Scheduler**.

   The **Task Scheduler** dialog box displays.
2. Select a deployment in the **Saved**, **Scheduled**, or **Log** tab.
3. Click **Snapshot**.

   The **Configuration Snapshot Search** dialog box displays.
4. Identify the targets you want to search.

   Select a target in the **Available Targets** list and click the right arrow to move the target to the **Selected Targets** list.
5. Define search criteria.

   You can specify whether the targets should contain or not contain specific text, and whether to display all configurations, the most recent configurations, or only those configurations that fall within a specific date range.

6. Click **Find**.

   The Management application displays the list of snapshots that match the search criteria you specified.

You can select configurations in the **Search Results** list to display details, view the snapshot report, and compare two configurations,

Fibre Channel Troubleshooting

FC troubleshooting

NOTE
FC troubleshooting is only available for Fabric OS devices.

You can perform the following operations using FC troubleshooting:

- Trace Route (Path Information and FC Ping) – Use to obtain the detailed routing information for any two selected device ports. The devices can exist in the same fabric or in two different fabrics shared through FC Routers.
- Device Connectivity Troubleshooting – Use to identify any problems that might be preventing communication between the two selected device ports. The device ports can be selected from the same fabric or from two different fabrics.
- Fabric Device Sharing Diagnosis (pure Fabric OS fabrics only) – Use to confirm that any two or more selected fabrics are capable of sharing devices between them.
- Diagnostic Port Testing (Fabric OS 10 Gbps-capable, 16 Gbps-capable ports, and 32 Gbps-capable ports only) – Use to run the following diagnostic port tests on the 10 Gbps-capable, 16 Gbps-capable ports, and 32 Gbps-capable ports: electrical (32 Gbps only), optical (32 Gbps only), measure link distance, and link traffic (16 Gbps and 32 Gbps only).

  NOTE
  Link Traffic Test is not supported for ports on Qlogic HBA devices with Firmware 8.0.0 or later.

Tracing FC routes

The Management application enables you to select a source port and a destination port and displays the detailed routing information from the source port or area on the local switch to the destination port or area on another switch.

NOTE
Trace route cannot be performed on offline devices.

NOTE
Trace route cannot be performed in a mixed (Fabric OS) fabric.

Fabric OS trace route requirements

- Fabric OS trace route is only supported in a pure-Fabric OS fabric.
- All Fabric OS switches in the fabric must be running Fabric OS 7.0 or later.

To trace routes, complete the following steps.

1. Select **Monitor > Troubleshooting > FC > FC Trace Route**.

   The **Trace Route** dialog box displays.

2. Choose from one of the following options:
   - Select a fabric from the **Fabric** list.

- Select a router from the **Routing** list. Requires Fabric OS 7.0 or later.

3. Select the source and destination ports by choosing one of the following:

   The source and destination ports must be on the same fabric; however, they cannot be connected to the same switch.

   - To enter the ports, select the **Enter port FC Address** option.
     a. Enter the source port FC address in the **Source** field.
     b. Enter the destination port FC address in the **Destination** field.
   - To select the ports, select the **Select two device ports** option.
     a. Right-click a fabric in the **Available Device Ports** table and select **Expand All**.
     b. Select the ports (two) for which you want to display the detailed routing information from the **Available Device Ports** table.

4. Click the right arrow button.
5. Click **OK**.

   The **Trace Route Summary** dialog box displays. This dialog box includes the following information:

   - **Trace Route Summary** — This table shows a brief summary of the trace including the following:
     - Port WWN
     - Port Name
     - FC Address
     - Switch Name
     - (Fabric OS only) Whether ping was successful (Fabric OS only)
     - (Fabric OS only) Round trip time (minimum, maximum, and average)
     - (Fabric OS only) Whether the device ports are in active zones.
   - **Forward Route** — This tab shows the path taken by data packets from the port belonging to the switch on which the trace route has been invoked (source port) to the port on the other switch (destination port). This tab includes the following path details:
     - Hop
     - In Port Address
     - In Port Slot/Port
     - Domain ID
     - Switch Name
     - Out Port Address
     - Out Port Slot/Port
     - Bandwidth (Gb/sec)
     - Cost
   - (Fabric OS only) **Reverse Route** — This tab shows the path from the destination port to the source port. This tab contains the same path details as the **Forward Route** tab.

     NOTE
     The reverse route may sometimes be different from the forward route.

   - (Fabric OS only) **FC Ping** — This tab shows the minimum, maximum and average round trip times between the selected device port WWNs and the domain controller. It details whether the selected device port WWNs are zoned or not. It also shows the number of frames sent to the device port, frames rejected, frames timed-out and frames received by the device port.
   - (Fabric OS only) **Add Flow** button — Click to launch the **Add Flow Definition** dialog box. For more information about Flow Vision, refer to "Flow Vision" on page 1013.

6. Click **Close** on the **Trace Route Summary** dialog box.

7. Click **Cancel** on the **Trace Route** dialog box.

"Confirming Fabric Device Sharing"

Troubleshooting device connectivity

To troubleshoot device connectivity, complete the following steps.

1. Select **Monitor > Troubleshooting > FC > Device Connectivity**.

   The **Device Connectivity Troubleshooting** dialog box displays.

2. Select the source and destination ports on which you want to troubleshoot device connectivity using one of the following options:

   - Enter the source and destination ports directly by selecting the **Enter port FC Address** option and completing the following steps.
     a. Enter the source port in the **Source** field.
     b. Enter the destination port in the **Destination** field.
     c. Click **Search and Add**.
   - Select the source and destination ports from a list by selecting the **Select two device ports** option and completing the following steps.
     a. Right-click a fabric in the **Available Device Ports** table and select **Expand All**.
     b. Select the ports (source and destination) for which you want to confirm device sharing from the **Available Device Ports** table.
     To add a detached device to troubleshoot device connectivity, refer to "Adding a detached device" on page 948.
     c. Click the right arrow button.

3. Click **OK**.

   The following diagnostic tests are performed:

   - Device Status
   - Switch port health status
   - Zone configuration in the fabric
   - LSAN zone configuration in edge fabrics
   - Edge fabric - FC router physical connection status.
   - Active ACL DCC policy check (Fabric OS only)

   The **Device Connectivity Troubleshooting Results** dialog box displays.

   If no problems are found, the diagnostic test is marked with a check mark. If problems are found, an alert icon appears next to the test, with a brief statement detailing the error as well as a suggested resolution.

4. Click **Re-run Diagnosis** to run the device connectivity on the same ports.
5. Click **Trace Route** to trace the route between the two selected ports.
6. Click **Close** on the **Device Connectivity Troubleshooting Results** dialog box.

Adding a detached device

To add a detached device to the **Selected Device Ports** table, complete the following steps.

1. Select **Monitor > Troubleshooting > FC > Device Connectivity**.

   The **Device Connectivity Troubleshooting** dialog box displays.

2. Click **Add Detached**.
3. Enter the port WWN of the detached device port in the **Port WWN** field.
4. Click **OK**.

Confirming Fabric Device Sharing

NOTE
Fabric device sharing is only available with Trial or Licensed version.

NOTE
Fabric device sharing is only available on pure Fabric OS fabrics.

To confirm that two or more fabrics have been configured to share devices, complete the following steps.

1. Select **Monitor > Troubleshooting > FC > Fabric Device Sharing**.

   The **Fabric Device Sharing Diagnosis** dialog box displays.

2. Select the fabrics (two or more) for which you want to confirm device sharing from the **Available Fabrics** table.
3. Click the right arrow button.
4. Click **OK**.

   The following checks are performed on the selected fabrics:

   - Are the selected fabrics configured with an FC Router?
   - Are the selected fabrics connected to the same backbone fabric?
   - Is sharing of devices between backbone and edge fabric supported?

   The **Fabric Device Sharing Diagnosis Results** dialog box displays with the details of the fabrics selected for diagnosis, the details of the tests performed, the results of the test, as well as short description of the test results.

5. Click **Close** on the **Fabric Device Sharing Diagnosis Results** dialog box.
6. Click **Cancel** on the **Fabric Device Sharing Diagnosis** dialog box.

Troubleshooting port diagnostics

This feature allows you to run a diagnostic port test and a link traffic test on the selected ports.

Port diagnostics requirements

- Only supported on devices with 8 Gbps-capable E-ports, ICL-ports, and AG N-ports running Fabric OS 7.2.1 or later.
- Only supported on devices with 10 Gbps-capable D-ports or E-ports running Fabric OS 7.0 or later. The source and destination ports must be the same.
- Only supported on devices with 16 Gbps-capable E-ports running Fabric OS 7.0 or later and with 32 Gbps-capable E-ports and F-ports running Fabric OS 8.0 or later.
- Only supported on devices with 16 Gbps-capable F-ports, ICL-ports, and AG N-ports running Fabric OS 7.1 or later.
- Only supported on devices with 32 Gbps-capable E-ports, F-ports, and AG N-ports running Fabric OS 8.0 or later.
- Beginning with Fabric OS 8.0, D-Port test is supported on QLogic HBA devices.
- Both the source and destination ports must be managed by the Management application.

ATTENTION
The Management application changes the port type for all selected ports and associated attached ports to a D port for the duration of the test. This may cause the fabric to segment. When the test is complete, the Management application changes the port type back to an E port.

ATTENTION
If you run more than one test per slot, the result may go wrong or the test may fail.

NOTE
Electrical and Optical loopback tests are not supported on QSFP except for 32 Gbps-capable QSFPs, that support E_WRAP and O_WRAP. To support Electrical and Optical loopback test, both the local and remote end must be running on the Fabric OS v8.1.0a and later.

TABLE 73 D-Port test support matrix

D-Ports Tests		Fabric OS 7.0	Fabric OS 7.1				HBA driver 3.2	QLogicHBA v8.0.0
		E-Port	E-Port	F-Port	AG N-Port	ICL-Port		D-Port
Electrical Test		Supported	Supported	Supported	Supported	Not supported	Supported	Supported
Optical Test		Supported	Supported	Supported	Supported	Not supported	Not supported	Supported
Link Traffic Test	Configuration	Not supported	Supported	Supported	Supported	Supported	Not supported	Not supported
	Test	Supported	Supported	Supported	Supported	Supported	Not supported	Not supported
Link Distance		Supported	Supported	Supported	Supported	Supported	Not supported	Not supported
Link Measurement		Not supported	Not supported	Supported	Supported	Supported	Not supported	Not supported
Dense Wave Division Multiple-xing	Configuration	Not supported	Not supported	Supported	Not supported	Not supported	Not supported	Supported
	Test	Not supported	Not supported	Supported	Not supported	Not supported	Not supported	Supported

To run a diagnostic port test, complete the following steps.

1. Select **Monitor > Troubleshooting > FC > Diagnostic Port Test**.

   The **Diagnostic Port Test** dialog box displays.

   FIGURE 455 Diagnostic Port Test dialog box

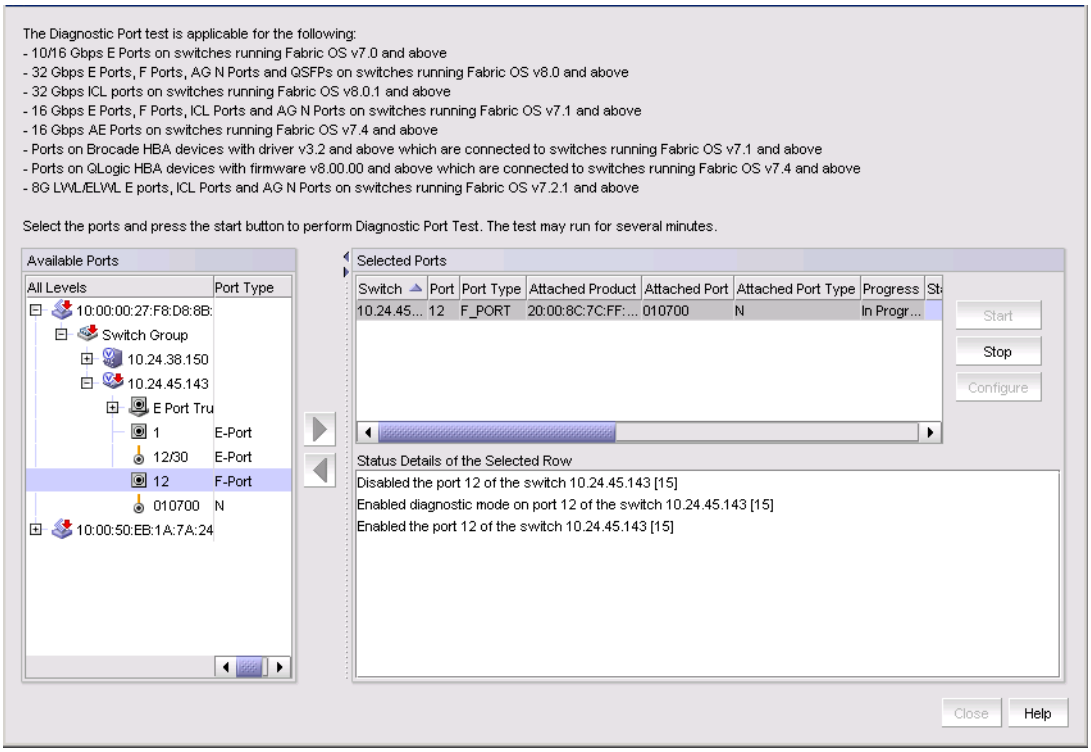

2. Select the ports or Qlogic devices for which you want to run a diagnostic port test from the **Available Ports** table.

   You can only run 10 diagnostic port tests at a time. If you select more than 10 ports, the Management application runs the first 10 diagnostic port tests and queues the rest. When the first test is completed, the next test in the queue begins and so on until all tests are completed.

3. Click the right arrow button to move the ports to the **Selected Ports** table.

4. To configure parameters for the link traffic test, select the ports for which you want to configure link traffic test in the **Selected Ports** table and click the **Configure** button.

   Refer to "Configuring test configuration parameters" on page 953.

5. Click **Start**.

   The Management application performs the following operations to enable diagnostic mode on the selected ports:

   1. Disable the source port.
   2. Disable the destination port.
   3. Enable the diagnostic mode on source E-port.
   4. Enable the diagnostic mode on destination E-port.

5. Enable the source port.
6. Enable the destination port.

The following tests are performed on the selected ports:

- Electrical loopback (32 Gbps)
- Optical loopback (32 Gbps)
- Link traffic (16 Gbps and 32 Gbps)
- Latency measurement
- Measure link distance

TABLE 74 Supported link distance measurements

SFP speed	Accuracy	Precision
10 Gbps	124 meters	+ or - 50meters
16 Gbps	5 meters	+ or - 5 meters
32 Gbps	3 meters	+ or - 3meters

If any of the tests fail, the Management application does not rollback to already executed operations.

When the test successfully completes, the Management application performs the following operations to change the port type back to E-port:

1. Disable the source port.
2. Disable the destination port.
3. Disable the diagnostic mode on source D-port.
4. Disable the diagnostic mode on destination D-port.
5. Enable the source port.
6. Enable the destination port.

The **Progress** column shows whether the test is not started, in progress, or completed.

The **Status** column shows the overall status (Success or Failed) of the test.

Or

6. Click **Stop** to stop the diagnostic port test. Only one port can be selected.

   The Management application performs the following operations to disable diagnostic mode on the selected ports:

   - **Status Details of the selected Row** section displays **Dport Test Stopped**.
   - Port state is reverted to its previous state

7. Select a port row in the Selected Ports table to display the detailed status in the **Status Details of the Selected Row** table. Select **Configure** > **Task Scheduler** > **Logs** tab if you want to view and generate the detailed status as a report.

   The **Status Details of the Selected Row** table displays with the details of the port selected for diagnosis, the details of the tests performed, the results of the test, as well as short description of the test results. The following table details the messages that display depending on the success or failure of the operations and tests.

TABLE 75 Status Detail messages

Operation/Test	Possible message
Disable the source or destination port	Disabled the port *slot_number*/*port_number* of the switch *switch_IP_address*.
	Failed to disable the port *slot_number*/*port_number* of the switch *switch_IP_address*. Reason: *CAL_error_message*
Enable the diagnostic mode on source or destination E ports	Enabled diagnostic mode on port *slot_number*/*port_number* of the switch *switch_IP_address.*
	Failed to enable diagnostic mode on port *slot_number*/*port_number* of the switch *switch_IP_address*. Reason: *CAL_error_message*
Enable the source or destination port	Enabled the port *slot_number*/*port_number* of the switch *switch_IP_address*.
	Failed to enable the port *slot_number*/*port_number* of the switch *switch_IP_address*. Reason: *CAL_error_message*
Disable the diagnostic mode on source or destination D ports	Disabled diagnostic mode on port *slot_number*/*port_number* of the switch *switch_IP_address*.
	Failed to disable diagnostic mode on port *slot_number*/*port_number* of the switch *switch_IP_address*. Reason: *CAL_error_message*
Lost connectivity to switch while test is in progress	Connection failed to the switch during the operation.
Diagnostic port test timed out The Management application waits 30 minutes to complete the test. If not completed, the test times out.	Diagnostic port test time-out. You may need to change the port configuration before retrying.
Electrical Loopback Test	Successfully completed Electrical Loopback Test on port *slot_number*/*port_number* of the switch *switch_IP_address*.
	Electrical Loopback Test failed on port *slot_number*/*port_number* of the switch *switch_IP_address.* Reason: *CAL_error_message*
	Electrical Loopback Test skipped on port *slot_number*/*port_number* of the switch *switch_IP_address*. Reason: *CAL_error_message*
Optical Loopback Test	Successfully completed Optical Loopback Test.
	Optical Loopback Test failed
	Optical Loopback Test skipped. Reason: *CAL_error_message* If configure on a DWDM link, the optical loopback test is skipped
Link Traffic Test	Successfully completed Link Traffic Test.
	Link Traffic Test failed.
Distance between ports	Approximate distance between the ports is *numerical_value* meters.
Reverse Optical Loopback Test	Successfully completed Reverse Optical Loopback Test.
	Reverse Optical Loopback Test failed.
Roundtrip link latency	Roundtrip link latency: *numerical_value* nano-seconds.
Buffers required	Buffers required: *numerical_value*

TABLE 75 Status Detail messages

Operation/Test	Possible message
Link Traffic Test Configuration used	No. of test frames: *numerical_value* Million Duration of test (HH:MM): hours:minutes Test frame size: *numerical_value* Bytes Payload Pattern: *pattern_name* or *frame_data* FEC (enabled/active): Yes/No CR (enabled/active): Yes/No Start time: day month date hour:minute:seconds year End time: day month date hour:minute:seconds year **NOTE:** If you do not set a payload pattern, results do not show payload pattern.
If any test fails, that test displays as failed and a Failure report displays.	Sample failure report : Errors detected (local): CRC, Bad_EOF, Enc_out Errors detected (remote): CRC, Bad_EO Run portstatsshow and porterrshow for more detail on the errors.
HBA Electrical test successful	Successfully completed Electrical Loopback Test on port *HBA_port_number* of the HBA *HBA_node*
HBA Electrical test failed	Electrical Loopback Test failed on port *HBA_port_number* of the HBA *HBA_node*
Stop D-port Test	DPort test stopped

8. Click **Close** on the **Diagnostic Port Test** dialog box.

Configuring test configuration parameters

You can configure these parameters to improve accuracy and to measure link performance on stress conditions with more traffic.

1. Enter the number of frames to use in the traffic test in the **Number of Frames** field.

   Minimum is 1 million (default). Maximum is 2,147,483,647 million.

2. Enter the duration you want the test to run in the **Duration Hours** and **Minutes** fields.

   This option is mutually exclusive of frames option.

   Minimum is 0 hours and 1 minute (default). Maximum is 99 hours and 59 minutes.

3. Enter the frame size you want for the traffic test in the **Frame Size** field.

   Minimum is 36 bytes. Maximum is 2236 bytes. Default is 1024 bytes.

4. Choose one of the following options in the **Payload Pattern** area to configure the payload pattern to use in the traffic test

   - Select the **Predefined** option and select a pre-defined payload patteren from the list.

     Options include BYTE_NOT, WORD_NOT, QUAD_NOT, BYTE_RAMP, WORD_RAMP, QUAD_RAMP, BYTE_LFSR, RANDOM, CRPAT, CSPAT, CHALF_SQ, CQTR_SQ, RDRAM_PAT, jCRPAT, jCJTPAT, jCSPAT, PRED_RAND, SMI_TEST, CJPAT, QUAD_NOTP, JSPAT, and JTSPAT.

   - Select the **User Defined** option and enter a pattern in the field.

     Minimum is 0. Maximum is 2,147,483,647.

   An example of the payload pattern displays in the **Example** field.

5. Select the **Forward Error Correction - Enable** check box to enable forward error correction (FEC) during the D-Port test.

   Clear to disable (default).

6. Select the **Credit Recovery - Enable** check box to enable credit recovery (CR) during the D-Port test.

   Clear to disable (default).

7. Select **Optical Loopback test** check box to enable DWDM and skip the test for E-Ports.

   Clear to disable (default).

8. Click **OK.**

   The **Diagnostic Port Test** dialog box displays. Return to step 5 of "Troubleshooting port diagnostics" on page 949.

   NOTE
   Beginning with Fabric OS 7.4.0, MAPS is used to monitor D-Port test and threshold will be configured for monitoring errors.

   NOTE
   Beginning with Fabric OS 7.4.0, during the D-Port test, a user can identify the RX and TX power loss occurring in the network path link.

FCIP troubleshooting

NOTE
FCIP troubleshooting is only available for Fabric OS devices.

You can perform the following operations using FCIP troubleshooting:

- **Ping.** Use to confirm that the configured FCIP tunnels are working correctly.
- **Trace Route.** Use to view the route information from a source port on the local device to a destination port on another device and determine where connectivity is broken.
- **Performance.** Select to view FCIP tunnel performance between two devices.

Configuring IP ping

NOTE
IP Ping only supported on Fabric OS devices.

NOTE
IP Perf is not supported on the Fabric OS 8 Gbps Extension Switch or Blade.

You can also verify IP connectivity when configuring an FCIP circuit. For more information, refer to "Adding an FCIP circuit".

To configure IP ping, complete the following steps.

1. Select **Monitor > Troubleshooting > FCIP > Ping**.

   The **IP Ping** dialog box displays.

2. Select a switch from the **Available Switches** table.
3. Select a port from the **GigE Port** list.

4. Select an IP address switch from the **IP Interface** list.
5. Enter the remote IP address in the **Remote IP Address** field.
6. Click **OK**.

   Ping sends four Internet Control Message Protocol (ICMP) Ping packets to the destination address and records the time until a response.

   The **IP Ping Result** dialog box displays with two tables.

   The top table (**FCIP IP Ping Response Details**) contains the following statistics:

   TABLE 76 FCIP IP Ping Response Details

Field or Component	Description
Status	Always displays 'Completed'. If there is a failure, an error message displays instead of the **IP Ping Result** dialog box.
Packets Sent	Always displays '4. This is not configurable.
Packets Received	The number of received responses.
Packets Lost	Equal to the number of packets sent minus the number of packets received.
Packet Lost percentage	The number of packets lost expressed as a percentage of the packets sent. This will be 0%, 25%, 50%, 75% or 100% for 0, 1, 2, 3, or all 4 packets lost.
Minimum Round Trip Time	The shortest time, in milliseconds, of any response. If no response, the round trip times is 0.
Maximum Round Trip Time	The longest time, in milliseconds, of any response. If no response, the round trip times is 0.
Average Round Trip Time	The average time, in milliseconds, of all responses. If no response, the round trip times is 0.

   The bottom table (**IP Ping Details**) provides details for each ping attempt.

   TABLE 77 IP Ping Details

Field or Component	Description
Reply From	The IP address of the device that sent the reply. For a normal response, this is the destination IP address. Some error responses (such as "destination unreachable") may come from an intermediate router.
Status	Displays either Success or an error message (such as request timed out or destination unreachable) from the switch.
Number of bytes	The number of bytes in the data portion of the response. Should be 64, matching the 64 bytes of data sent in the transmitted packet.
Round Trip Time (ms)	The time in milliseconds between sending the packet and receiving the response. This provides a rough indication of network congestion or latency. It is normal for the first packet to experience a higher round trip time than later packets, if the intermediate routers need to do ARP requests to locate the next hop.
Time To Live (hops)	The number of hops remaining in the received response. The time to live is decremented by each router that forwards the packet. The packet is dropped if the time to live reaches zero.

7. Click **Close** on the **IP Ping Result** dialog box.
8. Click **Cancel** on the **IP Ping** dialog box.

Tracing IP routes

The Management application enables you to select an source and a target and displays the detailed routing information from the source port or area on the local switch to the destination port or area on another switch.

Trace route cannot be performed on the offline devices or virtual devices.

To trace routes, complete the following steps.

1. Select **Monitor > Troubleshooting > FCIP > Trace Route**.

   The **IP Traceroute** dialog box displays.

2. Select a switch from the **Available Switches** table.
3. Select a port from the **GigE Port** list.
4. Select an IP address switch from the **IP Interface** list.
5. Enter the remote IP address in the **Remote IP Address** field.
6. Click **OK**.

   The **IP Traceroute Result** dialog box displays.

   Traceroute sends three ICMP Ping packets to the destination address with a time to live (TTL) of one hop, and expects a 'TTL Expired' error back from the first router to obtain the IP address of the first hop. Traceroute then repeats the operation with a TTL of two hops to get the IP address of the second hop. This process repeats for up to ten hops, or until a successful PING response is received.

   The IP Trace Details table displays the results of each attempt.

   TABLE 78 IP Trace Details

Field or Component	Description
Hop Number	The TTL inserted in the transmitted probe packet.
IP Address 1	The IP address of the system that responded to the first of the three probes, or 0.0.0.0 if there was no response.
IP Address 2	The IP address of the system that responded to the second of the three probes, or 0.0.0.0 if there was no response.
IP Address 3	The IP address of the system that responded to the third of the three probes, or 0.0.0.0 if there was no response.
RTT 1	The time in milliseconds for the first of the three responses to be received, or blank if there was no response. This value helps identify a congested or slow link in the path.
RTT 2	the time in milliseconds for the second of the three responses to be received, or blank if there was no response. This value helps identify a congested or slow link in the path.
RTT 3	the time in milliseconds for the third of the three responses to be received, or blank if there was no response. This value helps identify a congested or slow link in the path.

7. Click **Close** on the **IP Traceroute Result** dialog box.
8. Click **Cancel** on the **IP Traceroute** dialog box.

Viewing FCIP tunnel performance

NOTE
IP Performance is only supported on the 4 Gbps Router, Extension Switch and Encryption Blade.

NOTE
If you run IP Performance over a link also being used for production traffic, it will impact the production traffic performance.

To view FCIP tunnel performance, complete the following steps.

1. Select **Monitor > Troubleshooting > FCIP > Performance**.

   The **IP Performance** dialog box displays.

2. Select a switch from the **Available Switches** table.
3. Select a port from the **GigE Port** list.
4. Select an IP address from the **IP Interface** list.
5. Enter the remote IP address in the **Remote IP Address** field.
6. Click **OK**.

   The **IP Performance Result** dialog box displays.

   IP Performance sends dummy data as fast as possible to the remote IP address and measures how much data can be sent over a given interval. IP Performance attempts to saturate the network link to see how much bandwidth is available. It will display the media link bandwidth only if no other traffic is flowing. The remote IP address must belong to a managed switch so that IP Performance can set up the receiving end on the remote switch.

   For more information about IP Performance, refer to Chapter 20 in the *Fabric OS Administrator's Guide*.

   During the IP Performance test, data is sent continuously and statistics are sampled every 30 seconds. At the end of the period, the IP Performance results dialog box displays. The IP Performance results dialog contains a table with one row for each 30-second sample of the test. Columns in the perf results dialog are:

TABLE 79

Field/Component	Description
Available Bandwidth	The average bytes per second sent during the sample interval. This is a count of FC payload bytes; for example, the throughput seen by an FC application. It is slightly lower than the actual bytes-per-second on the wire since it does not include headers and acknowledgments.
Weighted Bandwidth	The weighted bandwidth represents what the FCIP tunnel / FC application sees for throughput rather than the Ethernet on-the-wire bytes.
Loss Percent	An estimate of the percentage of data packets lost during the sampling interval, based on TCP re-transmits.
DELAY	The average round trip time to send a packet of data and receive the acknowledgment.
PMTU (Path Maximum Transmission Unit)	The largest packet size that can be transmitted over the end path without fragmentation. This value is measured in bytes and includes the IP header and payload. IP Performance tries the configured Fabric OS Jumbo MTU value (anything over 15000, then 1500, then 1260. The value displayed in the table is the largest value that worked.

7. Click **Close** on the **IP Performance Result** dialog box.
8. Click **Cancel** on the **IP Performance** dialog box.

Performance Data

SAN performance overview

Performance monitoring provides details about the quantity of traffic and errors that a specific port or device generates on the fabric over a specific time. You can also use performance monitoring to indicate the devices that create the most traffic and identify the ports that are most congested.

Performance monitoring allows you to monitor your SAN using the following methods (requires a Licensed version):

- Display the connections which are using the most bandwidth on the selected device or one of the F_Ports on the device with a feature called Top Talkers.
- Gather and display real-time performance data (Switch Ports - FC, Switch Ports - GE, Switch Ports - 10 GE, ISL Ports, E_Port Trunks, end-to-end Monitors, FCIP Tunnels, and device Ports).

  The Professional version only allows you to monitor your SAN by gathering and displaying real-time performance data (Switch Ports - FC, Switch Ports - GE, Switch Ports - 10GE, ISL Ports, E_Port Trunks, end-to-end Monitors, FCIP Tunnels, and device Ports).
- Persist and display historical performance data (Switch Ports - FC ports, ISL ports, device Ports, FCIP tunnels, SFP, and Switch Ports - 10 GE Ports) for selected fabrics or the entire SAN.
- Create custom port and time data filters for historical performance data that can be saved as a favorite.
- Support end-to-end monitors for real-time and historical performance data.
- Enforce user-defined performance thresholds and notification when thresholds are exceeded.
- Display “insufficient resources” message when the is busy and you request performance statistics.
- Display percentage utilization for FC and FCIP links.
- Select a granularity for collecting data:
  - 5 minutes for last 8 days.
  - 30 minutes granularity for last 30 days
  - 2 hour granularity for last 30 days
  - 1 day granularity for last 730 days.
- Provide enhanced performance reports.

SAN performance measures

Performance measures enable you to select one or more measures to define the graph or report. The measures available to you depend on the object type from which you want to gather performance data.

NOTE
Devices with 10GE ports must be running Fabric OS 7.0 or later to obtain the correct TE_Port statistics (TX/RX).

NOTE
Devices with 10GE ports must have the RMON MIB enabled on the switch. For more information about the **rmon collection** command, refer to the ***Fabric OS Converged Enhanced Ethernet Command Reference***.

You can define a report or graph for the following performance data:

- Current — Available in mAmps for installed SFPs.
- Rx Power — Available in dBm for installed SFPs.
- Tx Power — Available in dBm for installed SFPs.
- Temperature — Available in Centigrade for installed SFPs.
- Voltage — Available in mVolts for installed SFPs.
- Tx % Utilization — Available for FC, GE, E_Port trunks, 10GE ports, and FCIP tunnels.
- Rx % Utilization — Available for FC, GE, 10GE ports, E_port trunks, and FCIP tunnels.
- Tx MB/Sec — Available for FC, GE, managed HBA ports, managed CNA ports, 10GE ports, E_Port trunks, FCIP tunnels, and end-to-end monitors.
- Rx MB/Sec — Available for FC, GE, managed HBA ports, managed CNA ports, 10GE ports, E_port trunks, FCIP tunnels, and end-to-end monitors.
- CRC Errors — Available for FC, managed HBA ports, managed CNA ports, 10GE ports and end-to-end monitors.
- Signal Losses — Available for managed HBA ports, managed CNA ports, and FC ports.
- Sync Losses — Available for managed HBA ports, managed CNA ports, and FC ports.
- Link Failures — Available for managed HBA ports, managed CNA ports, and FC ports.
- Sequence Errors — Available for FC ports.
- Invalid Transmissions — Available for FC ports.
- Rx Link Resets — Available for FC ports.
- Tx Link Resets — Available for FC ports.
- C3 Discard — Available for FC ports.
- C3 Discard RX Timeout — Available for FC ports
- C3 Discard Tx Timeout — Available for FC ports
- C3 Discard Unreachable — Available for FC ports.
- C3 Discard Others — Available for FC ports.
- Encode Error out — Available for FC ports.
- BB Credit Zero — Available for FC ports.
- Truncated Frames — Available for FC ports.
- FEC Corrected Blocks — Available for FC ports.

- FEC Uncorrected Blocks — Available for FC ports.
- PCS Error Block — Available for FC ports.
- Dropped Packets — Available for FCIP tunnels only.
- Cumulative Compression Ratio — Available for FCIP tunnels only.
- Current Compression Ratio — Available for FCIP tunnels only.
- Latency — Available for FCIP tunnels only.
- Link Retransmits — Available for FCIP tunnels only.
- Timeout Retransmits — Available for FCIP tunnels only.
- Fast Retransmits — Available for FCIP tunnels only.
- Duplicate Ack Received — Available for FCIP tunnels only.
- Window Size RTT — Available for FCIP tunnels only.
- TCP Out of Order Segments — Available for FCIP tunnels only.
- Slow Start Status — Available for FCIP tunnels only.
- Uncompressed Tx/Rx MB/sec - Available for FCIP tunnels only.
- Overflow Errors — Available for 10GE ports only.
- Runtime Errors — Available for 10GE ports only.
- Receive EOF — Available for 10GE ports only.
- Too Long Errors — Available for 10GE ports only.
- Underflow Errors — Available for 10GE ports only.
- Alignment Errors — Available for 10GE ports only.
- NOS Count — Available for managed HBA ports and managed CNA ports.
- Error Frames — Available for managed HBA ports and managed CNA ports.
- Under Sized Frames — Available for managed HBA ports and managed CNA ports.
- Over Sized Frames — Available for managed HBA ports and managed CNA ports.
- Primitive Sequence Protocol Errors — Available for managed HBA ports and managed CNA ports.
- Dropped Frames — Available for managed HBA ports and managed CNA ports.
- Bad EOF Frames — Available for managed HBA ports and managed CNA ports.
- Invalid Ordered Sets — Available for managed HBA ports, managed CNA ports, and FC ports.
- Non Frame Coding Error — Available for managed HBA ports and managed CNA ports.

NOTE

Under Sized Frames, Over Sized Frames, Bad EOF Frames, Dropped Frames, and Non Frame Coding Error measures are not supported for managed HBA and CNA ports on Emulex Adapters.

SAN performance management requirements

To collect performance data, make sure the following requirements have been met:

- Make sure the SNMP access control list for the device is empty or the Management application server IP is in the access control list.

  Example of default access control list

  ```
  FCRRouter:admin> snmpconfig --show accesscontrol
  SNMP access list configuration:
  Entry 0:  No access host configured yet
  Entry 1:  No access host configured yet
  Entry 2:  No access host configured yet
  Entry 3:  No access host configured yet
  Entry 4:  No access host configured yet
  Entry 5:  No access host configured yet
  ```

  Example of Management application Server IP address included in access control list

  ```
  FCRRouter:admin> snmpconfig --show accesscontrol
  SNMP access list configuration:
  Entry 0:  Access host subnet area 172.26.1.86 (rw)
  Entry 1:  No access host configured yet
  Entry 2:  No access host configured yet
  Entry 3:  No access host configured yet
  Entry 4:  No access host configured yet
  Entry 5:  No access host configured yet
  ```

  To add the Management application server IP address to the access control list, use the **snmpconfig --add accesscontrol** command.

  To set the default access control, use the **snmpconfig --default accesscontrol** command.

- Make sure that the SNMP credentials in the Management application match the SNMP credentials on the device.
  - To check the SNMP v1 credentials on the device, use the **snmpconfig --show snmpv1** command.

    Example of SNMP v1

    ```
    HCLSwitch:admin> snmpconfig --show snmpv1
    SNMPv1 community and trap recipient configuration:
    Community 1: Secret C0de (rw)
    Trap recipient: 10.103.4.63
    Trap port: 162
    Trap recipient Severity level: 4
    Community 2: OrigEquipMfr (rw)
    Trap recipient: 10.1.12.240
    Trap port: 162
    Trap recipient Severity level: 4
    Community 3: private (rw)
    Trap recipient: 10.103.5.105
    Trap port: 162
    Trap recipient Severity level: 4
    Community 4: public (ro)
    Trap recipient: 2.168.102.41
    Trap port: 162
    Trap recipient Severity level: 4
    Community 5: common (ro)
    Trap recipient: 10.32.150.116
    Trap port: 162
    Trap recipient Severity level: 4
    ```

```
Community 6: FibreChannel (ro)
Trap recipient: 1001:0:0:0:0:0:0:172
Trap port: 162
Trap recipient Severity level: 4
```

- To set the SNMP v1 credentials on the device, use the **snmpconfig --set snmpv1** command.

  Example of setting SNMP v1

```
HCLSwitch:admin> snmpconfig --set snmpv1
SNMP community and trap recipient configuration:
Community (rw): [test]
Trap Recipient's IP address : [172.26.1.183]
Trap recipient Severity level : (0..5) [4]
Trap recipient Port : (0..65535) [162]
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address : [172.26.24.26]
Trap recipient Severity level : (0..5) [4]
Trap recipient Port : (0..65535) [162]
Community (rw): [custom]
Trap Recipient's IP address : [172.26.1.158]
Trap recipient Severity level : (0..5) [4]
Trap recipient Port : (0..65535) [162]
Community (ro): [custom]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [common]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address : [172.26.1.145]
Trap recipient Severity level : (0..5) [4]
Trap recipient Port : (0..65535) [162]
```

- To check the SNMP v3 credentials on the device, use the **snmpconfig --show snmpv3** command.

  Example of SNMP v3

```
sw1:FID128:admin> snmpconfig --show snmpv3
SNMPv3 USM configuration:
User 1 (rw): snmpadmin1
Auth Protocol: noAuth
Priv Protocol: noPriv
User 2 (rw): snmpadmin2
Auth Protocol: noAuth
Priv Protocol: noPriv
User 3 (rw): snmpadmin3
Auth Protocol: noAuth
Priv Protocol: noPriv
User 4 (ro): snmpuser1
Auth Protocol: noAuth
Priv Protocol: noPriv
User 5 (ro): snmpuser2
Auth Protocol: noAuth
Priv Protocol: noPriv
User 6 (ro): admin
Auth Protocol: noAuth
Priv Protocol: noPriv
```

- To set the SNMP v3 credentials on the device, use the **snmpconfig --set snmpv3** command.

```
FM_4100_21:admin> snmpconfig --set snmpv3
SNMPv3 user configuration(SNMP users not configured in Fabric OS user database will have
physical AD and admin role as the default):
User (rw): [snmpadmin1] admin
```

```
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 1
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol [DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES2(5)/AES256(6)]): (1..6) [2] 1
New Priv Passwd:
Verify Priv Passwd:
User (rw): [snmpadmin2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES2(5)/AES256(6)]): (2..2) [2]
User (rw): [snmpadmin3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES2(5)/AES256(6)]): (2..2) [2]
User (ro): [snmpuser1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES2(5)/AES256(6)]): (2..2) [2]
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES2(5)/AES256(6)]): (2..2) [2]
User (ro): [snmpuser3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/3DES(3)/AES128(4)/AES2(5)/AES256(6)]): (2..2) [2]
SNMPv3 trap recipient configuration:
Trap Recipient's IP address : [2.168.71.32]
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [4]
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [1.1.1.1]
UserIndex: (1..6) [2]
Trap recipient Severity level : (0..5) [4]
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [10.64.209.171]
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [4]
Trap recipient Port : (0..65535) [162]
Trap Recipient's IP address : [0.0.0.0]
Trap Recipient's IP address : [0.0.0.0]
Trap Recipient's IP address : [0.0.0.0]
```

- To check SNMP credentials in the Management application, complete the following steps.

    1. Select **Discover > Fabrics**.
       The **Discover Fabrics** dialog box displays.
    2. Select an IP address from the **Available Addresses** list.
    3. Click **Edit**.
       The **AddFabric Discovery** dialog box displays.
    4. Select the **Manual** option to view SNMP credentials.
    5. Click the **SNMP** tab.
    6. Select **v1** or **v3** from the **SNMP Version** list.
    7. Make sure SNMP credentials match those on the device.
    8. Click **OK** on the **AddFabric Discovery** dialog box.
    9. Click **Close** on the **Discover Fabrics** dialog box.

- To set SNMP credentials in the Management application, refer to "Discovery" on page 33.

- Make sure that the SNMP security level is set to the appropriate level for the switch.
  - To check the SNMP security level, use the snmpconfig **--show secLevel** command.

    Example of checking SNMP security level

    ```
    snmpconfig --show secLevel
    GET security level = 0, SET level = 0
    SNMP GET Security Level: No security
    SNMP SET Security Level: No security
    ```

  - To set the SNMP security level, use the **snmpconfig --set secLevel** command.

    Example of checking SNMP security level

    ```
    snmpconfig --set secLevel 0
    Select SNMP GET Security Level
    (0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 = No Access): (0..3) [0]
    ```

- To collect performance data for GE ports and FCIP statistics, make sure that SNMP v3 credentials match and that FCIP-MIB capability is enabled.
  - To check FCIP-MIB capability, use the **snmpconfig --show mibcapability** command.

    Example of showing FCIP-MIB

    ```
    FCRRouter:admin> snmpconfig --show mibcapability
    FCIP-MIB: YES
    ```

  - To enable FCIP-MIB capability, use the **snmpconfig --set mibcapability** command.

    Example of enabling FCIP-MIB

    ```
    FCRRouter:admin> snmpconfig --set mibcapability
    FA-MIB (yes, y, no, n): [yes]
    FICON-MIB (yes, y, no, n): [yes]
    HA-MIB (yes, y, no, n): [yes]
    FCIP-MIB (yes, y, no, n): [yes]
    ISCSI-MIB (yes, y, no, n): [yes]
    ```

- To collect performance data on a Virtual Fabrics-enabled device, use the **userconfig --show** command to make sure the Fabric OS user has access to all the Virtual Fabrics. Make sure that the SNMPv3 user name is the same as the Fabric OS user name. Otherwise, the data is not collected for virtual switches with a non-default Virtual Fabric ID. By default, the **admin** user has access to all Virtual Fabrics.

  Example of Fabric OS user verification

  ```
  sw1:FID128:admin> userconfig --show
  Account name: admin
  Description: Administrator
  Enabled: Yes
  Password Last Change Date: Unknown
  Password Expiration Date: Not Applicable
  Locked: No
  Home LF Role: admin
  Role-LF List: admin: 1-128
  Chassis Role: admin
  Home LF: 128
  ```

  - Make sure I/O is running on the switch to obtain real statistics. To view switch statistics, use the **portperfshow** (FC ports) or **portshow fciptunnel** (FCIP tunnels) command.

  Example for FC ports

  ```
  Sprint-65:root> portperfshow 5
  ```

Example for FCIP tunnels

```
Sprint-65:root> portshow fciptunnel ge0 1 -perf
```

SAN real-time performance data

Real-time performance monitoring enables you to collect data from managed devices in your SAN. Real-time performance monitoring is only supported on the following managed objects: FC (E_Ports and F_Ports), GE_Ports, E_Port trunks, 10GE_Ports, and FCIP tunnels. You can use real-time performance monitoring to configure the following options:

- Select the polling rate from 10 seconds up to 1 minute.
- Select up to 100 ports total from a maximum of 20 devices for graphing performance.

  For E_Port trunks, you can select up to 25 trunks (the trunk member [port] count must be below 100) from a maximum of 20 devices for graphing performance.

  NOTE
  Virtual Fabric logical ISL ports are not included in performance collection.

- Choose to display the same Y-axis range for both the Tx MB/Sec and Rx MB/Sec measure types for easier comparison of graphs.

Generating a real-time performance graph

You can monitor the device performance through a performance graph that displays transmit and receive data. The graphs can be sorted by the column headers. You can create multiple real-time performance graph instances.

NOTE
To make sure that statistics for a switch do not fail, you must configure SNMP credentials for the switch. For step-by-step instructions, refer to "Discovery" on page 33.

To generate a real-time performance graph for a device, complete the following steps.

1. Select the fabric, device, or port for which you want to generate a performance graph. Right-click and select **Performance > Real-Time Graph/Table** or select **Monitor > Performance > Real-Time Graph**. The **Real Time Graphs/Tables** dialog box displays.

   If you selected a port, the **Real Time Performance Graphs** dialog box for the selected port displays. To filter real-time performance data from the **Real Time Performance Graphs** dialog box, refer to "Filtering real-time performance data" on page 968.

   If you selected a fabric or a device, the **Realtime Port Selector** dialog box displays. Continue with step 2.

FIGURE 456 Realtime Port Selector dialog box

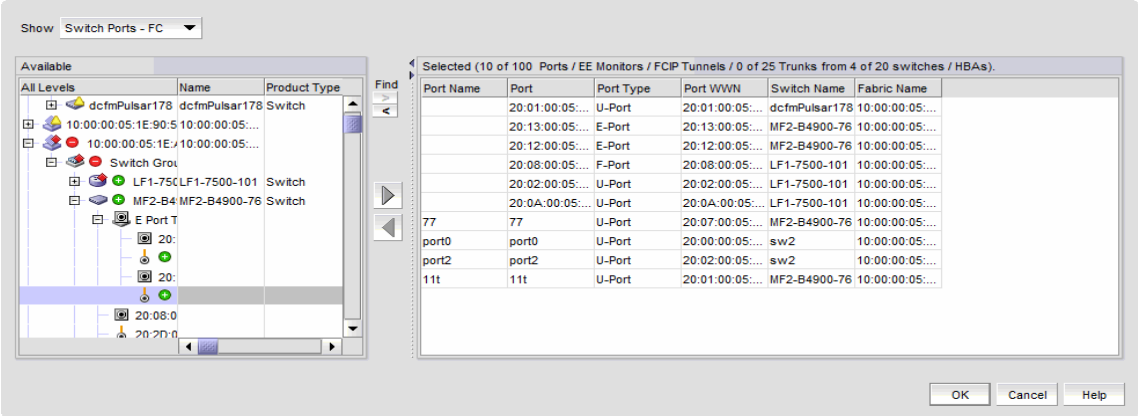

NOTE
You can set columns in the right side of the dialog box for FICON display using **Server** > **Options** > **SAN Display**. The first eight columns will display FC Address, Serial #, Tag, Product Type, Model, Vendor, Port Name, Port Type, and Port WWN.

If you selected a host adapter, the **Realtime Port Selector - Hosts** dialog box displays with all the managed HBA and CNA ports of the discovered hosts. Continue with step 3.

FIGURE 457 Realtime Port Selector - Hosts dialog box

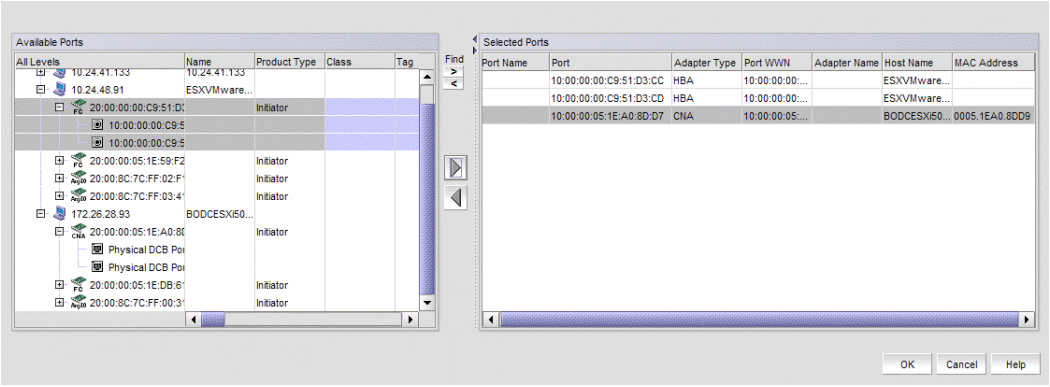

2. From the **Show** list, select the object type for which you want to performance graph.

NOTE
Devices with 10 GbE ports must be running Fabric OS 7.0 or later to obtain the correct TE_Port statistics (TX/RX).

NOTE
Devices with 10 GbE ports must have the RMON MIB enabled on the switch. For more information about the **rmon collection** command, refer to the *Fabric OS Converged Enhanced Ethernet Command Reference*.

If you selected a host adapter, the **Show** list of the **Realtime Port Selector - Hosts** dialog box displays Host ports and Host port measures of the discovered hosts.

3. Right-click anywhere in the **Available** list and select **Expand All** from the menu.
4. Select the ports or trunks you want to include in the performance graph in the **Available** list.

5. Click the right arrow to move the selected ports or trunks to the **Selected** list.

   NOTE
   For E_Port trunks, the port node can be moved to the **Selected** list along with the occupied and attached ports.

6. Click **OK**.

   The **Real Time Performance Graphs** dialog box displays.

Filtering real-time performance data

To filter real-time performance data from the **Real Time Performance Graphs** dialog box, complete the following steps.

1. Open the **Real Time Performance Graphs** dialog box.

   For step-by-step instructions, refer to "Generating a real-time performance graph" on page 966.

2. Select how the data is measured, in received frames, transmitted frames, or CRC errors.

   For a list of possible performance measures, refer to "SAN performance measures" on page 960.

3. To select more than one measure, click the **Additional Measures** expand arrows and select the check box for each additional measure. If the **Additional Measures** area is not shown, click the down arrow.

   For a list of possible performance measures, refer to "SAN performance measures" on page 960.

4. (Optional) Enter a value (percentage) in the **Reference Line** field to set a reference for the transmit and receive utilization.

   Note that this field is only available when you select **Tx % Utilization** or **Rx % Utilization** from the **Measures** list.

5. Select how detailed the data will display from the **Granularity** list. Options are in increments of 10 seconds, 15 seconds, 20 seconds, 25 seconds, 30 seconds, 45 seconds, or 1 minute.

6. Select **Plot Events** to display advanced monitoring service (AMS) violation events received during the chart time range and Master Log events logged on the same product as the measure being plotted.

7. Move the **Row Height** slider to the left to make the row height smaller or to the right to make it larger.

8. Select the **Display tabular data only** check box to show only text with no graphs or icons.

   The **Source** and **Destination** icons and the **Graph** column do not display.

9. Click **Apply**.

   The selected data automatically displays in the **Real Time Performance Graphs** dialog box.

10. Click the close button (**X**) to close the **Real Time Performance Graphs** dialog box.

Graph display

The columns in the graphical portion of the **Real Time Performance Graphs** dialog box display the following information:

- Additional Measures - Displays each measure selected in the **Measures** list and **Additional Measures** area.
- Measures - A list for each selected measure in the **Measures** list or **Additional Measures** area.
- Granularity - The selected granularity for collecting data.
- Source Fabric - The source fabric being monitored.
- Source - The source device being monitored.

- Source Port - The source port being monitored.
- Port Type - Type of port being monitored.
- Graph - Graph of data over time.
- Destination - The destination device.
- Destination Port - The port through which the selected device is connected to the destination device.
- Rx% Utilization - The bytes received.
- Tx% Utilization - The bytes transmitted.

  Select any of the columns and click **Add Flow** to create an add flow definition. For more information, refer to "Monitoring flows" on page 1021.

Exporting real-time performance data

To export real-time performance data, complete the following steps.

1. Generate a performance graph.

   To generate a performance graph, refer to "Generating a real-time performance graph" on page 966.

2. Right-click anywhere in the graph table and select **Export Table**.

   The **Save table to a tab delimited file** dialog box displays.

3. Browse to the file location where you want to save the performance data.
4. Enter a name for the file and click **Save**.

Clearing port counters

To reset all port statistics counters to zero on a selected device or fabric, complete the following steps.

1. Right-click a device or a fabric on the Connectivity Map or Product List and select **Monitor > Performance > Clear Counters**. An attention message displays.
2. Click **Yes** on the message.

   All the port statistics counters and port logs will be cleared and the audit events log generated by the switches is displayed in the Master Log.

   NOTE
   Clear counters is not initiated for Not Reachable switches.

   NOTE
   The audit log is supported in Fabric OS firmware version 7.3.0 and later only. Beginning with Fabric OS firmware version 7.3.0, clearing port counters and port logs will be initiated at the chassis level for Virtual Fabrics-enabled switches.

SAN historical performance data

Performance monitoring should be enabled constantly to receive the necessary historical data required for a meaningful report. The following options and features are available for obtaining historical performance data:

- Collect historical performance data from the entire SAN or from a selected fabric.

  NOTE
  Virtual Fabric logical ISL ports are not included in performance data collection.

- Persist data on every polling cycle (5 minutes).
- Store records for each port.
- Use the Round Robin Database (RRD) style aging scheme.
- Enable a granularity for data collection:
  - 5 minute granularity for last 8 days
  - 30 minutes granularity for last 30 days
  - 2 hour granularity for last 30 days
  - 1 day granularity for last 730 days
- Plot advanced monitoring service (AMS) violation events received during the chart time range and Master Log events logged on the same product as the measure being plotted.
- Generate reports. For instructions on generating reports, refer to "Generating SAN performance reports" on page 1285.
- Configure the graph display using right-click menu options. For more information refer to "Configuring the graph display" on page 973.

Enabling SAN-wide historical performance collection

To enable historical performance collection, select **Monitor > Performance > Historical Data Collection.**

The **Fabric Selector** dialog box displays with **Enable SAN Wide enabled by default. This enables historical performance data collection for all fabrics in the SAN.**

NOTE
After enabling historical data collection, information for switches, ports, and FCIP tunnels also displays in the IP **Historical Graph/Tables** dialog box. If available, click the **IP** tab, then select **Monitor > Performance > Historical Graphs/Tables**.

Enabling historical performance collection for selected fabrics

To enable historical performance collection for selected fabrics, complete the following steps.

1. Select **Monitor > Performance > Historical Data Collection.**

   The **Fabric Selector** dialog box displays.

2. Select **Enable Selected.**

   The **Historical Data Collection** dialog box displays, as shown in Figure 458 on page 971.

FIGURE 458 Historical Data Collection dialog box

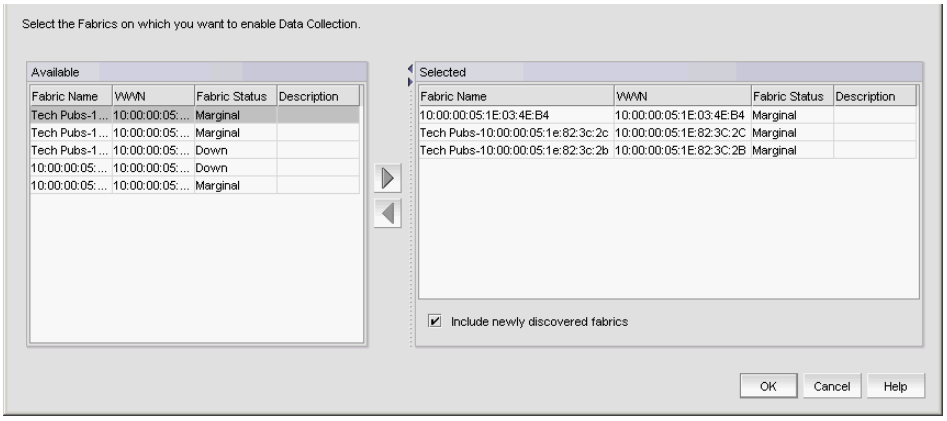

3. Select the fabrics for which you want to collect historical performance data in the **Available** list.

   NOTE
   Devices with 10GE ports must be running Fabric OS 7.0 or later to obtain the correct TE_Port statistics (TX/RX).

   NOTE
   Devices with 10GE ports must have the RMON MIB enabled on the switch. For more information about the **rmon collection** command, refer to the *Fabric OS Converged Enhanced Ethernet Command Reference*.

4. Click the right arrow to move the selected fabrics to the **Selected** list.
5. Select the **Include newly discovered fabrics** check box to automatically add all newly discovered fabrics to the **Selected** list.
6. Click **OK**.

   Historical performance data collection is enabled for all selected fabrics.

   NOTE
   After enabling historical data collection, information for switches, ports, and FCIP tunnels also displays in the IP **Historical Graph/Tables** dialog box. If available, click the **IP** tab, then select **Monitor > Performance > Historical Graphs/Tables**.

Disabling historical performance collection

Perform the following steps to disable historical performance collection on all fabrics.

1. Select **Monitor > Performance > Historical Data Collection.**

   The **Fabric Selector** dialog box displays.

2. **Select Disable All**.

Historical performance data collection is disabled for all fabrics in the SAN.

Generating and saving a historical performance graph

The **Historical Performance Graph** is available through the SAN tab or through the IP tab if you select SAN devices. If selecting through the IP tab, refer to ", as shown in Figure 707 on page 1710Configuring the performance graph display" on page 1002.

To generate a historical performance graph for a device, complete the following steps.

1. Select the device for which you want to generate a performance graph.
2. Select **Monitor > Performance > Historical Graph**.

   The **Historical Performance Graph** dialog box displays, as shown in Figure 459 on page 972.

   FIGURE 459 Historical Performance Graph dialog box

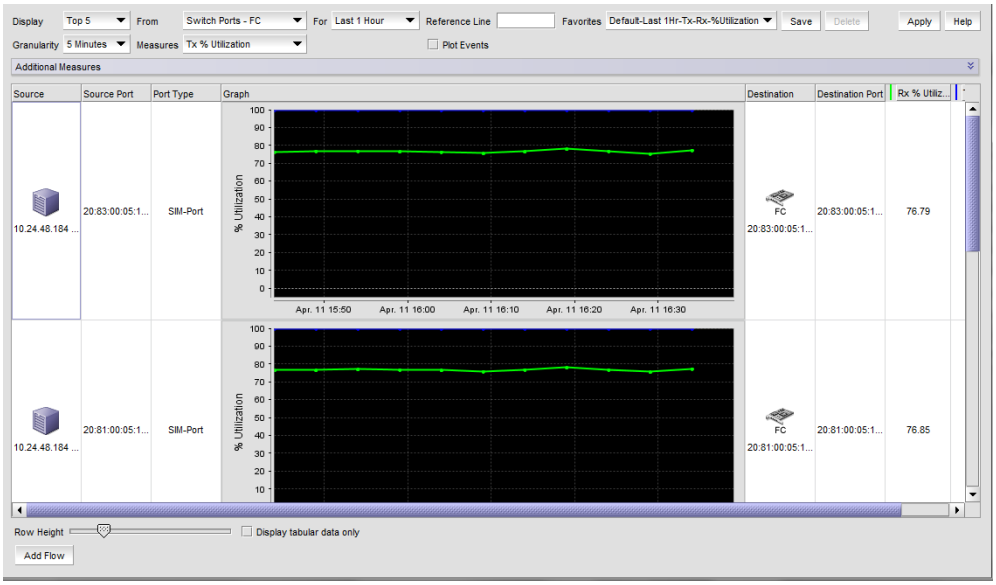

3. Select a default or custom-saved port and time from the **Favorites** list or filter the historical data by completing the following steps.
   a. Select the number of results to display from the **Display** list.
   b. Select the type of port from which you want to gather performance data from the **From** list.

      NOTE
      Devices with 10GE ports must be running Fabric OS 7.0 or later to obtain the correct TE port statistics (TX/RX).

      NOTE
      Devices with 10GE ports must have the RMON MIB enabled on the switch. For more information about the **rmon collection** command, refer to the ***Fabric OS Converged Enhanced Ethernet Command Reference***.

      If you select **Custom**, the **Custom Port Selector** dialog box displays where you can save selected ports as a favorite.

      If you select **Custom**, refer to "Filtering data by ports" on page 974.

   c. Select the historical period for which you want to gather performance data from the **For** list.

      If you select **Custom**, you can save selected time as a favorite.

      If you select **Custom**, refer to "Filtering data by time" on page 975.

d. Select the granularity at which you want to gather performance data from the **Granularity** list.
   - 5 minutes for last 8 days
   - 30 minutes granularity for last 30 days
   - 2 hour granularity for last 30 days
   - 1 day granularity for last 730 days

   NOTE
   The graph will not update dynamically if the granularity is 30 Minutes, 2 Hours, or 1 day. To update, click **Apply**. The graph will update dynamically when 5 Minutes is selected.

e. Select the measure by which you want to gather performance data from the **Measures** list.

   To select more than one measure, click the **Additional Measures** expand arrows and select the check box for each additional measure.

f. If selecting **Tx % Utilization** or **Rx % Utilization** from the **Measures** list, enter a percentage in **Reference Line.**

g. Select **Plot Events** to plot advanced monitoring service (AMS) violation events received during the chart time range and Master Log events logged on the same product as the measure being plotted.

h. Move the **Row Height** slider to the left to make the row height smaller or to the right to make it larger.

i. Select the **Display tabular data only** check box at the bottom of the graph to show only text with no graphs or icons.

   The **Source** and **Destination** icons and the **Graph** column do not display.

j. **Add flow** button - Select any of the columns and click **Add Flow** to create an Add flow definition. For more information, refer to "Monitoring flows" on page 1021.

k. Click **Apply**.

   The selected graph automatically displays in the **Historical Performance Graph** dialog box, if you do not select the **Display tabular data only** check box.

   To save a filtered graph, refer to "Generating and saving a historical performance graph" on page 972.

   To delete user-defined graph, refer to "Deleting a favorite graph configuration" on page 976.

   To configure graph display, right-click in the graph and select desired options. For details on these options, refer to "Configuring the graph display" on page 973.

4. Enter a name for the configuration in the **Favorites Name** field.
5. Save this configuration by selecting **Save**.

   The **Save Favorites** dialog box displays. This enables you to save the selected configuration so that you can use it to generate the same type of report at a later date.

6. Click the close button (**X**) to close the **Historical Performance Graph** dialog box.

Configuring the graph display

To configure the historical performance graph display, right-click in the graph and select the following options:

- Select **Zoom In** to zoom in on the graph.
- Select **Zoom Out** to zoom out on the graph.
- Select **Fit in window** to fit the graph in the window.
- Select **Go to Latest** to go to the latest data point on the graph.

- Select the **Use Logarithmic Axis** check box to present data on a logarithmic or non-logarithmic axis.
- Select the **Show Values** check box to annotate data point values in the graph.
- Select the **Enable Auto Scrolling** check box to automatically jump to display the new data when new data is collected while the graph is in view.
- Select the **Enable Transition Effect** check box to automatically adjust the range on the vertical axis so that all the data are contained within the view area when you drag the chart into a different time range on the SNMP monitoring graph.
- Select **Plot Min/Max** to plot minimum and maximum values along with the average data point. This option is not available if minimum interval granularity (5 minutes for a SAN historical graph) is selected. The width of the color band displayed on the graph indicates the variation during the time period.
- Select **Show Events** to display advanced monitoring service (AMS) violation events received during the chart time range and master log events logged on the same product as the measure being plotted.
- Select **Chart Styles** to display data as a line chart, area chart, or bar chart.
- Select **Export** to export to a spreadsheet (.csv) or an image (.png).
- Select **Print** to print the graph.

Filtering data by ports

To filter data for a historical performance graph by ports, complete the following steps.

1. Select **Custom** from the **From** list on the **Historical Performance Graph** dialog box.

   The **Custom Port Selector** dialog box displays.

2. Select the type of ports from the **Show** list, as shown in Figure 460 on page 974.

   FIGURE 460 Custom Port Selector dialog box

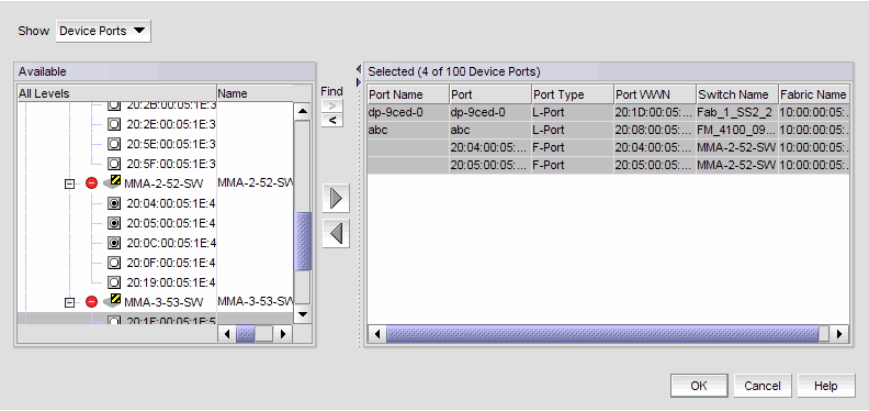

3. Right-click a device in the **Available** list and select **Expand All**.
4. Select the ports (press **Ctrl** or **Shift** and then click to select multiple ports) from which you want to gather performance data from the **Available** list and click the right arrow button.

   NOTE
   For E_Port trunks, the port node can be moved to the **Selected** list along with the occupied or connected ports and attached ports.

   NOTE
   Devices with 10 GbE ports must be running Fabric OS 7.0 or later to obtain the correct TE_Port statistics (TX/RX).

NOTE
Devices with 10GE ports must have the RMON MIB enabled on the switch. For more information about the **rmon collection** command, refer to the *Fabric OS Converged Enhanced Ethernet Command Reference*.

The selected ports move to the **Selected** list.

5. Click **OK**.

Filtering data by time

To filter data for a historical performance graph by time, complete the following steps.

1. Click **Monitor > Performance > Historical Graph**.

   The **Historical Performance Graph** dialog box displays.

2. Select **Custom** from the **For** list.

   The **Custom Time Frame** dialog box displays as shown in Figure 461 on page 975. Perform one of the following steps:

   - Select the **Last** option and enter the number of minutes, hours, or days that you want to monitor.
   - Select the **From** option and enter the start date and time (in MM DD YYYY HH MM AM/PM format) that you want to monitor.
   - Select the **To** option and enter the end date and time (in MM DD YYYY HH MM AM/PM format) that you want to monitor.

   FIGURE 461 Custom Time Frame dialog box

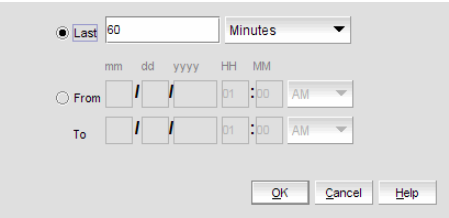

3. Click **OK**.

Exporting historical performance data

To export historical performance data, complete the following steps.

1. Generate a performance graph.

   To generate a performance graph, refer to "Generating and saving a historical performance graph" on page 972.

2. Right-click anywhere in the graph table and select **Export**.

   The **Save to a tab delimited file** dialog box displays.

3. Browse to the file location where you want to save the performance data.

4. Enter a name for the file and click **Save**.

Deleting a favorite graph configuration

To delete a favorite historical performance graph configuration, complete the following steps.

1. Select **Monitor > Performance > Historical Graph**.

   The **Historical Performance Graph** dialog box displays.

2. Select the configuration you want to delete from the **Favorites** list.

   You can only delete a user-defined historical performance graph. You cannot delete a default favorite historical performance graph.

3. Click **Delete**.
4. Click **Yes** on the confirmation message.
5. Click the close button (**X**) to close the **Historical Performance Graph** dialog box.

Performance database views

The following view names are used to extract data similar to the 11.3.0 database schema from the server with the version greater than or equal to 12.0.2.

NOTE
The FC_PORT_STATS and FCIP_STATS views definition are available under the tree view of *databases >dcmdb>Schemas>dcm>Views node hierarchy and* can be extracted from the 12.0.2 database schema from the server by the following ways:
- Search for the view definitions with the view names at *<Managment_Application install-home>\conf\schema\dcm-postgres-schema.sql location*
- Open the PostgreSQL user interface by double clicking on *<Managment_Application install-home>\bin\dbadmin.*

- FC_PORT_STATS_5MIN_INFO
- FC_PORT_STATS_30MIN_INFO
- FC_PORT_STATS_2HOUR_INFO
- FC_PORT_STATS_1DAY_INFO
- FCIP_STATS_5MIN_INFO
- FCIP_STATS_30MIN_INFO
- FCIP_STATS_2HOUR_INFO
- FCIP_STATS_1DAY_INFO

The following EE_MONITOR_STATS and TE_PORT_STATS view names are used to extract data similar to the 11.3.0 database schema from the server with the version greater than or equal to 12.0.0. Refer to "Database Fields" for view definitions.

How to extract performance statistics data from the database

Following are the steps used to extract any PM data from the database:

- Check PM_DATA_COLLECTOR table, to identify the collector database ID
- Check PM_COLLECTOR_TIME_SERIES_MAPPING table, to find the mapping table that contains the required data

- Construct the select query using the mapping table

Execute the following query to extract the FCIP tunnel statistics for last 1 day

```
Select * from TIME_SERIES_DATA_2 where COLLECTOR_ID = 13;
```

Execute the following query to extract the FC port statistics for last 3 days

```
Select * from TIME_SERIES_DATA_1_30MIN where COLLECTOR_ID = 11;
```

Execute the following query to extract the TE port statistics for last 30 days

```
Select * from TIME_SERIES_DATA_1_2HOUR where COLLECTOR_ID = 12;
```

Execute the following query to extract all SAN product statistics for last 730 days

```
select * from TIME_SERIES_DATA_2_1DAY where COLLECTOR_ID = 15;
```

Performance statistics counters

Table 80 details the formulas used to calculate performance statistics based on counter type and protocol.

To calculate FC, GE, FCIP and TE port statistics, the Management application uses SNMP to query the respective object identifiers (OID) (listed in Table 80).

To calculate HBA and CNA statistics, the Management application uses APIs provided by HCM.

To calculate end-to-end monitor (EE monitor) statistics, the Management application uses HTTP to obtain the TX, RX, and CRC error values.

The polling interval for historical graphs is 5 minutes. The polling interval for real-time graphs is based on the granularity value (configured in the Real Time Graph dialog box).

TABLE 80 Performance statistic counters

Counter name	Type	Protocol	Source OID value	Formula
TX	FC	SNMP	.1.3.6.1.3.94.4.5.1.6	TX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
RX	FC	SNMP	.1.3.6.1.3.94.4.5.1.7	RX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
TX	GE	SNMP	.1.3.6.1.2.1.31.1.1.1.10	TX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
RX	GE	SNMP	.1.3.6.1.2.1.31.1.1.1.6	RX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
TX	FCIP	SNMP	.1.3.6.1.2.1.31.1.1.1.10	TX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
RX	FCIP	SNMP	.1.3.6.1.2.1.31.1.1.1.6	RX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
Uncompressed Tx/Rx MB/sec	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.6	(delta value[1] / (1000 * 1000)) / (polling interval[2])
TX	EE Monitors	HTTP	PortRX (variable from the return html file)	TX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
RX	EE Monitors	HTTP	PortTX (variable from the return html file)	RX = (delta value[1] / (1000 * 1000)) / (polling interval[2])

TABLE 80 Performance statistic counters (Continued)

Counter name	Type	Protocol	Source OID value	Formula
TX	HBA, CNA	HCM API	N/A	TX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
RX	HBA, CNA	HCM API	N/A	RX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
TX	TE	SNMP	.1.3.6.1.2.1.31.1.1.1.10	TX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
RX	TE	SNMP	.1.3.6.1.2.1.31.1.1.1.6	RX = (delta value[1] / (1000 * 1000)) / (polling interval[2])
TX% / RX%	FC	SNMP	TX = .1.3.6.1.3.94.4.5.1.6 RX = .1.3.6.1.3.94.4.5.1.7	TX% or RX% for FC = ((delta value[1] of TX or RX) / ((Bytes transmitted * port speed)*(polling interval[2]))) * 100 where Bytes transmitted for 1G, 2G, 4G, 8G and 16G port speed is 106250000 and Bytes transmitted for 10G port speed is 127500000. If utilization is less than 1, the value is 0.0.
TX% / RX%	GE	SNMP	TX = .1.3.6.1.2.1.31.1.1.1.10 RX =.1.3.6.1.2.1.31.1.1.1.6	TX% or RX% for FC = ((delta value[1] of TX or RX) /((125000000* port speed)*(polling interval[2]))) * 100. If utilization is less than 1, the value is 0.0.
TX% / RX%	FCIP	SNMP	TX = .1.3.6.1.2.1.31.1.1.1.10 RX = .1.3.6.1.2.1.31.1.1.1.6	TX% or RX% for FCIP = ((delta value[1] of TX or RX) /(maximum bytes transmitted))*polling interval[2]))) * 100, where maximum bytes transmitted = tunnel speed * 125000000. If utilization is less than 1, the value is 0.0.
TX% / RX% (Pre-Fabric OS 6.4.1 release)	TE	SNMP	TX = .1.3.6.1.2.1.31.1.1.1.10 RX = .1.3.6.1.2.1.31.1.1.1.6	TX% or RX% for TE = ((delta value[1] of TX or RX) /((125000000* 10)*(polling interval[2])))* 100. If utilization is less than 1, the value is 0.0.
Cumulative Compression Ratio	FCIP	N/A	.1.3.6.1.4.1.1588.4.1.1.4	Compression Ratio = current value/ 1000 The compression ratio is the current compression ratio value.
Current Compression Ratio	FCIP	N/A	N/A	(ifHCInOctets + ifHCOutOctets) / fcipExtendedLinkCompressedBytes
Receive EOF	TE		.1.3.6.1.2.1.16.1.1.1.5	Receive EOF = delta value[1] / (1000 * 1000)
Other[3]				Other counters = delta value[1] / polling interval[2]

1. The difference of the value retrieved between two consecutive polling cycles.
2. The duration between two polling cycle in seconds.
3. Additional performance counters are detailed in Table 80.

Table 81 lists the additional counters for which you can obtain performance statistics.

TABLE 81 Performance counters

Counter name	Type	Protocol	Source OID value
CRC Errors	FC	SNMP	.1.3.6.1.3.94.4.5.1.40
Signal Losses	FC	SNMP	.1.3.6.1.3.94.4.5.1.43
Sync Losses	FC	SNMP	.1.3.6.1.3.94.4.5.1.44
Link Failures	FC	SNMP	.1.3.6.1.3.94.4.5.1.39

TABLE 81 Performance counters (Continued)

Counter name	Type	Protocol	Source OID value
Sequence Errors	FC	SNMP	.1.3.6.1.3.94.4.5.1.42
Invalid Transmissions	FC	SNMP	.1.3.6.1.3.94.4.5.1.41
Rx Link Resets	FC	SNMP	.1.3.6.1.3.94.4.5.1.33
Tx Link Resets	FC	SNMP	.1.3.6.1.3.94.4.5.1.34
C3 Discard	FC	SNMP	.1.3.6.1.3.94.4.5.1.28
C3 Discard Rx Timeout	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.27.1.25
C3 Discard Unreachable	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.27.1.26
C3 Discard Tx Timeout	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.27.1.27
C3 Discard Others	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.27.1.28
PCS Error Block	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.27.1.29
Temperature	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.28.1.1.1
Voltage	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.28.1.1.2
Current	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.28.1.1.3
Rx Power	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.28.1.1.4
Tx Power	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.28.1.1.5
Encode Error Out	FC	SNMP	.1.3.6.1.4.1.1588.2.1.1.1.6.2.1.26
Invalid Oredered Set	FC	SNMP	.1.3.6.1.3.94.4.5.1.45
BB Credit Zero	FC	SNMP	1.3.6.1.3.94.4.5.1.8
Truncated Frames	FC	SNMP	1.3.6.1.3.94.4.5.1.47
FEC Corrected Blocks	FC	SNMP	1.3.6.1.4.1.1588.2.1.1.1.27.1.31
FEC Uncorrected Blocks	FC	SNMP	1.3.6.1.4.1.1588.2.1.1.1.27.1.32
Latency	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.5
Dropped Packets	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.3
Link Retransmits	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.2
Timeout Retransmits	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.9
Fast Retransmits	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.10
Duplicate Ack Received	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.11
Window Size RTT	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.12
TCP Out of Order Segments	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.13
SlowStart Status	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.14
CRC Errors	EE Monitors	HTTP	PortCRC (variable from the return html file)

SAN end-to-end monitoring

Procedures in this section pertain to end-to-end monitoring using the legacy End-to-End Monitor feature instead of using Flow Vision to create end-to-end monitors.

NOTE
The **End-to-End Monitors** feature is disabled (grayed-out) and not supported in Fabric OS 7.4.0 and later. The Fabric OS 7.4.0 switches will be filtered and not monitored.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > End-to-End Monitors**, a message displays that you can use Flow Vision to provide End-to-End monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Set End-To-End Monitors** dialog box. To use the legacy End-to-End Monitor feature, you must deactivate existing flows defined for the switch for Flow Vision.

Refer to the following important notes when using this feature:

- For systems using Fabric OS version 7.2 or later, you can create end-to-end monitors using the Flow Vision feature. Refer to "Flow Vision" for details.
- End-to-end monitoring requires a Fabric OS device.
- An end-to-end monitor and a Top Talker monitor cannot be configured on the same fabric or external F_Port application-specific integrated circuit (ASIC). You must delete the Top Talker monitor before you configure the end-to-end monitor.
- End-to-end monitoring on an Access Gateway device requires Fabric OS 7.0 or later with an Advanced Performance Monitor license.

  Performance monitoring enables you to provision end-to-end monitors of selected target and initiator pairs. These monitors are persisted in the database and are enabled on one of the F_Ports on the connected device (the Management application server determines the port). You can use these monitors to view both real-time and historical performance data.

Configuring an end-to-end monitor pair

Procedures in this section pertain to configuring monitors on systems using the legacy End-to-End Monitor feature instead of using Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > End-to-End Monitors**, a message displays that you can use Flow Vision to provide End-to-End monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Set End-To-End Monitors** dialog box. To use the legacy End-to-End Monitor feature, you must deactivate existing flows defined for the switch for Flow Vision.

NOTE
Either the initiator device or the target device must have an Advanced Performance Monitor license configured to create an end-to-end monitor.

To configure an end-to-end monitor pair, complete the following steps.

1. Select **Monitor > Performance > End-to-End Monitors**.

   The **Set End-to-End Monitors** dialog box displays as shown in Figure 462 on page 981.

FIGURE 462 Set End-to-End Monitors dialog box

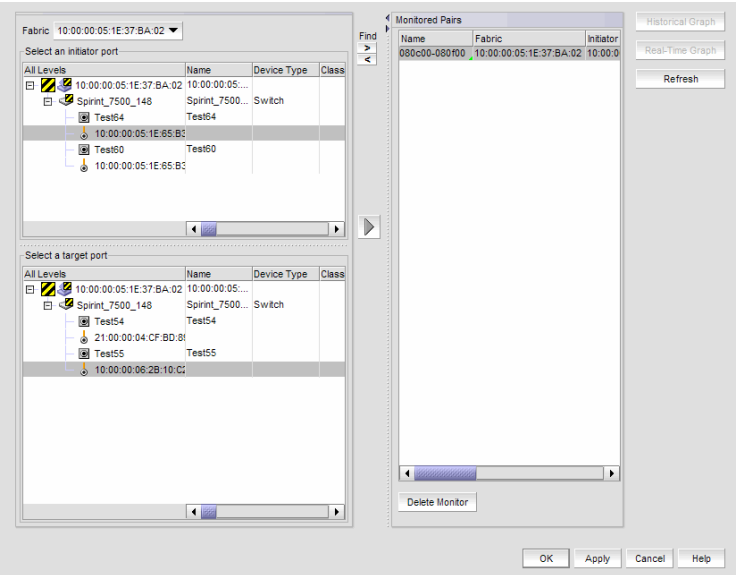

2. Select the fabric for which you want to configure end-to-end monitoring from the **Fabric** list.
3. Select an initiator port from the **Select an initiator port** list.
4. Select a target port from the **Select a target port** list.
5. Click the right arrow to move the selected initiator and target ports to the **Monitored Pairs** list.

   The system automatically determines the initiator SID and the target DID identifiers for the pair and displays them in the **Monitored Pairs** list.

6. Click **Apply**.

   Before you apply end-to-end monitoring to ports moved to the **Monitored Pairs list the Status** column displays "Not Configured." When you **Apply** the monitored pair, the **Status** column displays "Enabled". If the end-to-end monitored pair fails, the **Status** column displays "Failed:Reason".

   NOTE
   If the initiator or target port is part of a logical switch and you move it to another logical switch, the end-to-end monitor fails.

   Once you have created the end-to-end monitored pair, you can view both real-time and historical performance data. For step-by-step instructions, refer to "Displaying end-to-end monitor pairs in a real-time graph" on page 981 or "Displaying end-to-end monitor pairs in a historical graph" on page 982.

Displaying end-to-end monitor pairs in a real-time graph

Procedures in this section pertain to displaying monitors on systems using the legacy End-to-End Monitor feature instead of using Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > End-to-End Monitors**, a message displays that you can use Flow Vision to provide End-to-End monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.

- Clicking **OK** opens the legacy **Set End-To-End Monitors** dialog box. To use the legacy End-to-End Monitor feature, you must deactivate existing flows defined for the switch for Flow Vision.

To display an end-to-end monitor pair in a real-time graph, complete the following steps.

1. Select **Monitor > Performance > End-to-End Monitors**.

   The **Set End-to-End Monitor** dialog box displays.

2. Select one or more end-to-end monitor pairs you want to view from the **Monitored Pairs** list.

   You can select up to 100 monitored pairs.

3. Click **Real-Time Graph**.

   The **Real Time Performance Graphs** dialog box displays.

Displaying end-to-end monitor pairs in a historical graph

Procedures in this section pertain to configuring monitors on systems using the legacy End-to-End Monitor feature instead of using Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > End-to-End Monitors**, a message displays that you can use Flow Vision to provide End-to-End monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Set End-To-End Monitors** dialog box. To use the legacy End-to-End Monitor feature, you must deactivate existing flows defined for the switch for Flow Vision.

To display monitored pairs in a historical graph, data collection must be enabled for the selected fabric or enabled SAN-wide.

To display an end-to-end monitor pair in a historical graph, complete the following steps.

1. Select **Monitor > Performance > End-to-End Monitors**.

   The **Set End-to-End Monitor** dialog box displays.

2. Select one or more end-to-end monitor pairs you want to view from the **Monitored Pairs** list.

   You can select up to 100 monitored pairs.

3. Click **Historical Graph**.

   The **Historical Performance Graph** dialog box displays.

Refreshing end-to-end monitor pairs

Procedures in this section pertain to refreshing monitors on systems using the legacy End-to-End Monitor feature instead of using Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > End-to-End Monitors**, a message displays that you can use Flow Vision to provide End-to-End monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.

- Clicking **OK** opens the legacy **Set End-To-End Monitors** dialog box. To use the legacy End-to-End Monitor feature, you must deactivate existing flows defined for the switch for Flow Vision.

The Management application enables you to rewrite the end-to-end monitors (deleted through the CLI or an Element Manager) back to a device.

To refresh all end-to-end monitor pairs, complete the following steps.

1. Select **Monitor > Performance > End-to-End Monitors**.

   The **Set End-to-End Monitor** dialog box displays.

2. Click **Refresh**.

   All end-to-end monitor pairs are rewritten back to any devices where the end-to-end monitor pairs were deleted through the CLI or an Element Manager.

3. Click **OK**.

Deleting an end-to-end monitor pair

Procedures in this section pertain to deleting monitors on systems using the legacy End-to-End Monitor feature instead of using Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > End-to-End Monitors**, a message displays that you can use Flow Vision to provide End-to-End monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Set End-To-End Monitors** dialog box. To use the legacy End-to-End Monitor feature, you must deactivate existing flows defined for the switch for Flow Vision.

To delete an end-to-end monitor pair, complete the following steps.

1. Select **Monitor > Performance > End-to-End Monitors**.

   The **Set End-to-End Monitor** dialog box displays.

2. Select the end-to-end monitor pair you want to delete from the **Monitored Pairs** list.
3. Click **Delete Monitor**.
4. Click **OK**.

SAN Top Talker monitoring

Procedures in this section pertain to configuring the legacy Top Talkers feature instead of using Flow Vision.

> NOTE
> Top Talkers monitoring is disabled (grayed-out) and not supported in Fabric OS 7.4.0 and later. The Fabric OS 7.4.0 switches will be filtered and not monitored.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > Top Talkers**, a message displays that you can use Flow Vision to provide Top Talkers monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Top Talkers** dialog box. To use the legacy Top Talkers feature, you must deactivate existing flows defined for the switch for Flow Vision.

The following are some important notes for using the Top Talkers feature:

- To access Flow Vision, the Fabric Vision (FV) license or both the Fabric Watch (FW) and the Advanced Performance Monitor (APM) licenses must be installed on the hardware platform.
- Top Talkers cannot be enabled on a single-switch fabric.
- Top Talkers require Fabric OS version 7.0 or later.
- A Top Talker monitor and an end-to-end monitor cannot be configured on the same external F_Port application-specific integrated circuit (ASIC). You must delete the end-to-end monitor before you configure the Top Talker monitor.
- On the 8 Gbps 8-FC port, 10 GbE 24-DCB port Switch, Top Talkers is only supported on the 8 Gbps FC Ports.

You can create Top Talker monitors on selected devices. Use Top Talkers to display the connections which are using the most bandwidth on the selected device or port. Top Talkers can be enabled on the device or one of the F_Ports on the device. You can only use Top Talkers to view real-time performance data.

You can have multiple Top Talker monitors configured at the same time. You can monitor up to 10 switches for fabric mode Top Talkers and 32 ports and 10 switches for F_Port Top Talkers; however, you can only monitor one device or port for each Top Talker you configure.

NOTE
If the Fabric OS device is configured for Fibre Channel routing (FCR), you can only configure a Top Talker monitor on the following devices:
- 16 Gbps Backbone Chassis with a FC 16 Gbps 32-port or 48-port blade
- 16 Gbps 48-port switch

Configuring a fabric mode Top Talker monitor

Procedures in this section pertain to configuring the legacy Top Talkers feature instead of using Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > Top Talkers**, a message displays that you can use Flow Vision to provide Top Talkers monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Top Talkers** dialog box. To use the legacy Top Talkers feature, you must deactivate existing flows defined for the switch for Flow Vision.

Here are some important notes for using this feature.

NOTE
A fabric mode Top Talker and an end-to-end monitor cannot be configured on the same fabric. You must delete the end-to-end monitor before you configure the fabric mode Top Talker.

NOTE
A fabric mode Top Talker and an F_Port mode Top Talker cannot be configured on the same fabric. You must delete the F_Port mode Top Talker before you configure the fabric mode Top Talker.

NOTE
You cannot enable Top Talkers for a single-switch fabric.

To configure a fabric mode Top Talker monitor on systems using Fabric OS before v7.2, complete the following steps.

1. Select the fabric on which you want to monitor Top Talker data.

   NOTE
   On the 8 Gbps 8-FC port, 10 GbE 24-DCB port Switch, Top Talkers is only supported on the 8 Gbps FC Ports.

2. Select **Monitor > Performance > Top Talkers**.

   The **Top Talker Selector** dialog box displays, as shown in Figure 463 on page 985.

   FIGURE 463 Top Talker Selector dialog box

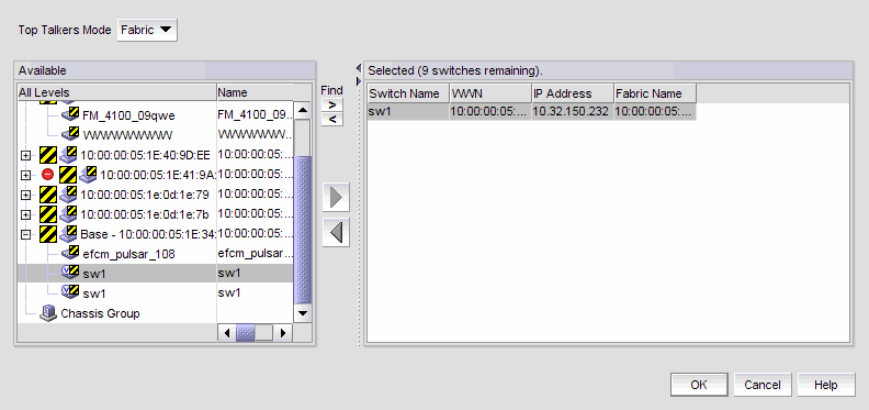

3. Select **Fabric** in the **Top Talker Mode** list to select a switch to monitor.
4. Select an available switch from a fabric in the left panel, and then the right arrow to move it to the right panel.

   You can select only one device on which to enable Top Talkers.

5. Click **OK** on the **Top Talker Selector** dialog box.

   Top Talkers is enabled on the selected device. The **Top Talkers - Fabric Mode for** *Device_Name* dialog box displays.

6. Select the number of Top Talkers (1 through 20) to display from the **Display** list.
7. Select how often you want the Top Talker to refresh (10, 20, 30, 40, or 50 seconds, or 1 minute) from the **Refresh Interval** list.
8. Click **Apply**.

   The top 20 conversations display in the **Current Top Talkers** list. The **Top Talkers Summary** list displays all Top Talkers that occurred since the **Top Talkers - Fabric Mode for** *Device_Name* dialog box was opened (displays a maximum of 360). When the maximum is reached, the oldest Top Talker drops as a new one occurs.

   The fabric mode Top Talker provides the following details:

- Tx+Rx Ave (MB/sec)
- Occurrences
- Source
- Last Occurred
- SID
- Source Port

- Source Switch/Port
- Destination
- Destination Switch/Port
- DID
- Destination Port

1. Click **Destination** to launch the **Port Properties** dialog box for the Destination port.
2. Click **Source** to launch the **Port Properties** dialog box for the Source port.

Configuring an F_Port mode Top Talker monitor

Procedures in this section pertain to configuring the legacy Top Talkers feature instead of using Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > Top Talkers**, a message displays that you can use Flow Vision to provide Top Talkers monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Top Talkers** dialog box. To use the legacy Top Talkers feature, you must deactivate existing flows defined for the switch for Flow Vision.

Here are some important notes for using this feature:

- An F_Port mode Top Talker and an end-to-end monitor cannot be configured on the same F_Port. You must delete the end-to-end monitor before you configure the F_Port mode Top Talker.
- A Top Talker monitor and an end-to-end monitor cannot be configured on the same fabric or external F_Port application-specific integrated circuit (ASIC). You must delete the end-to-end monitor before you configure the Top Talker monitor.
- Launching a fabric mode Top Talker monitor from the connectivity map or **Top Talker Selector** dialog box displays a "Top Talkers cannot be enabled for single switch fabric" warning.

To configure an F_Port mode Top Talker monitor on systems using Fabric OS before v7.2, complete the following steps.

1. Select a fabric that you want to monitor Top Talker data for an F_Port.
2. Select **Monitor > Performance > Top Talkers**.

   The **Top Talker Selector** dialog box displays.
3. Select **F Port** from the **Top Talkers Mode** list.
4. Select an available F_Port in the left panel, and then the right arrow to move it to the right panel.

   You can only select one F_Port on which to enable the Top Talker monitor.
5. Click **OK** on the **Top Talker Selector** dialog box.

   Top Talkers is enabled on the selected port. The **Top Talkers - F Port Mode for** *Port_Name* dialog box displays.
6. Select the number of Top Talkers (1 through 20) to display from the **Display** list.
7. Select how often you want the Top Talker to refresh (10, 20, 30, 40, or 50 seconds, or 1 minute) from the **Refresh Interval** list.
8. Select whether you want to monitor the receive (Rx) flow or the transmit (Tx) flow for the port from the **Flow** list.
9. Click **Apply**.

The top 20 conversations display in the **Current Top Talkers** list. The **Top Talkers Summary** list displays all Top Talkers that occurred since the **Top Talker Selector** dialog box was opened (displays a maximum of 360). When the maximum is reached, the oldest Top Talker drops as a new one occurs.

The F_Port mode Top Talker provides the following details:

- Rx Ave (MB/sec) or Tx Ave (MB/sec)
- Occurrences
- Source
- Source Switch/Port
- Destination
- Destination Switch/Port
- % Utilization
- Last Occurred
- SID
- Source Port
- DID
- Destination Port
- Port Speed

Deleting a Top Talker monitor

Procedures in this section pertain to deleting monitors created on systems using the legacy Top Talkers feature and not those created with Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > Top Talkers**, a message displays that you can use Flow Vision to provide Top Talkers monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Top Talkers** dialog box. To use the legacy Top Talkers feature, you must deactivate existing flows defined for the switch for Flow Vision.

To delete Top Talker monitors, use the following steps:

1. Select the dialog box of the Top Talker monitor you want to delete.

   Refer to steps 1-5 under "Configuring a fabric mode Top Talker monitor" on page 984 or "Configuring an F_Port mode Top Talker monitor" on page 986 to display this dialog box.

2. Click **Close**.
3. Click **Yes** on the "Do you want to delete this monitor?" message.

Pausing a Top Talker monitor

Procedures in this section pertain to pausing monitors created on systems using the legacy Top Talkers feature and not those created with Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > Top Talkers**, a message displays that you can use Flow Vision to provide Top Talkers monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Top Talkers** dialog box. To use the legacy Top Talkers feature, you must deactivate existing flows defined for the switch for Flow Vision.

To pause a Top Talker monitor on systems using Fabric OS before 7.2, complete the following steps.

1. Select the dialog box of the Top Talker monitor you want to pause.

   Refer to steps 1-5 under "Configuring a fabric mode Top Talker monitor" on page 984 or "Configuring an F_Port mode Top Talker monitor" on page 986 to display this dialog box.

2. Click **Pause at the top of the dialog box**.

Related topics

"Flow Vision"

"Configuring a fabric mode Top Talker monitor"

"Configuring an F_Port mode Top Talker monitor"

"Deleting a Top Talker monitor"

"Restarting a Top Talker monitor"

Restarting a Top Talker monitor

Procedures in this section pertain to restarting monitors created on systems using the legacy Top Talkers feature and not those created with Flow Vision.

For systems using Fabric OS version 7.2 or later, when you select a device or device port, and then select **Monitor > Performance > Top Talkers**, a message displays that you can use Flow Vision to provide Top Talkers monitoring. You have these options:

- To use Flow Vision, delete existing monitors, then use the **Add Flow Definition** dialog box to define an initiator and target port pair for monitoring. Refer to "Flow Vision" for more information.
- Clicking **OK** opens the legacy **Top Talkers** dialog box. To use the legacy Top Talkers feature, you must deactivate existing flows defined for the switch for Flow Vision.

To restart a top talker monitor, perform the following steps:

1. Select the dialog box of the Top Talker monitor you want to restart.

   Refer to steps 1-5 under "Configuring a fabric mode Top Talker monitor" on page 984 or "Configuring an F_Port mode Top Talker monitor" on page 986 to display this dialog box.

2. Click **Continue**.

Bottleneck detection

NOTE
Bottleneck Detection is not supported for switches running Fabric OS 8.0.0 or later. **MAPS FPI** monitoring must be enabled for detecting bottleneck. Bottlenecked Ports widget and Bottleneck Port Status are supported using MAPS FPI for switches running Fabric OS 8.0.0 or later.

A ***bottleneck*** is a port in the fabric where frames cannot get through as fast as they should. In other words, a bottleneck is a port where the offered load is greater than the achieved egress throughput. Bottlenecks can cause undesirable degradation in throughput on various links. When a bottleneck occurs at one place, other points in the fabric can experience bottlenecks as the traffic backs up.

The bottleneck detection feature detects two types of bottlenecks:

- Latency bottleneck

- Congestion bottleneck

A *latency bottleneck* is a port where the offered load exceeds the rate at which the other end of the link can continuously accept traffic, but does not exceed the physical capacity of the link. This condition can be caused by a device attached to the fabric that is slow to process received frames and send back credit returns. A latency bottleneck due to such a device can spread through the fabric and can slow down unrelated flows that share links with the slow flow.

A *congestion bottleneck* is a port that is unable to transmit frames at the offered rate because the offered rate is greater than the physical data rate of the line. For example, this condition can be caused by trying to transfer data at 8 Gbps over a 4 Gbps ISL.

You can set alert thresholds for the severity and duration of the bottleneck.

If a bottleneck is reported, you can then investigate and optimize the resource allocation for the fabric. Using the zone setup and Top Talkers, you can also determine which flows are destined to any affected F_Ports.

You configure bottleneck detection on a per-fabric or per-switch basis, with per-port exclusions.

NOTE
Bottleneck detection is disabled by default. The best practice is to enable bottleneck detection on all switches in the fabric, and leave it on to continuously gather statistics.

Supported configurations for bottleneck detection

Note the following configuration rules for bottleneck detection:

- The switch must be running Fabric OS 7.0 or later.
- Bottleneck detection is supported on Fibre Channel ports and FCoE F_Ports.
- Bottleneck detection is supported on the following port types:
  - E_Ports
  - EX_Ports
  - F_Ports
  - FL_Ports
- F_Port and E_Port trunks are supported.
- Long distance E_Ports are supported.
- FCoE F_Ports are supported.
- Bottleneck detection is supported on 4 Gbps, 8 Gbps, and 16 Gbps platforms.
- Bottleneck detection is supported in Access Gateway mode.
- Bottleneck detection is supported whether Virtual Fabrics is enabled or disabled. In VF mode, bottleneck detection is supported on all fabrics, including the base fabric.

How bottlenecks are reported

Bottlenecks are reported through alerts in the Master Log. A bottleneck cleared alert is sent when the bottleneck is cleared.

NOTE
A bottleneck cleared alert is sent if you disable bottleneck detection on a bottlenecked port, even though the port is still bottlenecked.

Bottlenecks can be highlighted in the Connectivity Map and Product List. Select **Monitor > Performance > View Bottlenecks**. If a port is experiencing a bottleneck, a bottleneck icon is displayed in the Connectivity Map for the switch and fabric, and in the Product List for the port, switch, and fabric, as shown in Figure 464. In the figure, port15 and port22 are bottlenecked.

FIGURE 464 Bottleneck port indications

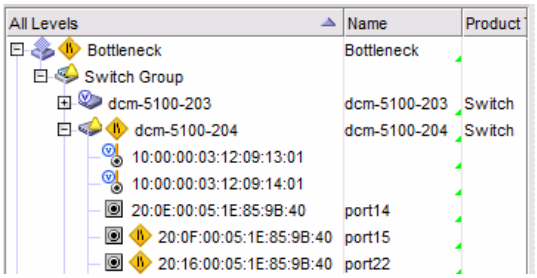

Limitations of bottleneck detection

The bottleneck detection feature for latency detection is not recommended for link utilizations above 85 percent.

The bottleneck detection feature detects latency bottlenecks only at the point of egress, not ingress. For example, for E_Ports, only the traffic egressing the port is monitored. For FCoE ports, bottleneck detection monitors traffic going from the FC side to the DCB side, and does not monitor traffic going in the reverse direction.

Enabling bottleneck alerts and configuring alert parameters

Bottleneck detection is enabled on a switch or fabric basis. It enables both latency and congestion detection. Consider these points when enabling bottleneck detection:

- If you enable bottleneck detection on a fabric, the feature is applied to all eligible switches in the fabric and all eligible ports on the switches.
- If you enable bottleneck detection on a switch, the feature is applied to all eligible ports on that switch.
- You can override switch configuration by changing parameters for specific ports.
- When changing switch-level parameters, such as time and severity threshold values, bottleneck detection will be disabled, then enabled.

If ineligible ports later become eligible or, in the case of a logical switch, if ports are moved to the logical switch, bottleneck detection is automatically applied to those ports.

If you add additional switches, including logical switches, to the fabric, bottleneck detection is not automatically applied, so be sure to enable bottleneck detection on those switches as well.

NOTE
It is recommended that you enable bottleneck detection on every switch in the fabric.

When you enable bottleneck detection, you also determine whether alerts are to be sent when the bottleneck conditions at a port exceed a specified threshold. The alert parameters include whether alerts are sent and the threshold, time, and quiet time options. These alert parameters apply to all ports in the switch, unless you override them later.

After you enable bottleneck detection, you can change the alert parameters on all eligible ports, switches, and fabrics.

NOTE

The best practice is to enable alerts and use the default values:

Congestion	80%
Latency	10%
Window	300 seconds
Quiet Time	300 seconds
Time threshold	0.8
Severity threshold	50

If you change the Window value, you should use a setting that is 300 seconds or higher.

If you change the alert parameters for a port, you can later cancel these settings and inherit the settings from the switch. Refer to "Inheriting alert parameters from a switch" on page 992 for instructions.

Use the following steps to enable bottleneck alerts and configure alert parameters.

1. Select **Monitor > Performance > Bottlenecks**.

   The **Bottlenecks** dialog box displays, as shown in Figure 465 on page 991.

   FIGURE 465 Bottlenecks dialog box

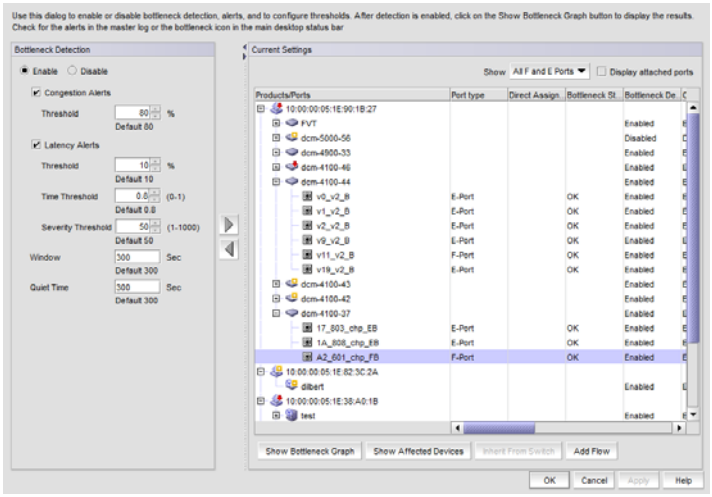

2. Select **Enable** if it is not already selected.
3. Select the **Congestion Alerts** check box to enable alerts for congestion bottlenecks. Clear this check box to disable alerts. If you enabled alerts, enter threshold values between 1 and 100, or use the default value for triggering a congestion alert.
4. Select the **Latency Alerts** check box to enable alerts for latency bottlenecks. If you enabled alerts, enter values for the following thresholds:

- **Threshold** - Enter values between 1 and 100, which is the percentage of one-second intervals affected by congestion conditions within a specified time window that will trigger a latency alert.
- **Time Threshold** - Enter the minimum fraction of a second (sub-second time) that must be affected by latency in order for that second to be considered affected by a latency bottleneck and trigger a latency alert. Values are in tenths of a second from 0 through 10 tenths, or 1 second. You can only configured Time Threshold for switches running Fabric OS v7.1.0 and later.
- **Severity Threshold** - Enter a severity threshold from 1 through 1000. This specifies the factor that throughput must drop in a second (sub-second severity) for that second to be considered affected by a latency bottleneck and trigger a latency alert. You can only configured Severity Threshold for switches running Fabric OS v7.1.0 and later.

  NOTE
  When setting time and severity threshold values the at switch level or fabric level, all values applied to individual ports are overridden and updated with the new values.

5. Enter a value for **Window** in seconds over which the percentage of seconds affected by bottleneck conditions is computed and compared with the threshold. Values can be from 1 through 10800 seconds (3 hours).
6. Enter a value for **Quiet Time**, which is the minimum number of seconds between consecutive alerts. Enter values from 1 through 31556926 (approximately 1 year).
7. Select one or more fabrics, switches, or ports from the **Products/Ports** list.

   You can select fabrics or switches or ports, but you cannot select a mix of fabrics, switches, and ports.
8. Click the right arrow to apply the settings in the **Bottleneck Detection** pane to the selected elements in the **Products/Ports** list.

   If you selected one or more ports, a right arrow displays in the **Direct Assigned** column for these ports, indicating that the alert parameters for the ports override the alert parameters for the switch.

   If you selected switches or fabrics, the alert parameters are changed for all of the eligible ports in those switches and fabrics except for the ports that had been directly assigned alert parameters previously.
9. Select the following options at the bottom of the dialog box as necessary:
   - **Show Bottleneck Graph**. This displays the **Bottleneck Graph Port Selector** dialog box for configuring a bottleneck graph. Refer to "Displaying bottleneck statistics" on page 993.
   - **Show Affected Devices**. Lists the hosts and targets that might be affected by the selected bottlenecked port.
   - **Inherit From Switch**. Restores the switch bottleneck parameters to a port that has direct assigned settings. Refer to "Inheriting alert parameters from a switch".
   - **Add Flow**. Displays the **Add Flow Definition** dialog box to define flows for flow monitoring. Applicable fields on the dialog box, such as source device and destination device IDs, will be populated according to the port selected. Refer to "Monitoring flows" on page 1021 for more information on using the **Add Flow Definition** dialog box and Flow Vision features. The **Add flow** button is enabled when you select a single bottleneck or non-bottleneck port.
10. Click **OK** or **Apply** to save your changes.

Inheriting alert parameters from a switch

When you enable bottleneck detection on a switch, all eligible ports on that switch inherit the same bottleneck parameters as the switch. You can then change the parameters for specific ports or exclude specific ports from bottleneck detection.

Use the following procedure if you want to restore the switch bottleneck parameters to a port that has direct assigned settings.

1. Select **Monitor > Performance > Bottlenecks**.

   The **Bottlenecks** dialog box displays.

2. Select a port that has directly assigned bottleneck settings, which is indicated by a right arrow in the **Direct Assigned** column.
3. Click **Inherit From Switch**.
4. Select the **Alerts** check box to enable alerts. Clear this check box to disable alerts.

   The bottleneck parameters that are specified for the switch are applied to the port.
5. Click **OK** or **Apply** to save your changes.

Copying alert parameters from one switch or port to another

1. Select **Monitor > Performance > Bottlenecks**.

   The **Bottlenecks** dialog box displays.
2. Select the switch or port from which you want to copy the bottleneck parameters.
3. Click the left arrow.

   The parameters display in the **Bottleneck Detection** pane.
4. Select one or more switches, ports, or fabrics to which you want to copy the bottleneck parameters.

   You can select fabrics or switches or ports, but you cannot select a mix of fabrics, switches, and ports.
5. Click the right arrow.

   The bottleneck parameters are applied to the selected items.
6. Click **OK** or **Apply** to save your changes.

Displaying bottleneck statistics

You can display a graph of bottleneck statistics for up to 32 ports at one time.

You can display a graph showing the history of bottleneck conditions, for up to the last 150 minutes.

1. Select **Monitor > Performance > Bottleneck Graph**.

   The **Bottleneck Graph Port Selector** dialog box displays with bottlenecked ports shown in the **Available** list.
2. (*Optional*) Select **All Ports** from the **Show** list to display all ports in the **Available** list.
3. Select one or more ports for which you want to display bottleneck statistics and click the right arrow to move them to the **Selected** list.

   You can select up to 32 ports.

   You can select a port on the **Available** list or **Selected** list to find and highlight the port on the alternate list.
4. Click **OK**.

   The **Bottleneck Graph** dialog box displays, showing bottleneck statistics for the selected ports. This dialog box has several options for displaying the data:

   - Change the display interval and the display range.

     Bottleneck port statistics is limited to a maximum of the last 150 minutes with display intervals of 10, 60, and 300 seconds.

- Click **Refresh** to update the displayed data with fresh data.

  If you change the display interval or display range, you must click **Refresh** for the changes to take effect.
- Display real-time and historical performance graphs.
- Select a bottlenecked F_Port or FL_Port and click **Show Affected Devices** to see the hosts and targets that may be affected by the bottleneck.

Displaying devices that could be affected by an F_Port or FL_Port bottleneck

The following procedure displays hosts and targets that could be affected because of a bottlenecked F_Port or FL_Port. These devices are determined based on zoning information and are not based on actual traffic flow.

Affected devices cannot be determined for bottleneck E_Ports.

1. Select **Monitor > Performance > Bottlenecks**.

   The **Bottlenecks** dialog box displays.
2. In the **Current Settings** list, select a bottlenecked port (a port with "Bottlenecked" in the **Bottleneck Status** column).
3. Click **Show Affected Devices**.

   The **Show Affected Devices** dialog box displays.
4. Select a port in the **Bottleneck Ports** list to display the affected hosts and targets in the table on the right side of the dialog box.
5. Select a device in the table, then click the **Show affected VM** button to identify virtual machines with the same target ports aa the device port attached to the bottlenecked F_Port or FL_Port.

Disabling bottleneck detection

Use this procedure to exclude specific ports from bottleneck detection or to disable bottleneck detection on entire switches or fabrics.

It is not recommended to disable bottleneck detection on a port except under special circumstances. For example, if a long-distance port is known to be a bottleneck because of credit insufficiency, you could disable bottleneck detection on that port.

1. Select **Monitor > Performance > Bottlenecks**.

   The **Bottlenecks** dialog box displays.
2. Select **Disable**.
3. Select one or more fabrics, switches, or ports from the **Products/Ports** list.

   You can select fabrics or switches or ports, but you cannot select a mix of fabrics, switches, and ports.
4. Click the right arrow to apply the settings in the **Bottleneck Detection** pane to the selected elements in the **Products/Ports** list.
5. Click **OK** or **Apply** to save your changes.

Thresholds and event notification

NOTE
This feature is only available for Fabric OS devices running 7.3.X and earlier. It not supported on devices running 7.4.0 or later.

Performance monitoring allows you to apply thresholds and event notification to real-time performance data. A performance monitor process (thread) monitors the performance data against the threshold setting for each port and issues an appropriate alert to notify you when the threshold is exceeded. For information about configuring event notification, refer to "Event notification" on page 1132.

NOTE
It is not necessary to configure event notification to receive events in the Master Log. If the threshold is exceeded for a threshold, an event is automatically generated and displayed in the Master Log.

NOTE
If you set the threshold for a particular critical event to 100 percent, by the time you are notified, it may be too late to prevent a failure. However, when you set the threshold to 85 percent, for example, you may be able to prevent the failure from occurring.

Example

The values at 1 second, 3 seconds, and 5 seconds generate events because they exceed boundaries. The value at 2 seconds does not generate an event because, although it crosses the boundary, it remains in the buffer zone. The value at 6 seconds generates an event because it crosses the lower boundary and returns to a value beyond the buffer zone. The example is shown as a graph in Figure 466 on page 995.

FIGURE 466 Threshold example

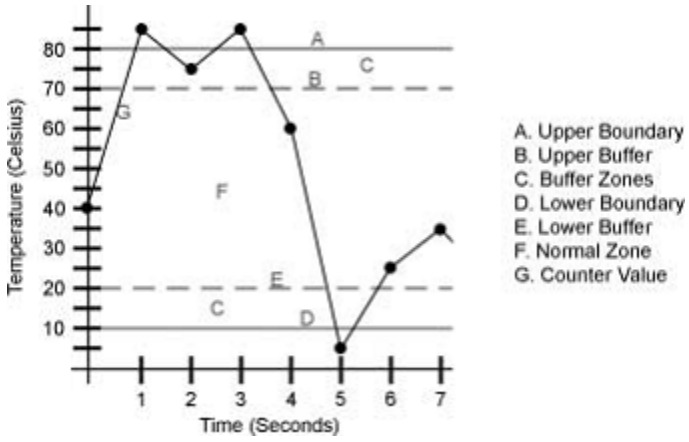

Creating and editing a threshold policy

To create or edit a threshold policy, complete the following steps.

1. Select **Monitor > Fabric Watch > Performance Thresholds**.

   The **Set Threshold Policies** dialog box displays, as shown in Figure 467 on page 996.

FIGURE 467 Set Threshold Policies dialog box

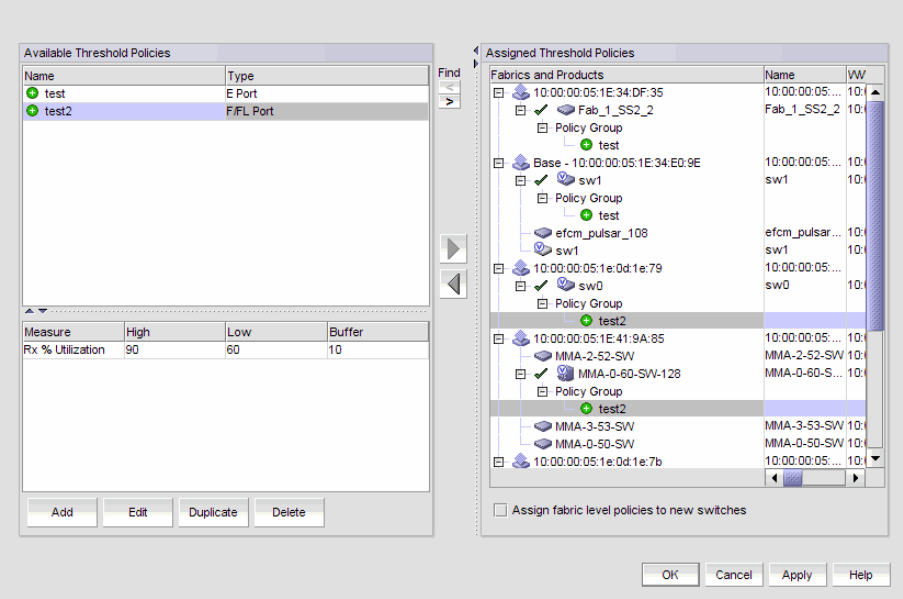

NOTE
Monitoring and Alerting Policy Suite (MAPS)-enabled switches will be filtered out.

NOTE
Performance threshold configurations are not supported in Fabric OS 7.4.0 and later.

2. To edit a current policy, select a policy form the **Available Threshold Policies** list and click **Edit**.

   The **Edit Threshold Policy** dialog box displays, as shown in Figure 468 on page 996.

FIGURE 468 Edit Threshold Policy dialog box

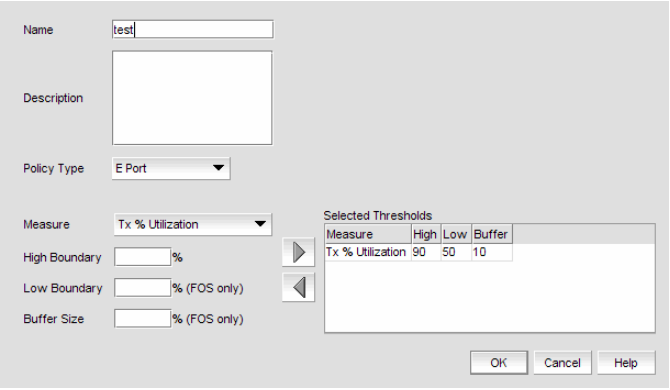

3. To add a new policy, perform the following steps.

   a. Click **Add**.

      The **New Threshold Policy** dialog box displays as shown in Figure 469 on page 997.

FIGURE 469 New Threshold Policy dialog box

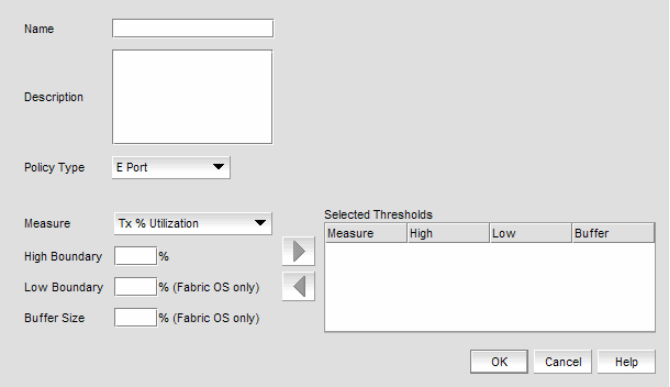

b. Enter a name for the policy (100 characters maximum) in the **Name** field.

4. Select a policy type from the **Policy Type** list.

   You can only define policies for E_Port and F_Port, and FL_Ports.

5. Select a measure from the **Measure** list.

   You can only define policies for the Tx % Utilization and Rx % Utilization measures. You cannot add the same measure more than once. If you try to add another threshold with the same measure, the new values overwrite the older threshold values in the **Selected Thresholds** list.

6. Enter a percentage for the upper boundary in the **High Boundary** field.

   When the counter value exceeds high boundary, an event is raised.

7. (Fabric OS only) Enter a percentage for the lower boundary in the **Low Boundary** field.

   When the counter value goes below the low boundary, an event is raised.

8. (Fabric OS only) Enter a percentage for the buffer in the **Buffer Size** field.

   Counters may fluctuate around the upper or lower boundary of a range threshold, and as a result cause numerous events in a short period of time. To reduce the number of events, configure a buffer (a range of values just below the upper boundary and just above the lower boundary) in which a counter does not register an event if it returns to a “normal” value. An event only registers if the counter returns to a “normal” value beyond the buffer.

9. Click the right arrow button to move the threshold to the **Selected Thresholds** list.

   If an error is detected, a message displays informing you to enter a valid value. Click **OK** to close this message. Fix any errors and repeat step 9.

10. Repeat steps 5 through 9 for each measure that you want to add to the policy.

11. Click **Assign fabric level polices to new switches** on the **Set Threshold Policies** dialog box if you want to assign these policies to new switches.

12. Click **OK**.

    The threshold policy displays in the **Available Threshold Policies** table with an added icon ( ).

13. Click **OK** on the **Set Threshold Policies** dialog box.

    The **Confirm Threshold Changes** dialog box displays as shown in Figure 470 on page 998.

FIGURE 470 Confirm Threshold Changes dialog box

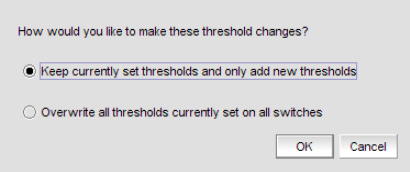

14. Make the threshold changes by selecting one of the following options:
    - To add only new thresholds, select the **Keep currently set thresholds and only add new thresholds** check box.
    - To overwrite all existing thresholds on all fabrics and devices, select the **Overwrite all thresholds currently set on all switches** check box.
15. Click **OK** on the **Confirm Threshold Changes** dialog box.

Duplicating a threshold policy

To duplicate a threshold policy, complete the following steps.

1. Select **Monitor > Fabric Watch > Performance Thresholds**.

   The **Set Threshold Policies** dialog box displays.
2. Select the threshold policy you want to copy in the **Available Threshold Policies** list.
3. Click **Duplicate**.

   The threshold policy displays in the **Available Threshold Policies** list with an added icon ( ) using "copy of *Threshold_Name*" as the naming format. To edit the threshold, refer to "Creating and editing a threshold policy" on page 995. To assign a threshold policy to a fabric or device, refer to "Assigning a threshold policy" on page 998.
4. Click **OK** on the **Set Threshold Policies** dialog box.

   The **Confirm Threshold Changes** dialog box displays.
5. Make the threshold changes by selecting one of the following options:
   - To add only new thresholds, select the **Keep currently set thresholds and only add new thresholds** check box.
   - To overwrite all existing thresholds on all fabrics and devices, select the **Overwrite all thresholds currently set on all switches** check box.
6. Click **OK** on the **Confirm Threshold Changes** dialog box.

Assigning a threshold policy

To assign a threshold policy to a fabric or device, complete the following steps.

1. Select **Monitor > Fabric Watch > Performance Thresholds**.

   The **Set Threshold Policies** dialog box displays.
2. Select one or more threshold policies you want to assign to a fabric or device in the **Available Threshold Policies** list.

   Press **Ctrl** or **Shift** and then click to select multiple policies.

3. Select one or more fabrics or devices to which you want to assign the policy in the **Available Threshold Policies** list.

   If you choose to assign the policy to a fabric and a M-EOS logical switch is present in the fabric, the policy is not assigned to the M-EOS logical switch. You must directly assign a policy to a M-EOS physical chassis.

   When you directly assign a policy to a M-EOS physical chassis, the policy is assigned to all logical switches in the physical chassis.

   Press **Ctrl** or **Shift** and then click to select multiple fabrics or devices.

4. Click the right arrow button to apply the selected policies to the selected fabrics and devices.

   If any of the selected devices do not have a Fabric Watch license, the threshold policies are not set on the device and a message displays listing the affected devices. You will need to upgrade the Fabric Watch license and then assign threshold policies to these devices. Click **OK** to close the message.

5. Click **OK** on the **Set Threshold Policies** dialog box.

   The **Confirm Threshold Changes** dialog box displays.

6. Make the threshold changes by selecting one of the following options:

   - To add only new thresholds, select the **Keep currently set thresholds and only add new thresholds** check box.
   - To overwrite all existing thresholds on all fabrics and devices, select the **Overwrite all thresholds currently set on all switches** check box.

7. Click **OK** on the **Confirm Threshold Changes** dialog box.

Deleting a threshold policy

To delete a threshold policy, complete the following steps.

1. Select **Monitor > Fabric Watch > Performance Thresholds**.

   The **Set Threshold Policies** dialog box displays.

2. Select the threshold policy you want to delete in the **Available Threshold Policies** list.

   When you delete a policy from the M-EOS physical chassis, the policy is deleted from all logical switches in the physical chassis.

3. Click **Delete**.

   The threshold policy displays in the **Available Threshold Policies** list with a removed icon ( ).

4. Click **Yes** on the confirmation message.
5. Click **OK** on the **Set Threshold Policies** dialog box.

   The **Confirm Threshold Changes** dialog box displays.

6. Make the threshold changes by selecting one of the following options:

   - To add only new thresholds, select the **Keep currently set thresholds and only add new thresholds** check box.
   - To overwrite all existing thresholds on all fabrics and devices, select the **Overwrite all thresholds currently set on all switches** check box.

7. Click **OK** on the **Confirm Threshold Changes** dialog box.

SAN connection utilization

NOTE
Connection utilization is only supported on the following managed objects: E_Ports, F_Ports, N_Ports, 10 GE_Ports and FCIP tunnels.

NOTE
Fabrics where performance data collection is not enabled display connections as thin black lines.

Performance connection utilization for device ports provides the following features:

- Turns the utilization display on and off from the menu and toolbar.
- Displays moving dotted colored lines that originate from a port.
- Displays two lines in the topology (when turned on); one represents percentage utilization for transmit and the other represents the percentage utilization for receive. The movement of the line determines if it is a transmit or a receive.
  - Receive (Rx) — Line moves into a port.
  - Transmit (Tx) — Line moves out of a port.
- Displays different colors to represent the percentage utilization range, as shown in Figure 471 on page 1000.

FIGURE 471 Utilization Legend

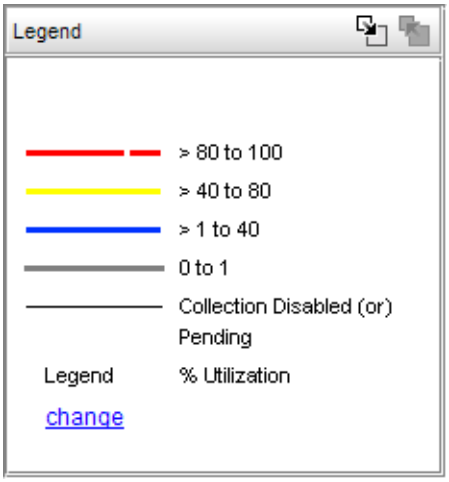

The colors and their meanings are outlined in Table 82 on page 1000.

TABLE 82 Utilization Legend

Line color	Utilization defaults
Red line	80% to 100% utilization
Yellow line	40% to 80% utilization
Blue line	1% to 40% utilization
Gray line	0% to 1% utilization
Black line	Utilization disabled

Enabling connection utilization

To display the connection utilization, complete the following steps.

1. Choose from one of the following options:
   - Select **Monitor > Performance > View Utilization**.
   - Press **CTRL + U**.
   - Click the Utilization icon ( ).

   If you have already enabled historical data collection, the Utilization Legend displays in the main interface window.

   If you have not already enabled historical data collection, a message appears informing you that you must enable historical data collection before you can view utilization, as shown in Figure 705 on page 1707.

2. Choose one of the following options:
   - Select **Enable SAN Wide** to enable data collection for the entire SAN.
   - Select **Enable Selected Fabrics** to enable data collection for specific fabrics.

     The **Historical Data Collection** dialog box displays. To select the fabrics on which you want to enable data collection, refer to "Enabling historical performance collection for selected fabrics" on page 970.

     If you click **Close** on the **Historical Data Collection** message, Historical Data Collection is not enabled; however, the Utilization Legend still displays in the main window.

   There is a 5-minute delay before the values are displayed.

Disabling connection utilization

To turn off the connection utilization, choose one of the following options while connection utilization is enabled:

- Select **Monitor > Performance > View Utilization**.
- Press **CTRL + U**.
- Click the Utilization icon ( ).

  The Utilization Legend is removed from the main interface window.

Changing connection utilization percentages

You can change the utilization percentages.

To change the utilization percentages, complete the following steps.

1. Click the **change** link in the **Utilization Legend**, as shown in Figure 472 on page 1002.

FIGURE 472 Utilization Legend in edit mode

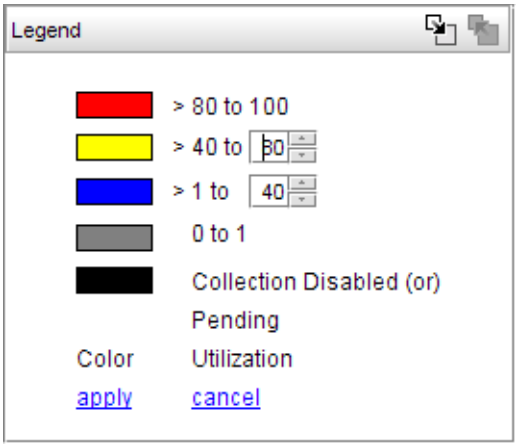

2. Enter or select the end percentage you want for the blue line.

   When you make a change to the end percentage of a utilization line, you also change the start percentage for the utilization line immediately above the one you changed when you click **apply**. For example, if you change the blue line end percentage to 60 the yellow line start percentage changes to 60 when you click **apply**.

3. Enter or select the end percentage you want for the yellow line.
4. Click the **apply** link.

   The new values appear in the **Utilization Legend**.

, as shown in Figure 707 on page 1710## Configuring the performance graph display

Use this procedure to configure the graph display for the **Real Time Graphs/Tables** dialog box and **Historical Graphs/Tables** dialog box as well as time series monitors on the **Dashboard** tab.

1. Right-click the graph and select one of the following options:
   - Select the **Show Controls** check box to show or hide additional display options on the graph (refer to step 2 through step 8 for more information).
   - Select the **Show Legend** check box to show or hide the measurements beneath the graph.
   - Select **Clear Graph** to clear the graph.
   - Select **Delete Selected Measures** to delete the selected measures from performance.
   - Select **Zoom In** to zoom in on the graph.
   - Select **Zoom Out** to zoom out on the graph.
   - Select **Fit in window** to fit the graph in the window.
   - Select **Go to Latest** to go to the latest data point on the graph.
   - Select the **Use Logarithmic Axis** check box to present data on a logarithmic or non-logarithmic axis.
   - Select the **Show Values** check box to annotate data point values in the graph.
   - Select the **Enable Auto Scrolling** check box to automatically jump to display the new data when new data is collected while the graph is in view.

- Select the **Enable Transition Effect** check box to automatically adjust the range on the vertical axis so that all the data are contained within the view area when you drag the chart into a different time range on the SNMP monitoring graph.
- Select **Plot Min/Max** to plot minimum and maximum values along with the average data point. This option is not available if minimum interval granularity (5 minutes for a SAN historical graph) is selected. The width of the color band displayed on the graph indicates the variation during the time period.
- Select **Show Events** to display advanced monitoring service (AMS) violation events received during the chart time range and master log events logged on the same product as the measure being plotted.
- Select **Chart Styles** to display data as a line chart, area chart, or bar chart.
- Select **Options** to launch the **Graph Options** dialog box. Refer to "Configuring graph options" on page 1004 for more information.
- Select **Export** to export to a spreadsheet (.csv) or an image (.png).
- Select **Print** to print the graph.

2. Click **Options** to launch the **Graph Options** dialog box. Refer to "Configuring graph options" on page 1004 for instructions on using this dialog box.
3. Select the **Graph** or **Table** option to display data in graphical or tabular format.
4. Select a time range relative to the present for the display of historical data from the **For** list.

   The options are incremental from the last 30 minutes to the last 24 hours.
5. (Historical graphs and monitors only) Select the **Plot Min/Max** check box to plot minimum and maximum values along with the average data point.

   The range between the minimum and maximum values will be represented in a color band surrounding the data points. The width of the color band indicates the variation during the time period. Note that this option is not available if you select **Minimum Interval** granularity.
6. (Historical graphs and monitors only) Select one of the following options from the **Granularity** list to set the granularity of the data point to display on the graph:

   - **5 minutes**
   - **30 minutes**
   - **2 hours**
   - **1 day**

   NOTE
   The graph will not update dynamically if the granularity is 30 minutes, 2 hours, or 1 day. To update, move from one granularity setting to another. The graph will update dynamically when **Minimum Interval** is selected.
7. Select the **Events** check box to display advanced monitoring service (AMS) violation events received during the chart time range.
8. (**Real Time Graphs/Tables** and **Historical Graphs/Tables** dialog boxes only) Click **Save as Widget** to create a performance monitor published widget on the active dashboard. For instructions, refer to "Configuring a monitor from a performance graph" on page 441.

Configuring graph options

Use the following steps to configure graph options for real-time performance graph display as well as time series monitors on the **Dashboard** tab.

1. Click **Options** on the graph.

   The **Graph Options** dialog box displays, as shown in Figure 473 on page 1004.

   FIGURE 473 Graph Options dialog box (Historical Graphs/Tables dialog box)

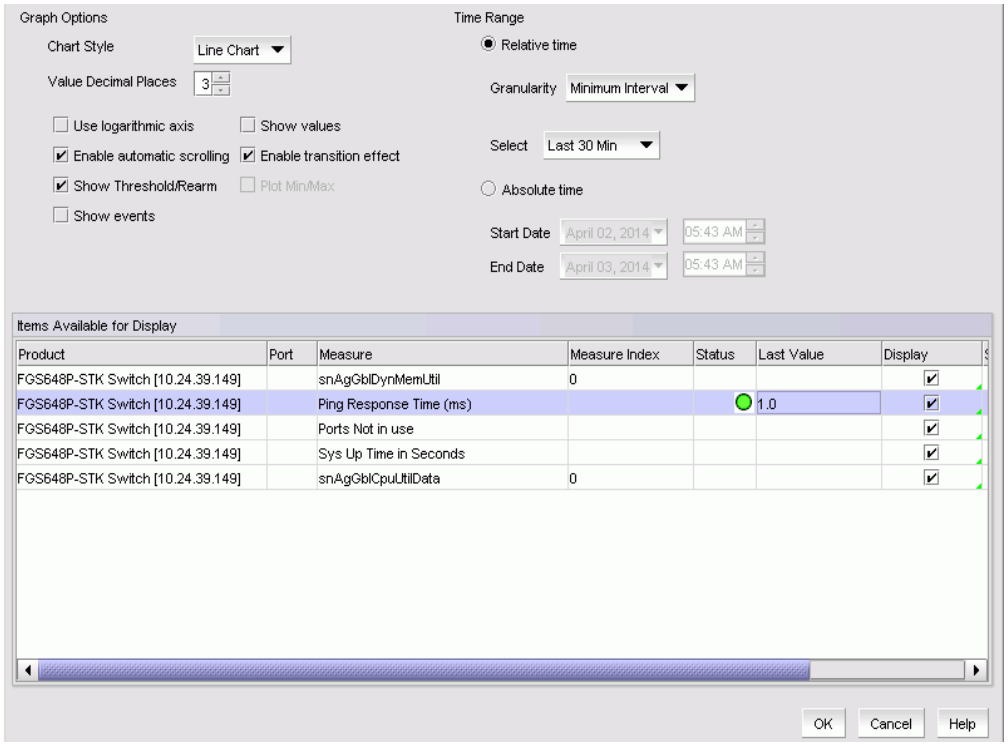

   NOTE
   Figure 473 illustrates the **Graph Options** dialog box available from the **Historical Graphs/Tables** dialog box. The **Graph Options** dialog box available from the **Real Time Graphs/Tables** dialog box is similar, but has fewer control options.

2. Select the type of chart style from the **Chart Style** list.

   Available chart styles include **Line Chart**, **Area Chart**, or **Bar Chart**.

3. Select the graph accuracy to up to three decimal places in the **Value Decimal Places** list.

4. Select from the following check boxes to define how polled data displays:

   - **Use logarithmic axis** — Data can be presented on a logarithmic or non-logarithmic axis. Each unit in a non-logarithmic axis presents the data in equal segments. However, logarithmic axis units are not equal and can increase exponentially by 10. Therefore, use a logarithmic axis if you have a large amount of data to view.
   - **Show values** — Annotates data point values in the graph.
   - **Enable automatic scrolling** — If new data is collected while the chart is in view, the chart will automatically jump to display the new data.

- **Enable transition effect** — The SNMP monitoring chart automatically adjusts the range on the vertical axis so that all the data are contained within the view area when you drag the chart into a different time range. Enabling this option provides an animated smooth transition between the adjustments while the monitoring chart is being dragged or during any action that may cause the range of vertical axis to change.
- **Show Threshold/Rearm** — Displays threshold and rearm events on the chart.
- (Historical graphs and monitors only) **Plot Min/Max** - Plots minimum and maximum values along with the average data. The range between the minimum and maximum values will be represented by the width of a color band surrounding the data points as shown in the following illustration. Note that this option is not available if you select **Minimum Interval** granularity. It also does not apply and is not available for real-time performance graphs.

FIGURE 474 Data points graph

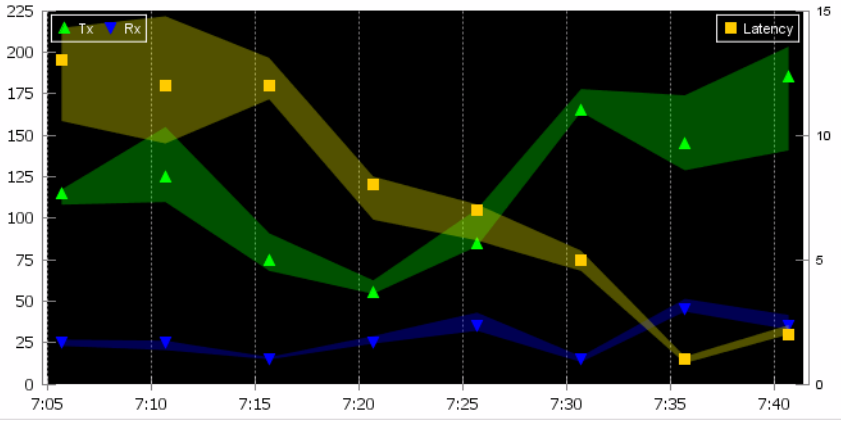

- **Show events** - Select to display advanced monitoring service (AMS) violation events received during the chart time range and master log events logged on the same product as the measure being plotted. Each event will be represented by the same severity icon that is shown in the Master Log (refer to the icons at the bottom of the following graph). Pausing the pointer over the icon displays details of the violation, such as violation time, switch/port information, violated rule name, and violated rule condition. Monitoring and Alerting Policy Suite (MAPS) violations are plotted for a product or port level measure (whichever is selected) during the plotted time range. The show events graph is shown in Figure 475 on page 1005.

FIGURE 475 Show events graph

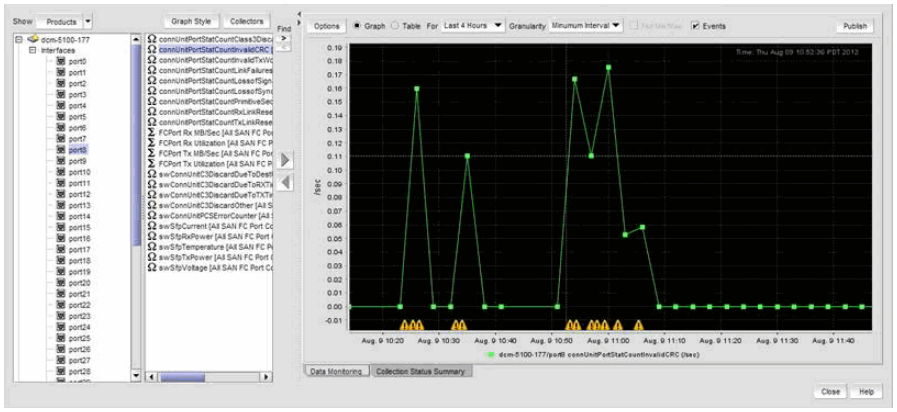

5. In the **Time Range** area, select one of the following options:
   - Select **Relative time** to set a time range relative to the present for the display of historical data.

a. (Historical graphs and monitors only) Select the granularity of the data points to display on the graph from the **Granularity** list. Options are 5 minutes, 30 minutes, 2 hours, or 1 day.

NOTE
The graph will not update dynamically if the granularity is 30 minutes, 2 hours, or 1 day. To update, click **Apply**. The graph will update dynamically when **Minimum Interval** is selected.

b. Select the duration of time for data display on the graph from the **Select** list.
Real-time options are incremental from the last 30 minutes to the last 6 hours.
Historical options are incremental from the last 30 minutes to the last 24 hours.

- (Historical graphs and monitors only) Select **Absolute time** to get a snapshot of data from a specific time range and complete the following steps.
  a. Select the start date from the **Start Date** list and start time in the spin box.
  b. Select the end date from the **End Date** list and end time in the spin box.

6. Include items in the graph by selecting the **Display** check box for each item in the **Items Available for Display** list.
7. Set the scale factor for each item by entering a value (an integer between -2147483648 and 2147483647) in the **Scaling Factor** column for each item in the **Items Available for Display** list.
8. Click **OK** on the **Graph Options** dialog box.

Viewing Historical Graphs/Tables

1. Right-click a row in a performance monitor on the dashboard and select **Show Graph/Table**.

   The **Historical Graphs/Tables** dialog box displays.

2. Select the **Data Monitoring** tab.

   The main features are a tree structure and a graph area. You can collapse the tree structure to expand the graph area.

3. Use the **Show** selector to toggle the tree structure display in the left panel between **Products** and **Collectibles**.

   - Select **Products** and the left panel displays the tree structure of devices and device interfaces on the network being polled for collectible data. The right panel displays measures currently being collected for the selected product or port in the left panel.

     Measures also display for SAN products, ports, and FCIP tunnels that appear in the device tree. Refer to Figure 476 on page 1007 and Figure 477 on page 1007 for examples.

FIGURE 476 SAN Fibre Channel port display

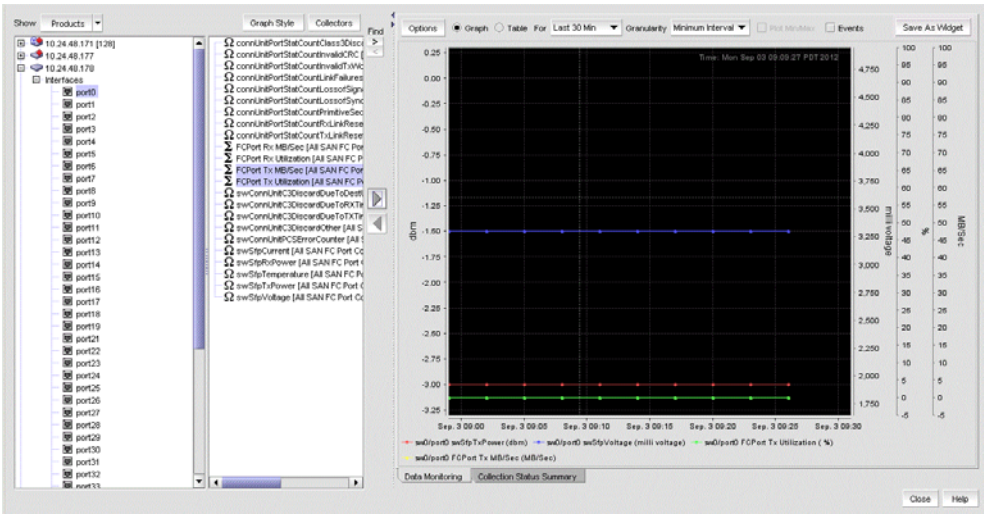

FIGURE 477 SAN FCIP tunnel display

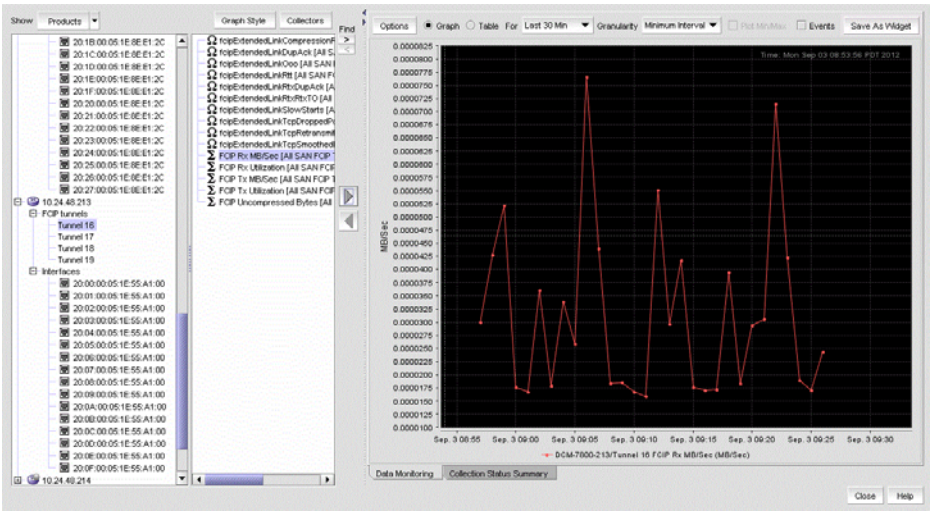

- Select **Collectibles** and the left panel displays measures (MIB objects and expressions) currently being collected. Select a measure, and the targets (products or ports) from which the measure was collected display in the right panel. If SAN historical data collection is enabled, corresponding SAN products and ports display.

  Measures also display for SAN products, ports, and FCIP tunnels that appear in the device tree. In addition, measures collected for attached wireless access point (AP) devices and controllers display. You can select these collectibles to create applicable historical graphs and tables.

4. Select **Collectors** to open the **Historical Data Collectors** dialog box.
5. *(Optional)*: To configure the look and feel of the performance graph from the **Historical Graphs/Tables** dialog box, refer to "Configuring the performance graph" on page 1717.
6. *(Optional)*: Once data collection begins, the data is presented on the chart (if **Graph** is selected) or table (if **Table** is selected).

   If a graph is displayed, the legend under the graph shows what data each color represents. Also, you see the following text:

   - *MIB*: Shows the name of the MIB object that is being used to collect the data and the device that is being polled. If the target is a port, then the interface ID is also displayed.

- *EXP*: Shows the name of the expression being used to collect the data and the device that is being polled. If the target is a port, then the interface ID is also displayed.

Each collectible is represented by a different color and the color for a collectible can change as new data is collected.

If a table is displayed, the first column displays the time of the collection. The remaining columns display the value of each collectible at the specified time. There is one column for every collectible you select to display.

7. *(Optional)*: To add the performance monitor published widget to the active dashboard, click **Save As Widget**.

   The **Performance Dashboard Monitor Title** dialog box displays.

   Select **Add the Widget to active dashboard (Product status and Traffic)** check box to enable the published widget. By default, the check box is enabled.

   Click **OK.**

8. Select the **Collection Status Summary** tab.

   The **Collection Status Summary** tab provides a high-level overview of all defined collectors. The information is displayed in the following columns:

   - **Product** - Shows the product name. There maybe multiple instances of the product name for each collectible assigned to the product.
   - **Port** - The port name when a port is selected.
   - **Collectible** - The MIB objects and expressions used by the data collector. When you select a collectible row, collectible information displays in the bottom portion of the panel, such as errors, error count, and messages.
   - **Collector** - The data collector name.
   - **Status** - The status field uses the following icons:.

     Failed. No value was ever collected for this collectible.

     Warning: Data collection failed in the last polling cycle.

     Successful: Last collection successful.

     Scheduled but not currently active.

   - **Last Value** - The last (most current) value collected.
   - **Last Time Polled** - The time that the collector was last polled.

When you use the **Show** selector to select **Products**, devices and ports display in a tree structure in the left-most column. If you select a device or port, the right collectibles column lists all the collectors that have been defined for the device or port.

If you use the **Show** selector to select **Collectible**, the left-most column shows all the collectibles (MIB objects or SNMP expressions) currently being collected. Select a collectible to display a tree structure in the right column of all products and ports from which the expression or MIB object are to be collected.

When a specific collectible is selected, collectible detail, error count, and error messages display in an area below the table.

Mouse functions for graphs

The following mouse functions can be used for graphs:

- Zoom: Use the mouse wheel to zoom in or zoom out of a graph.
- Graph panning: Hold down the left mouse button and move the mouse left and right to pan through the graph.
- Selective zooming: Select an area that you want to zoom by holding down the right mouse button at one edge of the area, then drag the mouse to the left or right to the other edge of the area. The area you selected changes color. Release the right mouse button to zoom the selected area.
- Highlighting: Place the mouse over a data point. Information about that data point appears in a tooltip-like format.
- Drag and drop from trees: If you want to monitor additional devices on the same graph, select the device from the device tree, then drag and drop it into the graph. You can monitor up to twenty entries in one graph. If you drag and drop a device node, all MIB variables and expressions collected from that device are included.

Performance collection configuration using batch files

You can use the following procedures to configure performance collection using batch files.

Enabling or disabling performance statistics collection

This batch file allows you to enable or disable SAN performance data collection at a global level.

To enable or disable performance statistics collection, complete the following steps.

- On Windows systems, complete the following steps.
  a. Open a command prompt and navigate to the *Install_Home*\utilities directory.
  b. Enable performance statistics collection by typing `sanperformancestatsenable.bat` *dbusername dbpassword* `enable` and pressing **Enter**.

     For example, `sanperformancestatsenable.bat dcmadmin passw0rd enable`.

     Disable performance statistics collection by typing `sanperformancestatsenable.bat` *dbusername dbpassword* `disable` and pressing **Enter**.
- On UNIX systems, complete the following steps.
  a. Open a terminal and navigate to the *Install_Home*\utilities directory.
  b. Enable performance statistics collection by typing `sanperformancestatsenable` *dbusername dbpassword* `enable|disable` and pressing **Enter**.

     For example, `sanperformancestatsenable dcmadmin passw0rd enable`.

     Disable performance statistics collection by typing `sanperformancestatsenable` *dbusername dbpassword* `disable` and pressing **Enter**.

Updating system threshold data

This batch file enables you to modify the default disk space threshold values and assign custom values.

To update the file system threshold data in the system property table, complete the following steps.

- On Windows systems, complete the following steps.
  a. Open a command prompt and navigate to the *Install_Home*\utilities directory.
  b. Update file system threshold data by typing `updatethresholddata.bat` *dbusername dbpassword THRESHOLD_WARN THRESHOLD_RISK THRESHOLD* and pressing **Enter**.

     For example, `updatethresholddata.bat dcmadmin passw0rd 80 90 95`.
- On UNIX systems, complete the following steps.
  a. Open a terminal and navigate to the *Install_Home*\utilities directory.
  b. Update file system threshold data by typing `updatethresholddata` *dbusername dbpassword THRESHOLD_WARN THRESHOLD_RISK THRESHOLD* and pressing **Enter**.

Configuring custom duration for performance aging

NOTE
Changes to the custom duration for performance aging automatically trigger a Management application server restart.

This batch file enables you to set a custom duration for performance aging. You cannot select the maximum limit which is as follows:

- 5 minute granularity for last 8 days
- 30 minutes granularity for last 30 days
- 2 hour granularity for last 30 days
- 1 day granularity for last 730 days (2 years)

To configure custom duration for performance aging, complete the following steps.

- On Windows systems, complete the following steps.
  a. Open a command prompt and navigate to the *Install_Home*\utilities directory.
  b. Set a custom duration for performance aging by typing `performanceaginglimit.bat dbusername dbpassword` and pressing **Enter**.
  c. Enter the number of days for the custom duration by typing *number_of_days* at the **Enter number of days** prompt and pressing **Enter**.
  d. Continue by typing `y` at the **Do you want to continue(Y/N)** prompt and pressing **Enter**.

     The Management application server restarts automatically.
- On UNIX systems, complete the following steps.
  a. Open a terminal and navigate to the *Install_Home*\utilities directory.
  b. Set a custom duration for performance aging by typing `performanceaginglimit dbusername dbpassword` and pressing **Enter**.
  c. Enter the number of days for the custom duration by typing *number_of_days* at the **Enter number of days** prompt and pressing **Enter**.
  d. Continue by typing `y` at the **Do you want to continue(Y/N)** prompt and pressing **Enter**.

     The Management application server restarts automatically.

Exporting configuration data

This batch file enables you to to export device configuration data from database tables to a text file.

To export configuration data from the CFG_backup_archive table, complete the following steps.

- On Windows systems, complete the following steps.
  a. Open a command prompt and navigate to the *Install_Home*\utilities directory.
  b. Export configuration data by typing `exportconfigdata.bat` *dbusername dbpassword* and pressing **Enter**.
     For example, `exportconfigdata.bat dcmadmin passw0rd`.
- On UNIX systems, complete the following steps.
  a. Open a terminal and navigate to the *Install_Home*\utilities directory.
  b. Export configuration data by typing `exportconfigdata` *dbusername dbpassword* and pressing **Enter**.
     For example, `exportconfigdata dcmadmin passw0rd`.

Clearing performance data

This batch file enables you to delete old performance data which was not deleted due to purge failure. It also enables to you delete performance data before you reach the purging limit to free disk space.

To clear performance data (all time series child table data), complete the following steps.

- On Windows systems, complete the following steps.
  a. Open a command prompt and navigate to the *Install_Home*\utilities directory.
  b. Clear performance data by typing `clear-performance-data.bat` *dbusername dbpassword* and pressing **Enter**.
     For example, `clear-performance-data.bat dcmadmin passw0rd`.
- On UNIX systems, complete the following steps.
  a. Open a terminal and navigate to the *Install_Home*\utilities directory.
  b. Clear performance data by typing `clear-performance-data` *dbusername dbpassword* and pressing **Enter**.
     For example, `clear-performance-data dcmadmin passw0rd`.

Clearing sFlow data

This batch file enables you to delete old sFlow data which was not deleted due to purge failure. It also enables to you delete sFlow data before you reach the purging limit to free disk space.

To clear sFlow data (all time series child table data), complete the following steps.

- On Windows systems, complete the following steps.
  a. Open a command prompt and navigate to the *Install_Home*\utilities directory.
  b. Clear sFlow data by typing `clear-sflow-data.bat` *dbusername dbpassword* and pressing **Enter**.
     For example, `clear-sflow-data.bat dcmadmin passw0rd`.
- On UNIX systems, complete the following steps.
  a. Open a terminal and navigate to the *Install_Home*\utilities directory.
  b. Clear sFlow data by typing `clear-sflow-data` *dbusername dbpassword* and pressing **Enter**.
     For example, `clear-sflow-data dcmadmin passw0rd`.

Flow Vision

VM Insight

NOTE
VM Insight support requires Fabric OS v8.1.0 and later.

VM Insight configures the Management application to query switches periodically to obtain the data pertaining to any virtual machines (VMs) logged into a fabric. The Management application collects the VM Name as well as the VM Instance UUID (which is registered by the HBA on the Application server). When the vCenter is discovered, the Management application automatically maps the VM Name to the VM Instance UUID. If the vCenter is not discovered, only the VM Instance UUID displays. The VM data displays in the **Virtual Machine** tab of the **vCenter Host** properties dialog box as well as other feature dialog boxes (such as the **Flow Vision** dialog box (**Flow Definitions** table)).

VM Insight enables you to associate any applications running on the VM by mapping the application name to the VM Instance UUID and VM Name. For more information, refer to "Adding an application name to a VM" on page 482.

NOTE
The vCenter must be discovered to add an application name to the VM and look into VM section for adding an application.

MAPS and Flow vision enhancements for VM Insight is available in respective section, For more information, refer to "Monitoring flows" on page 1021, "Creating a Flow Monitor flow definition" on page 1026, and "Importing Flow definitions" on page 1237.

NOTE
The Fabric OS 24-port, 32 Gbps and 64-port, 32 Gbps switches do not support VM-based IOS measures statistics.

Flow Vision overview

Flow Vision is a Fibre Channel (FC) SAN network diagnostic tool supported on all platforms supported by Fabric OS 7.2 and later that provides you with a comprehensive vision of fabric traffic flows and with the ability to non-disruptively create and capture copies of traffic flows for later analysis. Flow Vision also provides a test flow generation capability that you can use to pre-test a SAN infrastructure for robustness. This test flow generation capability is also useful for testing the internal connections on a switch before deploying the switch into a production environment.

NOTE
You cannot run Flow Vision and Advanced Performance Monitor (APM) or Port Mirroring at the same time on a chassis (across logical switches).

Supported hardware platforms

Flow Vision is supported on Fabric OS platforms using 8 Gbps-, 16 Gbps, and 32 Gbps-capable FC platforms; there are no platform exclusions. For a list of port types that support Flow Vision, refer to "Supported port types" on page 1014.

Flow Vision requires the Fabric Vision (FV) license ***or*** both the Fabric Watch (FW) and the Advanced Performance Monitor (APM) licenses must be installed on the hardware platform. For more information about Fabric OS licenses, refer to the ***Software Licensing Guide***.

For details on Flow Vision feature and parameter support on switch platforms, Access Gateway switches, and virtual fabrics, refer to "Supported port configurations for each feature" on page 1018.

Supported port types

Flow Vision is supported on the following 8 Gbps-, 16 Gbps-, or 32 Gbps-capable FC port types. For more information on support for a specific feature, such as Flow Monitor, Flow Mirror, or Flow Generator, refer to the appropriate sections on those features.

- E_Ports
- F_Ports
- VE_Ports (supported only on 16 Gbps 24-FC port, 18 GbE port switch and 32 Gbps, Router Extension Blade)
- EX_Ports
- E_Port trunk
- F_Port trunk
- Mirror Port
- SIM-Ports
- Base E-port

Flow Vision terminology

Table 83 explains the terms used in the discussion of Flow Vision.

TABLE 83 Flow Vision terminology

Term	Description
Defined flow	User-created flow; it can be active or inactive.
Local flow	Flow defined on the switch on which you run the flow.
Root flow	Instance of a static flow used to create learned flows.
Static flow	Flow created without using a learned flow.
Sub-flow	System auto-created flow based on a root flow. There can be more than one sub-flow.
Remote flow	Flow defined on a different switch from the one on which you are viewing it.
Learned flow	Flow defined using an asterisk (*) for the source and destination end devices, which enables Flow Vision to learn all of the source and destination device pairs passing through the switch through a specified ingress or egress port without having to identify all the devices.
Local switch	Switch on which you run the flow.
Remote switch	Switch other than the switch on which you run the flow.
ISL	An Inter-Switch Link (ISL) is a protocol that maintains VLAN information in Ethernet frames as traffic flows between switches and routers, or switches and switches.

TABLE 83 Flow Vision terminology

Term	Description
DISL	A Dynamic ISL (DISL) is a physically-connected link between two logical switches that belong to the same Fabric ID (FID). A DISL is dedicated to carry frames only related to the FIDs of connected logical switches.
LISL	A Logical ISL (LISL) is a logical link between two logical switches that is used for control frames. Depending on the fabric topology, a LISL may or may not map directly to a single physical ISL.
XISL	An eXtended ISL (XISL) is a logical link connecting base switches together to form the base fabric. It carries frames from the base fabric and other logical fabrics using the encapsulation and inter-fabric link (IFL) header as identifiers.
Backbone E_Port	The E_Port on a Fibre Channel Routing (FCR)-enabled switch.

Flow Vision features

Flow Vision has three features: Flow Monitor, Flow Generator, and Flow Mirror. The following sections describe each feature and provide links to a more detailed explanation and to step-by-step procedures for each feature.

Flow Monitor

Flow Monitor provides flow monitoring and the gathering of frame statistics for fabric application flows, including the ability to learn (discover) flows automatically. For a more detailed description, refer to "Flow Monitor" on page 1024. For a general step-by-step procedure and example procedures, refer to "Creating a Flow Monitor flow definition" on page 1026 and "Flow Monitor example procedures" on page 1038.

To monitor a Flow Monitor, refer to "Monitoring a Flow Monitor flow" on page 1035.

Flow Generator

Flow Generator simulates and generates test-load traffic in specific flows; this allows you to validate hardware components, connectivity, and verify performance. For a more detailed description, refer to "Flow Generator" on page 1058. For a general step-by-step procedure and example procedures, refer to "Creating a Flow Generator flow definition" on page 1061 and "Flow Generator example procedures" on page 1067.

To monitor a Flow Generator, refer to "Monitoring a Flow Generator flow" on page 1064.

Flow Mirror

Flow Mirror provides the ability to non-disruptively create copies of application flow frames that can be captured for deeper analysis. For a more detailed description, refer to "Flow Mirror" on page 1069. For a general step-by-step procedure and example procedures, refer to "Creating a Flow Mirror flow definition" on page 1071 and "Flow Mirror example procedures" on page 1076.

To monitor a Flow Mirror, refer to "Monitoring a Flow Mirror flow" on page 1073.

Flow Vision flows

A flow is a set of FC frames or packets that share similar traits, such as an ingress port or egress port identifier or any other data that can be used to differentiate one set of related frames or packets from a different set. You specify the following parameters as part of the flow definition:

- Port parameter — Also called the Point of Interest, or where the data you want to examine is from. This consists of an ingress port or an egress port. You can specify the ingress port or the egress port, but not both in a flow definition.

- Frame parameters — These include the following parameters: source device identifier (port ID or WWN), destination device identifier (port address or WWN), LUN, or frame type. At least one frame parameter must be present to define a flow (refer to "Supported frame type parameters" on page 1017 for details on frame types).
- Direction — A direction is implicitly defined from an ingress port to an egress port, or a source device to a destination device. For example, source device=x, destination device=y indicates traffic flowing from x to y. The bidirectional option causes the flow definition to be monitored in both directions.

Figure 478 illustrates how the frame and port parameters apply to a flow.

FIGURE 478 Frame and port parameters

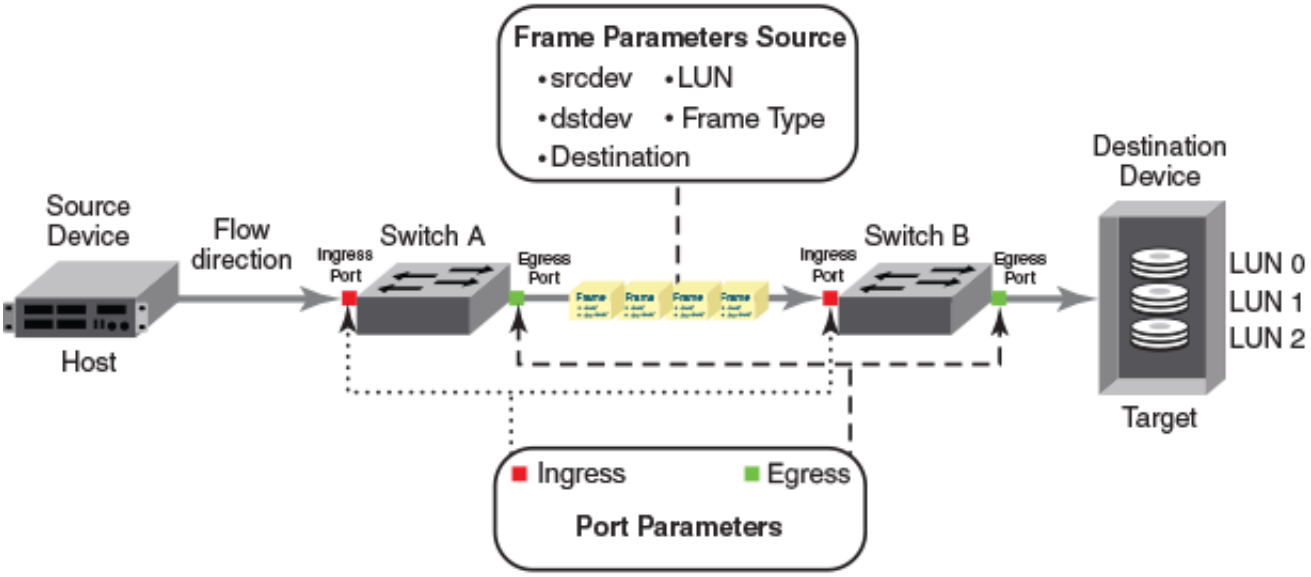

Flow definitions

To define a flow and configure Flow Vision to monitor that flow, you must provide a unique flow name and specify the flow parameters. These parameters identify the sets of related frames and can either be explicitly defined or Flow Vision can learn them through observation.

Flow definitions are stored on the switch on which you create the flow, and are not distributed across the fabric. This means that each switch knows only its own flows and does not know what flows exist on other switches.

You can create three types of flows: Monitor, Generator, and Mirror. For specific procedures for defining flows, refer to the following topics:

- "Creating a Flow Monitor flow definition" on page 1026
- "Creating a Flow Generator flow definition" on page 1061
- "Creating a Flow Mirror flow definition" on page 1071

Flow definition parameters and rules

The rules listed in Table 84 identify the parameters that can be used to define a flow.

TABLE 84 Flow definition rules

Parameters	Field names	Rules
Port	Ingress Egress	• Only one field must be specified • Values must be explicit
Frame	Source Destination LUN Frame Type[1]	• At least one field must be specified • Values for source and destination can be explicit or "*" ("*" indicates a learned flow) • Values for LUN and frame type must be explicit

1. Refer to Table 85 for more information on frame types.

Flow frame type parameters

Frame monitoring can be performed for a variety of frames using predefined frame type parameters. To access the **Frame Type Picker** dialog box, click the ellipsis button next to the **Frame Type** field on the **Add Flow Definition** dialog box.

Table 85 lists the frame type parameters and the frames counted for each.

TABLE 85 Supported frame type parameters

Frame type parameter	Frames counted
abts	Abort sequence frames
baacc	All frames accepted
barjt	All frames rejected
scsi	All SCSI frames (including both command and data frames)
scsiread	Only SCSI read command frames
scsiwrite	Only SCSI write command frames
scsirw	Both SCSI read and write command frames
scsi2reserve	Only SCSI 2 reserve command frames
scsi3reserve	Only SCSI 3 reserve command frames
scsi2release	Only SCSI 2 release command frames
scsi3release	Only SCSI 3 release command frames
scsiresrel2	Only SCSI 2 reserve-release command frames
scsiresrel3	Only SCSI 3 reserve-release command frames
scsitur	Only SCSI test unit ready frames
scsicmdsts	Only SCSI command status frames **NOTE:** This parameter is valid only for Flow Mirror. It implicitly assumes bidirectional flow and looks for both SCSI command and status frames.
scsigoodstatus	Only SCSI status frames with status marked as good (all 0s (zeros) in status byte)
scsicheckstatus	Only SCSI status frames with check status (Check Condition, Busy, Reservation Conflict, Task Full Set)
scsiinquiry	Only SCSI inquiry frames
scsiresvconflict	Only SCSI status frames with reservation conflict set

TABLE 85 Supported frame type parameters (Continued)

Frame type parameter	Frames counted
scsixferrdy	Only SCSI FCP XFER_RDY (transfer ready) frames
srr	Sequence retransmission request for F_Port. **NOTE:** This parameter is valid only for the devices running Fabric OS 8.0.1.

Supported port configurations for each feature

Table 86 lists the supported configurations for each Flow Vision feature using only the basic flow parameters (ingress port and source device, egress port and destination device).

TABLE 86 Port configurations supported in Flow Vision

Feature	Platforms		Switch Configuration Mode	
	16 Gbps- or 32 Gpbs-capable Fibre Channel[1]	8 Gbps-capable Fibre Channel	Access Gateway	Virtual Fabric
Flow Monitor	Supported (E_Ports, EX_Ports, F_Ports, VE_Ports, and XISL_Ports)	Supported (E_Ports, EX_Ports, F_Ports, and XISL_Ports)	Supported (F_Ports only)	Supported
Flow Generator	Supported (SIM-Ports only)	Supported (Destination SIM-Ports only)	Not Supported	Supported
Flow Mirror	Supported (F_Ports and F_Port trunks)	Not Supported	Not Supported	Supported

1. 16 Gbps-capable platforms include the FC8-32E and FC8-48E blades.

Notes on supported configurations

- If you are using at least one advanced parameter (LUN, frame type, or bidirectional), then feature-specific rules apply. Refer to the individual Flow Vision features for specific details.
- Neither ranges nor lists are supported for any parameter.
- Support for 16 Gbps and 32 Gbps F_Ports and F_Port trunks is provided on the following devices:
  - Switches: 24-port, 48-port, and 96-port 16 Gbps and 32 Gbps switches
  - Chassis: 4-slot and 8-slot 16 Gbps and 32 Gbps Backbone Chassis
  - Blades: FC8-32E, FC8-48E, FC16-32, FC16-48, FC16-64, FC32-48, and SX6
- Disabling a SIM-port that is receiving traffic may produce Class 3 discards for the simulated traffic; however, this will have no effect on other traffic flows.

Supported flow parameters

Table 87 lists the supported flow configuration parameters for the Flow Vision Generator, Monitor, and Mirror features.

TABLE 87 Supported flow parameters

Parameter	Flow Generator	Flow Monitor	Flow Mirror
Ingress switch port	Supported	Supported	Supported
Egress switch port	Supported	Supported	Supported
Source end device	Supported	Supported	Supported

TABLE 87 Supported flow parameters

Parameter	Flow Generator	Flow Monitor	Flow Mirror
Destination end device	Supported	Supported	Supported
Source FID/VFID	Supported	Supported	Supported
Destination FID/VFID	Supported	Supported	Supported
Bidirectional	Not applicable	Supported	Supported
Mirror port	Not applicable	Not applicable	Supported
LUN ID	Not applicable	Supported	Supported
Frame type	Not applicable	Supported	Supported

Number of supported flows

On chassis-based platforms, Flow Vision supports a maximum of 512 user-defined flows. On fixed-port platforms, Flow Vision supports a maximum of 128 user-defined flows. However, there is a combined limit from all features of 64 flows (including static flows, root flows, and sub-flows, whether active or inactive) for any one port. In addition, there are individual limits for each Flow Vision feature; Table 88 lists these limits. When you create a flow, Flow Vision verifies that there is no identical active flow.

Refer to the individual features for feature-specific restrictions.

TABLE 88 Feature-specific flow count restrictions in Flow Vision

Feature	Limit to number of flows
Flow Monitor	Up to 64 active flows per port, including static flows, root flows, and sub-flows.
Flow Generator	Up to 4 active flows per port, including static flows, root flows, and sub-flows.
Flow Mirror	One active flow per switch.

Learned flows

Flow Vision creates a learned flow by using an asterisk (*) for the source device, the destination device, or both devices. This allows Flow Vision to learn all of the source and destination device pairs passing through the switch using a specified ingress or egress port without having to identify all the devices.

- Flow Vision uses an asterisk (*) to indicate a learned flow.
- Learned source device or destination device values are only supported on 16 Gbps- and 32 Gbps-capable FC ports.
- Only supports one learning flow per ASIC.
- Each Flow Vision feature uses learned flows as follows:
  - Flow Monitor can learn all the source device and destination device pairs passing through the ingress or egress port defined in a flow. Does not support a learned flow for Flow Monitor flows defined using the LUN, Frame Type, or Bidirectional parameters. Only supports a learned flow on E_Ports, F_Ports, and VE_Ports.
  - Flow Generator can generate traffic to or from every source or destination device that shares the zone with the ingress or egress port defined in a flow. When you configure a learned Generator flow, it queries the Name Server database to identify source and destination devices that are zoned together. These pairings are not automatically changed if either member of the pair changes zones. If either member of the pair changes zones, you must deactivate the flow and then reactivate it to use the new zone values. Flow Generator allocates the first four flows per source ID to zoned destination IDs. The rest of the destination IDs are not tested. For learned flows, no zone enforcement is applied to either the source or destination SIM-Ports.

If the source or destination port for a sub-flow goes offline, the root flow is deactivated and traffic will be stopped on all sub-flows of that root flow.

- Flow Mirror can capture all the source device and destination device pairs passing through the ingress or egress port defined in a learned flow. Supports a learned flow for Flow Mirror flows defined using the LUN, Frame Type, or Bidirectional parameters. If you specify the frame type in the flow definition, the frame type value must be a fixed value for the flow to work. For a list of valid frame type values, refer to "Flow frame type parameters" on page 1017.

Creating a flow from a learned sub-flow

To create a new flow from a learned sub-flow, complete the following steps.

1. Select the device on which you defined a learned flow and select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows in the **Flow Definitions** table. For more information about the fields and components of the **Flow Definitions** table, refer to "Flows Definitions table fields and components" on page 1021.

2. Select one or more learned flows in the **Flow Definitions** table and click the right arrow button to display the sub-flows in the **Flows** table.
3. Select one learned sub-flow from which you want to create the new learned flow from the **Flows** table.
4. Click **Define Flow**.

   The **Create Flow Definition** dialog box displays.

5. Enter a name for the new flow in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

6. Select the **Active** check box to make this the active flow.

   NOTE
   Selecting the **Active** check box deactivates the existing learned flow definition.

   Clear the **Active** check box to create a new inactive flow.

7. Click **OK** to create the new learned flow.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. For more information, refer to "Monitoring a Flow Monitor flow" on page 1035, "Monitoring a Flow Generator flow" on page 1064, or "Monitoring a Flow Mirror flow" on page 1073.

Monitoring flows

Flow Vision enables you to monitor flows defined on devices through the Management application or CLI. Summary data for flows includes all measures supported by the Flow Vision features as well as MAPS violations on monitored flows. You can display summary data for multiple flows, set the time durations, hide and display SCSI and frame-related measures, launch performance graphs for flow data, and view flow data for any discovered fabric.

To view the summary data for a flow, complete the following steps.

1. Select **Monitor > Fabric Vision > Flow > Monitor** or click the Flow Vision icon on the toolbar.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows (active and inactive) in the **Flow Definitions** table, as shown is Figure 479.

FIGURE 479 Flow Vision dialog box (Flow Definitions table)

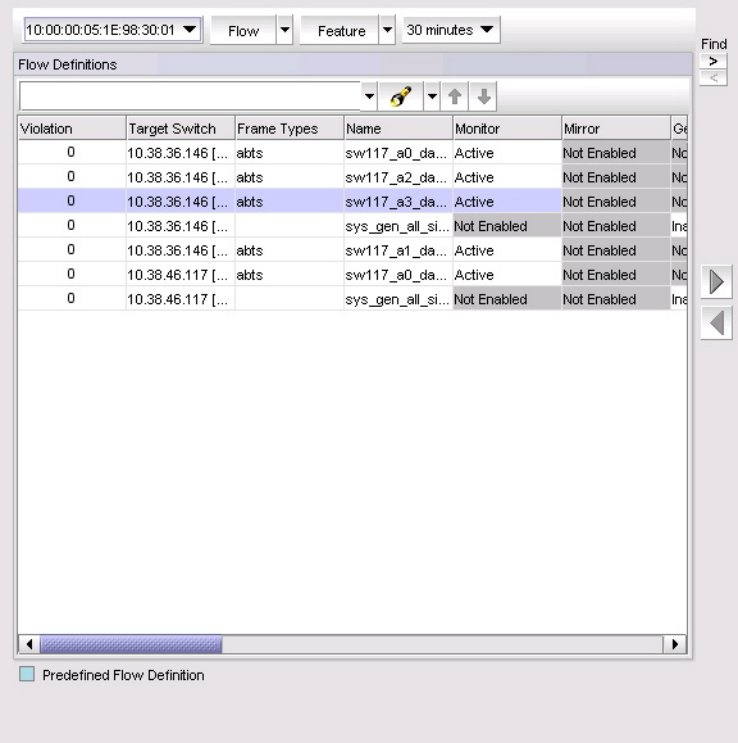

The **Flow Definitions** table has the following general characteristics and functions:

- Data updates dynamically every 5 minutes.
- You can sort the table by clicking any column head. You can reverse the sort order by clicking the column head again.

Table 89 describes the data displayed in the **Flow Definitions** table.

TABLE 89 Flows Definitions table fields and components

Field and components	Description
Violation	The number of Monitoring and Alerting Suite (MAPS) violations for the flow over the selected time duration. This field updates dynamically every 5 minutes with the violation count received from MAPS.
Target Switch	The switch on which you created the flow definition.
Frame Type	All frame types defined in the flow definition.

TABLE 89 Flows Definitions table fields and components (Continued)

Field and components	Description
Flow Name	The user-defined name for the flow definition.
Monitor **Mirror** **Generator**	Displays the status of the flow type. Possible values include: • Active — This flow type is defined and active. • Not Enabled (gray cell) — This flow type is not defined. • Inactive — This flow type is defined, but the feature is not active.
Source	The source identifiers defined in the flow definition.
Source Info	The icon and name for the source device. The device name is a hyperlink. Click to launch the device's property sheet. This field is empty if the source device is not defined in the flow definition.
Destination	The port number of the destination device defined in the flow definition. An * (asterisk) indicates learned flows.
Destination Info	The icon and name for the destination device. The device name is a hyperlink. Click to launch the device property sheet. This field is empty if the destination device is not defined in the flow definition.
Ingress Port	The ingress port defined in the flow definition. The port name is a hyperlink. Click to launch the port property sheet. A yellow icon indicates a bottlenecked port. Refer to "SAN port icons" on page 305 for details.
Egress Port	The egress port defined in the flow definition. The port name is a hyperlink. Click to launch the port property sheet. A yellow icon indicates a bottlenecked port. Refer to "SAN port icons" on page 305 for details.
Source Fabric ID	The fabric identifier of the source defined in the flow definition.
Destination Fabric ID	The fabric identifier of the destination defined in the flow definition.
LUN	Any LUN values defined in the flow definition.
Bi-Directional	Whether the flow is bidirectional (**yes**) or not (**no**).
Flow Definition Persistence	Whether the flow is configured to persist over switch reboots (**yes**) or not (**no**).
Size	The payload size defined in the flow.
Pattern	The payload pattern defined in the flow.
Mirror Port	The mirror port identifier defined in the flow definition.
Source Entity ID	The Instance UUID/Source Entity ID of the virtual machine.
VM Name	The name of the virtual machine.
Application Name	The Application name, which is running in the virtual machine.

2. **Change the fabric view by selecting a different fabric from the Fabric list.**

   Note that the fabric must contain at least one Flow Vision-capable device.

3. Define a new flow by selecting **Flow > Add**.

   For step-by-step instructions, refer to "Creating a Flow Monitor flow definition" on page 1026, "Creating a Flow Generator flow definition" on page 1061, or "Creating a Flow Mirror flow definition" on page 1071.

4. Clear the flow statistics by selecting the flow that you want to reset and selecting **Flow > Reset**.

   For step-by-step instructions, refer to "Resetting flow statistics" on page 1023.

5. Delete a flow by selecting the flow that you want to delete and selecting **Flow > Delete**.

   For step-by-step instructions, refer to "Deleting flows" on page 1024.

6. Configure a MAPS monitor for a flow by selecting the flow and selecting **Flow > MAPS > Configure**.

   For step-by-step instructions, refer to "Configuring a MAPS policy" on page 1232.

7. View MAPS viloations for a flow by selecting the flow and selecting **Flow > MAPS > Violations**.

   For step-by-step instructions, refer to "Viewing MAPS violations" on page 1255.

8. Customize the payload and pattern for a Generator flow by selecting the Generator flow and selecting **Feature > Generator > Configure**.

   For step-by-step instructions, refer to "Customizing Flow Generator flows" on page 1063.

9. Activate a flow by selecting one or more inactive flows and selecting **Feature >** ***feature_name*** **> Activate** (where ***feature_name*** is Monitor, Generator, or Mirror).

10. Deactivate a flow by selecting one or more active flows and selecting **Feature >** ***feature_name*** **> Deactivate** (where ***feature_name*** is Monitor, Generator, or Mirror).

11. Reset a flow by selecting one or more flows and selecting **Feature >** ***feature_name*** **> Reset** (where ***feature_name*** is Monitor, Generator, or Mirror).

12. Set the time interval for monitoring the flow in the **Time duration** list.

    Possible values are 30 minutes, 1 hour, 6 hours, 12 hours, 1 day, 3 days, 1 week, and 1 month.

13. Search for a flow definiton by entering text in the search field and pressing **Enter**.

    For more information about searching for content, refer to "Search" on page 313.

14. Highlights system-defined flows in light blue colour in the **Flow Definitions** table and diplays the **Predefined Flow Definition** legend at the bottom of the Flow Vision dialog box.

15. Select the flow that you want to monitor in the **Flow Definitions** table.

16. Click the right arrow button to display the selected flow in the **Flows** table.

    For more information about the summary data that displays for each flow type, refer to "Monitoring a Flow Monitor flow" on page 1035, "Monitoring a Flow Generator flow" on page 1064, or "Monitoring a Flow Mirror flow" on page 1073

Resetting flow statistics

Flow Vision allows you to clear (reset) the flow statistics record for any defined flow.

To clear all statistics for a flow, complete the following steps.

1. Select the device on which you defined the flow and select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows in the **Flow Definitions** table.

2. Select the flow that you want to reset in the **Flow Definitions** table.

3. Select **Reset** from the **Flow** list.

4. Select ***feature_name*** **> Reset** (where ***feature_name*** is Monitor, Generator, or Mirror) from the **Feature** list.

   This clears the flow feature data for the selected device.

Activating flows

Activating a flow automatically clears any existing flow statistics for that flow.

To activate a flow definition, complete the following steps.

1. Select the device on which you defined the flow and select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows in the **Flow Definitions** table.

2. Select one or more inactive flows that you want to activate in the **Flow Definitions** table.
3. Select ***feature_name*** **> Activate** (where ***feature_name*** is Monitor, Generator, or Mirror) from the **Feature** list.

   This activates the selected flow definitions.

Deactivating flows

You can deactivate Flow Monitor flows without deleting them. This allows you to create and store a "library" of flows that you can activate when needed without having to recreate them every time they are needed.

To deactivate a flow definition, complete the following steps.

1. Select the device on which you defined the flow and select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows in the **Flow Definitions** table.

2. Select the flow that you want to deactivate in the **Flow Definitions** table.
3. Select ***feature_name*** **> Deactivate** (where ***feature_name*** is Monitor, Generator, or Mirror) from the **Feature** list.

   This deactivates the selected flow.

Deleting flows

To delete a flow definition, complete the following steps.

1. Select the device on which you defined the flow and select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows in the **Flow Definitions** table.

2. Select the flow that you want to delete in the **Flow Definitions** table.
3. Select **Delete** from the **Flow** list.

Flow Monitor

Flow Monitor enables you to monitor all the traffic passing through E_Ports, EX_Ports, F_Ports, VE_Ports, and XISL_Ports using hardware-supported flow parameters. It also lets you define your own monitoring flows using combinations of ingress or egress ports, source and destination devices, logical unit numbers (LUNs), and frame types to create a flow definition for a specific use case.

> NOTE
> Beginning with Fabric OS 8.0.1 release, you can configure flow definition for VE_Ports and monitor the traffic only for the 16 Gbps 24-FC port, 18 GbE port switch and 32 Gbps, Router Extension blade. Pre-defined flows are not supported on VE_Ports.

Flow Monitor provides support for monitoring the following flows and traffic:

- Learned and static flows for traffic passing through E_Ports, F_Ports, and VE_Ports
- Learned and static flows monitoring edge-to-edge traffic, edge-to-backbone traffic, and backbone-to-edge traffic passing through EX_Ports
- Learned and static flows monitoring traffic inside logical fabrics and inter-fabric (routed) traffic passing through XISL_Ports
- Learned and static flows monitoring inter-fabric traffic and backbone traffic passing through backbone E_Ports
- Monitoring Ironware OS measures in Realtime graph for the F_Ports in Gen 6 chassis running Fabric OS 8.0.1 or later

In Fabric OS 7.1.x and earlier, the Advanced Performance Monitor (APM) provided the following monitors: End-to-End, Frame-based, ISL, and Top Talker. In Fabric OS 7.3.0, Flow Monitor provides you with the following abilities in addition to those provided by the APM:

- Monitoring of application flows (for example, a flow within a fabric from a host to a target or LUN) at a specified port.
- Comprehensive visibility into application flows in a fabric, including the ability to learn (discover) flows automatically.
- When N_Port ID Virtualization (NPIV) is used on the host, you can monitor Virtual Machine (VM)-to-LUN level performance.
- Capturing statistics for specified flows, which provides insights into application performance. These statistics include transmitted and received frame counts, transmitted and received frame throughput rates, SCSI Read and SCSI Write frame counts, the number of SCSI Reads and Writes per second, as well as others.

  A sample use case would be to monitor throughput statistics for inbound traffic between a source device and a destination device. For an example procedure for this use case, refer to "Monitoring LUN level statistics" on page 1038.
- Monitoring of various frame types at a switch port to provide deeper insights into storage I/O access patterns at a LUN, reservation conflicts, and I/O errors. Examples of the frame types that can be monitored include SCSI Aborts, SCSI Read, SCSI Write, SCSI Reserve, all rejected frames, and many others. For a list and description of the frame types that can be monitored, refer to "Flow frame type parameters" on page 1017.
- SCSI Read/Write Frame Count and SCSI Read/Write Data statistics are supported only for F_Ports for any flow configuration where either the source device or destination device exists on the switch, and the flow is defined using a combination of source device, destination device, ingress port, or egress port (with or without bidirectional), or a combination of source device, destination device, LUN, ingress port, or egress port.
- Integration with the Monitoring and Alerting Policy Suite (MAPS) service to enable threshold-based monitoring and alerting based on flows. For more information on integration with MAPS, refer to the "Monitoring and Alerting Policy Suite integration with Flow Vision" on page 1092.

Flow Monitor limitations

The following limitations apply to all Flow Monitor flows:

- Only one active learned flow is supported per ASIC.
- Only supports learned flows on 16 Gbps- and 32 Gbps-capable FC platforms.
- For scsiread, scsiwrite, and scsirdwr frame type parameters, only monitors SCSI 6-, 10-, 12-, and 16-bit Read and Write values. Flow Monitor cannot monitor Read Long and Write Long values.
- Flow Monitor is not supported on ports with Encryption or Compression enabled.
- Flow Monitor can monitor IFL flows only on EX_Ports in an FC router.
- Flow Monitor cannot monitor Inter-Fabric Link (IFL) flows on E_Ports or F_Ports.
- Flow Monitor cannot monitor flows that are using frame redirection for encryption.
- Flow Monitor can monitor flows of VE_Ports of shared Gige.
- You cannot convert Flow Monitor flows to Fabric OS 7.1.x flow performance monitors.
- The calculated Rx and Tx frame size values displayed in the output are accurate within a range of -4 through +8 bytes. For example, a frame size value of 256 bytes may actually be anywhere from 252 to 264 bytes.

The following restrictions apply to parameter usage in Flow Monitor flow definitions:

- You cannot use the LUN and Bidirectional parameters in a flow definition.
- You cannot use the Frame Type, LUN, and Bidirectional parameters for learned flows.
- Flow creation is not allowed if Advanced Performance Monitor (APM) or Port Mirroring is enabled. Similarly, APM and Port Mirroring-related operations are not allowed if any flow (active or defined) is present on the switch.

The following limitations apply to 16 Gbps and 32 Gbps-capable FC platforms and 8 Gbps enhanced blades:

- They support a maximum of 2 flows defined using a combination of ingress port and frame type parameters per ASIC chip.
- Each port supports a maximum of 12 flows defined using both egress port and frame type parameters.
- Flow Monitor can only monitor flows that are using EX_Ports.

The following limitations apply to 8 Gbps-capable FC platforms and blades:

- You cannot create a flow definition using both ingress port and frame type parameters.
- Each port supports a maximum of 12 flows defined using both egress port and frame type parameters, except for the following Fabric OS devices, which only support a maximum of 8 flows per port.
  - 24-port, 8 Gbps FC Switch
  - 40-port, 8 Gbps FC Switch
  - 80-port, 8 Gbps FC Switch
  - Embedded 12-port, 8 Gbps Switch
  - Embedded 24-port, 8 Gbps Switch
  - Embedded 16-port, 8 Gbps Switch
  - Embedded 24-port, 8 Gbps Switch
  - 8 Gbps Extension Switch
- They cannot show statistics for SIM-ports.
- They do not support learned flows.

Creating a Flow Monitor flow definition

This procedure provides step-by-step instructions for configuring a Monitor flow definition. For more specific example procedures, refer to "Flow Monitor example procedures" on page 1038.

1. Select the object on which you want to configure a flow monitor from the Connectivity Map or Product List. Options include:
   - Switch port
   - Initiator port
   - Target port
   - Switch that supports Flow Vision
2. Select **Monitor > Fabric Vision > Flow > Add**.

   NOTE
   You can also right-click on any of these objects in the Connectivity Map or Product List and select **Flow > Add**.

   The **Add Flow Definition** dialog box displays populated with criteria and flow identifiers (such as source and destination identifiers) for the appropriate flow definition based on the selected object. For example, if you selected a target port, the **Add Flow Definition** dialog box is populated with the source device port identifier, ingress switch port identifier, monitor type, and destination device (as an asterisk). An asterisk (*) denotes any selected device in the same zone as the source device.

FIGURE 480 Add Flow Definition dialog box

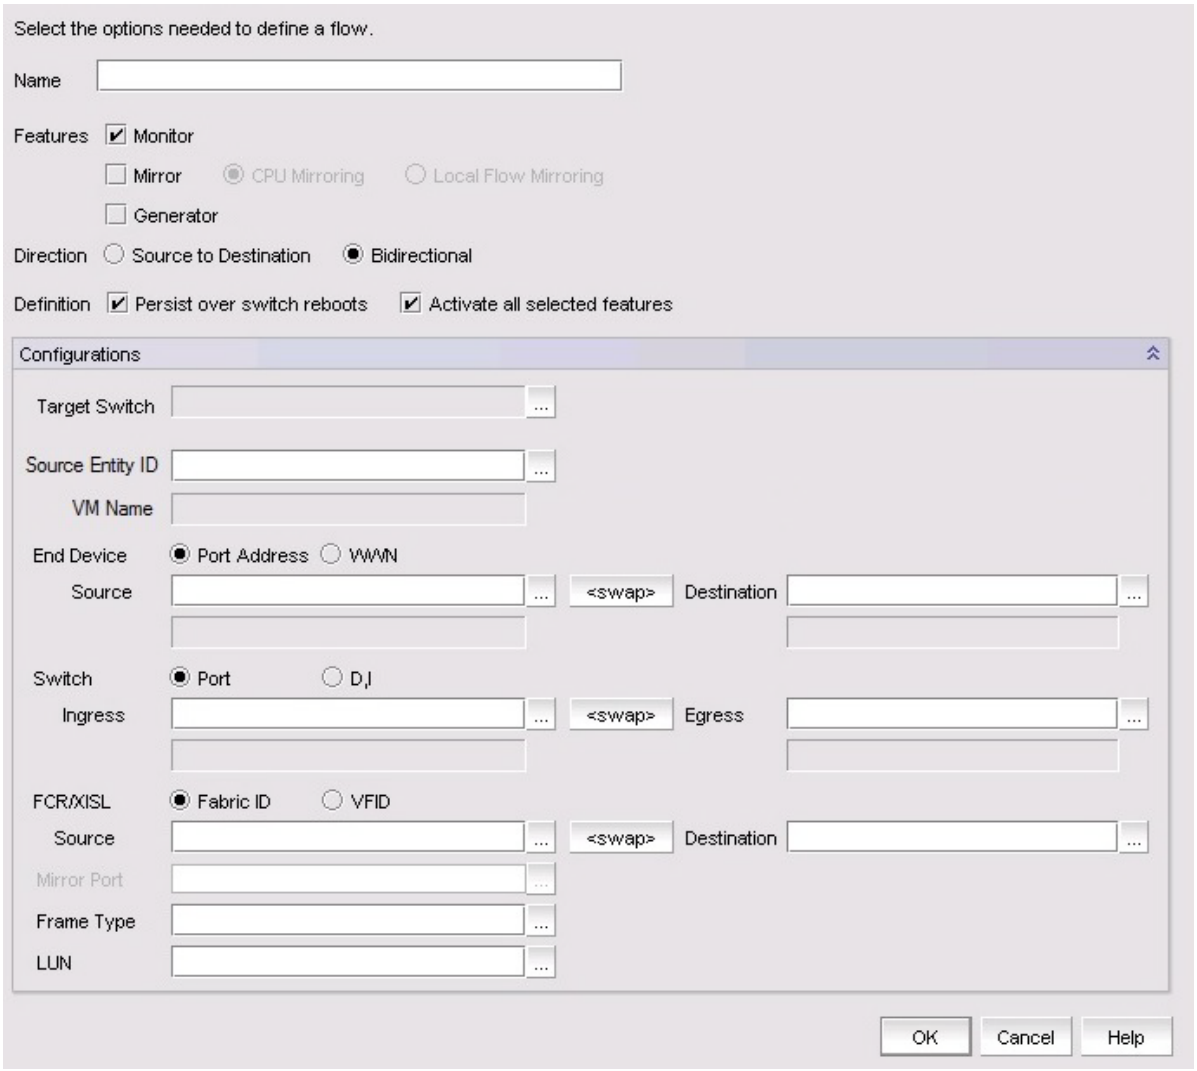

3. Enter a name for the flow in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

4. Select the **Monitor** check box.
5. Select one of the following options to define the flow direction:
   - **Source to Destination** — Select to define the flow direction from an ingress port to an egress port, or a source device (in traffic) to a destination device (out traffic).
   - **Bidirectional** — Select to define the flow direction to be monitored in both directions.

     NOTE
     You cannot use the LUN and Bidirectional parameters in a flow definition.

6. If you want to persist this flow definition over reboots, select the **Persist over switch reboots** check box.

7. If you want to immediately activate the flow after creation, select the **Activate all selected features** check box.
8. Change target switch for the flow definition by clicking the ellipsis button to the right of the **Target Switch** field.

   The **Select Switch** dialog box displays. To manually set the **Target Switch**, refer to "Selecting the target switch from a list of Fabric Vision-capable switches" on page 1030.
9. Enter the entity ID in the Source Entity ID field or click the ellipsis button to select an entity ID from the list.

   The Select VM picker dialog will display the VM Name and Entity ID only when a vCenter is discovered. The selected VM Name is displayed in the VM Name field of Add flow Definition dialog.
10. Select one of the following format options for **End Device** mode:
    - **Port Address** (port ID) — Select to display the source and destination device address using the port ID.
    - **WWN** (world wide name) — Select to display the source and destination device address using the port WWN.
11. Enter the port ID or WWN of the source port in the **Source** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the source port from a list, refer to "Selecting an end device port from a list of available device ports" on page 1030.

    When you enter a port ID or WWN in the **Source** or **Destination** fields, a port information field displays beneath with the port label based your topology layout settings (refer to "Changing the port label" on page 324). If you enter an asterisk (*) or no value in the **Source** or **Destination** fields, this field remains blank.
12. Enter the port ID or WWN of the source port in the **Destination** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the destination port from a list, refer to "Selecting an end device port from a list of available device ports" on page 1030.
13. (Optional) If you want to swap source and destination device port information, click **<swap>**.
14. Select one of the following format options for **Switch** mode:
    - **Port (slot/port)** — Select to display the switch ingress or egress port using the slot and port number.
    - **D,I** (domain ID, port number) — Select to display the switch ingress or egress port using the domain ID and port number.
15. Enter the ingress port in Port (slot/port) or D,I (domain ID,port number) format in the **Ingress** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the ingress port from a list, refer to "Selecting an ingress or egress port from a list of available switch ports" on page 1032.

    NOTE
    You can enter only the port number. For Backbone chassis, the slot number cannot be 0 (zero). For switches, you must enter 0 (zero) as the slot number. For logical (virtual) switches, follow the rule for the physical chassis (either Backbone chassis or switch) from which you created the logical switch.

    When you enter a value in the **Ingress** or **Egress** fields, a port information field displays beneath with the port label based your topology layout settings (refer to "Changing the port label" on page 324). If you enter an asterisk (*) or no value in the **Ingress** or **Egress** fields, this field remains blank.

16. Enter the egress port data in Port (slot/port) or D,I (domain ID,port number) format in the **Egress** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the egress port from a list, refer to "Selecting an ingress or egress port from a list of available switch ports" on page 1032.

    NOTE
    You must enter the slot number and the port number for the chassis, the slot number cannot be 0 (zero). For switches, you must enter the port number without any slot number.

17. For FCR or virtual fabrics, select one of the following format options for **FCR/XISL** mode:

    - **Fabric ID** — Select to configure a flow definition for a FCR fabric.
    - **VFID** (Virtual Fabric ID) — Select to configure a flow definition for virtual fabrics.

18. For FCR or virtual fabrics, enter the fabric ID or virtual fabric ID in the **Source** field.

    Enter an asterisk (*) to use any port. To select the source from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

19. For FCR or virtual fabrics, enter the fabric ID or virtual fabric ID in the **Destination** field.

    Enter an asterisk (*) to use any port. To select the destination from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

20. (Local Flow Mirroring only) Enter the mirror port number in the **Mirror** field or click the ellipsis button to select a port from the list.

    Only enabled when you select the **Local Flow Mirroring** option for the **Mirror** feature. Only supported on 16 Gbps-capable ports on devices running Fabric OS 7.3 or later.

    To select a port from a list, refer to "Selecting a mirror port from a list of available ports" on page 1033.

21. Enter a frame type in the **Frame Type** field or click the ellipses button select a frame type from a list.

    NOTE
    You cannot combine frame type and bidirectional parameters for Mirror flows.

    To select a frame type from a list, refer to "Selecting a mirror port from a list of available ports" on page 1033.

22. Enter a LUN ID in the **LUN** field or click the ellipses button to select a LUN ID from the list.

    NOTE
    You cannot use LUN IDs for Bidirectional flows.

    NOTE
    You cannot combine frame type, LUN, and bidirectional parameters for learning flows.

    LUN IDs can be from 0 through 65535. To select a LUN ID from a list, refer to "Selecting a LUN ID from a list of available LUNs" on page 1034.

23. Select **OK** to save the definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow (refer to "Monitoring a Flow Monitor flow" on page 1035).

Selecting the target switch from a list of Fabric Vision-capable switches

To change target switch for the flow definition, complete the following steps.

1. Click the ellipses button to the right of the **Target Switch** field.

   The **Select Switch** dialog box displays with all Fabric Vision-capable switches listed in the **Available Switches** list.

2. Select the **Expand All** check box to expand the **Group/Product** tree to show all Fabric Vision-capable devices.
3. Select a switch in the **Available Switches** list and click the right arrow to move to the **Selected Switch** list.

   Move the switch back to the **Available Switches** list by clicking the left arrow.

4. Click **OK** on the **Select Switch** dialog box.

   The selected switch displays in the **Target Switch** field.

Selecting an end device port from a list of available device ports

To select the port for the source or destination device, complete the following steps.

1. Click the ellipses button to the right of the **Source** or **Destination** field.

   The **Select Device Port** dialog box displays expanded to display all supported device ports for the target switch in the **Available Device Ports** list.

   Unsupported ports do not display in the **Available Device Ports** list.

2. Select one of the following options to display either ports in specific zones or available fabrics, devices, and ports for selection.
   - Select **Products and Ports** from the **Show** list to display available fabrics, devices, and ports as in Figure 481.
   - Select **Zones** from the **Show** list to display available ports under specific zones as shown in Figure 482.

FIGURE 481 Select Device Ports dialog box. (Products and Ports selected)

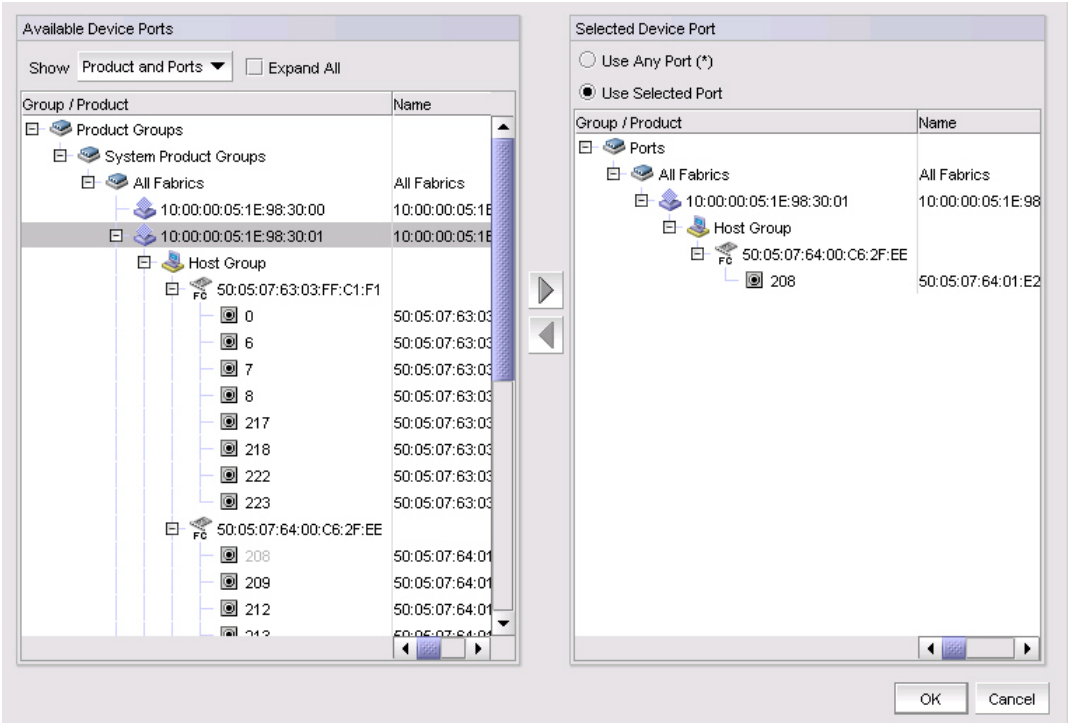

FIGURE 482 Select Device Port dialog box (Zones selected)

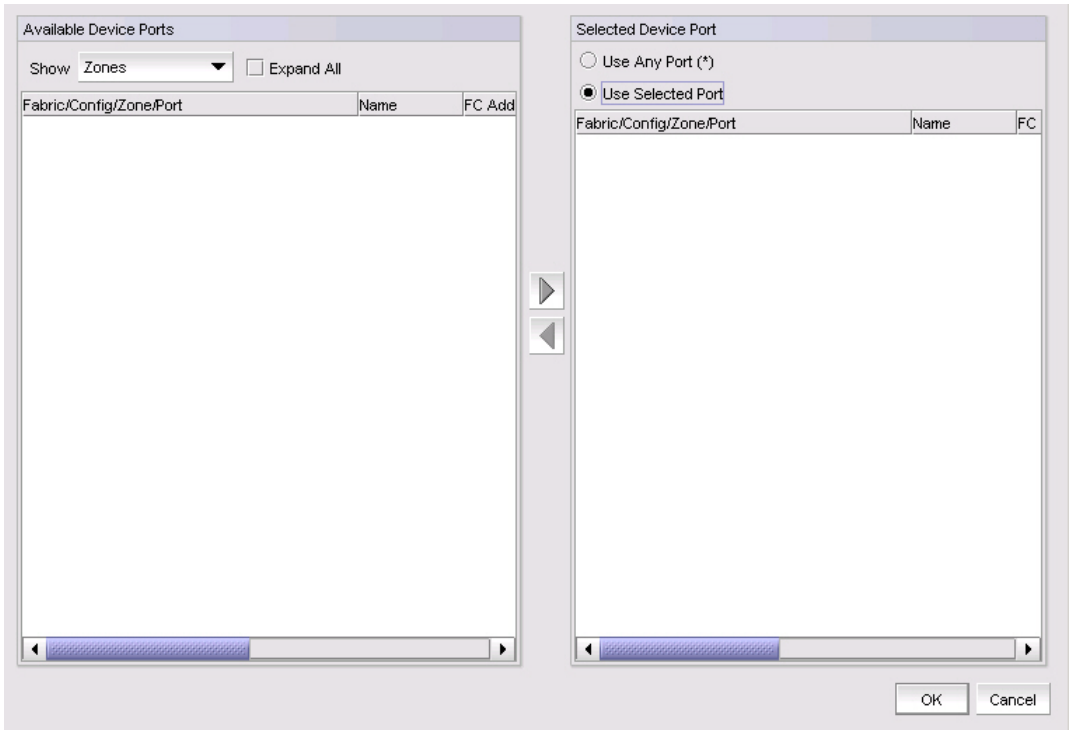

3. Select the **Expand All** check box to expand the **Group/Product** tree to show supported device ports for all discovered devices.

4. Select a device port in the **Available Device Ports** list and click the right arrow to move to the **Selected Device Ports** list.

   Select the **Use Any Port (*)** option to use any supported device port. If you already selected a device port, selecting this option clears the selection.

   Note that moving a device port to the **Selected Device Ports** list automatically selects the **Use Selected Port** option.

   Move the device port back to the **Available Device Ports** list by clicking the left arrow.

5. Click **OK** on the **Select Device Port** dialog box.

   The selected device port displays in the **Source** or **Destination** field.

Selecting an ingress or egress port from a list of available switch ports

To select the ingress or egress port for the switch, complete the following steps.

1. Click the ellipses button to the right of the **Ingress** or **Egress** field.

   The **Select Switch Port** dialog box displays expanded to display all supported device ports for the target switch in the **Available Switch Ports** list.

   Unsupported ports do not display in the **Available Switch Ports** list.

2. Select the **Expand All** check box to expand the **Group/Product** tree to show supported switch ports for all discovered devices.
3. Select a switch port in the **Available Switch Ports** list and click the right arrow to move to the **Selected Switch Ports** list.

   Move the switch port back to the **Available Switch Ports** list by clicking the left arrow.

4. Click **OK** on the **Select Switch Port** dialog box.

   The selected switch port displays in the **Ingress** or **Egress** field. Note that the **Target Switch** field updates dynamically based on the ingress port selection.

Selecting a fabric or virtual fabric ID from a list of available products

To select the fabric or virtual fabric ID, complete the following steps.

1. Click the ellipses button to the right of the **Source** or **Destination** field.

   The **Select Fabric ID** dialog box displays expanded to display all supported fabrics in the **Available Products** list.

   Unsupported products do not display in the **Available Products** list.

FIGURE 483 Select Fabric ID dialog box

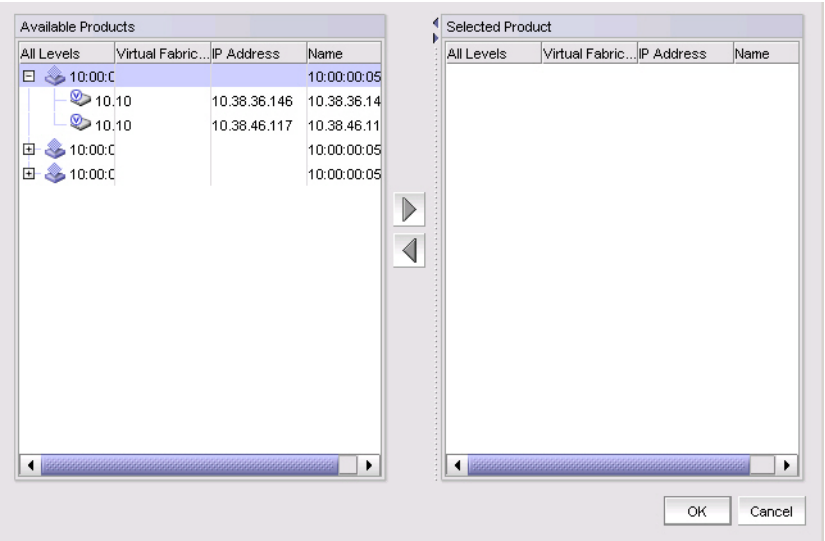

2. Select a fabric or virtual fabric ID in the **Available Products** list and click the right arrow to move to the **Selected Products** list.

   Move the fabric or virtual fabric ID back to the **Available Products** list by clicking the left arrow.

3. Click **OK** on the **Select Fabric ID** dialog box.

   The selected fabric or virtual fabric ID displays in the **Source** or **Destination** field.

Selecting a mirror port from a list of available ports

Only supported on 16 Gbps-capable ports on devices running Fabric OS 7.3 or later.

To select the mirror port, complete the following steps.

1. Click the ellipses button to the right of the **Mirror Port** field.

   The **Select Mirror Port** dialog box displays expanded to display all unoccupied, loopback, and mirror ports for the target switch in the **Available** list.

   Unsupported ports do not display in the **Available** list.

FIGURE 484 Select Mirror Port dialog box

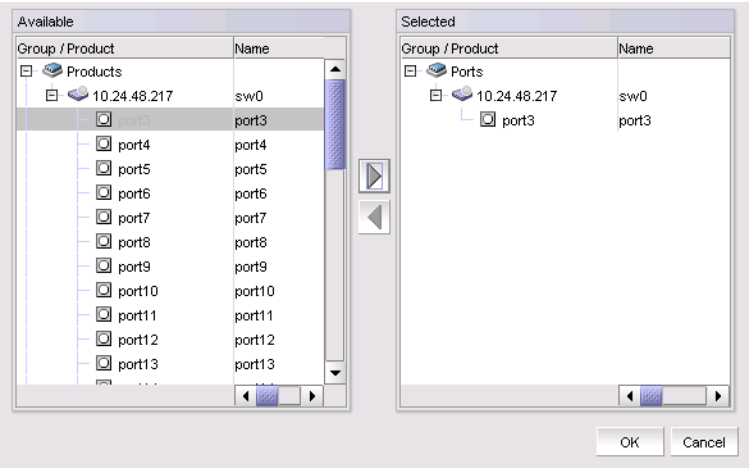

2. Select a switch port in the **Available** list and click the right arrow to move to the **Selected** list.

   Move the switch port back to the **Available** list by clicking the left arrow.

3. Click **OK** on the **Select Mirror Port** dialog box.

   The selected port displays in the **Mirror Port** field.

Selecting a LUN ID from a list of available LUNs

4. Use this dialog box to select the LUNs associated with discovered fabrics.

To select the LUN ID, complete the following steps.

1. Click the ellipses button to the right of the **LUN** field.

   The **Select LUN** dialog box displays with a list of fabrics with associated LUNs.

FIGURE 485 Select LUN dialog box

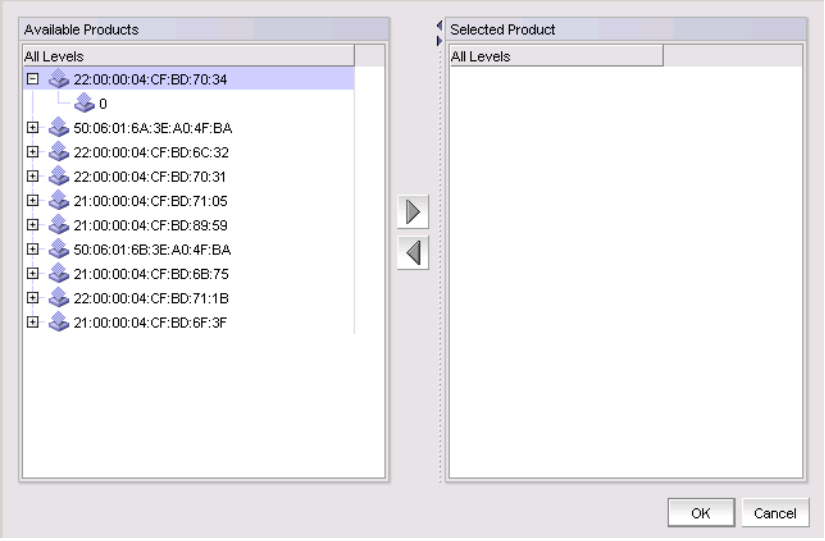

2. Select a LUN ID in the **Available Products** list and click the right arrow to move to the **Selected Products** list.

   Move the switch port back to the **Available Products** list by clicking the left arrow.

3. Click **OK** on the **Select LUN** dialog box.

   The selected LUN ID displays in the **LUN** field.

Monitoring a Flow Monitor flow

To view the summary data for a Flow Monitor flow, complete the following steps.

1. Select the device on which you defined the flow and select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows in the **Flow Definitions** table. For more information about the fields and components of the **Flow Definitions** table, refer to "Flows Definitions table fields and components" on page 1021.

2. Select a time interval for monitoring the flow in the **Time duration** list.

   Possible values are 30 minutes, 1 hour, 6 hours, 12 hours, 1 day, 3 days, 1 week, and 1 month.

3. Select the flow that you want to monitor in the **Flow Definitions** table.

4. Click the right arrow button to display the selected flow in the **Flows** table.

   You can sort the **Flows** table by clicking any column head. You can reverse the sort order by clicking the column head again.

5. Select the **SCSI** check box to display SCSI-related measures.

   SCSI-related measures include SCSI read count, write count, read rate, write rate, read data, write data, and read and write frame data.

   Clear the check box to hide SCSI-related measures.

6. Select the **Frame** check box to display frame-related measures.

   Frame-related measures include transmit (Tx) and receive (Rx) frame count, Transmit frame and receive frame rate, Transmit and receive word count, and transmit and receive throughput.

   Clear the check box to hide frame-related measures.

7. Review the sub-flow data for the selected Flow Monitor flow.

   The **Flows** table, as shown is Figure 497, displays statistics and data for the selected flows. The **Flows** table has the following general characteristics and functions:

   - Data updates dynamically every 5 minutes.
   - Sort the table by clicking any column head. You can reverse the sort order by clicking the column head again.
   - Locate an ingress port, egress port, source device, and destination device in the Product List or Topology Map by right-clicking a sub-flow in the **Flows** table and selecting **Locate >** *port_type* (where *port_type* is Ingress port, Egress port, Source device, or Destination device).
   - Highlights inactive sub-flows in yellow by selecting the **Show inactive flows** check box.

     Inactive sub-flows indicate no updates to the sub-flow statistics for over 15 minutes.

     Clear the check box to hide inactive sub-flows.

- Display a flow in a performance graph. Select a row in the **Flows table and select Performance Graph >** *graph_or_report* (where *graph_or_report* is **Real Time Graph**, **Historical Graph**, **Historical Report**).

  Refer to "Performance integration with Flow Vision" on page 1098 **for additional details. Note that the Performance Graph list is only available when there is at least one sub-flow selected in the Flows table.**

- A single flow definition might yield data in multiple rows in the **Flows** table. For example, if you defined the flow definition source ID (SID) and destination ID (DID) as *, this might result in five rows if the source is communicating with five destination IDs. In the case of a learning flow, a root flow also displays to summarize all sub-flows.
- Each unique sub-flow for the flow definition displays in the **Flows** table if it was reported within the selected time duration. If the last data point did not report that flow, the reported values may be 0.
- The measures that display are based on the flow definition. Therefore, not all columns may be populated for the selected flows.

FIGURE 486 Flow Vision dialog box (Monitor Flows table)

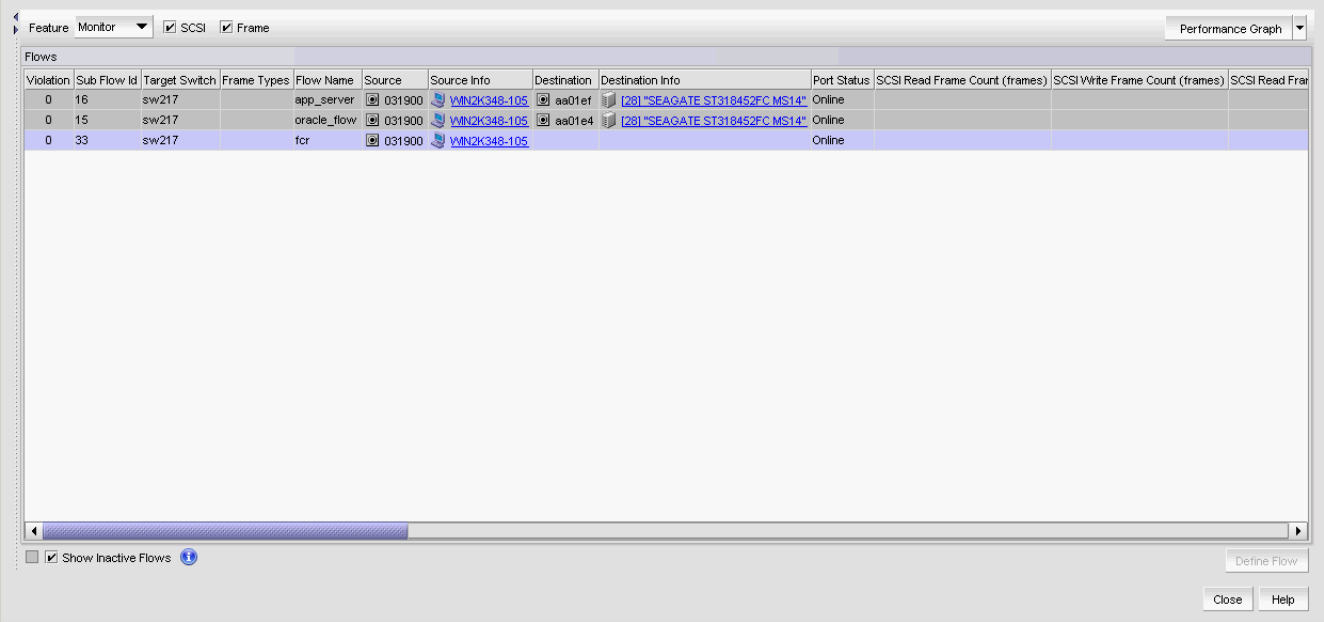

Table 90 describes information on sub-flows displayed in the **Flows** table when you select **Monitor** from the Feature list above the **Flows** table.

TABLE 90 Flows table fields and components (Monitor flow)

Fields and components	Description
Violation	MAPS violation for the flow over the time duration selected. This column gets updated dynamically for every 5 mins with the violation count receiving from MAPS.
Sub Flow ID	The sub-flow database identifier.
Source Entity ID	The Instance UUID/Source Entity ID of the virtual machine.
VM Name	The name of the virtual machine.
Application Name	The Application name, which is running in the virtual machine.
Target Switch	The switch on which you created the flow definition.
Frame Type	All frame types defined in the flow definition.

TABLE 90 Flows table fields and components (Monitor flow) (Continued)

Fields and components	Description
Flow Name	The user-defined name for the flow definition.
Source	The source identifiers defined in the flow definition.
Source Info	The icon and name for the source device. The device name is a hyper link. Click to launch the device's property sheet. This field is empty if the source device is not defined in the flow definition.
Destination	The port number of the destination device defined in the flow definition. An * (asterisk) indicates learned flows.
Destination Info	The icon and name for the destination device. The device name is a hyper link. Click to launch the device's property sheet. This field is empty if the destination device is not defined in the flow definition.
Port Status	The operational status (online or offline) for the port.
SCSI Read Frame Count (frames)	The number of SCSI read frames as reported in the last data point for the flow.
SCSI Write Frame Count (frames)	The number of SCSI write frames as reported in the last data point for the flow.
SCSI Read Frame Rate (f/s)	The SCSI read frame rate in frames per second as reported in the last data point for the flow.
SCSI Write Frame Rate (f/s)	The SCSI write frame rate in frames per second as reported in the last data point for the flow.
SCSI Read Data Rate (MB)	The SCSI read data rate in megabytes as reported in the last data point for the flow.
SCSI Write Data Rate (MB)	The SCSI write data rate in megabytes as reported in the last data point for the flow.
SCSI Read Data Rate (Mbps)	The SCSI read data rate in megabits per second as reported in the last data point for the flow.
SCSI Write Data Rate (Mbps)	The SCSI write data rate in megabits per second as reported in the last data point for the flow.
Transmit Frame Count (frames)	The number of frames transmitted as reported in the last data point for the flow.
Receive Frame Count (frames)	The number of frames received as reported in the last data point for the flow.
Transmit Frame Rate (f/s)	The transmit frame rate in frames per second as reported in the last data point for the flow.
Receive Frame Rate (f/s)	The receive frame rate in frames per second as reported in the last data point for the flow.
Transmit Word Count (MB)	The transmitted word count in megabytes as reported in the last data point for the flow.
Receive Word Count (MB)	The received word count in megabytes as reported in the last data point for the flow.
Transmit Throughput (Mbps)	The transmit throughput in megabits per second as reported in the last data point for the flow.
Receive Throughput (Mbps)	The receive throughput in megabits per second as reported in the last data point for the flow.
Average Transmit Frame Size (bytes)	The average transmit frame size in bytes size as reported in the last data point for the flow.
Average Receive Frame Size (bytes)	The average receive frame size in bytes as reported in the last data point for the flow.
Ingress Port	The ingress port defined in the flow definition. The port name is a hyper link. Click to launch the port's property sheet. A yellow icon indicates a bottlenecked port.
Egress Port	The egress port defined in the flow definition. The port names is a hyper link. Click to launch the port's property sheet. A yellow icon indicates a bottlenecked port.
Source Fabric ID	The fabric identifier of the source defined in the flow definition.
Destination Fabric ID	The fabric identifier of the destination defined in the flow definition.
LUN	Any LUN values defined in the flow definition.
Bi-Directional	Whether the flow is bidirectional (**yes**) or not (**no**).

TABLE 90 Flows table fields and components (Monitor flow) (Continued)

Fields and components	Description
Flow Definition Persistence	Whether the flow is configured to Persist over switch reboots (**yes**) or not (**no**).
Last Updated Time	The date and time the sub-flow was last updated.

Flow Monitor example procedures

The following examples provide step-by-step instructions for using the Flow Monitor feature.

"Monitoring LUN level statistics" on page 1038

"Viewing summary flow data for a specific device pair" on page 1039

"Configuring a learned flow" on page 1040

"Configuring an end-to-end monitor flow" on page 1041

"Configuring a frame monitor flow" on page 1041

"Configuring a Top Talker monitor flow" on page 1042

"Configuring a Flow Monitor flow for a trunk group" on page 1043

"FC router fabrics Flow Monitor flow example procedures" on page 1044

"FC router fabric monitors using virtual port IDs" on page 1049

Monitoring LUN level statistics

A common use of flow monitors is to monitor traffic flowing from a particular ingress port to a specified LUN.

The following example creates a flow named "lunFlow11" that monitors traffic ingressing on port 5 between device 010502 and device 030700 using LUN 4, and then displays the results of that flow.

Figure 487 provides an illustration of what occurs in the example.

FIGURE 487 LUN monitoring flow example

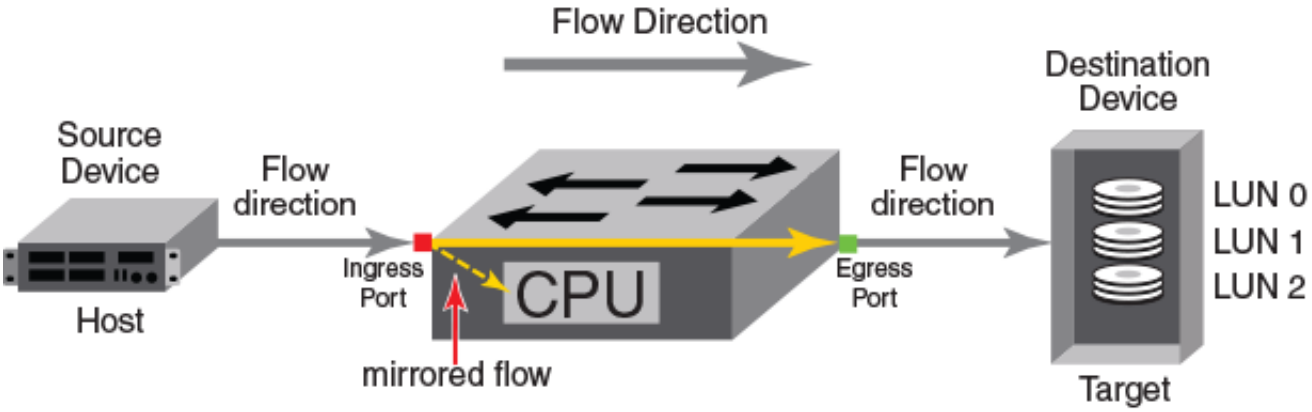

1. Right-click an initiator port and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

- Feature — Monitor
- Direction — Bidirectional
- Source Device — Source identifier (0x010502)
- Destination Device — * (an asterisk allows you to use any port)
- Ingress port — port number (5)

2. Enter a name (lunFlow11) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumerics or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
4. Select the **Activate all selected features** check box to immediately activate the flow after creation.
5. Enter the address or WWN of the source port in the **Destination** field.
6. Enter the LUN identifier (4) in the **LUN** field.
7. Click **OK**.

   The **Flow Vision** dialog box displays with the new flow definition highlighted.

8. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. For more information, refer to "Monitoring a Flow Monitor flow" on page 1035.

Viewing summary flow data for a specific device pair

The following example creates a Flow Monitor flow gathering statistics for frames ingressing through port 30 between device 010000 and device 010100, and then displays the results. The point of interest in this example is port 30; it can be either an E, EX, or F_Port.

1. Right-click an initiator port (30) on the source device and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier (010000)
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (30)

2. Enter a name (summlow1) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumerics or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option to gather flow data from the source device to the destination device.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.

5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Enter the address or WWN of the source port (010100) in the **Destination** field.
7. Click **OK**.

   The **Flow Vision** dialog box displays with the new flow definition highlighted.
8. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow (refer to "Monitoring a Flow Monitor flow" on page 1035).

Configuring a learned flow

The following example illustrates using the learning functionality for flow monitoring. The defined flow monitors for frames ingressing on port 30 between all devices.

1. Right-click an initiator port (30) on the source device and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (30)
2. Enter a name (ingressTT) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumerics or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.
3. Select the **Source to Destination** option to gather flow data from the source device to the destination device.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Enter * in the **Source** and **Destination** fields to monitor frames on the ingress port (30) between all devices.
7. Click **OK**.

   The **Flow Vision** dialog box displays with the new flow definition highlighted.
8. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow (refer to "Monitoring a Flow Monitor flow" on page 1035).

Configuring an end-to-end monitor flow

You can use the **Bidirectional** option to create the equivalent to an end-to-end monitor. The following example creates a bidirectional Flow Monitor flow between device 02d8c0 and device 023a00 egressing port 4/10 of the switch on which the flow is running.

1. Right-click an egress port (4/10) on the device and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier (02d8c0)
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (30)

2. Enter a name (EndtoEndMonitor) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumerics or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
4. Select the **Activate all selected features** check box to immediately activate the flow after creation.
5. Enter the address or WWN of the source port (023a00) in the **Destination** field.
6. Click **OK**.

   The **Flow Vision** dialog box displays with the new flow definition highlighted.

7. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring a frame monitor flow

You can use the **Frame Type** field to create a frame monitor flow. For more information about frame monitoring, refer to "Frame Monitor" on page 1101. The following example creates a Flow Monitor flow that counts SCSI Read-Write (scsirw) frames egressing port 2 of the switch on which the flow is running.

1. Right-click a switch port (2) on which you want to monitor frames and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Egress port — port number (2)

2. Enter a name (scsirw) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

NOTE
For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
4. Select the **Activate all selected features** check box to immediately activate the flow after creation.
5. Enter an asterisk (*) in the **Destination** field.
6. Enter the frame type (scsirw) in the **Frame Type** field.

   To select a frame type from a list, click the ellipsis button. Refer to "Flow frame type parameters" on page 1017 for a list of the supported frame type parameters.

7. Click **OK**.

   The **Flow Vision** dialog box displays with the new flow definition highlighted.

8. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring a Top Talker monitor flow

You can use the learn flow parameter ("*") to create the equivalent to a Top Talker monitor. Use a Top Talker monitor to identify high volume flows passing a port.

NOTE
The Top Talker monitor is only supported on F_Ports.

The following example creates an ingress Top Talker monitor. This procedure creates a Flow Monitor learning flow named "ingresstt" for all frames between any devices ingressing through port 41 of the switch on which the flow is running.

1. Right-click a switch port (41) on which you want to monitor Top Talkers and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (41)

2. Enter a name (ingresstt) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
4. Select the **Activate all selected features** check box to immediately activate the flow after creation.
5. Enter an asterisk (*) in the **Destination** field.

6. Click **OK**.

   The **Flow Vision** dialog box displays with the new flow definition highlighted.

7. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring a Flow Monitor flow for a trunk group

NOTE
You cannot create a learned flow in a trunk group.

Flow Monitor supports monitoring trunk ports subject to the following conditions:

- Although the flow created on the trunk master port monitors data traffic for the entire port, you must create the same flow on all trunk member ports.
- After a switch initialization or a recovery (cold or warm), existing flows are re-created on both master and slave ports, but only those flows associated with the master port are activated.
- If you create an active flow on a slave port, the flow is automatically activated when the slave port becomes the master port.
- If you create an inactive flow on a slave port, you must activate the flow manually when the slave port becomes the master port.

The following example creates four flows, one for each member of the trunk group.

1. Right-click the trunk master or slave port on which you want to monitor a flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)

2. Enter a name (trunkflow1 (master), trunkflow2, trunkflow3, and trunkflow4) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option to gather flow data from the source device to the destination device.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Enter an asterisk (*) in the **Destination** field.
7. Enter the egress port data in Port (slot/port) or D,I (domain ID,port number) format in the **Egress** field.

   NOTE
   You must enter the slot number and the port number. For Backbone chassis, the slot number cannot be 0 (zero). For switches, you must enter 0 (zero) as the slot number. For logical (virtual) switches, follow the rule for the physical chassis (either Backbone chassis or switch) from which you created the logical switch.

8. Click **OK**.

   The **Flow Vision** dialog box displays with the new flow definition highlighted.

9. Repeat step 1 through step 8 for each trunk port in the group.

10. Select **OK** to save the definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table.

    To view the Flow Monitor statistical data for the entire trunk group, select the master port flow definition in the **Flow Definitions** table and click the right arrow button to display it in the **Flows** table.

    The accumulated Flow Monitor statistical data for the entire trunk group is stored on the master port. If the master port changes, the data is transferred to the new master port.

    You cannot display Flow statistics for slave trunk ports.

    To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

FC router fabrics Flow Monitor flow example procedures

You can create Flow Monitor flow definitions using port WWNs, port IDs, and fabric IDs (FC router). When creating flow monitors on EX_Ports, you can use either a WWN or a port ID for the source device and destination device or you can use the fabric ID for the FCR/XISL source fabric and destination fabric. Inter-Fabric Link (IFL) flows can be monitored only on 16 Gbps-capable EX_Ports in an FC router. IFL flows are not supported on E_Ports or F_Ports.

When monitoring an FC router fabric, you may find it simpler to use port WWNs rather than virtual port IDs in your flow definitions. This is because you do not need to locate and map the virtual port IDs for the actual source and destination devices.

Even though you always create a flow definition in the backbone fabric, the perspective of the flow is from the edge fabric. In the following examples, the flow definitions are based on the perspective of the Edge 1 fabric .

FC router fabric monitors using WWNs

The following examples present the flow definitions with the **End Device** mode set to **WWN**.

In Figure 488, the physical devices are A, B, and C, and have the port WWNs a, b, and c, respectively.

FIGURE 488FC router fabric example

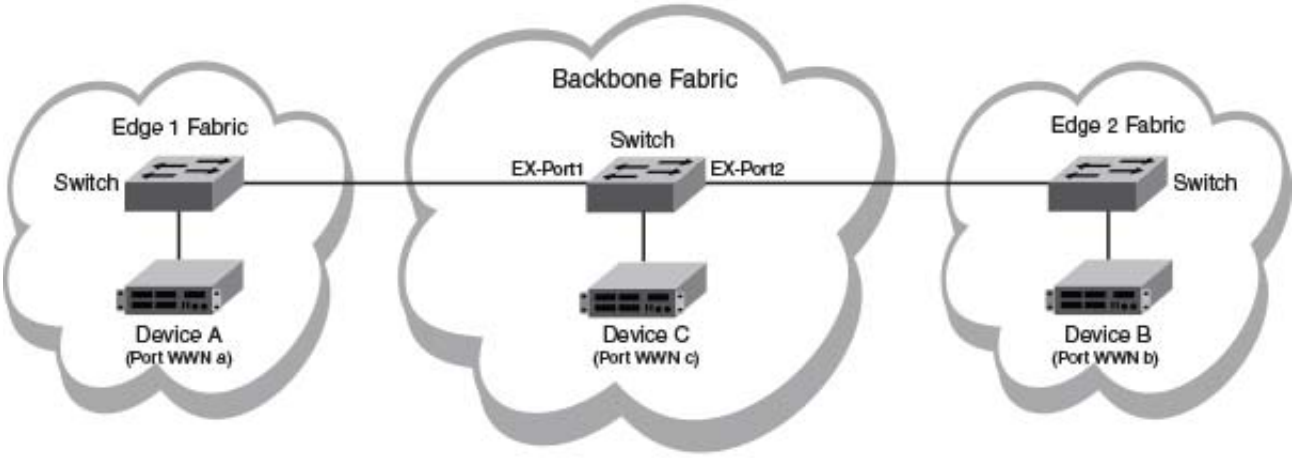

In Figure 489, provides the port WWN values for the following example procedures.

FIGURE 489FC router fabric annotated with port WWN values

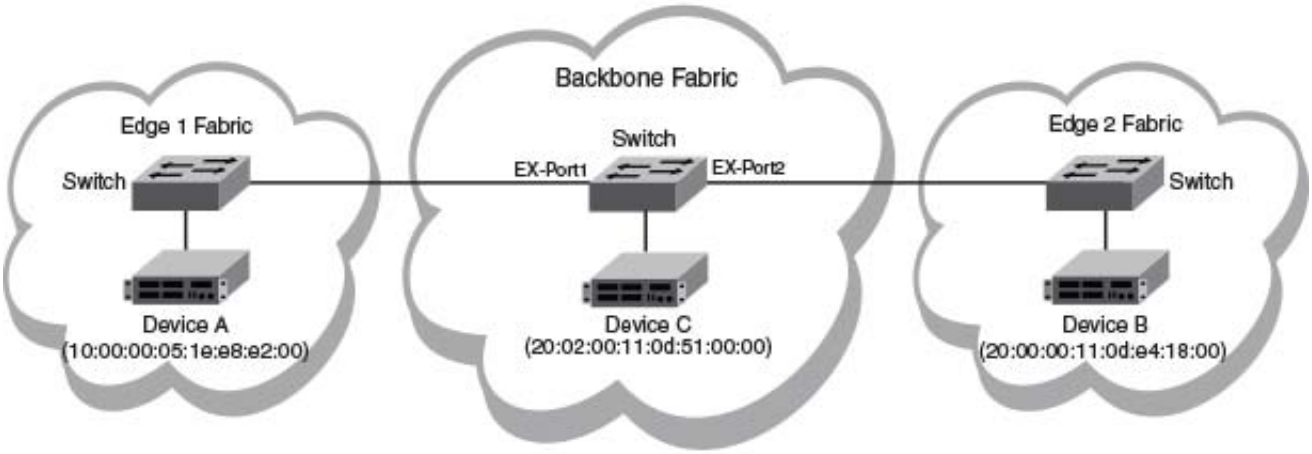

Configuring an edge-to-edge flow through an ingress port using WWNs

Figure 488, provides the values for and a diagram of the following example procedure which monitors a flow from Device A to Device B ingressing through EX_Port1 (traffic is running from left to right).

The following procedure creates a flow (e2e_src_dcx_wwn) that filters frames passing from one edge fabric (10:00:00:05:1e:e8:e2:00) to another edge fabric (20:00:00:11:0d:e4:18:00) using a specific ingress port (219) on the backbone.

1. Right-click the port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor

- Direction — Bidirectional
- Source Device — Source identifier
- Destination Device — * (an asterisk allows you to use any port)
- Ingress port — port number (219)

2. Enter a name (e2e_src_dcx_wwn) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **WWN** option for **End Device**.
7. Enter the port WWN (10:00:00:05:1e:e8:e2:00) in the **Source** field.
8. Enter the port WWN (20:00:00:11:0d:e4:18:00) in the **Destination** field.
9. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring an edge-to-edge flow through an egress port using WWNs

Figure 488, provides the values for and a diagram of the following example procedure which monitors a flow from Device B to Device A egressing through EX_Port1 (traffic is running from right to left).

The following procedure creates a flow (e2e_dst_dcx) that filters out frames passing from one edge fabric (20:00:00:11:0d:e4:18:00) to another edge fabric (10:00:00:05:1e:e8:e2:00) using a specific egress port (219) on the backbone.

1. Right-click the port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (219)

2. Enter a name (e2e_dst_dcx) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.

5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **WWN** option for **End Device**.
7. Enter the port WWN (20:00:00:11:0d:e4:18:00) in the **Source** field.
8. Enter the port WWN (10:00:00:05:1e:e8:e2:00) in the **Destination** field.
9. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring a backbone-to-edge flow through an egress port using WWNs

Figure 488, provides the values for and a diagram of the following example procedure which monitors a flow from Device C to Device A egressing through EX_Port1 (traffic is running from right to left).

The following procedure creates a flow (b2e_dst_dcx) that filters out frames passing from the backbone fabric (20:02:00:11:0d:51:00:00) to an edge fabric (10:00:00:05:1e:e8:e2:00) using a specific egress port (219).

1. Right-click the port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (219)

2. Enter a name (b2e_dst_dcx) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **WWN** option for **End Device**.
7. Enter the port WWN (20:02:00:11:0d:51:00:00) in the **Source** field.
8. Enter the port WWN (10:00:00:05:1e:e8:e2:00) in the **Destination** field.
9. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring an edge-to-backbone flow through an ingress port using WWNs

Figure 488, provides the values for and a diagram of the following example procedure which monitors a flow from Device A to Device C ingressing through EX_Port1 (traffic is running from left to right).

The following procedure creates a flow (e2b_src_dcx) that filters out frames passing from the backbone fabric (10:00:00:05:1e:e8:e2:00) to an edge fabric (20:02:00:11:0d:51:00:00) using a specific egress port (219).

1. Right-click the port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (219)

2. Enter a name (e2b_src_dcx) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **WWN** option for **End Device**.
7. Enter the port WWN (10:00:00:05:1e:e8:e2:00) in the **Source** field.
8. Enter the port WWN (20:02:00:11:0d:51:00:00) in the **Destination** field.
9. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

FC router fabric monitors using virtual port IDs

The following examples present flow definitions configured with the **End Device** mode set to **Port Address** (port ID).

In Figure 490, the physical devices are A, B, and C. The virtual devices are Device A?, B?, C?, A?, B?, and C?, representing the physical devices A, B, and C, respectively.

FIGURE 490FC router fabric

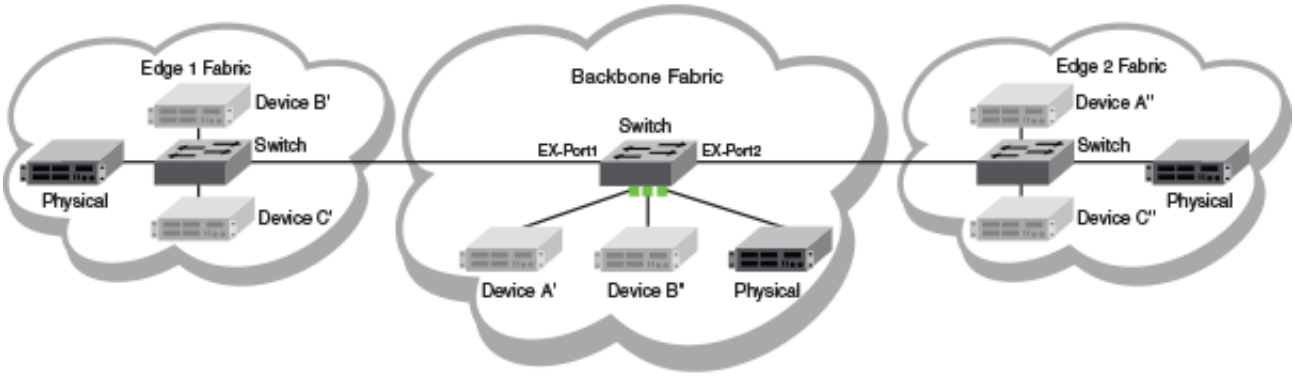

Figure 491 provides the physical port address (PID), fabric identifier (FID), and virtual port address (virtual PID) values for the following examples.

FIGURE 491FC router fabric annotated with port address, FID, and virtual port address values

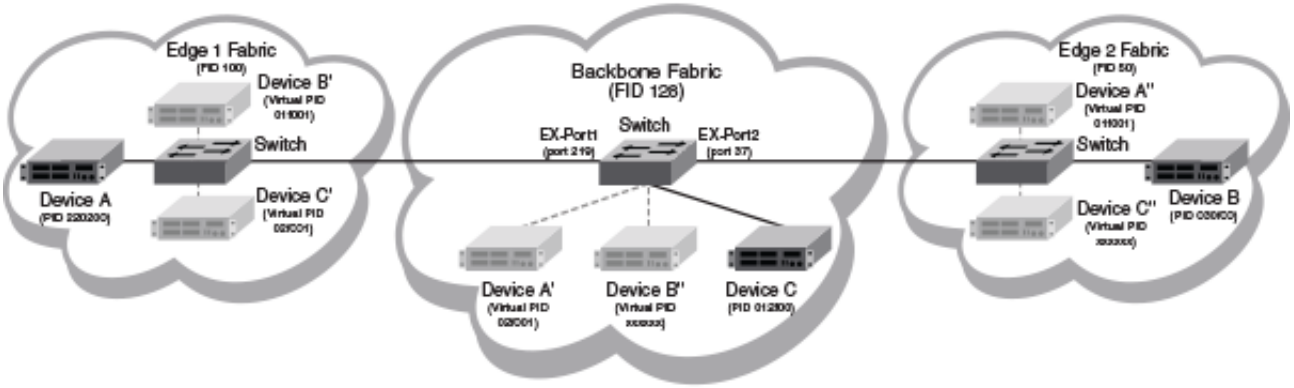

NOTE
The virtual port address values for devices B? and C? were not generated for the following examples, they are indicated by "xxxxxx" in Figure 491 .

Configuring an edge-to-edge flow through an ingress port using virtual port IDs

Figure 491, provides the values for and a diagram of the following example procedure which monitors a flow from Device A to Device B ingressing through EX_Port1, the source device is Device A, the destination device is Device B?, and the ingress port is EX_Port1 (traffic is running from left to right).

The following example creates a flow (e2e_src_dcx_pid) that filters frames passing from one edge fabric (220200) to another edge fabric (01f001) using a specific ingress port (219) on the backbone.

1. Right-click the port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (219)

2. Enter a name (e2e_src_dcx_pid) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **Port Address** option for **End Device**.
7. Enter the source port address (220200) in the **Source** field.
8. Enter the destination port address (01f001) in the **Destination** field.
9. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring an edge-to-edge flow through an egress port using virtual port IDs

Figure 491, provides the values for and a diagram of the following example procedure which monitors a flow from Device B to Device A egressing through EX_Port1, the source device is Device B?, the destination device is Device A, and the egress port is EX_Port1 (traffic is running from right to left).

The following example creates a flow (e2e_src_dcx) that filters out frames passing from one edge fabric (01f001) to another edge fabric (220200) using a specific egress port (219) on the backbone.

1. Right-click the port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Egress port — port number (219)

2. Enter a name (e2e_src_dcx) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric character,or underscores.

NOTE
For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **Port Address** option for **End Device**.
7. Enter the source port address (01f001) in the **Source** field.
8. Enter the destination port address (220200) in the **Destination** field.
9. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring a backbone-to-edge flow through an egress port using virtual port IDs

Figure 491, provides the values for and a diagram of the following example procedure which monitors a flow from Device C to Device A egressing through EX_Port1, the source device is Device C?, the destination device is Device A, and the egress port is EX_Port1 (traffic is running from right to left).

The following example creates a flow (b2e_dst_dcx) that filters out frames passing from the backbone fabric (01f001) to an edge fabric (220200) using a specific egress port.

1. Right-click the port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Egress port — port number (219)

2. Enter a name (b2e_dst_dcx) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **Port Address** option for **End Device**.
7. Enter the source port address (01f001) in the **Source** field.
8. Enter the destination port address (220200) in the **Destination** field.
9. Select **OK** to save the definition.

When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring a edge-to-backbone flow through an ingress port using virtual port IDs

Figure 491, provides the values for and a diagram of the following example procedure which monitors a flow from Device A to Device C ingressing through EX_Port1, the source device is Device A, the destination device is Device C?, and the ingress port is EX_Port1 (traffic is running from left to right).

The following example creates a flow (e2b_src_dcx) that filters out frames passing from an edge fabric (220200) to the backbone fabric (01f001) using a specific ingress port (219) .

1. Right-click the port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (219)

2. Enter a name (e2b_src_dcx) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **Port Address** option for **End Device**.
7. Enter the source port address (220200) in the **Source** field.
8. Enter the destination port address (01f001) in the **Destination** field.
9. Select **OK** to save the definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

FC router fabric monitors using fabric IDs

The following examples present flow definitions configured with the **FCR/XISL** mode set to **Fabric ID** (FC router fabric ID).

NOTE
This feature is only supported on 16 Gbps-capable platforms running Fabric OS 7.3 or later.

Configuring an edge-to-backbone flow through an ingress port using fabric IDs

The following example creates a flow (e2b_src_dcx_fid) that filters frames passing from an edge fabric (100) to the backbone fabric (128) using a specific ingress port (219).

1. Right-click the EX_Port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (219)

2. Enter a name (e2b_src_dcx_fid) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **Port Address** option for **End Device**.
7. Enter the source port ID (220200) in the **Source** field.
8. Enter an asterisk (*) for the destination port ID in the **Destination** field.
9. Select the **Fabric ID** option for **FCR/XISL** mode.
10. Enter the edge fabric ID (100) in the **Source** field.

    To select the source from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

11. Enter the backbone fabric ID (128) in the **Destination** field.

    To select the destination from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

12. Select **OK** to save the definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring an edge-to-edge flow through an ingress port using fabric IDs

The following example creates a flow (e2e_src_dcx_fid) that filters frames passing from one edge fabric (100) to the another edge fabric (50) using a specific ingress port (219).

1. Right-click the EX_Port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)

- Ingress port — port number (219)

2. Enter a name (e2e_src_dcx_fid) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.
4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **Port Address** option for **End Device**.
7. Enter the source port ID (220200) in the **Source** field.
8. Enter an asterisk (*) for the destination port ID in the **Destination** field.
9. Select the **Fabric ID** option for **FCR/XISL** mode.
10. Enter the edge fabric ID (100) in the **Source** field.

    To select the source from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

11. Enter the other edge fabric ID (50) in the **Destination** field.

    To select the destination from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

12. Select **OK** to save the definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring a backbone-to-edge flow through an egress port using fabric IDs

The following example creates a flow (b2e_src_dcx_fid) that filters frames passing from a backbone fabric (128) to the edge fabric (100) using a specific ingress port (219).

1. Right-click the EX_Port (219) on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Egress port — port number (219)

2. Enter a name (b2e_src_dcx_fid) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Source to Destination** option for **Direction**.

4. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
5. Select the **Activate all selected features** check box to immediately activate the flow after creation.
6. Select the **Port Address** option for **End Device**.
7. Enter the source port ID (01f001) in the **Source** field.
8. Enter an asterisk (*) for the destination port ID in the **Destination** field.
9. Select the **Fabric ID** option for **FCR/XISL** mode.
10. Enter the backbone fabric ID (100) in the **Source** field.

    To select the source from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.
11. Enter the edge fabric ID (128) in the **Destination** field.

    To select the destination from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.
12. Select **OK** to save the definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

XISL and backbone E_Port monitors

Flow Monitor supports both static and learned monitoring of fabric-wide statistics on both XISL_Ports and backbone E_Ports. You can use this data to estimate the logical fabric or inter-fabric utilization of an XISL_Port or a Backbone E_Port.

The following examples present flow definitions configured with the **FCR/XISL** mode set to **VFID** (virtual fabric ID).

NOTE
This feature is only supported on 16 Gbps-capable platforms running Fabric OS 7.3 or later.

Table 91 details the learned flow support for XISL_Ports and backbone E_Ports

TABLE 91 Learned flow support for Backbone E_port and XISL_port

Learned traffic	Backbone E_Port	XISL_Port
Intra-fabric traffic	Only learns backbone fabric traffic.	Learns logical fabric traffic.
Inter-fabric traffic	Learns edge-to-edge, backbone-to-edge, and edge-to-backbone traffic.	Only learns edge-to-edge traffic.

Configuring a specified fabric flow using an XISL E_Port

The following example creates a flow (e2e_src_dcx_xisl) that monitors frames passing from the source fabric to the destination fabric using a specific port (ingress or egress).

1. Right-click the XISL E_Port on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — * (an asterisk allows you to use any port)
   - Destination Device — * (an asterisk allows you to use any port)

- Ingress/Egress port — XISL E_Port number

2. Enter a name (e2e_src_dcx_xisl) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
4. Select the **Activate all selected features** check box to immediately activate the flow after creation.
5. Clear the source port ID in the **Source** field.
6. Clear the destination port ID in the **Destination** field.
7. Select the **VFID** option for **FCR/XISL** mode.
8. Enter the source edge fabric ID in the **Source** field.

   To select the source from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

9. Enter the destination edge fabric ID in the **Destination** field.

   To select the destination from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

10. Select **OK** to save the definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring an unspecified fabric flow using an XISL E_Port

The following example creates a flow (all_e2e_src_dcx_xisl) that monitors frames passing from all source fabrics to all destination fabrics using a specific port (ingress or egress).

1. Right-click the XISL E_Port on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — * (an asterisk allows you to use any port)
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress/Egress port — XISL E_Port number

2. Enter a name (all_e2e_src_dcx_xisl) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumerics or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
4. Select the **Activate all selected features** check box to immediately activate the flow after creation.
5. Select the **Port Address** option for **End Device**.

6. Clear the source port ID in the **Source** field.
7. Clear the destination port ID in the **Destination** field.
8. Select the **VFID** option for **FCR/XISL** mode.
9. Enter an asterisk (*) for the edge fabric ID in the **Source** field.
10. Enter an asterisk (*) for the edge fabric ID in the **Destination** field.
11. Select **OK** to save the definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Configuring a learned flow using an XISL E_Port

The following example creates a flow (learned_dcx_xisl) that monitors frames passing from all the source and destination fabrics as well as all devices associated with the traffic using a specific port (ingress or egress).

1. Right-click the XISL E_Port on which you want to monitor the flow and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following default criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — * (an asterisk allows you to use any port)
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress/Egress port — XISL E_Port number

2. Enter a name (learned_dcx_xisl) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
4. Select the **Activate all selected features** check box to immediately activate the flow after creation.
5. Enter an asterisk (*) for the source port ID in the **Source** field.
6. Enter an asterisk (*) for the destination port ID in the **Destination** field.
7. Select the **VFID** option for **FCR/XISL** mode.
8. Enter an asterisk (*) for the edge fabric ID in the **Source** field.
9. Enter an asterisk (*) for the edge fabric ID in the **Destination** field.
10. Select **OK** to save the definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Monitor flow" on page 1035.

Flow Generator

Flow Generator is a test traffic generator for pre-testing the SAN infrastructure (including internal connections) for robustness before deploying it. Flow Generator provides you with the ability to:

- Configure a 16 Gbps FC-capable port as a simulated device that can transmit frames at full 16 Gbps line rate.
- Emulate a 16 Gbps SAN without actually having any 16 Gbps hosts or targets or SAN-testers.
- Pre-test the entire SAN fabric at the full line rate, including optics and cables on ISLs as well as internal connections within a switch.

Flow Generator achieves this using simulation mode (SIM) ports. SIM-Ports behave like standard ports, but are used only for testing. By using SIM-Ports, Flow Generator traffic is terminated at the destination port and does not leave the switch. Refer to "SIM-Ports" on page 1059 for more information on SIM-Ports.

Flow Generator can generate standard frames or create custom frames with sizes and patterns you specify. A sample use case would be to create a traffic flow from a source ID to a destination ID to validate routing and throughput. "Creating a flow from a specific source to a specific destination" on page 1067 provides an example of this flow and the results for this use case.

You should not use Flow Generator in an active production environment, as the Flow Generator traffic can saturate the links and will impact the production traffic sharing the same links.

Flow Generator setup

Flow Generator generates and receives traffic only from simulated ingress and egress ports (SIM-Ports) which emulate device entries in the Name Server database, so that they are treated as real devices and can be used to evaluate various switch and fabric operations such as QoS and Traffic Isolation. For more information on working with SIM-Ports, refer to "SIM-Ports" on page 1059.

Flow Generator offers several flow control options that you can configure, including the ability to specify both the frame size and the frame payload pattern. Header parameters and other control parameters can also be added as part of the definition. The 0XID value for frames is random and cannot be specified.

Flow Generator flows are defined using a combination of the source device, destination device, ingress port, and egress port parameters. All of these must be SIM-Ports. The source device is the origination point of the test traffic. The destination device is the destination of the test traffic; for Flow Generator flows, it may be remote from the switch. The port that transmits (egress) the simulation traffic must be a 16 Gbps-capable FC port. The port that receives (ingress) the simulated traffic can be either an 8 Gbps- or a 16 Gbps-capable FC port. Figure 492 illustrates this concept.

FIGURE 492 Flow Generator flow structure

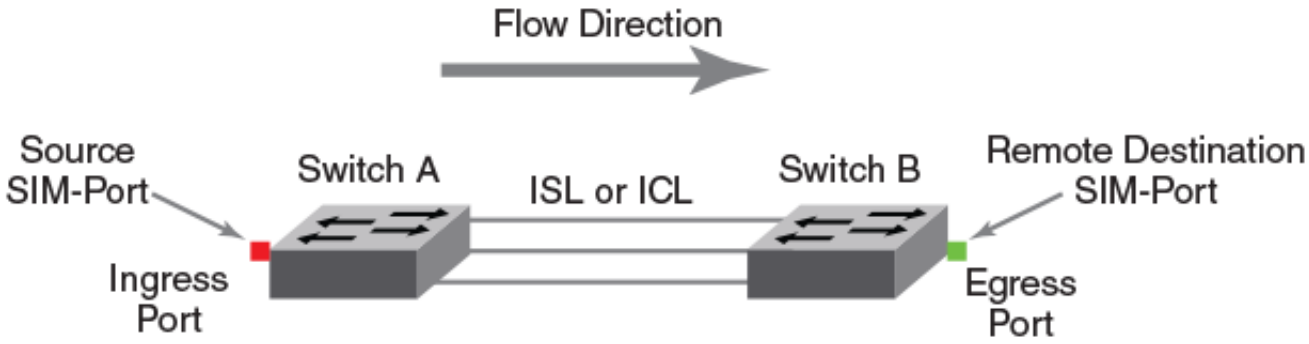

Flow Generator limitations

The following limitations apply specifically to Flow Generator:

- If used on a live production system, Flow Generator traffic will compete with any existing traffic. Consequently, E_Ports and FCIP links can become congested when using Flow Generator, leading to throughput degradation. FCIP links are more prone to congestion than E_Ports.
- Flow Generator is only supported on 16 Gbps- and 32 Gbps-capable FC devices.
- 8 Gbps-capable FC ports cannot transmit simulated traffic; however, they can receive simulated traffic. Therefore, source devices and ingress ports must be 16 Gbps- or 32 Gbps-capable FC ports, but destination devices and egress ports can be either 8 Gbps- or 16 Gbps-capable FC ports.
- The source device, ingress port, and egress port must all be local. The destination device can be local or remote; however, if you define the destination device and egress port, both must refer to the same destination and must be local. The source device and ingress port must refer to the same source.
- Flow Generator only supports four active flows per ingress port.
- Flow Generator is not supported in Access Gateway mode, on a base switch, or across a Fibre Channel Router (FCR) backbone fabric.
- Frame redirection is not supported on SIM-Ports.
- Zoning is not enforced. sources and destinations can be in different zones.
- Flow Generator gathers source and destination pairs from the zoning database for learning flows only at the time the flow is activated. Subsequent changes to this database will not be registered until a flow is reactivated.
- Flow creation is not allowed if Advanced Performance Monitor (APM) or Port Mirroring is enabled. Similarly, APM and Port Mirroring-related operations are not allowed if any flow (active or defined) is present on the switch.

SIM-Ports

Before you create and activate a Flow Generator flow, you must set the source device and destination device ports in simulation mode (SIM-Port mode). Setting the port to SIM-Port mode sets an internal loopback on the port and creates a filter that discards all incoming Flow Generator frames. This ensures that test flows are not unintentionally transmitted to real devices.

A SIM-Port simulates an F_Port on the switch using the port WWN or virtual WWN. Once you enable SIM-Port mode, the SIM-Port registers itself into the Name Server database. For learning mode, SIM-Ports must be part of a zoning database.

Flow Generator SIM-Ports have the following restrictions:

- Flow Generator supports up to four active flows per ingress SIM-Port and takes 52 credits per SIM-Port from the ASIC.

- Zoning is bypassed on SIM-Ports. Traffic will reach its destination regardless of zoning configuration.
- Flow Generator only uses zones that gather the Source ID-Destination ID pairs for learning flows.

On a switch that has live traffic passing through E_Ports or long distance ports, Flow Generator allows you to enable ("Enabling SIM-Port mode" on page 1061) and disable ("Disabling SIM-Port mode" on page 1061) SIM-Ports that are on the same ASIC as those active ports. Flow Generator generates traffic matching the flow definition on the SIM-Port. Traffic is generated between the local or remote ports at the speed configured on the source SIM-Port. Flow Generator can also activate and deactivate the artificial traffic, allowing you to verify the impact of the testing flow on existing traffic.

SIM-Port requirements and restrictions

Flow Generator SIM-Ports must meet the following criteria to be valid:

- SIM-Ports are supported on ASICs that support either 8 Gbps- or 16 Gbps-capable FC ports. Source devices or ingress ports can only be on 16 Gbps-capable FC ports. Destination devices or egress ports can be on either 8 or 16 Gbps-capable FC ports.
- SIM-Ports cannot be in the base switch or Access Gateway.
- SIM-Ports cannot be configured on a port that is online and connected to a real device.

  NOTE
  If a port is connected to a real device, you can disable the port, configure the SIM-Port, and then re-enable the port. The port will be a SIM-Port; the real device will not join the fabric.

- Existing SIM-Ports are added to Device Connection Control (DCC) policies when created with a wildcard (*) but not adhered to. These SIM-Port entries must be deleted if a new WWN is connected.
- SIM-Ports cannot be configured as any of the following port types; these restrictions also apply at the time a SIM-Port is enabled:
  - Any port running encryption or compression
  - Any F_Port connected to a real device (unless the port is disabled)
  - D_Port (Diagnostic Port)
  - E_Port
  - EX_Port
  - F_Port trunk
  - Fastwrite port
  - FCoE port
  - ICL port
  - L_Port
  - M_Port (Mirror Port)
  - VE port
- The following features of a SIM-Port are persistent across a reboot:
  - Each SIM-Port is assigned a PID.
  - Each SIM-Port's Port Worldwide Name by default is the switch PWWN, unless a user-defined Virtual Port Worldwide Name is assigned to it.
  - Each SIM-Port registers itself into the Name Server database.
- If a port is configured as a SIM-Port:
  - You cannot enable QoS.
  - You cannot enable CSCTL_mode.
  - You can set an Ingress Rate Limit.
- If a port is configured with QoS enabled, you cannot configure it as a SIM-Port.

- If a port is configured with CSCTL_mode enabled, you cannot configure it as a SIM-Port.
- If a port has an Ingress Rate Limit set, you can configure it as a SIM-Port.

Enabling SIM-Port mode

Prior to creating and activating flows using the Flow Generator feature, your must enable SIM-Port mode on the switch ports connected to the source and destination devices for your flow. For more information on SIM-Ports and SIM-Port mode, refer to "SIM-Ports" on page 1059.

To enable SIM-Port mode, complete the following steps.

1. Select a port on the local switch for the source device in the Product List, and then select **Monitor > Fabric Vision > Flow > SIM Mode > Enable.**
2. Select a port on the local switch for the destination device in the Product List, and then select **Monitor > Fabric Vision > Flow > SIM Mode > Enable**.

   NOTE
   You can enable and disable multiple SIM-Ports, and a combination of U_Ports and SIM-Ports.

   When you enable SIM-Port mode, the SIM-Port registers itself into the Name Server database. The SIM-Port for the source device registers itself as an initiator. Once the initiator is in the Name Server database, you must zone this virtual initiator to one or more destination SIM-Ports. For information about zoning, refer to "Zoning configuration" on page 783.

Disabling SIM-Port mode

To disable SIM-Port mode, use the following steps.

1. Select the SIM-Port for the source device in the Product List, and then select **Monitor > Fabric Vision > Flow > SIM Mode > Disable**.
2. Select the SIM-Port for the destination device in the Product List, and then select **Monitor > Fabric Vision > Flow > SIM Mode > Disable**.

Creating a Flow Generator flow definition

This procedure provides step-by-step instructions for configuring a Generator flow definition. For more specific example procedures, refer to "Flow Generator example procedures" on page 1067.

1. Create two SIM-Ports ("Enabling SIM-Port mode" on page 1061).
2. Right-click a SIM-Port and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — SIM-Port number
3. Enter a name for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

NOTE
For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

4. Clear the **Monitor** check box.
5. Select the **Generator** check box.
6. Select the **Source to Destination** option to gather flow data from the source device to the destination device.
7. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
8. Select the **Activate all selected features** check box to immediately activate the flow after creation.
9. Change the target switch for the flow definition by clicking the ellipsis button to the right of the **Target Switch** field.

   The **Select Switch** dialog box displays. To manually set the **Target Switch**, refer to "Selecting the target switch from a list of Fabric Vision-capable switches" on page 1030.
10. Select one of the following format options for **End Device** mode:
    - **Port Address** (port ID) — Select to display the source and destination device address using the port ID.
    - **WWN** (world wide name) — Select to display the source and destination device address using the port WWN.
11. Enter the address or WWN of the source port in the **Source** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the source port from a list, refer to "Selecting an end device port from a list of available device ports" on page 1030.

    When you enter a port ID or WWN in the **Source** or **Destination** fields, a port information field displays beneath with the port label based on your topology layout settings (refer to "Changing the port label" on page 324). If you enter an asterisk (*) or no value in the **Source** or **Destination** fields, this field remains blank.
12. Enter the address or WWN of the source port in the **Destination** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the destination port from a list, refer to "Selecting an end device port from a list of available device ports" on page 1030.
13. (Optional) If you want to swap source and destination device port information, click **<swap>**.
14. Select one of the following format options for **Switch** mode:
    - **Port (slot/port)** — Select to display the switch ingress or egress port using the slot and port number.
    - **D,I** (domain ID, port number) — Select to display the switch ingress or egress port using the domain ID and port number.
15. Enter the ingress port in Port (slot/port) or D,I (domain ID,port number) format for the SIM-Port in the **Ingress** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the ingress port from a list, refer to "Selecting an ingress or egress port from a list of available switch ports" on page 1032.

    NOTE
    You must enter the slot number and the port number. For Backbone chassis, the slot number cannot be 0 (zero). For switches, you must enter 0 (zero) as the slot number. For logical (virtual) switches, follow the rule for the physical chassis (either Backbone chassis or switch) from which you created the logical switch.

    When you enter a value in the **Ingress** or **Egress** fields, a port information field displays beneath with the port label based on your topology layout settings (refer to "Changing the port label" on page 324). If you enter an asterisk (*) or no value in the **Ingress** or **Egress** fields, this field remains blank.

16. Enter the egress port in Port (slot/port) or D,I (domain ID,port number) format for the SIM-Port in the **Egress** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the egress port from a list, refer to "Selecting an ingress or egress port from a list of available switch ports" on page 1032.

    NOTE
    You must enter the slot number and the port number. For Backbone chassis, the slot number cannot be 0 (zero). For switches, you must enter 0 (zero) as the slot number. For logical (virtual) switches, follow the rule for the physical chassis (either Backbone chassis or switch) from which you created the logical switch.

17. For FCR or virtual fabrics, select one of the following format options for **FCR/XISL** mode:

    - **Fabric ID** — Select to configure a flow definition for an FCR fabric.
    - **VFID** (Virtual Fabric ID) — Select to configure a flow definition for virtual devices.

18. For FCR or virtual fabrics, enter the fabric ID or virtual fabric ID in the **Source** field.

    Enter an asterisk (*) to use any port. To select the source from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

19. For FCR or virtual fabrics, enter the fabric ID or virtual fabric ID in the **Destination** field.

    Enter an asterisk (*) to use any port. To select the destination from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.

20. Enter a LUN ID in the **LUN** field or click the ellipsis button to select a LUN ID from the list.

    LUN IDs can be from 0 through 65535. To select a LUN ID from a list, refer to "Selecting a LUN ID from a list of available LUNs" on page 1034.

21. Click **OK** to save the flow definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Generator flow" on page 1064. To configure the frame payload size and pattern for the flow, refer to "Customizing Flow Generator flows" on page 1063.

Customizing Flow Generator flows

After you define a Generator flow, you can further customize the flow by specifying the frame payload size and pattern to be used for the flow.

If you customize an active flow, first the Management application deactivates the flow, then reconfigures the frame payload size and pattern, and finally reactives the flow. Inactive flows are reconfigured, but not reactivated.

To specify the frame payload size and pattern for a Generator flow, complete the following steps.

1. Select the device on which you defined the flow and select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows in the **Flow Definitions** table. For more information about the fields and components of the **Flow Definitions** table, refer to "Flows Definitions table fields and components" on page 1021.

2. Select the **Generator** flow that you want to customize in the **Flow Definitions** table.
3. Select **Generator > Configure** from the **Feature** list.

The **Configure Generator** dialog box displays.

FIGURE 493 Configure Generator dialog box

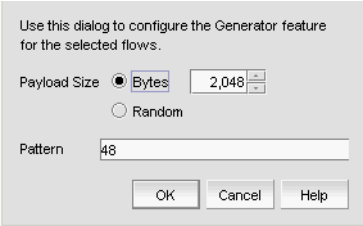

4. Define the frame payload size by choosing one of the following options:
   - To configure a specific frame payload size in bytes, select the **Bytes** option and enter the size you want in the **Bytes** field.

     The frame payload size value must be a multiple of 4 in the range from 64 through 2048. (64, 120, 140, 320, 512 and so on). The default payload size value is 2048.
   - To configure a random frame payload size, select the **Random** option.
5. Define the pattern to be used as the frame payload by entering the pattern you want to use in the **Pattern** field.

   The frame payload pattern must be an alphanumeric ASCII string between 1 and 32 characters in length. The default frame payload pattern value is 0, which produces a random pattern of alphanumeric ASCII characters with a variable string length.
6. Click **OK** to save your changes and return to the **Flow Vision** dialog box.

Monitoring a Flow Generator flow

To monitor the summary data for a Flow Generator flow, complete the following steps.

1. Select the device on which you defined the generator flow and select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows in the **Flow Definitions** table. For more information about the fields and components of the **Flow Definitions** table, refer to "Flows Definitions table fields and components" on page 1021.
2. Select the generator flow that you want to monitor in the **Flow Definitions** table.
3. Select a time interval for monitoring the flow in the **Time duration** list.

   Possible values are 30 minutes, 1 hour, 6 hours, 12 hours, 1 day, 3 days, 1 week, and 1 month.
4. Click the right arrow button to display the selected flow in the **Flows** table.

   You can sort the **Flows** table by clicking any column head. You can reverse the sort order by clicking the column head again.
5. Select the **SCSI** check box to display SCSI-related measures.

   SCSI-related measures include SCSI read count, write count, read rate, write rate, read data, write data, and read and write frame data.

   Clear the check box to hide SCSI-related measures.

6. Select the **Frame** check box to display frame-related measures.

   Frame-related measures include transmit (Tx) and receive (Rx) frame count, transmit frame and receive frame rate, transmit and receive word count, and transmit and receive throughput.

   Clear the check box to hide frame-related measures.

7. Review the sub-flow data for the selected Flow Generator flow.

   The **Flows** table, as shown in Figure 494, displays statistics and data for the selected flows. The **Flows** table has the following general characteristics and functions:

   - Data updates dynamically every 5 minutes. Sort the table by clicking any column head. You can reverse the sort order by clicking the column head again.
   - Locate an ingress port, egress port, source device, and destination device in the Product List or Topology Map by right-clicking a sub-flow in the **Flows** table and selecting **Locate >** ***port_type*** (where ***port_type*** is Ingress port, Egress port, Source device, or Destination device).
   - Highlight inactive sub-flows in gray by selecting the **Show inactive flows** check box.

     Inactive sub-flows indicate no updates to sub-flow statistics for over 15 minutes.

     Clear the check box to hide inactive sub-flows.
   - Display a flow in a performance graph. Select **Performance Graph >** ***graph_or_report*** (where ***graph_or_report*** is **Real Time Graph**, **Historical Graph**, or **Historical Report**).

     Refer to "Performance integration with Flow Vision" on page 1098 **for additional details. Note that the Performance Graph list is only available when there is at least one sub-flow in the Flows table.**
   - A single flow definition might yield data in multiple rows in the **Flows** table. For example, if you defined the flow definition source ID (SID) and destination ID (DID) as *, this may result in five rows if the source is communicating with five destination IDs. In the case of a learning flow, a root flow also displays to summarize all sub-flows.
   - Each unique sub-flow for the flow definition displays in the **Flows** table if it was reported within the selected time duration. If the last data point did not report that flow, the reported values may be 0.
   - The measures that display are based on the flow definition. Therefore, not all columns may be populated for the selected flows.

FIGURE 494 Flow Vision dialog box (Generator Flows table)

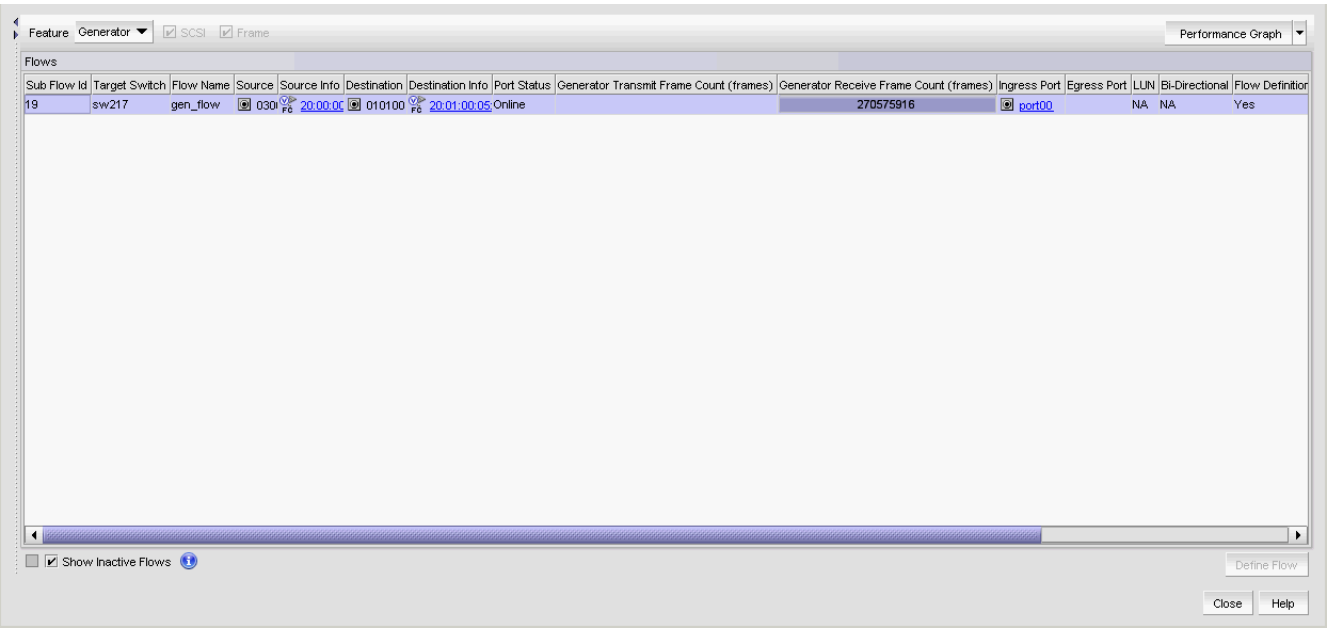

Table 90 describes information on sub-flows displayed in the **Flows** table when you select **Generator** from the **Feature** list above the **Flows** table.

TABLE 92 Flows panel information (Generator flow)

Column	Information Displayed
Sub Flow Id	The sub-flow database identifier.
Target Switch	The switch on which you created the flow definition.
Frame Type	All frame types defined in the flow definition.
Flow Name	The user-defined name for the flow definition.
Source	The source identifiers defined in the flow definition.
Source Info	The icon and name for the source device. The device name is a hyperlink. Click to launch the device's property sheet. This field is empty if the source device is not defined in the flow definition.
Destination	The port number of the destination device defined in the flow definition. An * (asterisk) indicates learned flows.
Destination Info	The icon and name for the destination device. The device name is a hyperlink. Click to launch the device's property sheet. This field is empty if the destination device is not defined in the flow definition.
Port Status	The operational status (online or offline) for the port.
Generator Transmit Frame Count (frames)	The number of frames transmitted for the Generator flow as reported in the last data point for the flow.
Generator Receive Frame Count (frames)	The number of frames received for the Generator flow as reported in the last data point for the flow.
Generator Number of complete runs (count)	The number of complete runs reported in the last data point received for the flow.

TABLE 92 Flows panel information (Generator flow) (Continued)

Column	Information Displayed
Generator Percent Complete of the Curent run (percentage)	The percentage completion of the current run reported in the last data point received for the flow.
Ingress Port	The ingress port defined in the flow definition. The port name is a hyperlink. Click to launch the port's property sheet. A yellow icon indicates a bottlenecked port.
Egress Port	The egress port defined in the flow definition. The port name is a hyperlink. Click to launch the port's property sheet. A yellow icon indicates a bottlenecked port.
Source Fabric ID	The fabric identifier of the source defined in the flow definition.
Destination Fabric ID	The fabric identifier of the destination defined in the flow definition.
LUN	Any LUN values defined in the flow definition.
Bi-Directional	Whether the flow is bidirectional (**yes**) or not (**no**).
Flow Definition Persistence	Whether the flow is configured to persist over switch reboots (**yes**) or not (**no**).
Last Updated Time	The date and time the sub-flow was last updated.

Flow Generator example procedures

The following examples describe how to work with Flow Generator flows.

Creating a flow from a specific source to a specific destination

The following example creates a generator flow, as shown in Figure 495.

FIGURE 495 Generator flow from a specific source to a specific destination

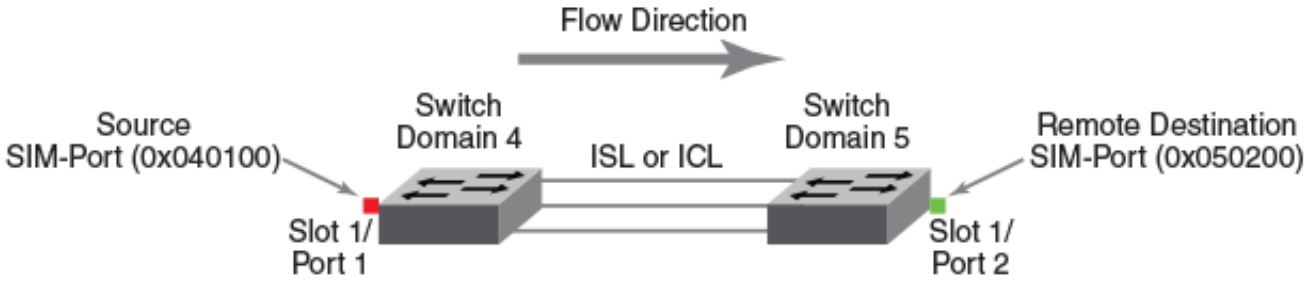

To create a generator flow between a specific source and a specific destination, complete the following steps.

1. Create two SIM-Ports ("Enabling SIM-Port mode" on page 1061).

   In this example, SIM-Port 1/1 is the source port and SIM-Port 1/2 is the destination device.

2. Right-click an initiator SIM-Port in the Product List and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier (0x040100)
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — SIM-Port number (1/1)

3. Enter a name (flowcase1) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

4. Clear the **Monitor** check box.
5. Select the **Generator** check box.

   The **Source to Destination** option is selected by default.

6. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
7. Select the **Activate all selected features** check box to immediately activate the flow after creation.
8. Enter the address or WWN of the destination SIM-Port in the **Destination** field.
9. Click **OK** to save the flow definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Generator flow" on page 1064. To configure the frame payload size and pattern for the flow, refer to "Customizing Flow Generator flows" on page 1063.

Integrating Flow Generator with Flow Monitor

Flow Generator flows can be monitored using Flow Monitor. For example, you can use a combination of Flow Generator flows and Flow Monitor flows to verify per-flow throughput at an ingress or egress port. This can be useful when more than one Flow Generator flow shares the same ingress or egress port. To do this, you must create a flow using both the Flow Generator and Flow Monitor features that share the ingress or egress port. The following example illustrates this integration.

To create a generator flow and a monitor flow, complete the following steps.

1. Create two SIM-Ports ("Enabling SIM-Port mode" on page 1061).

   In this example, SIM-Port 1/1 is the source port and SIM-Port 1/2 is the destination device.

2. Right-click an initiator SIM-Port in the Product List and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier (0x010100)
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — SIM-Port number (1/1)

3. Enter a name (flowCase3Src) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

4. Select the **Generator** check box.

The **Source to Destination** option is selected by default.

5. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
6. Select the **Activate all selected features** check box to immediately activate the flow after creation.
7. Enter an asteril (*) in the **Destination** field.
8. Click **OK** to save the flow definition.

   When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Generator flow" on page 1064. To configure the frame payload size and pattern for the flow, refer to "Customizing Flow Generator flows" on page 1063.

Flow Mirror

NOTE
Flow Mirror is only supported on 16 Gbps- or 32 Gbps-capable FC platforms.

As storage networks grow and become more complicated, it is becoming increasingly important to have non-intrusive diagnostic tools which can help identify problems without disturbing the existing fabric. Flow mirror is a diagnostic feature within Flow Vision that addresses this need.

Flow Mirror provides you with the ability to:

- Non-disruptively create copies of application flows that can be captured for deeper analysis.
- Conduct in-depth analysis of flows of interest, such as SCSI Reservation frames, ABTS frames, flows going to a bottlenecked device, frames during link bring-up, and others.
- Select the type of frames you want to be mirrored.
- Select a traffic pattern and create a real-time copy of this traffic, allowing you to debug a live system without disturbing existing connections. You can also use this feature as a way to view traffic passing through a port.

Flow Mirror duplicates the specified frames in a user-defined flow, and sends them to a mirror port. There are two types of mirrors:

- CPU Mirroring — Sends the duplicated frames to the local switch Control Processor Unit (CPU); however, CPU Mirroring has a limit of 256 frames per second.
- Local Flow Mirroring (LFM) — Sends the duplicated frames to an unoccupied, loopback, or mirror port on the same physical switch on which you defined the flow. LFM mirrored data can then be analyzed through an external Analyzer/Frame sniffer connected to the port. LFM requires that a loopback SFP be plugged in at the other end of the analyzer, or on the port configured as a mirror port.

  Note that if you select an unoccupied or loopback port, the Management application first reconfigures the port as a mirror port, and then creates the flow. This enables you to mirror a flow to a specific port, which is present on the same physical switch on which you defined the flow.

You can create a Flow Mirror flow by creating a copy of the flow you want to examine. Flow Mirror flows can be in an active state or an inactive state. If the mirror flow is "active", mirroring starts immediately; if the flow is "inactive", the flow must be activated (refer to "Activating flows" on page 1024) for mirroring to start. Mirrored flows can be unidirectional or bidirectional.

A sample use case would be to mirror the traffic flow from a slow-draining F_Port to see what is causing this condition. "Diagnosing a slow-draining F_Port" on page 1078 provides an example of this type of mirrored flow and the results for this use case.

Figure 496 provides a diagram of a mirrored flow with the ingress port mirroring the traffic flow to the CPU. Flow Mirror can also mirror the egress port, but only one port (ingress or egress) can be mirrored per flow. To mirror from one port in both flow directions (left to right and right to left), you must use the Bidirectional option.

FIGURE 496 Mirror flow

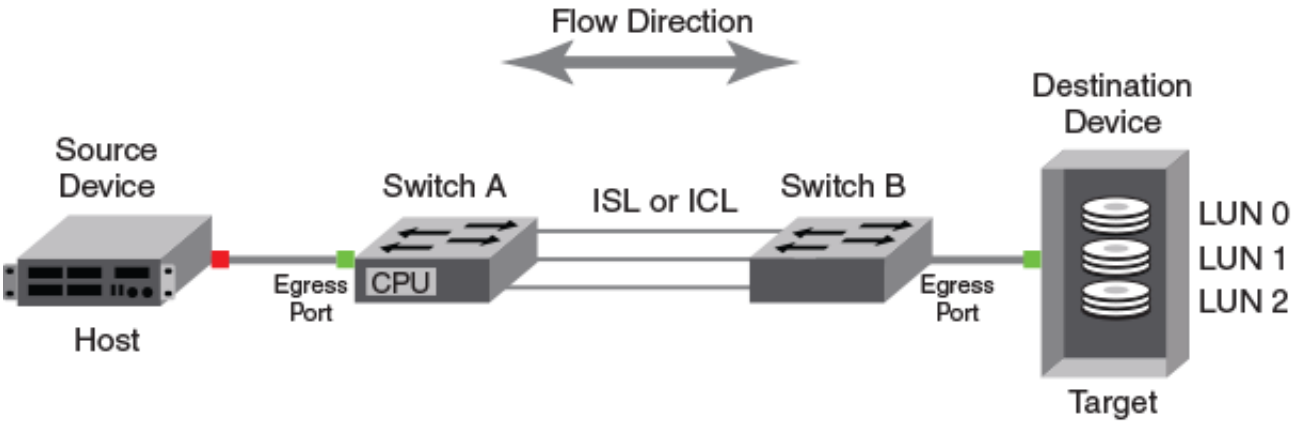

Flow Mirror limitations

- Flow Mirror is only supported on 16 Gbps and 32 Gbps and above capable FC devices.
- Flow Mirror requires a Fabric Vision license on the device where you are creating a mirrored flow.
- Flow Mirror is only supported on F_Ports or F_Port trunks.
- Flow Mirror only supports one active flow per device.
- A Flow Mirror port can be either an ingress or egress port of the flow definition.
- The Flow Mirror ingress or egress port must refer to a port on the current switch.
- Flow Generator flows can only be mirrored at the ingress port.
- Flow Mirroring is supported in Virtual Fabric and non-Virtual Fabric mode.
- Flow Mirror is not supported in Access Gateway mode.
- Flow Mirror is not supported across a Fibre Channel Routing (FCR) backbone fabric.
- Flow Mirror is not supported on SIM ports that are specified as ingress or egress ports in the flow definition.
- Flow Mirror flows can only mirror up to 256 frames (maximum) to the embedded port (CPU) per second. If a higher number of frames match the mirror criteria within a second, only the first 256 frames are mirrored. For chassis, 5120 frames are stored. For switches, 1280 frames are stored.
- If a Flow Monitor flow defined using the "-frametype" keyword as part of the flow command is installed on an ingress port, and a matching Flow Mirror flow is installed on an egress port, then traffic egressing through the egress port is not mirrored.
- If a flow is created for both Flow Monitor and Flow Mirror that uses a combination of either the "-frametype" and "-ingrport" keywords, or a combination of "-frametype", "-egrport", and "bidir" keywords, frames matching these definitions will be monitored but will not be mirrored.
- If a Flow Monitor flow is created on a blade or fixed-port switch that uses a combination of either the "-frametype" and "-ingrport" keywords, or a combination of "-frametype", "-egrport", and "bidir" keywords, then Flow Mirror frames matching any flow definitions defined using the "-ingrport" keyword that use ports on the same blade or fixed-port switch are not mirrored.

- Flow Mirror cannot mirror:
  - Frames belonging to device-switch communication (for example, a FLOGI or PLOGI action)
  - Link Primitives, discarded frames, frames from a remote Control Unit Port (CUP), Link Control Frames, or frames containing domain controller addresses used as source IDs
- Active Flow Mirror flows and the In-Flight Encryption/Compression feature are mutually exclusive on fixed-port switches or a single blade. However, they can co-exist in a chassis if they are enabled on different blades in that chassis.
- Flow creation is not allowed if Advanced Performance Monitor (APM) or Port Mirroring is enabled. Similarly, APM and Port Mirroring-related operations are not be allowed if any flow (active or defined) is present on the switch.

Creating a Flow Mirror flow definition

This procedure provides step-by-step instructions for configuring a Mirror flow definition. For more specific example procedures, refer to "Flow Mirror example procedures" on page 1076.

1. Right-click a switch and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Target Switch — Name of the target switch
   - Source Device — * (an asterisk allows you to use any port)
   - Destination Device — * (an asterisk allows you to use any port)

2. Enter a name for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Clear the **Monitor** check box.
4. Select **Mirror** check box.
5. Select one of the following options:

   - **CPU Mirroring**

     Select to mirror traffic to the switch CPU, which enables you perform debugging without disturbing existing connections. By default, **CPU Mirroring** is enabled.

   - **Local Flow Mirror**

     Select to mirror traffic to an unoccupied, loopback, or mirror port on the selected target switch, which enables you to mirror a flow to a specific port, which is present on the same physical switch on which you defined the flow.

     LFM mirrored data can then be analyzed through an external Analyzer/Frame sniffer connected to the port. LFM requires that a loopback SFP be plugged in at the other end of the analyzer, or on the port configured as a mirror port.

     Note that if you select an unoccupied or loopback port, the Management application first reconfigures the port as a mirror port, and then creates the flow.

6. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
7. Select the **Activate all selected features** check box to immediately activate the flow after creation.

8. Change the target switch for the flow definition by clicking the ellipsis button to the right of the **Target Switch** field.

   The **Select Switch** dialog box displays. To manually set the **Target Switch**, refer to "Selecting the target switch from a list of Fabric Vision-capable switches" on page 1030.

9. Select one of the following format options for **End Device** mode:

   - **Port Address** (port ID) — Select to display the source and destination device address using the port ID.
   - **WWN** (world wide name) — Select to display the source and destination device address using the port WWN.

10. Enter the address or WWN of the source port in the **Source** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the source port from a list, refer to "Selecting an end device port from a list of available device ports" on page 1030.

    When you enter a port ID or WWN in the **Source** or **Destination** fields, a port information field displays beneath with the port label based on your topology layout settings (refer to "Changing the port label" on page 324). If you enter an asterisk (*) or no value in the **Source** or **Destination** fields, this field remains blank.

11. Enter the address or WWN of the source port in the **Destination** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the destination port from a list, refer to "Selecting an end device port from a list of available device ports" on page 1030.

12. (Optional) If you want to swap source and destination device port information, click **<swap>**.

13. Select one of the following format options for **Switch** mode:

    - **Port (slot/port)** — Select to display the switch ingress or egress port using the slot and port number.
    - **D,I** (domain ID, port number) — Select to display the switch ingress or egress port using the domain ID and port number.

14. Enter the ingress port in Port (slot/port) or D,I (domain ID,port number) format in the **Ingress** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the ingress port from a list, refer to "Selecting an ingress or egress port from a list of available switch ports" on page 1032.

    NOTE
    You must enter the slot number and the port number. For Backbone chassis, the slot number cannot be 0 (zero). For switches, you must enter 0 (zero) as the slot number. For logical (virtual) switches, follow the rule for the physical chassis (either Backbone chassis or switch) from which you created the logical switch.

    When you enter a value in the **Ingress** or **Egress** fields, a port information field displays beneath with the port label based on your topology layout settings (refer to "Changing the port label" on page 324). If you enter an asterisk (*) or no value in the **Ingress** or **Egress** fields, this field remains blank.

15. Enter the egress port data in Port (slot/port) or D,I (domain ID,port number) format in the **Egress** field or click the ellipsis button to select a port from the list.

    Enter an asterisk (*) to use any port. To select the egress port from a list, refer to "Selecting an ingress or egress port from a list of available switch ports" on page 1032.

    NOTE
    You must enter the slot number and the port number. For Backbone chassis, the slot number cannot be 0 (zero). For switches, you must enter 0 (zero) as the slot number. For logical (virtual) switches, follow the rule for the physical chassis (either Backbone chassis or switch) from which you created the logical switch.

16. For FCR or virtual fabrics, select one of the following format options for **FCR/XISL** mode:
    - **Fabric ID** — Select to configure a flow definition for an FCR fabric.
    - **VFID** (Virtual Fabric ID) — Select to configure a flow definition for virtual devices.
17. For FCR or virtual fabrics, enter the fabric ID or virtual fabric ID in the **Source** field.

    Enter an asterisk (*) to use any port. To select the source from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.
18. For FCR or virtual fabrics, enter the fabric ID or virtual fabric ID in the **Destination** field.

    Enter an asterisk (*) to use any port. To select the destination from a list, refer to "Selecting a fabric or virtual fabric ID from a list of available products" on page 1032.
19. (Local Flow Mirroring only) Enter the mirror port number in the **Mirror** field or click the ellipsis button to select a port from the list.

    Only enabled when you select the **Local Flow Mirroring** option for the **Mirror** feature. Only supported on 16 Gbps-capable ports on devices running Fabric OS 7.3 or later.

    To select a port from a list, refer to "Selecting a mirror port from a list of available ports" on page 1033.
20. Enter a frame type in the **Frame Type** field or click the ellipsis button to select a frame type from a list.

    To select a frame type from a list, refer to "Selecting a mirror port from a list of available ports" on page 1033.
21. Enter a LUN ID in the **LUN** field or click the ellipsis button to select a LUN ID from the list.

    NOTE
    You cannot use LUN IDs for Bidirectional flows.

    NOTE
    You cannot combine frame type, LUN, and bidirectional parameters for learning flows.

    LUN IDs can be from 0 through 65535. To select a LUN ID from a list, refer to "Selecting a LUN ID from a list of available LUNs" on page 1034.
22. Click **OK** to save the flow definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Mirror flow" on page 1073.

Monitoring a Flow Mirror flow

To monitor the summary data for a Flow Mirror flow, complete the following steps.

1. Select the device on which you defined the mirror flow and select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays pre-populated with a list of all defined flows in the **Flow Definitions** table. For more information about the fields and components of the **Flow Definitions** table, refer to "Flows Definitions table fields and components" on page 1021.
2. Select the mirror flow that you want to monitor in the **Flow Definitions** table.
3. Select a time interval for monitoring the flow in the Time duration list.

Possible values are 30 minutes, 1 hour, 6 hours, 12 hours, 1 day, 3 days, 1 week, and 1 month.

4. Click the right arrow button to display the selected flow in the **Flows** table.

   You can sort the **Flows** table by clicking any column head. You can reverse the sort order by clicking the column head again.

5. Select the **SCSI** check box to display SCSI-related measures.

   SCSI-related measures include SCSI read count, write count, read rate, write rate, read data, write data, and read and write frame data.

   Clear the check box to hide SCSI-related measures.

6. Select the **Frame** check box to display frame-related measures.

   Frame-related measures include transmit (Tx) and receive (Rx) frame count, transmit frame and receive frame rate, transmit and receive word count, and transmit and receive throughput.

   Clear the check box to hide frame-related measures.

7. Review the sub-flow data for the selected Flow Mirror flow.

   NOTE
   Flow Mirror statistic counts greater than zero imply that the mirrored flow is functioning, but should not be inferred as accurate counts at this time.

The **Flows** table, as shown is Figure 497, displays statistics and data for the selected flows. The **Flows** table has the following general characteristics and functions:

- Data updates dynamically every 5 minutes. Sort the table by clicking any column head. You can reverse the sort order by clicking the column head again.
- Locate an ingress port, egress port, source device, and destination device in the Product List or Topology Map by right-clicking a sub-flow in the **Flows** table and selecting **Locate >** ***port_type*** (where ***port_type*** is Ingress port, Egress port, Source device, or Destination device).
- Highlight inactive sub-flows in gray by selecting the **Show inactive flows** check box.

  Inactive sub-flows indicate no updates to sub-flow statistics for over 15 minutes.

  Clear the check box to hide inactive sub-flows.
- Display a flow in a performance graph. Select a row in the **Flows table and select Performance Graph >** ***graph_or_report*** (where ***graph_or_report*** is **Real Time Graph**, **Historical Graph**, or **Historical Report**).

  Refer to "Performance integration with Flow Vision" on page 1098 **for additional details. Note that the Performance Graph list is only available when there is at least one sub-flow selected in the Flows table.**
- A single flow definition may yield data in multiple rows in the **Flows** table. For example, if you defined the flow definition source ID (SID) and destination ID (DID) as *, this might result in five rows if the source is communicating with five destination IDs. In the case of a learning flow, a root flow also displays to summarize all sub-flows.
- Each unique sub-flow for the flow definition displays in the **Flows** table if it was reported within the selected time duration. If the last data point did not report that flow, the reported values may be 0.
- The measures that display are based on the flow definition. Therefore, not all columns may be populated for the selected flows.

FIGURE 497 Flow Vision dialog box (Mirror Flows table)

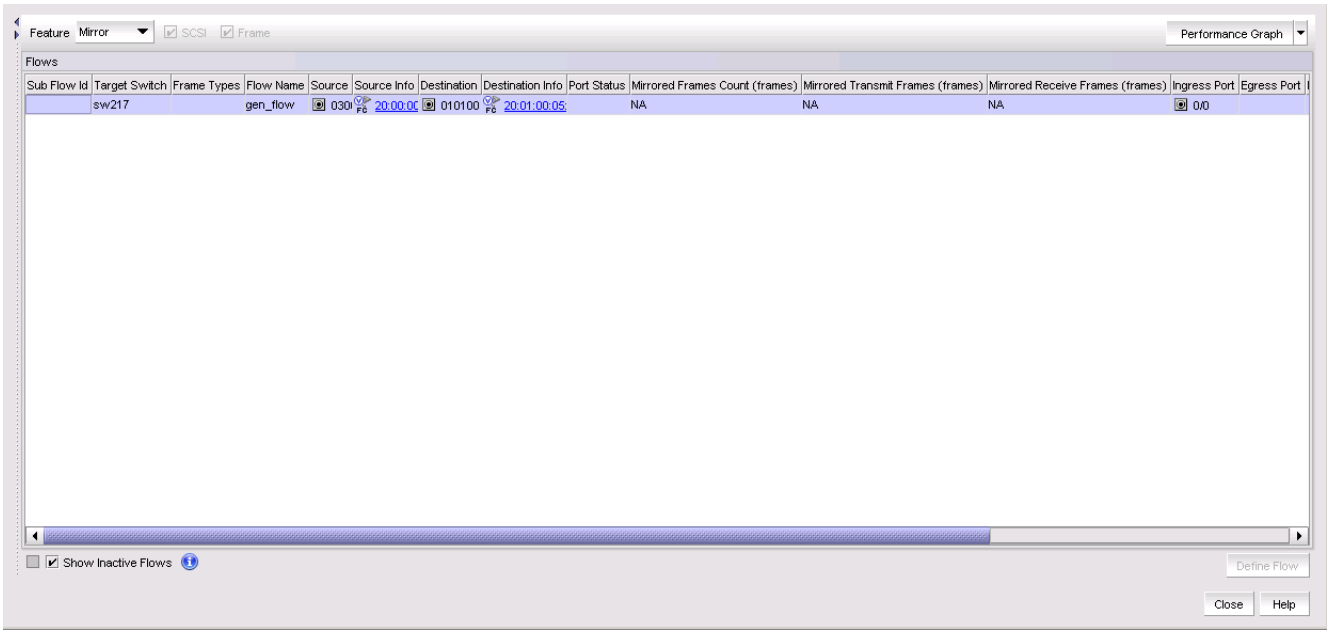

Table 93 describes the sub-flow data displayed in the **Flows** table when you select **Mirror** from the **Feature** list above the **Flows** table.

TABLE 93 Flows table fields and components (Mirror Flow)

Field and components	Description
Sub Flow Id	The sub-flow database identifier.
Target Switch	The switch on which you created the flow definition.
Frame Type	All frame types defined in the flow definition.
Flow Name	The user-defined name for the flow definition.
Source	The source identifiers defined in the flow definition.
Source Info	The icon and name for the source device. The device name is a hyperlink. Click to launch the device's property sheet. This field is empty if the source device is not defined in the flow definition.
Destination	The port number of the destination device defined in the flow definition. An * (asterisk) indicates learned flows.
Destination Info	The icon and name for the destination device. The device name is a hyperlink. Click to launch the device's property sheet. This field is empty if the destination device is not defined in the flow definition.
Port Status	The operational status (online or offline) for the port.
Mirrored Frames Count (frames)	The number of mirrored frames as reported in the last data point received for the flow.
Mirrored Transmit Frames (frames)	The number of transmitted mirrored frames as reported in the last data point received for the flow.
Mirrored Receive Frames (frames)	The number of received mirrored frames as reported in the last data point received for the flow.
Ingress Port	The ingress port defined in the flow definition. The port name is a hyperlink. Click to launch the port's property sheet. A yellow icon indicates a bottlenecked port.

TABLE 93 Flows table fields and components (Mirror Flow) (Continued)

Field and components	Description
Egress Port	The egress port defined in the flow definition. The port name is a hyperlink. Click to launch the port's property sheet. A yellow icon indicates a bottlenecked port.
Mirror Port	The mirror port identifier defined in the flow definition.
LUN	Any LUN values defined in the flow definition.
Bi-Directional	Whether the flow is bidirectional (**yes**) or not (**no**).
Flow Definition Persistence	Whether the flow is configured to persist over switch reboots (**yes**) or not (**no**).
Last Updated Time	The date and time the sub-flow was last updated.

8. Click **Close** to close the **Flow Vision** dialog box.

Flow Mirror example procedures

The following use cases describe how to use Flow Mirror to troubleshoot typical fabric performance problems.

Diagnosing SCSI reserve and SCSI release performance

If there is excessive SCSI reserve and release activity in a virtualized environment, you can use Flow Mirror to identify the affected LUNs. The following example creates a flow to mirror all the SCSI release frames from multiple servers to LUNs on the target port 1/20. You can then analyze the mirrored frames to determine the impacted LUNs.

1. Right-click a target port (1/20) and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Egress port — port number (1/20)

2. Enter a name (flow_scsi) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Clear the **Monitor** check box.
4. Select the **Mirror** check box.
5. Select the **CPU Mirroring** or **Local Flow Mirror** option.

   By default, **CPU Mirroring** is enabled.

6. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
7. Select the **Activate all selected features** check box to immediately activate the flow after creation.
8. Enter an asterisk (*) in the **Destination** field.

9. Enter the frame types (SCSI2Release,SCSI2Reserve) in the **Frame Type** field.

   To select a frame type from a list, click the ellipsis button. Refer to "Flow frame type parameters" on page 1017 for a list of the supported frame type parameters.

10. Click **OK** to save the flow definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Mirror flow" on page 1073.

Troubleshooting protocol errors

You can use Flow Mirror to mirror protocol error frames. The following example mirrors only ABTS frames egressing through port 1/20 to identify the ABTS protocol condition.

NOTE
This example can also be set up to mirror frames based on the total ABTS count provided by Flow Monitor.

1. Right-click a target port (1/20) and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Egress port — port number (1/20)

2. Enter a name (flow_protocol) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Clear the **Monitor** check box.
4. Select the **Mirror** check box.
5. Select the **CPU Mirroring** or **Local Flow Mirror** option.

   By default, **CPU Mirroring** is enabled.

6. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
7. Select the **Activate all selected features** check box to immediately activate the flow after creation.
8. Enter an asterisk (*) in the **Destination** field.
9. Enter the frame type (abts) in the **Frame Type** field.

   To select a frame type from a list, click the ellipsis button. Refer to "Flow frame type parameters" on page 1017 for a list of the supported frame type parameters.

10. Click **OK** to save the flow definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Mirror flow" on page 1073.

Diagnosing a slow-draining F_Port

The following example creates a flow to mirror traffic passing in both directions from device 0x010200 to F_Port 15 on device 0x040500. The collected frame data may help you diagnose the slow-draining device. Figure 498 provides a diagram of what is happening in this example.

FIGURE 498 Flow Mirror revealing a slow drain

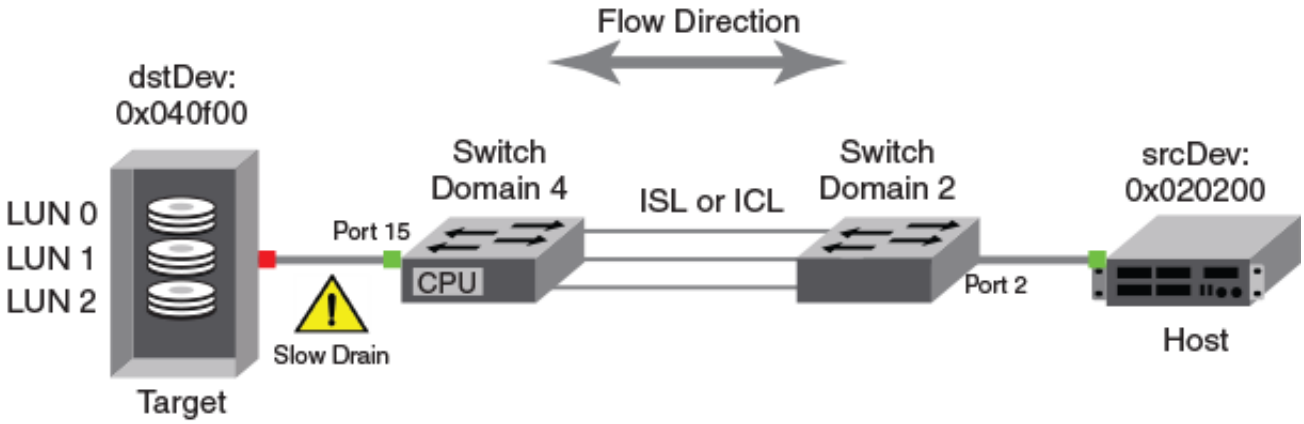

1. Right-click a target port (15) and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Egress port — port number (1/20)

2. Enter a name (flow_slowdrain) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Clear the **Monitor** check box.
4. Select the **Mirror** check box.
5. Select the **CPU Mirroring** or **Local Flow Mirror** option.

   By default, **CPU Mirroring** is enabled.

6. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
7. Select the **Activate all selected features** check box to immediately activate the flow after creation.

8. Enter the port address (0x010200) in the **Source** field.
9. Enter the port address (0x040500) in the **Destination** field.
10. Click **OK** to save the flow definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Mirror flow" on page 1073.

Tracking SCSI commands

You can use Flow Mirror to track the SCSI commands being initiated by a host. For example, you can use Flow Mirror to find all the targets with which the host is communicating. You can use this data to identify the favorite targets of a host, which then allows you to provide additional privileges, such as creating Quality of Service (QoS) or Traffic Isolation (TI) paths between those devices.

The following example mirrors scsicmd frames ingressing through port F1 to identify the SCSI commands being initiated by a host.

1. Right-click a port (F1) and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (F1)

2. Enter a name (mirror_scsismd) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Clear the **Monitor** check box.
4. Select the **Mirror** check box.
5. Select the **CPU Mirroring** or **Local Flow Mirror** option.

   By default, **CPU Mirroring** is enabled.

6. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
7. Select the **Activate all selected features** check box to immediately activate the flow after creation.
8. Enter an asterisk (*) in the **Destination** field.
9. Enter the frame type (scsicmd) in the **Frame Type** field.

   To select a frame type from a list, click the ellipsis button. Refer to "Flow frame type parameters" on page 1017 for a list of the supported frame type parameters.

10. Click **OK** to save the flow definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Mirror flow" on page 1073.

Tracking latency between a host and all connected targets

You can use Flow Mirror to track the latency of SCSI Initiator-Target pairs so that you can load balance them, which enables you to optimize application performance.

The following examples mirror all SCSI commands and their status frames initiated by device H1 ingressing through port F1, which enables you to determine the latency from the captures.

Tracking latency using CPU Mirroring

1. Right-click a port (F1) and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor
   - Direction — Bidirectional
   - Source Device — Source identifier
   - Destination Device — * (an asterisk allows you to use any port)
   - Ingress port — port number (F1)

2. Enter a name (cpu_mirror_scsicmdsts) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Clear the **Monitor** check box.
4. Select the **Mirror** check box.
5. Select the **CPU Mirroring** option. By default, **CPU Mirroring** is enabled.
6. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
7. Select the **Activate all selected features** check box to immediately activate the flow after creation.
8. Enter the source device (H1) in the **Source** field.
9. Enter an asterisk (*) in the **Destination** field.
10. Enter the frame type (scsicmdsts) in the **Frame Type** field.

    To select a frame type from a list, click the ellipsis button. Refer to "Flow frame type parameters" on page 1017 for a list of the supported frame type parameters.

11. Click **OK** to save the flow definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Mirror flow" on page 1073.

Tracking latency using Local Flow Mirror

1. Right-click a port (F1) and select **Fabric Vision > Flow > Add**.

   The **Add Flow Definition** dialog box displays with the following criteria and flow identifiers pre-populated:

   - Feature — Monitor

- Direction — Bidirectional
- Source Device — Source identifier
- Destination Device — * (an asterisk allows you to use any port)
- Ingress port — port number (F1)

2. Enter a name (lfm_mirror_scsicmdsts) for the flow definition in the **Name** field.

   The name cannot be over 20 characters and can only include alphanumeric characters or underscores.

   NOTE
   For a physical switch, the name must be unique. However, for logical switches, the name does not have to be unique.

3. Clear the **Monitor** check box.
4. Select the **Mirror** check box.
5. Select the **Local Flow Mirror** option. By default, **CPU Mirroring** is enabled.
6. Select the **Persist over switch reboots** check box to persist this flow definition over reboots.
7. Select the **Activate all selected features** check box to immediately activate the flow after creation.
8. Enter the source device (H1) in the **Source** field.
9. Enter an asterisk (*) in the **Destination** field.
10. Enter the mirror port number in the **Mirror** field or click the ellipsis button to select a port from the list.

    Only supported on 16 Gbps-capable ports on devices running Fabric OS 7.3 or later.

    To select a port from a list, refer to "Selecting a mirror port from a list of available ports" on page 1033.

11. Enter the frame type (scsismd) in the **Frame Type** field.

    To select a frame type from a list, click the ellipsis button. Refer to "Flow frame type parameters" on page 1017 for a list of the supported frame type parameters.

12. Click **OK** to save the flow definition.

    When the flow definition activates, the **Flow Vision** dialog box displays with the new flow selected (highlighted) in the **Flow Definitions** table. To review the sub-flow data for the selected flow, refer to "Monitoring a Flow Mirror flow" on page 1073.

Predefined flow definition templates

The Management application provides predefined flow templates that are available when you select a switch port, an initiator port, or a target port.

NOTE
The predefined flow templates are not supported on switches, storage arrays, or host groups.

Table 94 lists the predefined flow templates available by port type.

TABLE 94 Predefined templates

Port type	Template name	Platform support		Required Fabric OS version
		16 Gbps-capable port blade	8 Gbps-capable port blade	
Host Port Node	**Learn Flows from/to Host Port**	Supported	Not Supported	v7.2.0 or later
	Monitor All Flows from Host	Supported	Supported	v7.2.0 or later
	Monitor **Flows** to LUN	Supported	Supported	v7.2.0 or later
	Local **Flows M**irroring	Supported	Not Supported	v7.3.0 or later
	Ingress Monitoring	Supported	Supported	v7.2.0 or later
	Egress Monitoring	Supported	Supported	v7.2.0 or later
Storage Port	**Learn Flows from/to Storage Port**	Supported	Not Supported	v7.2.0 or later
	Monitor All Flows to Storage	Supported	Supported	v7.2.0 or later
	Monitor **Flows** to LUN	Supported	Supported	v7.2.0 or later
	Ingress Monitoring	Supported	Supported	v7.2.0 or later
	Egress Monitoring	Supported	Supported	v7.2.0 or later
	Local Flows Mirroring	Supported	Not Supported	v7.3.0 or later
Storage/Storage Array/Host Node/Host Group	• Static Flow • Aggregated Flow • Frame TYpe Monitoring • LUN Monitoring	Not Applicable		
SIM-Ports	**Learn Flows from/to Host Port** (connected to an initiator)	Supported	Not Supported	v7.2.0 or later
	Learn **Flows** from/to Storage Port (connected to a target)	Supported	Not Supported	v7.2.0 or later
	Monitor flows to LUN	Supported	Not Supported	v7.2.0 or later
	Monitor **All Flows** from Host	Supported	Supported	v7.2.0 or later
	Monitor **All Flows** to Storage	Supported	Supported	v7.2.0 or later
	Ingress Monitoring	Supported	Supported	v7.2.0 or later
	Egress Monitoring	Supported	Supported	v7.2.0 or later
	Local Flows Mirroring	Supported	Not Supported	v7.3.0 or later
Backbone E_Ports	Learn Routed **Flows**	Supported	Not Supported	v7.3.0 or later
E_Ports	Learn **Flows** on Port	Supported	Not Supported	v7.3.0 or later
	Monitor **Flows** by Frame Type	Supported	Supported	v7.2.0 or later
	Ingress Monitoring	Supported	Supported	v7.2.0 or later
	Egress Monitoring	Supported	Supported	v7.2.0 or later
	Learn ISL Flows	Supported	Not Supported	v7.3.0 or later
	Learn XISL **Flows** (Base Switch)	Supported	Not Supported	v7.3.0 or later

TABLE 94 Predefined templates (Continued)

Port type	Template name	Platform support		Required Fabric OS version
		16 Gbps-capable port blade	8 Gbps-capable port blade	
F_Ports	Learn **Flows** from/to Host Port (connected to an initiator)	Supported	Not Supported	v7.2.0 or later
	Monitor **All Flows** from Host	Supported	Supported	v7.2.0 or later
	Monitor **Flows** to LUN	Supported	Supported	v7.2.0 or later
	Local **Flows** Mirroring	Supported	Not Supported	v7.3.0 or later
	Learn **Flows** from/to Storage Port (connected to a target)	Supported	Not Supported	v7.2.0 or later
	Monitor **Flows** by Frame Type	Supported	Supported	v7.2.0 or later
	Monitor **All Flows** to Storage	Supported	Supported	v7.2.0 or later
VE_Ports	Learn Tunnel **Flows**	Supported	Not Supported	v8.0.1 or later
	Monitor **Flows** by Frame Type	Supported	Not Supported	v8.0.1 or later
	Ingress Monitoring	Supported	Not Supported	v8.0.1 or later
	Egress Monitoring	Supported	Not Supported	v8.0.1 or later

Creating a flow definition from a template

The predefined flow templates is displayed only when you select a Flow Vision-capable switch port, initiator port, or target port. Lists only the templates that are supported by the selected port.

1. Right-click a switch port, an initiator port, or a target port from the Connectivity Map or Product List and select **Flow Vision > Flow**.

   A list of the flow templates supported on the selected port displays.

2. Select the predefined flow template you want from the list, as shown in (Figure 499).

FIGURE 499 List of preconfigured flow templates

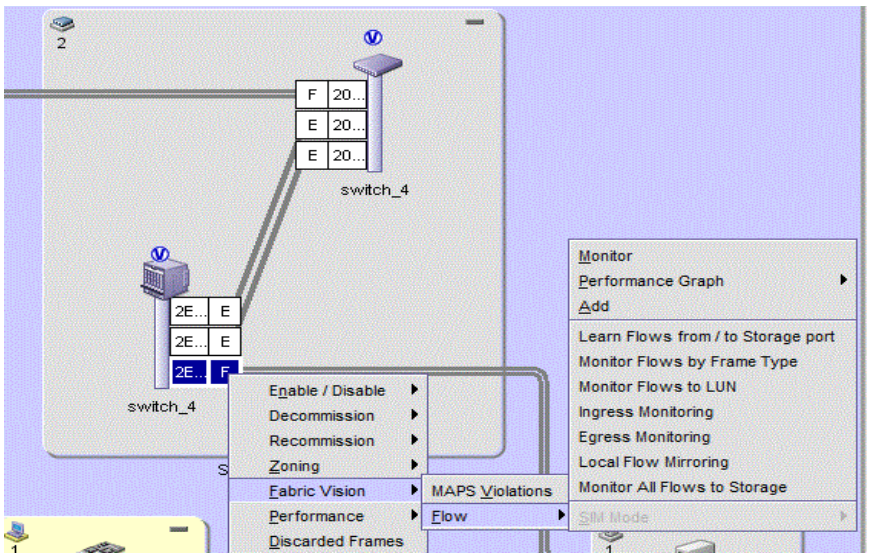

The **Add Flow - *predefined_flow*** dialog box displays, as shown in (Figure 500).

FIGURE 500 Sample Add Flow - predefined_monitor dialog box

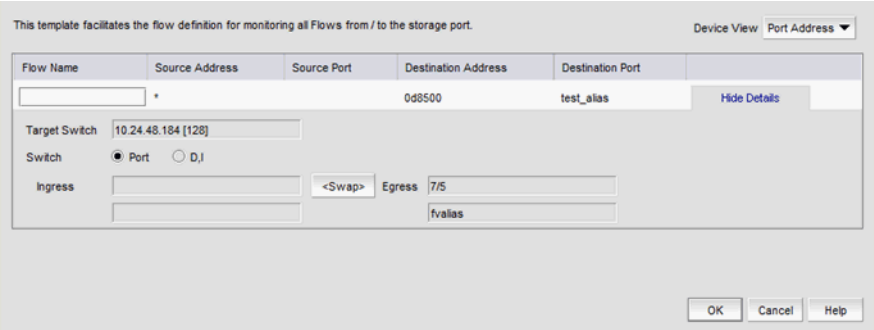

3. Enter a name for the flow definition in the **Flow Name** field.
4. Select one of the following format options from the **Device View** list:
   - **Port Address** (port ID) — Select to display the source and destination device address using the port ID.
   - **WWN** (world wide name) — Select to display the source and destination device address using the port WWN.
5. Enter the port ID or WWN of the source port in the **Source** field or click the ellipsis button to select a port from the list.

   Enter an asterisk (*) to use any port. To select the source port from a list, refer to "Selecting an end device port from a list of available device ports" on page 1030.
6. Enter the port ID or WWN of the destination port in the **Destination** field or click the ellipsis button to select a port from the list.

   Enter an asterisk (*) to use any port. To select the destination port from a list, refer to "Selecting an end device port from a list of available device ports" on page 1030.
7. Click **Show Details** to view information about the Target Switch, Ingress, and Egress ports based on the selected port.

   This data is not editable. Click **Hide Details** to hide the Target Switch, Ingress, and Egress port information.
8. Click **OK** to deploy the flow definition.

   The **Deployment Status** dialog box displays with a list of all defined flows in the **Flow Definitions Status** table, as shown is Figure 501.

FIGURE 501 Deployment Status dialog box

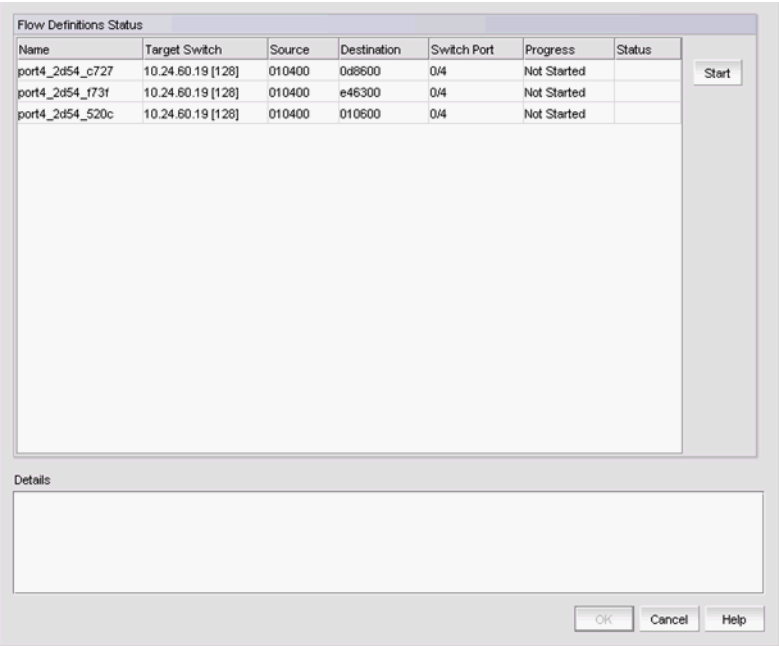

Table 89 describes the data displayed in the **Flow Definitions Status** table.

TABLE 95 Flows Definitions table fields and components

Field and components	Description
Name	The user-defined name for the flow definition.
Target Switch	The switch on which you created the flow definition.
Source	The source identifiers defined in the flow definition.
Destination	The port number of the destination device defined in the flow definition.
Switch Port	The switch port defined in the flow definition.
Progress	The progress of the flow (Not Started, In Progress, or Completed).
Status	The status of the flow (empty if not started; otherwise, Sucess or Failed).

9. **Select one or more flow definitions in the Flow Definitions Status** table and click **Start**.
10. View additional details for a deployed flow definition by selecting the flow defintion in the **Flow Definitions Status** table.

    The reason for sucess or failure displays in the **Details** area.

    To review the sub-flow data for the selected flow, refer to "Monitoring flows" on page 1021.

Predefined flow definition templates for initiator group and storage array

The predefined flow templates only display when you select a Flow Vision-capable initiator, or initiator group, or storage, or storage array.

1. Right-click an initiator, or initiator group, or storage, or storage array from the Connectivity Map or Product List and select **Flow** > **Static Flow** or **Aggregated Flow** or **Frame Type Monitoring** or **LUN Monitoring**.
2. Select the predefined flow template you want from the list, as shown in Figure 502.

FIGURE 502 List of preconfigured flow templates for initiator/initiator group, storage/storage array

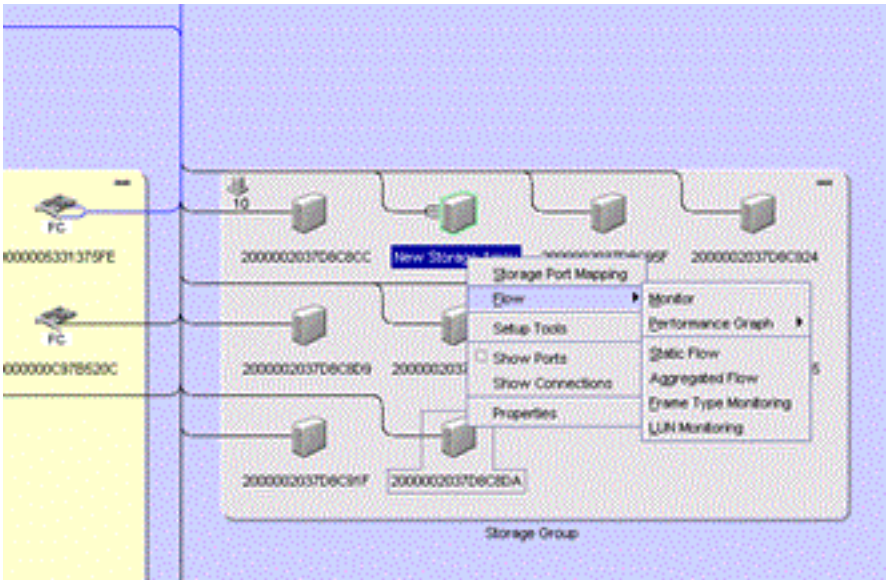

FIGURE 503 The Add Flow - predefined_flow dialog box displays, as shown in Figure 500 (where predefined_flow is the predefined monitor type).

- Static Flow — Use this **Add Flow Definition-Static Flow** dialog box to define bidirectionnal Monitor flow between a source and a destination device. Refer to "Configuring Static Flow template" on page 1087 for more information.
- Aggregated Flow — Use this **Add Flow Definition-Aggregated Flow** dialog box to define bidirectional Monitor flow for a storage port. Refer to "Configuring Aggregated Flow template" on page 1089 for more information.
- Frame Type Monitoring — Use this **Add Flow Definition-Frame Type Monitoring** dialog box to define unidirectional Monitor flow between a source and a destination device including a frame type from the list. Refer to "Configuring Frame Type Monitoring template" on page 1090 for more information.
- LUN Monitoring — Use this **Add Flow Definition-LUN Monitoring** dialog box to define unidirectional Monitor flow between a source and a destination device with the LUN details. Refer to "Configuring LUN Monitoring predefined template" on page 1091 for more information.

Configuring Static Flow template

The Static Flow predefined template receives the source device details from the storage array, searches all the zones where the source device is a member, iterates through all the zone members and receives each initiator or initiator with target port details, and creates flow definition row using the device port and source port details as shown in Figure 502.

FIGURE 504 Add Flow Definition - Static Flow template dialog box

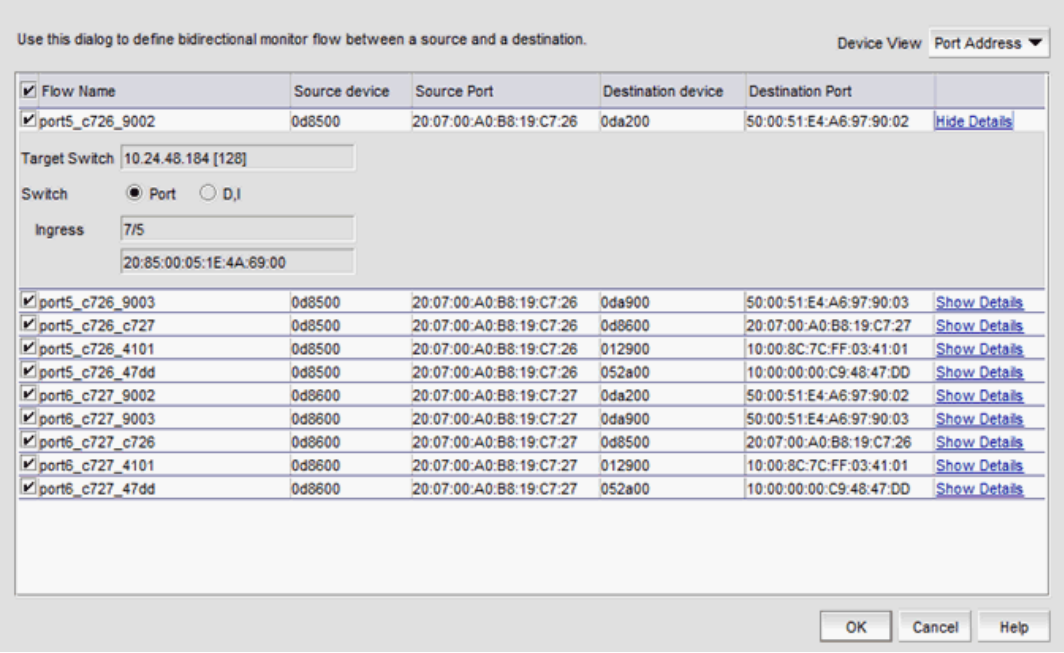

Perform the following steps, in the **Add Flow Definition - Static Flow** dialog box:

The default **Device View** is Port Address.

1. The **Flow Name** field is filled with the predefined flow name and the format is *<portname>_<last 4 digits of the source device WWN>_<last 4 digits of the destination device WWN>.*
2. Select one of the following format options for **Device View** list:
   - **Port Address** (port ID) — Select to display the source and destination device address using the port ID. The default **Device View** is Port Address.
   - **WWN** (world wide name) — Select to display the source and destination device address using the port WWN.
3. Click **Show Details** to view information about the Target Switch, Ingress, and Egress ports based on the selected device details.

   This data is not editable. Click **Hide Details** to hide the Target Switch, Ingress, and Egress port information.
4. Click **OK** to deploy the flow definition.

   The **Deployment Status** dialog box displays with a list of all defined flows in the **Flow Definitions Status** table, as shown is Figure 501. Refer to Table 89 for the data displayed in the **Flow Definitions Status** table.

FIGURE 505 Deployment Status dialog box

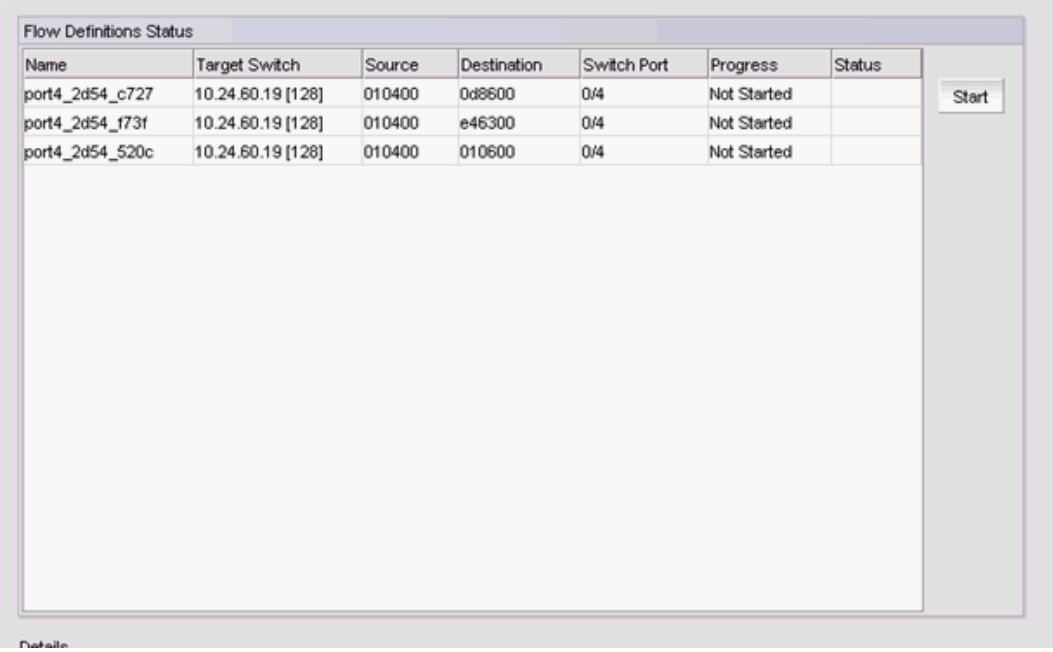

5. Click **Start** to deploy the flow definitions.
6. View additional details for a deployed flow definition by selecting the flow defintion in the **Flow Definitions Status** table.

   The reason for sucess or failure displays in the **Details** area.

   To review the sub-flow data for the selected flow, refer to "Monitoring flows" on page 1021.

Configuring Aggregated Flow template

The Aggregated Flow predefined template creates flow rows for all the ports in the storage array with the source device and empty destination device details as shown in Figure 506.

FIGURE 506 Add Flow Definition - Aggregated Flow template dialog box

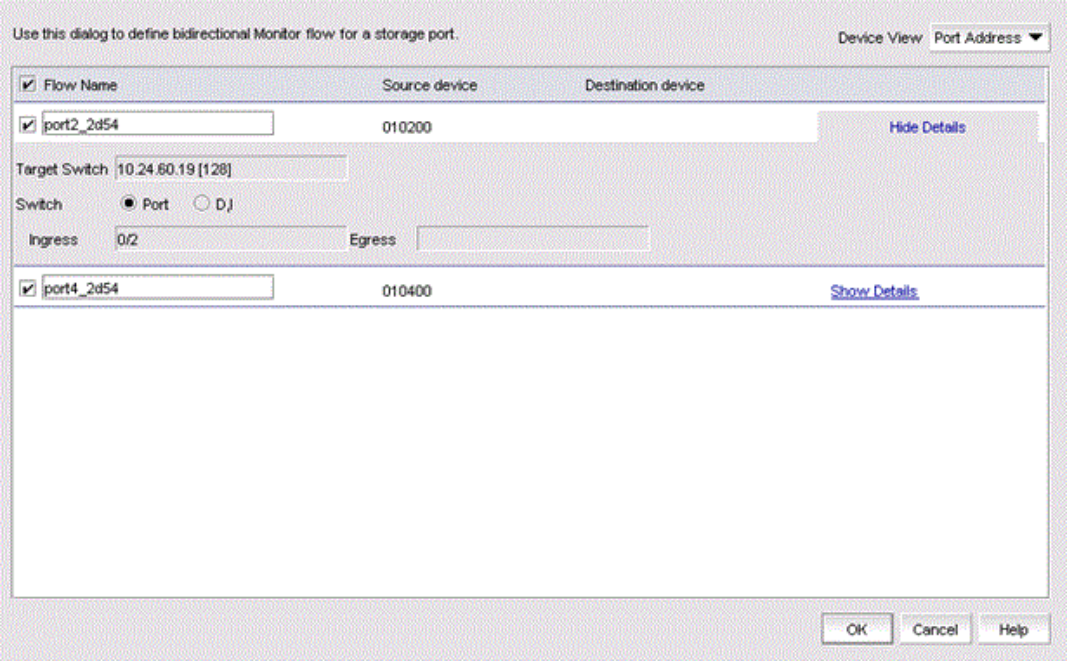

Perform the following steps, in the **Add Flow Definition - Aggregated Flow** dialog box:

The default **Device View** is Port Address.

1. The **Flow Name** field is filled with the predefined flow name and the format is *<portname>_<last 4 digits of the source device WWN>.*
2. Continue from step 2 through step 6 in "Configuring Static Flow template" on page 1087 to deploy the flow definition.

Configuring Frame Type Monitoring template

The Frame Type Monitoring template creates flow rows like Static Flow template with frame types as shown in Figure 507.

FIGURE 507 Add Flow Definition - Frame Type Monitoring template dialog box

Perform the following steps, in the **Add Flow Definition - Frame Type Monitoring** dialog box:

The default **Device View** is Port Address.

1. The **Flow Name** field is filled with the predefined flow name and the format is *<portname>_<last 4 digits of the source device WWN>_<last 4 digits of the destination device WWN>.*
2. Select **Frame Type** from the list.
3. Continue from step 2 through step 6 in "Configuring Static Flow template" on page 1087 to deploy the flow definition.

Configuring LUN Monitoring predefined template

The LUN Monitoring template creates flow rows like Static Flow template including the LUN ID field. You can configure LUN ID for the flows as shown in Figure 508.

FIGURE 508 Add Flow Definition - LUN Monitoring template dialog box

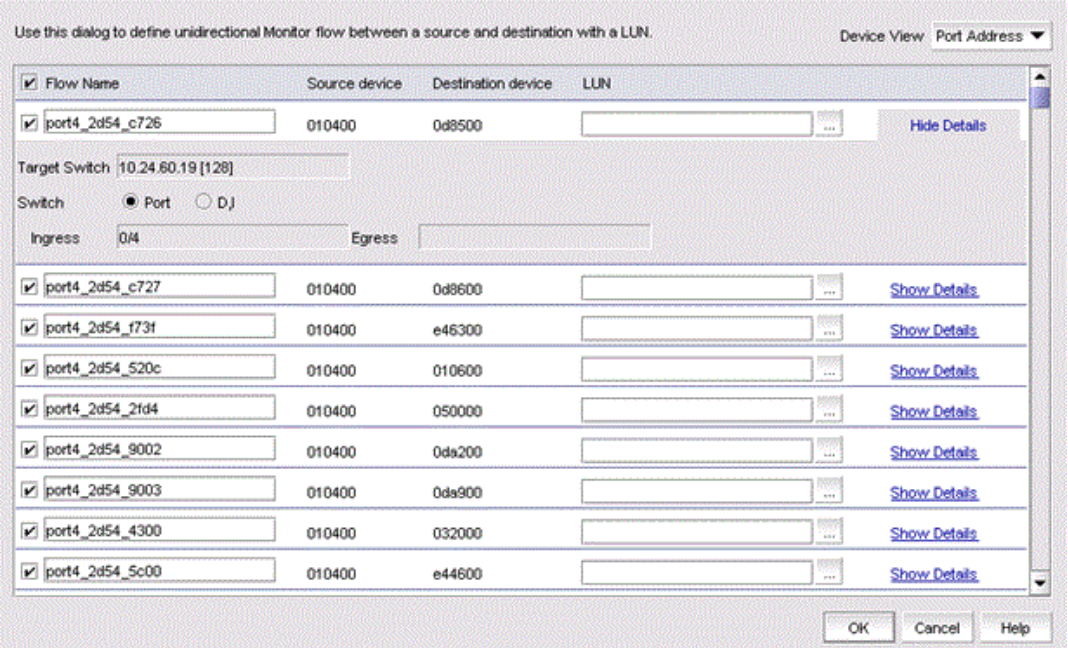

Perform the following steps, in the **Add Flow Definition - LUN Monitoring** dialog box:

The default **Device View** is Port Address.

1. The **Flow Name** field is filled with the predefined flow name and the format is *<portname>_<last 4 digits of the source device WWN>_<last 4 digits of the destination device WWN>.*
2. Click the ellipsis button next to the **LUN Monitoring** field, to configure LUN ID.

   Continue from step 2 through step 6 in "Configuring Static Flow template" on page 1087 to deploy the flow definition.

Flow Vision interoperability with other features

Flow Vision integration with other Management application features enables you to create flows from discarded frames, bottlenecked ports, routing information (Trace Route), and port connectivity. You can also monitor flows from the Dashboard and the Performance graphs.

Frame Viewer integration with Flow Vision

You can create a flow definition for a single discarded frame. For more information about Frame Viewer and viewing discarded frames, refer to "Frame viewer" on page 440.

Creating a Frame Viewer flow definition

1. Select a frame from the **Discarded Frame Log for the Selected Product** list in the **Discarded Frames** dialog box.
2. Click **Add Flow**.
3. The **Add Flow Definition** dialog box displays with the following fields populated:
   - Source Device — Source ID
   - Destination Device — Destination ID
   - Ingress port — Transmit port
   - Egress Port — * (an asterisk) if port is on a 16 Gbps-capable FC device; otherwise empty
   - Direction — Bidirectional
4. Complete the flow definition using one of the following procedures:
   - "Creating a Flow Monitor flow definition" on page 1026
   - "Creating a Flow Generator flow definition" on page 1061
   - "Creating a Flow Mirror flow definition" on page 1071

Monitoring and Alerting Policy Suite integration with Flow Vision

You can create a Monitoring and Alerting Policy Suite (MAPS) policy for events related to flows and sub-flows (refer to "Configuring a MAPS policy" on page 1232 and "Importing Flow definitions" on page 1237). You can also view a filtered list of MAPS violations for imported flows or sub-flows ("MAPS violations" on page 1255).

Bottleneck Detection integration with Flow Vision

You can create a flow definition for a single bottlenecked port. For more information about Bottleneck Detection, refer to "Bottleneck detection" on page 988.

Creating a bottlenecked port flow definition

1. Select a port from the **Products/Ports** list in the **Bottlenecks** dialog box.
2. Click **Add Flow.**

   The **Add Flow Definition** dialog box displays with fields and options populated based on the port selected on the **Bottlenecks** dialog box.

   Table 96 details the fields and options populated based on the port selected on the **Bottlenecks** dialog box.

TABLE 96 Add Flow Definition dialog box fields and options populated per Bottlenecks selection

Port selected	Options and fields populated
E_Port	• Target Switch = Selected switch • Source Device = * if port is on a 16-Gbps switch, otherwise empty • Destination Device = * if port is on a 16-Gbps switch, otherwise empty • Ingress Port = Selected port number • Direction = Bidirectional
Initiator	• Target Switch = Selected port's connected switch • Source Device = Initiator port ID • Destination Device = * if port is on a 16-Gbps switch, otherwise empty • Ingress Port = Selected port's connected switch port number • Direction = Bidirectional
Target	• Target Switch = Selected port's connected switch • Source Device = * if port is on a 16-Gbps switch, otherwise empty • Destination Device = Target port ID • Ingress Port = Selected port's connected switch port number • Direction = Bidirectional
Any port other than E_Port, initiator, or target port	• Target Switch = Selected switch • Source Device = * if port is on a 16-Gbps switch, otherwise empty • Destination Device = * if port is on a 16-Gbps switch, otherwise empty • Ingress Port = Selected switch port number • Direction = Bidirectional

FC Trace Route integration with Flow Vision

You can create a flow definition based on trace route data (forward routes, reverse routes, and FC ping). For more information about trace route and ping, refer to "Tracing FC routes" on page 945.

Creating a Forward Route flow definition

You can create a flow definition based on forward route data.

1. Select a row on the **Forward Route** tab in the **Trace Route Summary** dialog box.
2. Click **Add Flow.**

   The **Add Flow Definition** dialog box displays with fields and options populated based on the selected row.

   Table 97 describes how options are populated on the **Add Flow Definition** dialog box according to the row selected on the **Forward Route** tab.

TABLE 97 Add Flow Definition dialog box options populated per Forward Route row selection

Row selected	Options populated
First row (where **In Port Address** is the source device port's connected switch port address)	Target Switch = Switch of selected row Source Device = Source ID of source device port Destination Device = * if port is on a 16-Gbps switch, otherwise empty Ingress Port = "In" port slot/port of selected row Direction = Bidirectional
Last row (where **Out Port Address** is the destination device port's connected switch port address)	Target Switch = Switch of selected row Source Device = * if port is on a 16-Gbps switch, otherwise empty Destination Device = * if port is on a 16-Gbps switch, otherwise empty Ingress Port = "In" port slot/port of selected row Direction = Bidirectional
None	Target Switch = Switch in selected row Source Device = Source ID from source device port Destination Device = Destination ID from destination device port Direction = Bidirectional

3. Complete the flow definition using one of the following procedures:
   - "Creating a Flow Monitor flow definition" on page 1026
   - "Creating a Flow Generator flow definition" on page 1061
   - "Creating a Flow Mirror flow definition" on page 1071

Creating a Reverse Route flow definition

You can create a flow definition based on reverse route data. For more information about trace route and ping, refer to "Tracing FC routes" on page 945.

1. Select a row on the **Reverse Route** tab in the **Trace Route Summary** dialog box.
2. Click **Add Flow.**

   The **Add Flow Definition** dialog box displays with fields and options populated based on the selected row.

   Table 98 describes how options are populated on the **Add Flow Definition** dialog box according to the row selected on the **Reverse Route** tab.

TABLE 98 Add Flow Definition dialog box options populated per Reverse Route row selection

Row selected	Options populated
First row (where **In Port Address** is the destination device port's connected switch port address)	Target Switch = Switch of selected row Source Device = * if port is on a 16-Gbps switch, otherwise empty Destination Device = Device ID from destination device port Ingress Port = "In" port slot/port of selected row Direction = Bidirectional
Last row (where **Out Port Address** is the source device port's connected switch port address)	Target Switch = Switch of selected row Source Device = Source ID from source device port Destination Device = * if port is on a 16-Gbps switch, otherwise empty Ingress Port = "Out" port slot/port of selected row Direction = Bidirectional
Row other than the first or last row	Target Switch = Switch of selected row Source Device = * if port is on a 16-Gbps switch, otherwise empty Destination Device = * if port is on a 16-Gbps switch, otherwise empty Ingress Port = "In" port slot/port of selected row Direction = Bidirectional
None	Target Switch = Switch in selected row Source Device = Source ID from source device port Destination Device = Destination ID from destination device port Direction = Bidirectional

3. Complete the flow definition using one of the following procedures:
   - "Creating a Flow Monitor flow definition" on page 1026
   - "Creating a Flow Generator flow definition" on page 1061
   - "Creating a Flow Mirror flow definition" on page 1071

Creating a FC Ping flow definition

You can create a flow definition based on FC ping data. For more information about trace route and ping, refer to "Tracing FC routes" on page 945.

1. Select a row on the **FC Ping** tab in the **Trace Route Summary** dialog box.
2. Click **Add Flow.**

   The **Add Flow Definition** dialog box displays with fields and options populated based on the selected row:
   - Target Switch — Switch for source WWN
   - Source Device — Source ID
   - Destination Device — Destination ID
   - Direction — Bidirectional
3. Complete the flow definition using one of the following procedures:
   - "Creating a Flow Monitor flow definition" on page 1026
   - "Creating a Flow Generator flow definition" on page 1061

- "Creating a Flow Mirror flow definition" on page 1071

Port connectivity integration with Flow Vision

You can create a flow definition based on port connectivity data. For more information about port connectivity, refer to "Viewing port connectivity" on page 443.

Creating a port connectivity flow definition

1. Select a row on the **Port Connectivity View** dialog box.
2. Click **Add Flow.**

   The **Add Flow Definition** dialog box displays with fields and options populated based on the selected row.

   Table 99 describes how options are populated on the **Add Flow Definition** dialog box according to the selected row.

   TABLE 99 Add Flow Definition dialog box options populated per port connectivity row selection

Row selected	Options populated
E_Port	Target Switch = Selected switch Source Device = * if port is on a 16-Gbps switch, otherwise empty Destination Device = * if port is on a 16-Gbps switch, otherwise empty Direction = Bidirectional
Initiator	Target Switch = Selected port's connected switch Source Device = Initiator port ID Destination Device = * if port is on a 16-Gbps switch, otherwise empty Ingress Port = Selected port's connected switch port number Direction = Bidirectional
Target	Target Switch = Selected port's connected switch Source Device = * if port is on a 16-Gbps switch, otherwise empty Destination Device = Target port ID Ingress Port = Selected port's connected switch port number Direction = Bidirectional
Row other than the first or last row	Target Switch = Selected switch Source Device = * if port is on a 16-Gbps switch, otherwise empty Destination Device = * if port is on a 16-Gbps switch, otherwise empty Ingress Port = Selected port number Direction = Bidirectional

3. Complete the flow definition using one of the following procedures:
   - "Creating a Flow Monitor flow definition" on page 1026
   - "Creating a Flow Generator flow definition" on page 1061
   - "Creating a Flow Mirror flow definition" on page 1071

Dashboard integration with Flow Vision

You can use the Dashboard to monitor flows.

Monitoring flows from the Dashboard

You can create a flow monitor from a performance graph ("Monitoring flows using a performance graph" on page 1098) or by creating a user-defined performance monitor ("Configuring a user-defined traffic flow performance monitor" on page 288).

Figure 509 shows a flow performance monitor.

FIGURE 509 Flow performance monitor

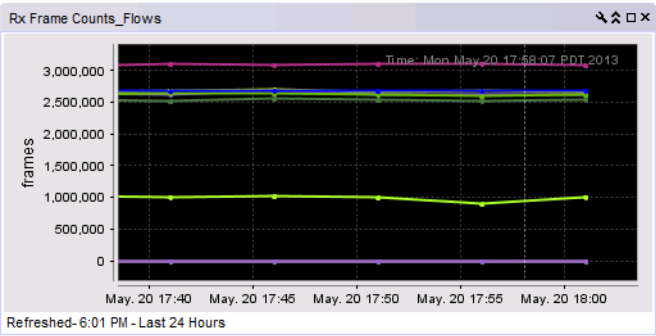

Flow widget for VE_Ports

Beginning with the Management application 14.0.1 release, you can monitor the traffic passing through the VE_Ports. You can create a custom widget for flows with ingress or egress port as VE_Ports and monitor the Flow statistics in custom dashboard. You can create Top N or Bottom N or Time series Traffic flows for VE_Port specific flows.

Performance integration with Flow Vision

You can use the performance graphs (historical or real-time) to monitor flows.

For platforms running Fabric OS 7.2 or later, Flow Vision replaces Top Talker and End-to-End monitoring.

Monitoring flows using a performance graph

1. Select **Monitor > Fabric Vision > Flow > Monitor**.

   The **Flow Vision** dialog box displays.

2. Select a row in the **Flows list.**

3. Select **Performance Graph > Historical Graph** or **Real-Time Graph**.

   You can also access performance graphs directly from the main menu (**Monitor > Fabric Vision > Flow > Historical Graph** or **Real-Time Graph**) or a shortcut menu (right-click a device or port and select **Fabric Vision > Flow > Historical Graph** or **Real-Time Graph**).

   If you select **Real-Time Graph** using a port that does not have any flows, the following message displays: “There are no active flows with the monitor feature. Do you want to creates a flow for the selected port?”. Click **OK** to create a flow for the selected port. For instructions, refer to “Creating a Flow Monitor flow definition” on page 1026, “Creating a Flow Generator flow definition” on page 1061, or “Creating a Flow Mirror flow definition” on page 1071.

   The **Historical** or **Real Time Graphs/Tables** dialog box displays.

   NOTE
   Beginning with the 14.2.0 release, devices running Fabric OS version 8.0.1 supports Ironware OS measures in both Historical and Real Time Graphs. Only F_Ports are supported with Ironware OS measures.

FIGURE 510 Historical Graph (Flow Measures selected)

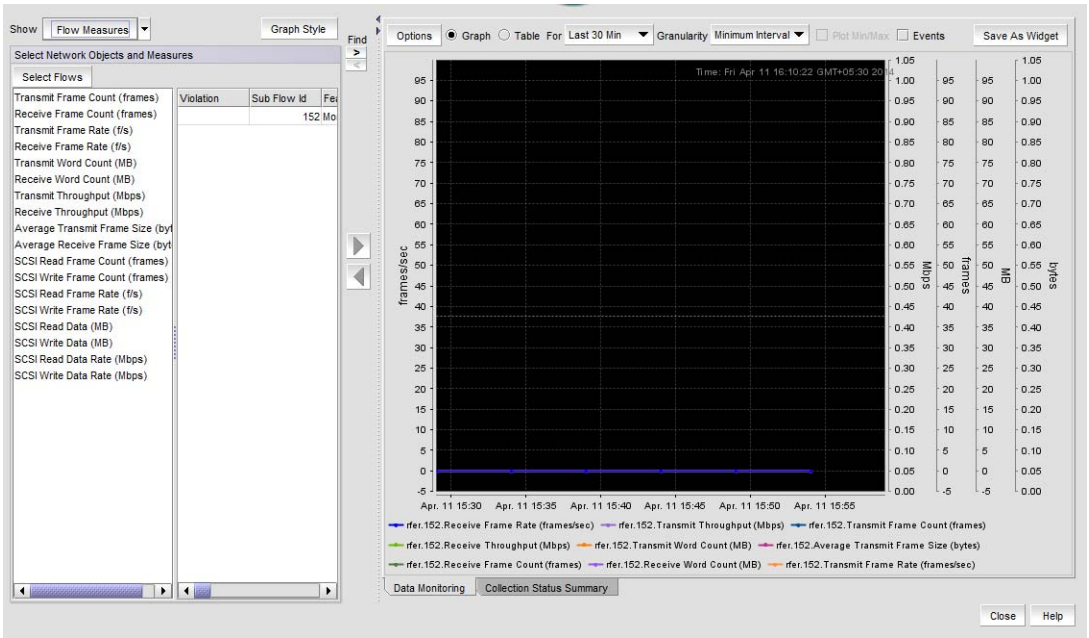

4. Select **Flows** or **Flow Measures** from the **Shows list.**

   When you select **Flows**, the tree structure shows flows for the selected device. When you select **Flow Measures**, the tree structure shows the measures associated with the selected flow.

5. Select the measures you want to display in the graph in the **Select Network Objects and Measures** list and click the right arrow button.

   Select multiple measures by holding down the CTRL key and clicking more than one measure. Select all measures for a flow by selecting the flow.

6. Remove measures from the graph by selecting the measure beneath the graph and clicking the left arrow button.
7. Configure the look and feel of the graph using the procedure in **"Configuring the performance graph"** on page 1717**.**
8. Create a flow performance monitor on the Dashboard by clicking **Save As Widget**.

   For step-by-step instructions, refer to "Configuring a traffic flows monitor from a performance graph" on page 288.

9. Click **Close** to close the dialog box.

Defining flows using Historical/Real-time performance graph

1. Select the port for which you want to generate a performance graph.
2. Select **Monitor > Performance > Historical Graph** or **Real-Time Graph**. The Historical or Real-time performance graph displays.
3. Make changes to the fields as required.
4. Select any of the column in **Additional Measures** table and click **Add Flow** button to launch **Add flow** dialog box to define flows.

Refer to "Creating a Flow Monitor flow definition" on page 1026 for more information.

Replacing Top Talker monitors

For platforms running Fabric OS 7.2 or later, Flow Vision replaces Top Talker monitoring. To use Flow Vision to configure a Top Talker monitor, complete the following steps.

1. Note any existing Top Talker monitors (refer to "SAN Top Talker monitoring" on page 983).
2. Delete existing monitors (refer to "Deleting a Top Talker monitor" on page 987).
3. Create a flow to define an initiator and target port pair for monitoring (refer to "Configuring a Top Talker monitor flow" on page 1042).

To continue using the legacy Top Talker feature, you must deactivate existing flows defined for the switch (refer to "Deactivating flows" on page 1024).

Replacing End-to-End monitors

For platforms running Fabric OS 7.2 or later, Flow Vision replaces End-to-End monitoring. To use Flow Vision to configure an End-to-End monitor, complete the following steps.

1. Note any existing End-to-End monitors (refer to "SAN end-to-end monitoring" on page 979).
2. Delete existing monitors (refer to "Deleting an end-to-end monitor pair" on page 983).
3. Create a flow to define an initiator and target port pair for monitoring (refer to "Configuring an end-to-end monitor flow" on page 1041).

To continue using the legacy End-to-End feature, you must deactivate existing flows defined for the switch (refer to "Deactivating flows" on page 1024).

Frame Monitor

Frame Monitor

NOTE
This feature is only available for Fabric OS devices running 7.3.X and earlier. It not supported on devices running 7.4.0 or later.

NOTE
Frame Monitoring is supported in Professional Plus and Enterprise Editions only. It is not supported in the Professional Edition.

Frame monitors count the number of frames transmitted through a port that match specific values in the first 64 bytes of the frame. Since the entire Fibre Channel frame header and many upper protocol (for example, SCSI) headers fall within the first 64 bytes of a frame, frame monitors can detect different types of traffic transmitted through a port. Each frame monitor keeps a timestamp of its last refresh. It also keeps a generation count, which is incremented each time the monitor is cleared.

Frame monitors generate alerts whenever the frame count for a certain frame type crosses the threshold configured for that frame type. You can configure high thresholds for every frame type, specify actions to be taken when the threshold is exceeded, and configure how often the data are sampled.

Virtual Fabrics considerations: You can assign frame monitors to ports in a logical switch. If a port is moved from one logical switch to another, however, all monitors that were assigned to the port are cleared in the new logical switch.

Trunking considerations: For trunked ports, the frame monitor is configured on the trunk master.

Frame types

The frame type can be a standard type (for example, a SCSI read command filter that counts the number of SCSI read commands that have been transmitted by the port) or a user-defined frame type customized for your particular use.

Pre-defined frame types

Pre-defined frame types include the following:

- ABTS (Abort Sequence Basic Link Service command)
- BA_ACC (Abort Accept)
- IP
- SCSI
- SCSI Read
- SCSI Write

- SCSI RW
- SCSI-2 Reserve
- SCSI-3 Reserve

Custom frame types

In addition to the standard frame types, you can create custom frame types to gather statistics that fit your needs. To define a custom frame type, you must specify a series of *offsets*, *bitmasks*, and *values*. For all transmitted frames, the switch performs these tasks:

- Locates the byte found in the frame at the specified *offset.*
- Applies the *bitmask* to the byte found in the frame.
- Compares the new value with the given *value*.
- Increments the filter counter if a match is found.

You can specify up to four values to compare against each offset. If more than one offset is required to properly define a filter, the bytes found at each offset must match one of the given values for the filter to increment its counter. If one or more of the given offsets does not match any of the given values, the counter does not increment. The value of the offset must be between 0 and 63, in decimal format. Byte 0 indicates the first byte of the Start of Frame (SOF), byte 4 is the first byte of the frame header, and byte 28 is the first byte of the payload. Thus only the SOF, frame header, and first 36 bytes of payload can be selected as part of a filter definition. Offset 0 is a special case, which can be used to monitor the first 4 bytes of the frame (SOF). When the offset is set to 0, the values 0–7 that are checked against that offset are predefined as shown in Table 100.

TABLE 100 Predefined values at offset 0

Value	SOF	Value	SOF
0	SOFf	4	SOFi2
1	SOFc1	5	SOFn2
2	SOFi1	6	SOFi3
3	SOFn1	7	SOFn3

Frame Monitoring requirements

To configure Frame Monitoring, the following requirements must be met:

- The switch must be running Fabric OS 7.0.0 or later.
- Frame Monitoring requires the Advanced Performance Monitoring license and the Fabric Watch license.

  NOTE
  The Advanced Performance Monitoring license is required to configure frame monitors. The monitoring functionality requires the Fabric Watch license.

The maximum number of frame monitors and offsets per port is platform-specific. Refer to the ***Flow Vision Administrator's Guide*** for more information.

Creating a custom frame monitor

Pre-defined frame monitors are already installed on switches that support Frame Monitoring. Use this procedure if you want to create a custom frame monitor.

1. Select **Monitor > Fabric Watch > Frame Monitor**.

   The Frame Monitor dialog box displays (Figure 511).

   FIGURE 511 Frame Monitor dialog box

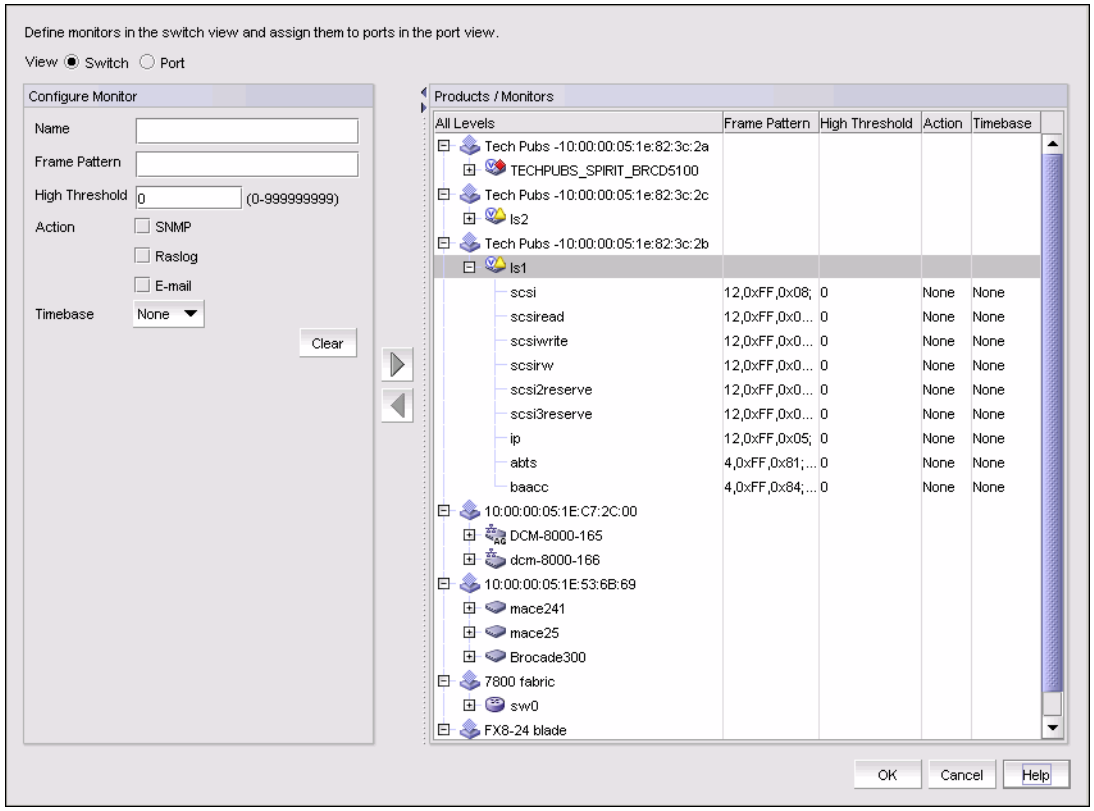

2. Select the **Switch** option.

   The Products / Monitors list displays the switches that support Frame Monitoring.

3. Enter the monitor data in the Configure Monitor area.

4. Select one or more switches in the Products / Monitors list, and click the right arrow button to assign the frame monitor to those switches.

5. Select the **Port** option.

6. Expand the switch in the Products / Ports list.

   The Monitors list displays all of the frame monitors defined for that switch.

7. Select one or more ports.

   You must select only ports belonging to the same switch.

8. Select one or more frame monitors in the Monitors list.

9. Click the right arrow button to move the frame monitor to the selected ports.

   The Monitor Details list displays the monitors that are assigned to a selected port. If no monitors are assigned, or if more than one port is selected, the Monitor Details list does not display.

10. Click **OK**.

    The Frame Monitor Configuration Status dialog box displays (Figure 512).

    FIGURE 512 Frame Monitor Configuration Status dialog box

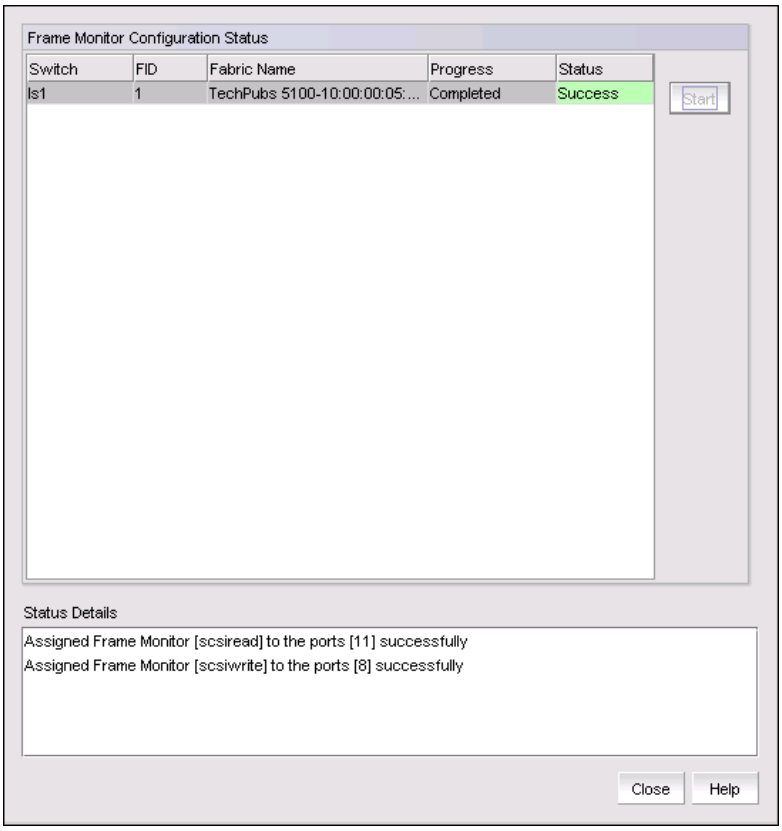

11. Click **Start**.

    The frame monitor configuration is applied to the switches.

12. Click **Close** after configuration is complete (indicated by "Completed" in the Progress column).

Editing a frame monitor

1. Select **Monitor > Fabric Watch > Frame Monitor**.

   The Frame Monitor dialog box displays.

2. Select the **Switch** option.
3. Expand the Products / Monitors list to display the frame monitors for each switch.

4. Select a frame monitor and click the left arrow button.

   The frame monitor is removed from the switch and the Configure Monitor area is populated with the values for that frame monitor.

5. Make changes to the monitor data in the Configure Monitor area.
6. Select one or more switches in the Products / Monitors list, and click the right arrow button to assign the frame monitor to those switches.

   If the frame monitor already exists on the switches, the frame monitor is modified. If the frame monitor does not exist on the switch, it is added.

7. Click **OK**.

   The Frame Monitor Configuration Status dialog box displays.

8. Click **Start**.

   The frame monitor configuration is applied to the switches and ports.

9. Click **Close** after configuration is complete (indicated by "Completed" in the Progress column).

Assigning a frame monitor to a port

1. Select **Monitor > Fabric Watch > Frame Monitor**.

   The Frame Monitor dialog box displays.

2. Select the **Port** option.
3. Expand the switch in the Products / Ports list.

   The Monitors list displays all of the frame monitors defined for that switch.

4. Select one or more ports.

   You must select only ports belonging to the same switch.

5. Select one or more frame monitors in the Monitors list.
6. Click the right arrow button to move the frame monitor to the selected ports.

   The Monitor Details list displays the monitors that are assigned to a selected port. If no monitors are assigned, or if more than one port is selected, the Monitor Details list does not display.

7. Click **OK**.

   The Frame Monitor Configuration Status dialog box displays.

8. Click **Start**.

   The frame monitor configuration is applied to the ports.

9. Click **Close** after configuration is complete (indicated by "Completed" in the Progress column).

Finding frame monitor assignments

Using the following procedure, you can select a frame monitor on a switch and see the ports to which it is assigned.

1. Select **Monitor > Fabric Watch > Frame Monitor**.

   The Frame Monitor dialog box displays.

2. Select the **Port** option.
3. Select a switch in the Products / Ports list.

   The Monitors list displays all of the frame monitors defined for that switch.

4. Select a frame monitor in the Monitors list.
5. Click the **Find** arrow.

   The ports to which the frame monitor is assigned are highlighted.

Removing a frame monitor from a port

1. Select **Monitor > Fabric Watch > Frame Monitor**.

   The Frame Monitor dialog box displays.

2. Select the **Port** option.
3. Expand the switch in the Products / Ports list.

   The Monitors list displays all of the frame monitors defined for that switch.

4. Select the port from which you want to remove the frame monitor.

   The Monitor Details list displays all of the frame monitors assigned to that port.

5. Select one or more frame monitors in the Monitor Details list.
6. Click **Remove**.
7. Click **OK**.

   The Frame Monitor Configuration Status dialog box displays.

8. Click **Start**.

   The frame monitor configuration is applied to the ports.

9. Click **Close** after configuration is complete (indicated by “Completed” in the Progress column).

Removing a frame monitor from a switch

When you remove a frame monitor from a switch, the frame monitor is automatically removed from all assigned ports in the switch.

You can remove only custom frame types; you cannot remove the pre-defined frame types.

1. Select **Monitor > Fabric Watch > Frame Monitor**.

   The Frame Monitor dialog box displays.

2. Select the **Switch** option.

   The Products / Monitors list displays the switches that support Frame Monitoring.

3. Expand the Products / Monitors list to display the frame monitors for each switch.
4. Select a frame monitor and click the left arrow button.

   The frame monitor is removed from the switch and the Configure Monitor area is populated with the values for that frame monitor.

5. Click **OK**.

   The Frame Monitor Configuration Status dialog box displays.

6. Click **Start**.

   The frame monitor configuration is applied to the switches and ports.

7. Click **Close** after configuration is complete (indicated by "Completed" in the Progress column).

Configuration Policy Manager

Configuration policy manager overview

Use the Configuration Policy Manager feature to provide best practice guidelines for network setup at the fabric, switch, port, and device level, as well as software configurations at the Fabric OS and the Management application level.

Configuring policy managers enables you to perform the following:

- Provide selectable and configurable built-in rules to check for best practices
- Schedule policies to run periodically
- Run a policy manually (on demand)
- Generate a report that will detail any issues found by the policy

Fabric configuration policy manager

Fabric configuration policy managers enable you to set the following configuration policy manager on SAN (refer to "Adding a configuration policy manager" on page 1117):

- **Check zoning status —** This fabric configuration policy manager enables you to determine if zoning is enabled or disabled on the fabric.

  Zoning plays a key role in the management of device communication. When you enforce zoning, devices not in the same zone cannot communicate. Zoning provides protection from disruption in the fabric (putting bounds on the scope of RSCNs). The best practice is always to enable zoning.

  Rule Violation Fix — If the configuration policy manager report shows a violation, the Administrator can use the **Zoning** dialog box (**Configure > Zoning > Fabric**) to fix the violation. Refer to "Zoning" on page 779.

  For example, if you use the configuration policy manager to make sure that the zoning status is enabled, you can fix the violation through the **Zoning** dialog box by locating the target fabric, defining a zone configuration, and activating the zone configuration.

- **Check that all zones belong to at least one zone config —** This fabric configuration policy manager enables you to determine if there are any orphaned zones in the fabric zone database.

  Too many orphaned zones can fill up the fabric zone database and complicate other ongoing administrative tasks.

  Rule Violation Fix — If the configuration policy manager report shows a violation, the Administrator can use the **Zoning** dialog box (**Configure > Zoning > Fabric**) to fix the violation. Refer to "Zoning" on page 779.

  For example, the Administrator can fix the violation through the **Zoning** dialog box using one of the following methods:

  - Defining a new zone configuration and moving the orphaned zones to the new zone configuration.
  - Moving the orphaned zones to an existing zone configuration.
  - Cleaning up unused orphaned zones.
- **Check the number of initiator ports zoned to each storage port —** This fabric configuration policy managerconfiguration policy managerenables you to determine the total number of initiator ports zoned to each storage port.

  When too many initiators share the same connection (share the bandwidth of the storage port), congestion can occur.

  There are four possible zone member types: device port WWN, device node WWN, (D,I), and Fabric Assigned WWN.

  - Device port WWN — The application counts the connected device ports and uses them for the ratio calculation.
  - Device node WWN zone member — The application finds the corresponding device ports and uses them for the ratio calculation.
  - D,I — If the switch port is connected to a device, the application finds the connected device ports and uses them for the ratio calculation.
  - Fabric Assigned WWN — If the switch or Access Gateway (AG) port has a connected device port, the application finds the connected device ports and uses them for the ratio calculation.

  Some devices can function as both initiator and target. If the application finds this type of device as one of the active zone members, this device port is treated as both initiator and target:

  - Target (storage port) — The application counts the number of initiator ports zoned to this storage port.
  - Initiator — The application counts this device as an initiator port for other storage ports in the same zone.

  Rule Violation Fix — If the configuration policy manager report shows a violation, the Administrator must make sure the initiator port limit is under the recommended number.
- **Check zones that do not contain any online member —** This fabric configuration policy manager enables you to identify zones in which all zone members are offline.

  NOTE
  The application does not count end devices which are missing from the fabric and D,PI zone members (online or offline) as online zone members. The application only counts zones with online WWN members as online zone members.

  Rule Violation Fix — If the configuration policy manager report shows a violation, you can use the **Zoning** dialog box (**Configure > Zoning > Fabric**) to bring the devices back online (refer to "Zoning" on page 779).

  For example, if you use the configuration policy manager to determine when all WWN members in a zone are offline, you can fix the violation through the **Zoning** dialog box by locating the target fabric and bringing the devices back online.

Switch and router configuration policy managers

Switch and router configuration policy managers enable you to set the following configuration policy managers on switches and routers.

- **Check connections: redundant connections to neighboring switches (SAN only)** — This switch and router configuration policy manager enables you to determine if there are at least the minimum number of configured inter-switch links (ISLs) between each switch pair.

The resiliency and redundancy of the fabric is an important aspect of the SAN topology. To remove any single point of failure, SAN fabrics have resiliency built into the Fabric OS.

For example, when a link between two switches fails, routing is recalculated and traffic is assigned to a new route. Therefore, to provide redundancy and enable resiliency, using ISLs, the best practice is to make sure that there are at least two ISLs between each switch pair.

The redundant link refers to both the physical connection and the logical ISL. No matter how many physical connections exist between the two base switches, there is only one logical ISL between two logical switches. A logical ISL counts as one connection between the source and destination switches; therefore, when a logical ISL is present, the connection count may be inaccurate. To pass this monitor, the total number of logical ISL and physical connections must be greater than the minimum connection.

For FCIP tunnels, one tunnel counts as one connection. This rule does not check circuits within the FCIP tunnel. The total number of trunk ISLs, single ISLs, and the number of tunnels is compared with the minimum number settings to decide if the redundant ISL check is a success or a failure.

Rule Violation Fix — If the configuration policy manager report shows a violation, the SAN Administrator can add redundant ISLs between the source and the target switch.

- **Check for HTTPS (secure HTTP) configuration** — This switch and router configuration policy manager enables you to check each target to see if HTTPS is active for device data transmission.

  The preferred Management application product communication must be HTTPS for this check to pass.

  For Fabric OS products, verifies the IP ACL active policy rules. You should verify that the IP ACL active rules deny HTTP access to all.

  For Fabric OS products, if the IPv6 interface is enabled, verifies both IPv4 and IPv6 IP ACL active policies.

  Rule Violation Fix — If the configuration policy manager report shows a violation, enable HTTPS on the device. Disable HTTP settings on the device, if enabled.

- **Check if the product is configured to send events to this server** — This switch and router configuration policy manager enables you to determine if the Management application server is registered as an SNMP recipient and Syslog recipient.

  If the server has multiple NICs, the server uses an IP address reachable from the switch for event registration. This policy cannot determine if the server is using a reachable IP address for the event registration.

  If the Management application server fails to register as a listener for SNMP, Syslog, and other events, the Management application server cannot notify you of changes to the fabric or device. If a fabric or switch fails, the Management application cannot provide notification, log, or support data. Therefore, you may not realize that there is an inconsistency between the physical device status and the device status in the Management application for some time. This policy cannot determine if the SNMP trap or syslog listener ports are available or working.

  Rule Violation Fix — If the configuration policy manager report shows an "SNMP not registered as recipient" violation, the Administrator can register the Management server as an SNMP recipient through the **SNMP Trap Recipients** dialog box (**Monitor > SNMP Setup > Product Trap Recipients**). Refer to "Fault Management" on page 1131.

  If the configuration policy manager report shows an "Syslog not registered as recipient" violation, the Administrator can register the Management server as an Syslog recipient through the **Syslog Recipients** dialog box (**Monitor > Syslog Configuration > Product Syslog Recipients**). Refer to "Fault Management" on page 1131.

  NOTE
  This check is not supported on front or Xlate domains.

- **Check if the product is configured to send Upload Failure Data Capture to an FTP server (SAN only) —** This switch and router configuration policy manager enables you to determine if Upload Failure Data Capture is enabled on the selected switches, that the configured FTP Server is accessible, and that you have write permission to the directory.

  Upload Failure Data Capture enables you to collect switch data periodically. This assists you to troubleshoot switch failure.

  Rule Violation Fix — If the report shows a violation, the SAN Administrator can change the Upload Failure Data Capture configuration through the **Upload Failure Data Capture** dialog box (**Monitor > Technical Support > Upload Failure Data Capture**). Refer to "Enabling upload failure data capture" on page 1272.

- **Check for SSH (secure Telnet) configuration** — This switch and router configuration policy manager enables you to check each target to see if SSH is enabled for device data transmission.

  The preferred Management application product communication must be SSH for this check to pass.

  For Fabric OS products, verifies SSH access is enabled and telnet access is disabled through the IP ACL active or applied policy rules. You should verify that the IP ACL active rules deny telnet access to all.

  For Fabric OS products, if the IPv6 interface is enabled, verifies both IPv4 and IPv6 through the active IP ACL policy.

  Rule Violation Fix — If the configuration policy manager report shows a violation, enable SSH on the device. Disable Telnet settings on the device, if enabled.

- **Check for SNMPv3 (secure SNMP) configuration** — This switch and router configuration policy manager enables you to check each target to see if SNMPv3 is active for device data transmission and SNMPv1 and SNMPv2 are not configured.

  NOTE
  For this check to pass, you must discover the products using SNMPv3 credentials.

  Rule Violation Fix — If the configuration policy manager report shows a violation, configure SNMPv3 on the device. Remove SNMPv1 and SNMPv2 settings on the device, if configured.

- **Check for MAPS actions enabled-** This switch and router configuration policy manager enables you to determine whether the chosen MAPS actions are enabled on the selected switches.

  Rule Violation Fix — If the configuration policy manager report shows a violation, the SAN Administrator can use the **MAPS Configuration** dialog box (**Monitor > Fabric Vision > MAPS** > **Configure > Actions**) to enable the required MAPS actions.

  NOTE
  For this check to pass, you must enable MAPS in the switches and discover Fabric OS switches running Fabric OS 7.2.0 or later.

- **Check to compare remote SFP metrics (SAN only)** - This switch and router configuration policy manager enables you to determine whether the Tx/Rx power loss between connected ports exceeds the provided Tx/Rx power threshold in percentage.

  Rule Violation Fix — If the configuration policy manager report shows a violation, the SAN Administrator can check and replace the faulty SFPs or links to reduce the power loss.

- **Configuration Rules** — This switch and router configuration policy manager enables you to use predefined rules or create your own rules to compare content against a baseline (such as a product's backup configuration file). A configuration rule is a logical expression built with configuration conditions and blocks. For more information, refer to "Viewing a predefined configuration rule" on page 1924.
  - Predefined rules — The predefined rules include the following:
    - **No Interface Shutdown Rule** — Fails if any of the interfaces in the device are shut down.
    - **Port Profile Interface Rule** — Fails if any of the interfaces on the device do not have a port profile.
  - User-defined rules — You can configure your own configuration rules using predefined conditions and blocks (refer to "Adding a configuration rule" on page 1927).

Host configuration policy managers

Host configuration policy managers enable you to set the following checks on host devices:

- **Check for multiple fabrics connections —** This host configuration policy manager enables you to determine if each host is connected to multiple fabrics to prevent a single point of failure.

  Available hosts include both automatic hosts and manual hosts. Automatic hosts are those hosts discovered through Host or VM Manager discovery. Manual hosts are those host enclosures that are manually created through host port mapping in the fabric topology.

  The Management application determines if the host has redundant connections to different fabrics based on discovery type and connection knowledge that the Management application collects; however, there is no guarantee that redundant paths exist to the same storage target.

  Depending on how you discover the hosts, there are recommended configurations you should complete to avoid inaccuracy:

  - Fabric discovery for manual host enclosures to fabric connections (refer to "Discovery" on page 33)

    Make sure there are supported HBAs (refer to "Supported Fibre Channel HBA models" on page 478) on the host.

    Make sure to configure the host port mapping. (refer to "Host port mapping overview" on page 465)

  - Host adapter discovery with 2.1 or later driver for host to unmanaged fabric connections (refer to "Host discovery" on page 51)

    Make sure there are supported HBAs (refer to "Supported Fibre Channel HBA models" on page 478) on the host.

  - Fabric plus Host adapter discovery with 2.1 or earlier driver (refer to "Host discovery" on page 51)

    Make sure there are supported HBAs (refer to "Supported Fibre Channel HBA models" on page 478) on the host.

  - Fabric plus VM Manager for hosts discovered through vCenter (refer to "VM Manager discovery" on page 63)

    Make sure there are supported HBAs (refer to "Supported Fibre Channel HBA models" on page 478) on the host.

    Make sure you discover the associated fabrics.

  - VM Manager plus Host adapter discovery (refer to "VM Manager discovery" on page 63)

    Make sure there are supported HBAs (refer to "Supported Fibre Channel HBA models" on page 478) on the host.

    Make sure you discover the associated fabrics.

  Rule Violation Fix — If the configuration policy manager report shows a violation, the Administrator can add a host connection to additional fabrics.

- **Check for connections through two fabrics to each target LUN** — This host configuration policy manager enables you to determine if there are redundant connections between the host group and the target LUN.

  To prevent a single point of failure, the host should have a redundant connection to the target LUN. Available hosts include both automatic hosts and manual hosts. An automatic host is a host discovered through Host adapter discovery or VM Manager discovery. A manual host is a host enclosure manually created through host port mapping in the fabric topology.
  Depending on how you discover the hosts, there are recommended configurations you should complete to avoid inaccuracy:

  - Host adapter discovery (refer to "Host discovery" on page 51)

    Make sure there are supported HBAs (refer to "Supported Fibre Channel HBA models" on page 478) on the host.

  - Fabric plus Host discovery (refer to "Discovery" on page 33)

    Make sure there are supported HBAs (refer to "Supported Fibre Channel HBA models" on page 478) on the host connected to the fabric.

Make sure to configure the host port mapping (refer to "Host port mapping overview" on page 465).

- Fabric plus VM Manager discovery (refer to "Discovery" on page 33)

  Make sure there are supported HBAs (refer to "Supported Fibre Channel HBA models" on page 478) on the host connected to the fabric.

- VM Manager plus Host discovery (refer to "VM Manager discovery" on page 63)

  Make sure there are supported HBAs (refer to "Supported Fibre Channel HBA models" on page 478) on the host.

  Make sure you discover the associated fabrics.

Rule Violation Fix — If the configuration policy manager report shows a violation, the Administrator can add redundant connections (either a host to attached fabrics or attached fabrics to a target LUN or more inter-fabric routes) to establish a complete path from host to target LUN.

Management configuration policy manager

The management configuration policy manager enables you to set a configuration policy manager on the Management application.

Check to see if the server backup is enabled and working — This management configuration policy manager enables you to determine if backup is enabled for the Management application server and if the backup output directory is accessible and writable.

Server backup automatically backs up the Management application database on a user-defined schedule.

Rule Violation Fix — If the configuration policy manager report shows a violation, the Administrator can edit the backup configuration through the **Options** dialog box, **Server Backup** pane (**Server > Options**). Refer to "Management server backup" on page 71.

Preconfigured configuration policy managers

The Management application provides preconfigured configuration policy managers. The preconfigured configuration policy managers include the following:

Default SAN Policy — Available for SAN products and contains the following values:

- **Name** — Default SAN Policy
- **Description** — Default policy to run on all SAN targets
- **Frequency** — Daily
- **Next Run** — Next time the policy will run using the format: *<Day_of_Week><Month><Date><Time_in_24_Hour_Format><Time_Zone><Year>*. For example, Fri Jun 08 08:00:00 PST 2014.
- **Last Run** — Empty
- **Result** — Empty
- **Rule** — The default SAN policy is configured with the following rules:
  - Zoning status
  - Fan in Ratio
  - Event registration
  - Redundant connection
  - Management application backup enabled

- **Targets** — The default SAN policy is configured with the following targets:
  - Fabric Checks — All Fabrics
  - Switch/Router Checks — All SAN Switches product group

Viewing configuration policy manager status

You can view configuration policy manager status from the main Management application window or from the **Configuration Policy Manager** dialog box.

The Management application enables you to view the configuration policy manager status at a glance by providing a configuration policy manager status icon on the status bar. Table 101 illustrates and describes the icons that indicate the current status of the configuration policy manager function.

TABLE 101 Configuration policy manager icons

Icon	Description
	Passed — Displays when all configuration policy managers, excluding un-alerted and acknowledged monitors, pass. Pause on icon to display flyover detail: Configuration policy manager is OK.
	Failed — Displays when at least one configuration policy manager failed. Pause on icon to display flyover detail: The last run of *number* configuration policy manager(s) has one or more failures.

To view more detail regarding configuration policy manager status, click the **Configuration Policy Manager** icon. The **Configuration Policy Manager** dialog box displays. For more information, refer to "Viewing existing configuration policy managers" on page 1116.

Viewing existing configuration policy managers

To view existing configuration policy managers, complete the following steps.

1. Select **Monitor > Configuration Policy Manager** (Figure 513).

   The **Configuration Policy Manager** dialog box displays.

   FIGURE 513 Configuration Policy Manager dialog box

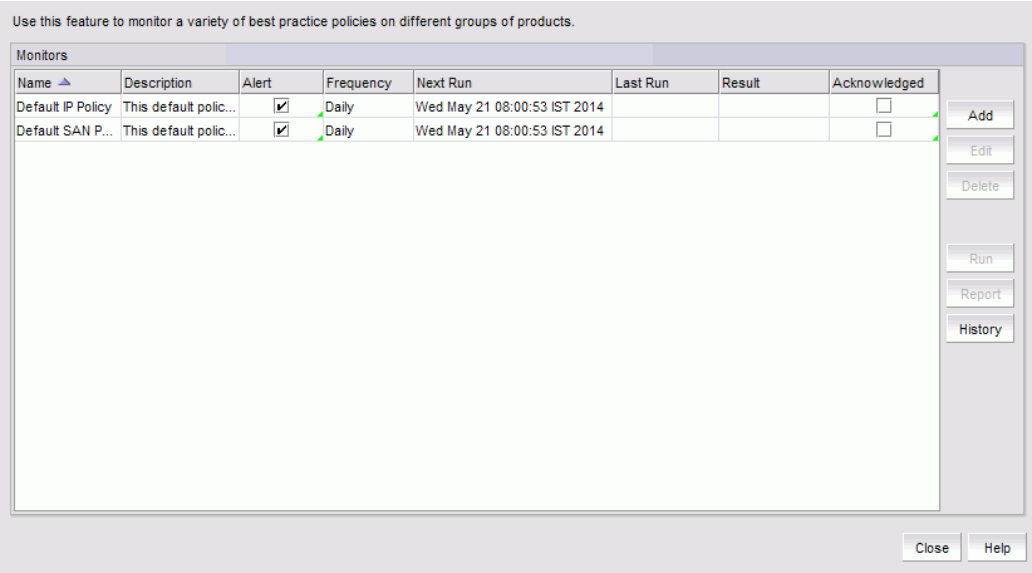

2. Review the configuration policy manager details:

   - **Name** — The user-defined name of the policy.
   - **Description** — A description of the policy.
   - **Alert** — Select to receive e-mail alerts and have the configuration policy manager status icon display in the status bar when the monitor fails or partially fails.
   - **Frequency** — The frequency (one time, hourly, daily, weekly, or monthly) at which the policy is scheduled.
   - **Next Run** — The time the policy will run again.
   - **Last Run** — The time the policy ran last.
   - **Result** — The result of the last configuration policy manager run. There are four possible results: Passed, Partially Failed, Failed, and Not Applicable.
   - **Acknowledged** — Whether the policy is acknowledged or not. Select the check box to acknowledge the policy. Disabled when the associated **Acknowledged** check box is cleared.

3. To add a configuration policy manager, click **Add** (refer to "Adding a configuration policy manager" on page 1117).
4. To edit the selected configuration policy manager, click **Edit** (refer to "Editing a configuration policy manager" on page 1123).
5. To delete the selected configuration policy manager, click **Delete** (refer to "Deleting a configuration policy manager" on page 1124).
6. To run the selected policy and view the report, click **Run** (refer to "Running a configuration policy manager" on page 1124).
7. To open the last executed report for a selected configuration policy manager, select a configuration policy manager and click **Report** (refer to "Viewing a configuration policy manager report" on page 1125).

8. To view the report history for all configuration policy managers, click **History** (refer to "Viewing historical reports for a configuration policy manager" on page 1129).
9. To view the report history for a selected configuration policy manager, select a configuration policy manager and click **History** (refer to "Viewing historical reports for a configuration policy manager" on page 1129).
10. Click **Close** on the **Configuration Policy Manager** dialog box.

Adding a configuration policy manager

To add a configuration policy manager, complete the following steps.

1. Select **Monitor > Configuration Policy Manager**.

   The **Configuration Policy Manager** dialog box displays.

2. Click **Add**.

   The **Add Configuration Policy Manager** dialog box displays (Figure 514).

   FIGURE 514 Add Configuration Policy dialog box, Fabric Checks tab

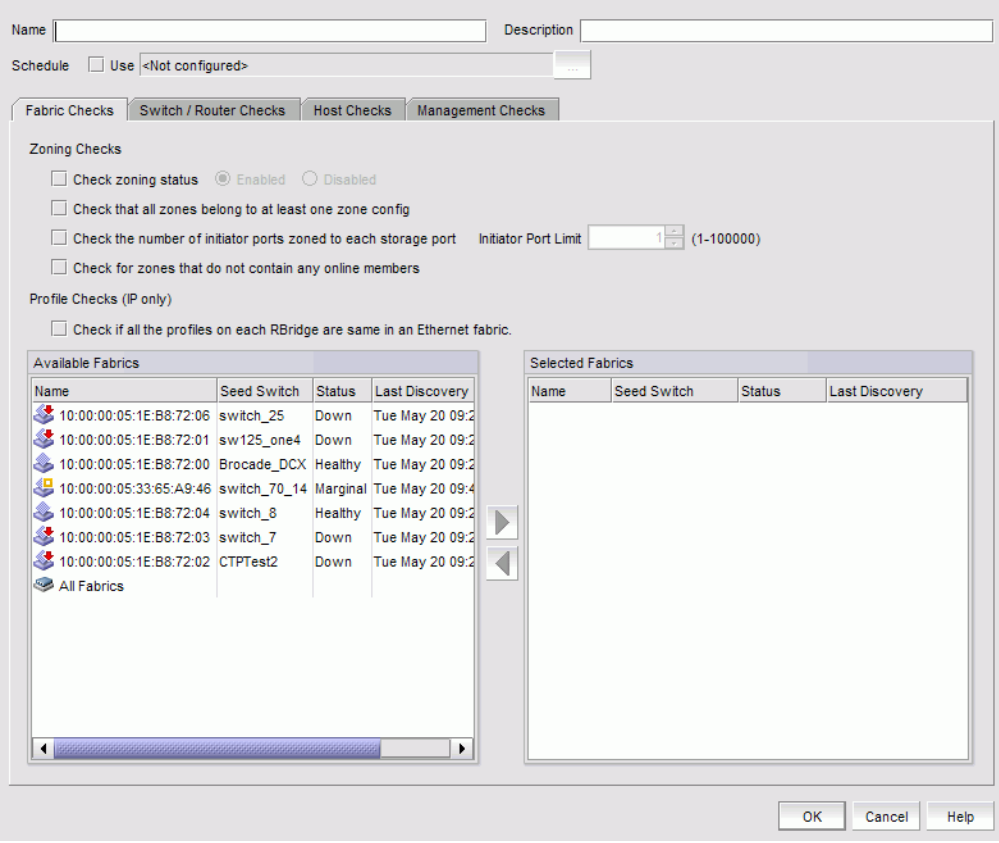

3. Enter a user-defined name for the policy in the **Name** field.

   The name must be unique. It cannot be over 64 characters, nor can the field be empty. It cannot include asterisks.

4. Enter a description of the policy in the **Description** field.

   The description cannot be over 128 characters. It cannot include asterisks.

5. Click the **Schedule Use** check box and choose one of the following options:

   - To use the default frequency (one time, runs at current system time plus fifteen minutes), go to step 6.
   - To configure the frequency, click the ellipsis button and choose one of the following options to configure the frequency at which deployment runs for the configuration policy manager:
     - To configure deployment to run only once, refer to "Configuring a one-time configuration policy manager schedule" on page 1121.
     - To configure hourly deployment, refer to "Configuring an hourly configuration policy manager schedule" on page 1122.
     - To configure daily deployment, refer to "Configuring a daily configuration policy manager schedule" on page 1122.
     - To configure weekly deployment, refer to "Configuring a weekly configuration policy manager schedule" on page 1122.
     - To configure monthly deployment, refer to "Configuring a monthly configuration policy manager schedule" on page 1123.

6. To set configuration policy managers for fabrics, select the **Fabric Checks** tab and complete the following steps.

   a. Select the **Check zoning status** check box to determine if zoning is enabled or disabled on the fabric.

      - Select the **Enabled** option to determine if zoning is enabled.
      - Select the **Disabled** option to determine if zoning is disabled.

      For more information about this check and a fix for rule violations, refer to "Fabric configuration policy manager" on page 1109.

   b. Select the **Check that all zones belong to at least one zone config** check box to determine if there are orphaned zones in the fabric zone database.

      For more information about this check and a fix for rule violations, refer to "Fabric configuration policy manager" on page 1109.

   c. Select the **Check the number of initiator ports zoned to each storage port** check box to determine the total number of initiator ports zoned to each storage port.

      For more information about this check and a fix for rule violations, refer to "Fabric configuration policy manager" on page 1109.

   d. Select the **Check zones that do not contain any online member** check box to identify zones in which all zone members are offline.

      For more information about this check and a fix for rule violations, refer to "Fabric configuration policy manager" on page 1109.

   e. Enter the initiator port limit in the **Initiator Port Limit** field.

      The default recommended threshold ratio is 1:1 (one initiator port to one target port). Therefore, if the ratio for the storage port is equal to or higher than 1:1, the configuration policy manager considers it as a violation and logs it in the report.

   f. Select the fabrics to which you want to apply this policy in the **Available Fabrics** list and click the right arrow button.

      NOTE
      You can use the All Fabrics target in the **Available Fabrics** list for future provisioning. Select **All Fabrics** and click the right arrow button to apply this policy to all discovered fabrics.

      The selected fabrics display in the **Selected Fabrics** list.

7. To set configuration policy managers for switches, select the **Switch/Router Checks** tab (Figure 515) and complete the following steps.

FIGURE 515 Add Configuration Policy Manager dialog box, Switch/Router Checks tab

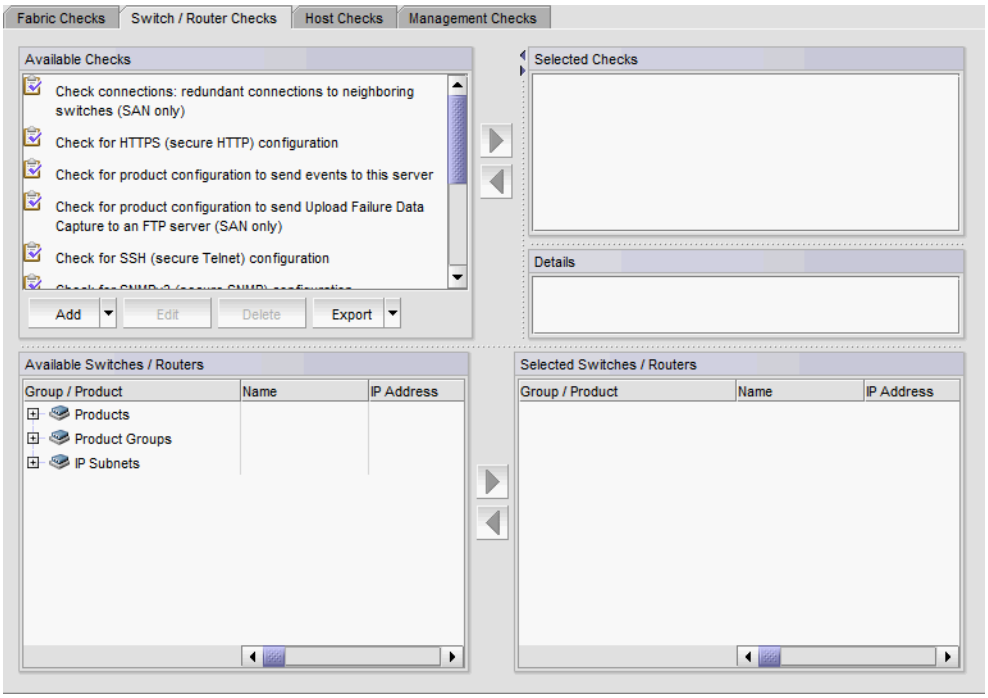

a. Select one or more of the following checks in the **Available Checks** list to include them in the configuration policy manager.

   For more information about these checks and fixes for rule violations, refer to "Switch and router configuration policy managers" on page 1110.

   - Select the **Check if the product is configured to send events to this server** to determine if the Management application server is registered as an SNMP recipient and syslog recipient.
   - Select the **Check for redundant connections to neighboring switches (SAN only)** check box to determine if there are at least the minimum number of configured ISLs between each switch pair.
   - Select the **Check for HTTPS (secure HTTP) configuration** check box to check each target to see if HTTPS is active for device data transmission.
   - Select the **Check if the product is configured to send Upload Failure Data Capture to an FTP server (SAN only)** check box to determine the following configurations:
     – Upload Failure Data Capture is enabled on the selected switches.
     – A configured FTP Server is accessible.
     – You have write permission to the directory.
   - Select the **Check for SSH (secure Telnet) configuration** check box to check each target to see if SSH is enabled for device data transmission.
   - Select the **Check for SNMPv3 (secure SNMP) configuration** check box to check each target to see if SNMPv3 is active for device data transmission and SNMPv1 and SNMPv2 are not configured.
   - Select the **Check for MAPS action enabled (SAN only)** check box to determine whether the MAPS actions are enabled over the selected switches.
   - Select the **Check to compare remote SFP metrics (SAN only)** check box to report the deviations found by comparing the Tx and Rx power values of E_Ports and F_Ports in a switch with its corresponding remote ports.

b. Click the right arrow button to move the selected checks to the **Selected Checks** list.

c. If you selected the **Check for redundant connections to neighboring switches (SAN only)** check box, enter the minimum number of connections allowed between a switch pair in the **Minimum Connections** field.

   The recommended default is 2. Valid values are from 2 through 512.

d. If you selected the **Check for MAPS action parameter dialog** check box, you must select the necessary actions that need to be verified in the **MAPs action parameter** dialog box.

   NOTE
   You can use the All SAN Switches targets (under the Product Groups > System Product Groups node) in the **Available Switches/Routers** list) for future provisioning. Select **All SAN Switches** and click the right arrow button to apply this policy to all discovered switches.

   The selected switches display in the **Selected Switches/Routers** list.

8. To set configuration policy managers for hosts, select the **Host Checks** tab (Figure 516) and complete the following steps.

FIGURE 516 Add Configuration Policy Manager dialog box, Host Checks tab

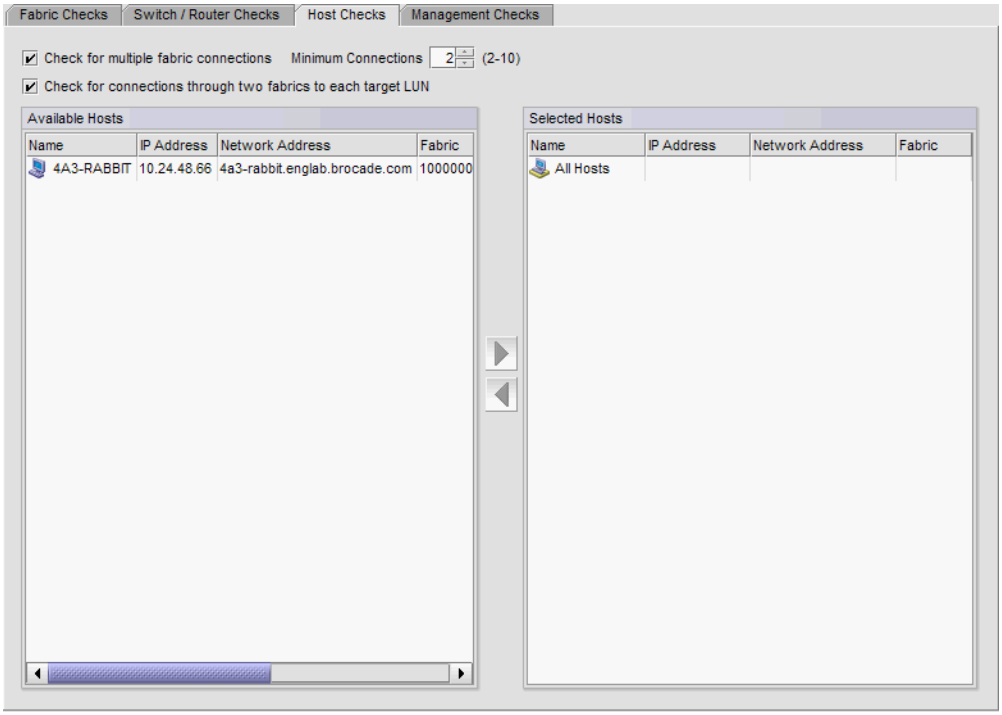

a. Select the **Check for redundant connections to attached fabrics** check box to determine if there are at least the minimum number of configured physical connections between the host and the attached fabric.

   The default is 2. For more information about this check and a fix for rule violations, refer to "Host configuration policy managers" on page 1113.

b. Enter the minimum number of connections between the host and the attached fabric in the **Minimum Connections** field.

   The default is 2.

c. Select the **Check for connections through two fabrics to each target LUN** check box to determine if there are redundant connections between the host group and the target LUN.

For more information about this check and a fix for rule violations, refer to "Host configuration policy managers" on page 1113.

d. Select the hosts to which you want to apply this policy in the **Available Hosts** list and click the right arrow button.

NOTE
You can use the All Hosts target in the **Available Hosts** list for future provisioning. Select **All Hosts** and click the right arrow button to apply this policy to all discovered hosts.

The selected hosts display in the **Selected Hosts** list.

9. To set configuration policy managers for the Management application (Figure 517), complete the following steps.

a. Select the **Management Checks** tab.

FIGURE 517 Add Configuration Policy Manager dialog box, Management Checks tab

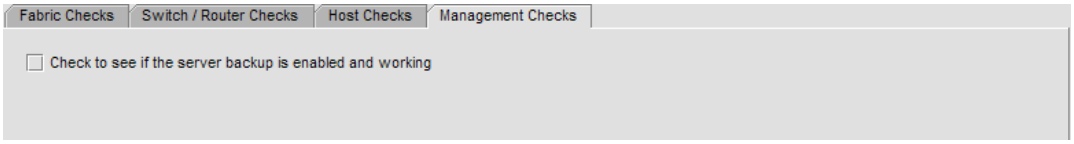

b. Select the **Check to see if the server backup is enabled and working** check box to determine the following configurations:

- Backup enabled for the Management application server.
- Backup output directory is accessible and writable.

This policy only applies to scheduled backup, not manual (on-demand) backup.

For more information about this check and a fix for rule violations, refer to "Management configuration policy manager" on page 1114.

10. Click **OK** on the **Add Monitor** dialog box.

The **Configuration Policy Manager** dialog box displays with the new configuration policy manager in the **Monitors** list.

11. Click **Close** on the **Configuration Policy Manager** dialog box.

Configuration policy manager scheduling

You can schedule a configuration policy manager to run automatically. For step-by-step instructions, refer to the following procedures:

- "Configuring a one-time configuration policy manager schedule" on page 1121
- "Configuring an hourly configuration policy manager schedule" on page 1122
- "Configuring a daily configuration policy manager schedule" on page 1122
- "Configuring a weekly configuration policy manager schedule" on page 1122
- "Configuring a monthly configuration policy manager schedule" on page 1123

Configuring a one-time configuration policy manager schedule

To configure a one-time schedule, complete the following steps.

1. Select **One Time** from the **Frequency** list.
2. Select the time of day you want deployment to run from the **Time (hh:mm)** lists.

   Where the hour value is from 1 through 12, the minute value is from 00 through 59, and the day or night value is AM or PM.
3. Click the **Date** list to select a date from the calendar.
4. Click **OK** on the **Schedule Properties** dialog box.

   To finish configuring the configuration policy manager, return to step 6 of "Adding a configuration policy manager" on page 1117.

Configuring an hourly configuration policy manager schedule

To configure an hourly schedule, complete the following steps.

1. Select **Hourly** from the **Frequency** list.
2. Select the minute past the hour you want deployment to run from the **Minutes past the hour** list.

   Where the minute value is from 00 through 59.
3. Click **OK** on the **Schedule Properties** dialog box.

   To finish configuring the configuration policy manager, return to step 6 of "Adding a configuration policy manager" on page 1117.

Configuring a daily configuration policy manager schedule

To configure a daily deployment schedule, complete the following steps.

1. Select **Daily** from the **Frequency** list.
2. Select the time of day you want deployment to run from the **Time (hh:mm)** lists.

   Where the hour value is from 1 through 12, the minute value is from 00 through 59, and the day or night value is AM or PM.
3. Click **OK** on the **Schedule Properties** dialog box.

   To finish configuring the configuration policy manager, return to step 6 of "Adding a configuration policy manager" on page 1117.

Configuring a weekly configuration policy manager schedule

To configure a weekly schedule, complete the following steps.

1. Select **Weekly** from the **Frequency** list.
2. Select the time of day you want deployment to run from the **Time (hh:mm)** lists.

   Where the hour value is from 1 through 12, the minute value is from 00 through 59, and the day or night value is AM or PM.
3. Select the day you want deployment to run from the **Day of the Week** list.
4. Click **OK** on the **Schedule Properties** dialog box.

To finish configuring the configuration policy manager, return to step 6 of "Adding a configuration policy manager" on page 1117.

Configuring a monthly configuration policy manager schedule

To configure a monthly schedule, complete the following steps.

1. Select **Monthly** from the **Frequency** list.
2. Select the time of day you want deployment to run from the **Time (hh:mm)** lists.

   Where the hour value is from 1 through 12, the minute value is from 00 through 59, and the day or night value is AM or PM.
3. Select the day you want deployment to run from the **Day of the Month** list (1 through 31).
4. Click **OK** on the **Schedule Properties** dialog box.

   To finish configuring the configuration policy manager, return to step 6 of "Adding a configuration policy manager" on page 1117.

Editing a configuration policy manager

To edit an existing configuration policy manager, complete the following steps.

1. Select **Monitor > Configuration Policy Manager**.

   The **Configuration Policy Manager** dialog box displays.
2. Select the policy you want to edit in the **Monitors** list and click **Edit**.

   The **Edit Configuration Policy Manager** dialog box displays. The **Edit Configuration Policy Manager** dialog box has the same fields and components as the **Add Configuration Policy Manager** dialog box.
3. Change the user-defined name for the policy in the **Name** field.

   The name must be unique. It cannot be over 64 characters, nor can the field be empty. It cannot include asterisks.
4. Change the description of the policy in the **Description** field.

   The description cannot be over 128 characters. It cannot include asterisks.
5. To edit the configuration policy manager checks, repeat step 5 through step 9 of "Adding a configuration policy manager" on page 1117.
6. Click **OK** on the **Edit Monitor** dialog box.

   The updated configuration policy manager displays in the **Monitors** list of the **Configuration Policy Manager** dialog box.
7. Click **Close** on the **Configuration Policy Manager** dialog box.

Deleting a configuration policy manager

To delete an existing configuration policy manager, complete the following steps.

1. Select **Monitor > Configuration Policy Manager**.

   The **Configuration Policy Manager** dialog box displays.

2. Select the policy you want to delete in the **Monitors** list.
3. Click **Delete**.
4. Click **Yes** on the confirmation message.
5. Click **Close** on the **Configuration Policy Manager** dialog box.

Running a configuration policy manager

Before you run a configuration policy manager, make sure your configuration policy managers are valid. Valid configuration policy managers must have at least one policy selected with one or more targets. Management checks do not require a target.

To run an existing configuration policy manager, complete the following steps.

1. Select **Monitor > Configuration Policy Manager**.

   The **Configuration Policy Manager** dialog box displays.

2. Select the policy you want to run in the **Monitors** list.
3. Click **Run**.

When the configuration policy manager check is complete, the *Policy_Name* - **Configuration Policy Manager Report** displays (Figure 518) in a web browser.

FIGURE 518 *Policy_Name* - Configuration Policy Manager Report

Export | Email

Status Summary 30 Failed / 9 Passed / 6 Not Applicable **Trigger: Manual** **Run Time: Tue Oct 30 2012 15:19:41 PDT**

Management

Name	Status
Check to see if the server backup is enabled and working	Failed No write permission on the directory D:/Backup

Fabric Checks - Check zoning is Enabled

Name	Status
10:00:00:05:33:5B:8E:A8	Failed Zoning is disabled in the fabric.
10:00:00:05:1E:53:6B:69	Passed
10:00:00:05:1E:53:89:CF	Passed
10:00:00:05:33:5B:8E:A6	Passed
10:00:00:05:33:5B:8E:A7	Failed Zoning is disabled in the fabric.
10:00:00:05:33:52:A0:A0 [10.24.60.56]	Failed Zoning is disabled in the fabric.
10:00:00:05:1E:0A:73:0D	Passed
10:00:00:05:1E:53:8A:1A	Passed

SAN Switch - Check for at least 2 connections to neighboring switches

Name	Status
IBM.5100.45.251 (10.24.45.251)	Not Applicable The switch does not have any neighboring switches.

4. Review the report details (refer to "Viewing a configuration policy manager report" on page 1125).

   To export a report, refer to "Exporting a configuration policy manager report" on page 1128.

   To e-mail a report, refer to "Exporting IP reports to e-mail recipients" on page 2152.

5. Click the close button (**X**) on the *Policy_Name* - **Configuration Policy Manager Report** browser window.
6. Click **Close** on the **Configuration Policy Manager** dialog box.

Viewing a configuration policy manager report

NOTE
You must run the configuration policy manager at least once before you can view a report.

To view an existing configuration policy manager report, complete the following steps.

1. Select **Monitor > Configuration Policy Manager**.

   The **Configuration Policy Manager** dialog box displays.

2. Select the policy for which you want to view a report in the **Monitors** list.
3. Click **Report**.

NOTE
If you have run this policy more than once, the latest report displays.

The *Policy_Name* - **Configuration Policy Manager Report** displays (Figure 518) in a web browser.

4. Review the report details:

- **Name** — Name of the configuration policy manager report.
- **Date** — Date and time the report was finished.
- **Export** button — To export a report, refer to "Exporting a configuration policy manager report" on page 1128.
- **E-Mail** button — To e-mail a report, refer to "Exporting IP reports to e-mail recipients" on page 2152.
- **Status Summary —** Number of checks that passed, partially failed, failed, not applicable, or unknown.

  When a policy status fails or partially fails, the status is highlighted in pink.
- **Trigger** — Trigger for the report. Valid results include Manual, Event Action, and Scheduled.
- **Run Time** — Date and time the report was triggered.
- *Individual_Policy_Checks* — Name of the policy check and a table displaying the results of the check. The following information is included in the report data for each policy check:

  Management Check — Displays the status of the management check. The management check provides the following information:

  - **Name** — Name of the management check.
  - **Status** — Result of the check and reason for failure if known. Valid results include Passed, Partially Failed, Failed, Not Applicable, and Unknown.

  Fabric Checks — Fabric checks provide the following information for each selected check:

  - **Name** — Fabric name.
  - **Status** — Result of the check and reason for failure if known. Valid results include Passed, Partially Failed, Failed, Not Applicable, and Unknown.

  Fabric checks include the following options:

  - Check zoning is enabled
  - Check that all zones belong to at least one zone configuration
  - Check the number of initiator ports zoned to each storage port is less than *Configured_Value*. This check provides the following additional detail for this check:
    - **Storage Port** — WWN of the storage port.
    - **Initiator Count** — Number of initiator ports zoned to the storage port.
    - **Initiator Port** — WWN of the initiator port.
    - **Zone** — Zone name containing the initiator/storage port zoning pair.
  - Check zones that do not contain any online member. This check lists the zones that contain only offline members.

  Switch Checks — Switch checks provide the following information for each selected check:

  - **Name** — Product name.
  - **Status** — Result of the check and reason for failure if known. Valid results include Passed, Partially Failed, Failed, Not Applicable, and Unknown.

Switch Checks include the following options:

- Switch — Check for at least *Configured_Minimum_Value* connections to neighboring switches. This check provides the following additional detail for this check:
  - **Neighboring Switch** — Name of the neighboring switch.
  - **Connection Count** — Number of connections to the neighboring switch.
- Switch — Check for HTTPS (secure HTTP) configuration. This check provides the following additional detail for this check:
  - **HTTPs Status —** Whether HTTPS is enabled or disabled on the product.
  - **HTTP Status —** Whether HTTP is enabled or disabled on the product.
- Switch — Check if the product is configured to send events to this server.
- Switch — Check if the product is configured to send Upload Failure Data Capture information to an FTP server.
- Switch - Check for SSH (secure Telnet) configuration. This check provides the following additional detail for this check:
  - **SSH Status —** Whether SSH is enabled or disabled on the product.
  - **Telnet Status —** Whether Telnet is enabled or disabled on the product.
- Switch - Check for SNMPv3 (secure SNMP) configuration. This check provides the following additional detail for this check:
  - **SNMPv3 Status —** Whether SNMPv3 is enabled or disabled on the product.
  - **SNMP Status —** Whether SNMP is enabled or disabled on the product.
- Switch - Check for MAPS actions enabled (SAN only). This check provides the following details:
  - **RAS Log Event** — Whether the RAS log event action is enabled or disabled on the switch.
  - **Port Decommission** — Whether the Port Decommission event action is enabled or disabled on the switch
  - **SNMP Trap** — Whether the SNMP Trap event action is enabled or disabled on the switch.
  - **Fence** — Whether the Fence event action is enabled or disabled on the switch.
  - **E-mail** — Whether the e-mail event action is enabled or disabled on the switch.
  - **SFP Status Marginal** — Whether the SFP Status Marginal event action is enabled or disabled on the switch.
  - **Switch Status Marginal** — Whether the Switch Status Marginal event action is enabled or disabled on the switch.
  - **Switch Status Critical** — Whether the Switch Status Critical event action is enabled or disabled on the switch.
- Configuration Rule Checks — Switch checks provide the following information for each selected check:
  - **Block/Condition Name —** Name of the block or condition.
  - **Matched Block —** Name of the matched block.
  - **Status** — Whether the configurations matched (Passed) or did not match (Failed).
  - **Failed Condition —** Name of the failed condition.
  - **Match/Not Match —** Whether the configurations matched (Match) or did not match (Not Match).
  - **Condition Details —** Details about the condition.
  - **Remediation —** Details how to correct the failure, if the condition fails.

Host Checks — Host checks provide the following information for each selected check:

- **Name** — Product name.
- **Status** — Result of the check and reason for failure if known. Valid results include Passed, Partially Failed, Failed, Not Applicable, and Unknown.

Displays the Host name and status of the policy check for the following option:

- Host — Check for at least *Configured_Minimum_Value* connections to attached fabrics
- Host — Check for connections through two fabrics to each target LUN. This check provides the following additional detail for this check:
  - **LUN Serial #** — LUN serial number.
  - **Adaptor Port** — Host adapter port number.
  - **Fabric** — Fabric name.
  - **Storage Port** — Storage port number.

5. Click the close button (X) on the *Policy_Name* - **Configuration Policy Manager Report** browser window.
6. Click **Close** on the **Configuration Policy Manager** dialog box.

Exporting a configuration policy manager report

1. Click **Export**.

   The **File Download** dialog box displays.

2. Click **Save**.

   The **Save** dialog box displays.

3. Browse to the file location where you want to save the report and click **Save**.
4. Click the close button (X) on the *Policy_Name* - **Configuration Policy Manager Report** browser window.

Viewing historical reports for all configuration policy managers

1. Select **Monitor > Configuration Policy Manager**.

   The **Configuration Policy Manager** dialog box displays.

2. Click **History.**

   The **Report History** dialog box displays the last 10 reports run for all monitors. The **Report History** dialog box retains up to 10 reports for each configuration policy manager.

   - **Name** — Name of the configuration policy manager.
   - **Date** — Date and time the report was finished.
   - **Result —** Result of the configuration policy manager run. Valid results include Passed, Partially Failed, Failed, Not Applicable, and Unknown.

3. Select the report you want to view and click **Display**.

   The *Policy_Name* - **Configuration Policy Manager Report** displays in a web browser. For detailed information about reports, refer to "Viewing a configuration policy manager report" on page 1125.

4. Click the close button (X) on the *Policy_Name* - **Configuration Policy Manager Report** browser window.
5. Click **Close** on the **Report History** dialog box.

Viewing historical reports for a configuration policy manager

1. Select **Monitor > Configuration Policy Manager**.

   The **Configuration Policy Manager** dialog box displays.

2. Select the policy for which you want to view the report history and click **History.**

   The **Report History** dialog box displays. The **Report History** dialog box displays up to 10 reports for the selected configuration policy manager.

   - **Name** — Name of the configuration policy manager.
   - **Date** — Date and time the report was finished.
   - **Result —** Result of the configuration policy manager run. Valid results include Passed, Partially Failed, Failed, Not Applicable, and Unknown.

3. Select the report you want to view and click **Display**.

   The *Policy_Name* **- Configuration Policy Manager Report** displays in a web browser. For detailed information about reports, refer to "Viewing a configuration policy manager report" on page 1125.

4. Click the close button (X) on the *Policy_Name* **- Configuration Policy Manager Report** browser window.

5. Click **Close** on the **Report History** dialog box.

Fault Management

Fault management overview

Fault management enables you to monitor your managed SAN and IP networks using the following methods:

- Listen, forward, and process SNMP traps for SAN and IP devices, which eliminates the need to poll devices for events.
- Receive and forward Syslog messages from Fabric OS switches, IP devices, and Brocade adapters — HBAs and CNAs are managed using the Host Connectivity Manager (HCM) Agent.
- Manage pseudo events.
- Configure the following event actions:
  - Logging policy
  - E-mail alerts
  - Scripts
  - Broadcast to clients
  - Special events handling
  - Run supportSave (SAN only)
- Monitor audit logs and event logs for specified conditions.
- Support application events.

Restrictions

The following items affect Fault Management operation.

Supported IP address types

The Management application receives traps and syslog messages for physical IP addresses only.

Event Purging

The default maximum number of days that historical events are stored is 365. You can select a different default (from 1 to 365) in the **Options** dialog box under **Event Storage**.

Event Archiving

The default number of days that purged events are archived is 30. This value cannot be changed.

Event notification

The Management application records the SAN and IP events in the Master Log. You can configure the application to send event notifications to e-mail addresses at certain time intervals. This is a convenient way to keep track of events that occur on the SAN and IP networks. You can also configure products to "call home" for certain events, notifying the service center of product problems. For instructions about configuring call home for events, refer to "Call Home" on page 327.

Configuring e-mail notification

To send e-mail notification of events to users, complete the following steps.

1. Select **Monitor > Event Notification > E-mail**.

   The **E-mail Event Notification Setup** dialog box (shown in Figure 519) displays.

   FIGURE 519 E-mail Event Notification Setup dialog box

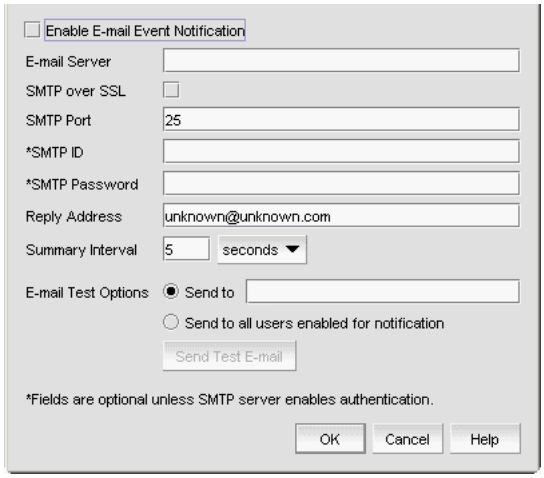

2. Select the **Enable E-mail Event Notification** check box to enable the application to send e-mail messages in case of event notifications.
3. Enter the IP address or the name of the SMTP mail server that the server can use to send the e-mail notifications in the **E-mail Server** field.

   The Management application accepts IP addresses in IPv4 and IPv6 formats. The IPv4 format is valid when the operating system has IPv4 mode only or dual stack mode. The IPv6 format is valid when the operating system has IPv6 mode only or dual stack mode.

4. Select the **SMTP over SSL** check box to enable secure communication.
5. Enter the port number of the SMTP mail server in the **SMTP Port** field.

   If SMTP over SSL is not enabled, the default is 25.

   If SMTP over SSL is enabled, the default is 465.

6. Enter the authentication ID of the SMTP mail server in the **SMTP ID** field.

NOTE
The **SMTP ID** field is optional unless the SMTP server enables authentication.

7. Enter the authentication password of the SMTP mail server in the **SMTP Password** field.

   NOTE
   The **SMTP Password** field is optional unless the SMTP server enables authentication.

8. Enter the sender's e-mail address in the **Reply Address** field.
9. Enter the length of time the application should wait between notifications in the **Summary Interval** field and list.

   Notifications are combined into a single e-mail message and sent at each interval setting. An interval setting of zero causes notifications to be sent immediately.

   ATTENTION
   Setting too short an interval can cause the recipient's e-mail inbox to fill ***very*** quickly.

10. Select one of the following e-mail test options:
    - Select **Send to** and enter an e-mail address for a user to send a test e-mail message to a specific user.
    - Select **Send to all users enabled for notification** to send a test e-mail message to all users already set to receive notification.
11. Click **Send Test E-mail** to test the e-mail server.

    A message displays whether the server was found. If the server was not found, verify that the server address was entered correctly and that the server is running. If you are using an SMTP mail server, also verify that the SMTP ID and password information was entered correctly.

12. Click **OK** to save your work and close the **E-mail Event Notification Setup** dialog box.

Defining filters

The **Define Filter** dialog box, shown in Figure 520, allows you to define event filters by product, event category, and severity. You can define event filters on SAN products, IP products, or hosts.

Setting up basic event filtering

To set up basic event filtering on the selected events for a user, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select a user in the **Users** list and click **Edit**.

   The **Edit User** dialog box displays.

3. Select the **E-mail Notification Enable** check box and click the **Filter** link.

   The **Define Filter** dialog box, shown in Figure 520, displays.

   FIGURE 520 Define Filter dialog box

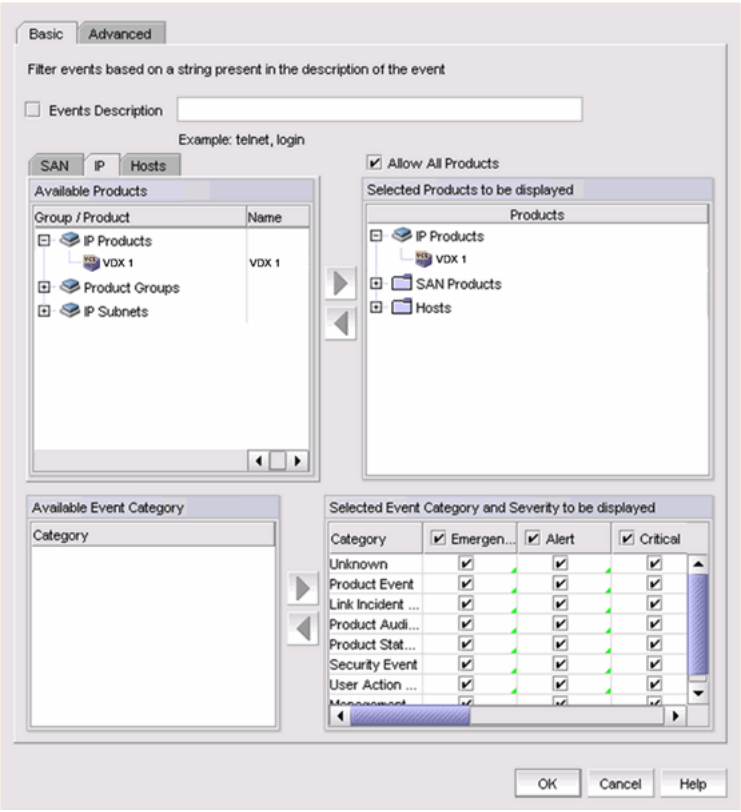

4. Select which product type you are defining (SAN, IP, or Hosts) and click the appropriate tab.
5. Select the **Events Description** check box and enter a description of the event in the field.
6. Select the **Allow All Products** check box to control whether or not all products are always displayed.
   - When selected (the default), all products, even newly-added products, are added to the **Selected Products to be displayed** list.
   - If the check box is cleared, only the products listed in the **Selected Products to be displayed** list are shown in the Master Log and all newly-added products are added to the **Available Products** list. You can include or exclude individual VCS cluster and the node members while defining Master Log filter.
7. Select one or more event categories from the **Available Event Category** list and click the right arrow button to move the event categories to the **Selected Event Category and Severity to be displayed** list. You can move any or all event categories.
8. Select at least one severity for each event. Severity options include Emergency, Alert, Critical, Error, Warning, Notice, Debug, Info, and Unknown.

NOTE
If you delete event actions that are part of the filtering criteria, they will not display in the Master Log, which displays in the lower left area of the main window, and lists all events and alerts that have occurred on the managed networks.

Setting up advanced event filtering

To set up advanced event filtering on the selected events for a user, complete the following steps.

1. Select **Server > Users**.

   The **Users** dialog box displays.

2. Select a user in the **Users** list and click **Edit**.

   The **Edit User** dialog box displays.

3. Select the **E-mail Notification Enable** check box and click the **Filter** link.

   The **Define Filter** dialog box displays.

4. Click **Advanced**.

   The **Advanced** tab of the **Define Filter** dialog box, shown in Figure 521, displays.

   FIGURE 521 Define Filter dialog box - Advanced tab

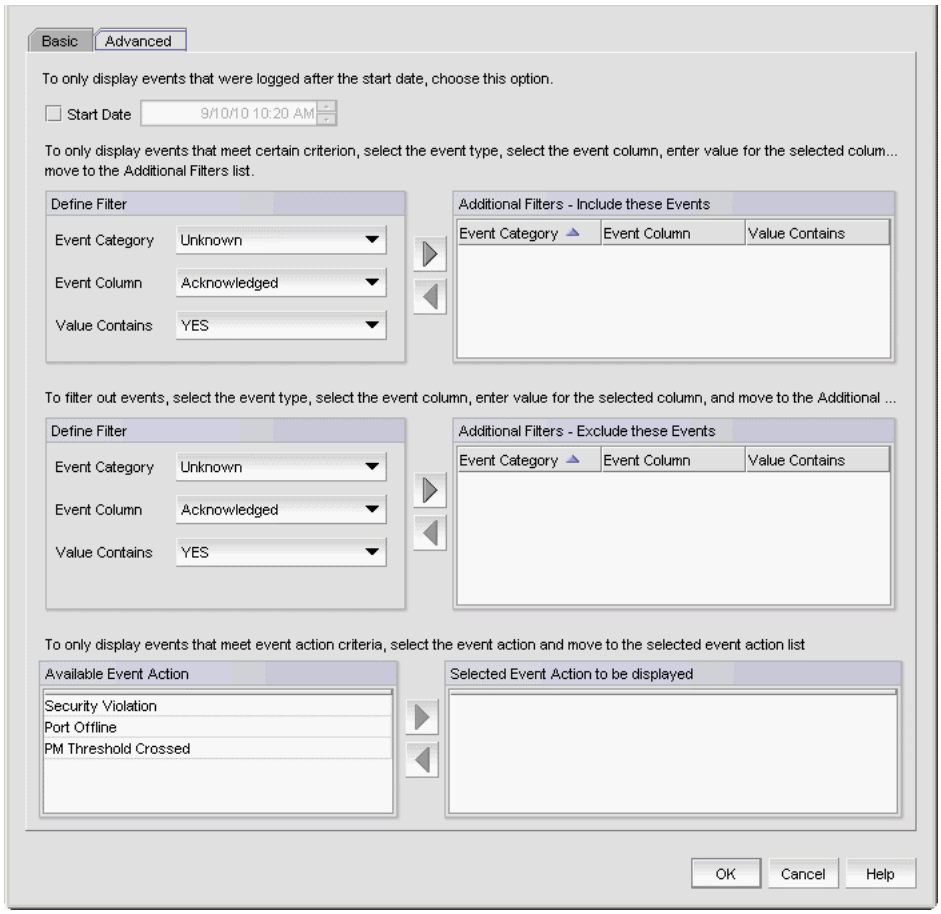

5. Select the **Start Date** check box to display only the events that were logged after the specified start date. The default start date and time is the current date and time.

6. To include events in the event filter, complete the following steps.
    a. Select the event type you want to include from the **Event Category** list.

       All event types are listed in alphabetical order.
    b. Select the event column for the event from the **Event Column** list.

       All event columns are listed in alphabetical order.
    c. Enter all or part of the event type value in the **Value Contains** field.
    d. Click the right arrow button to move the event type to the **Additional Filters - Include these Events** list.
    e. To add additional filters, repeat step a through step d.
7. To exclude events from the event filter, complete the following steps.

   NOTE
   You can configure a maximum of ten filters to be included.

    a. Select the event type you want to remove from the **Event Category** list.

       All event types are listed in alphabetical order.
    b. Select the event column for the event from the **Event Column** list.

       All event columns are listed in alphabetical order.
    c. Enter all or part of the event type value in the **Value Contains** field.
    d. Click the right arrow button to move the event type to the **Additional Filters - Exclude these Events** list.
    e. To remove additional filters, repeat step a through step d.
8. To display events generated by an event action, select the event action from the **Available Event Action** list and click the right arrow button to move it to the **Selected Event Action to be displayed** list.
9. Click **OK** to close the **Define Filter** dialog box.

Viewing events

The **All Events** dialog box enables you to view all events that have occurred on the selected switch, even events that were filtered using advanced filtering criteria.

To view events for a selected device, complete the following steps.

1. Right-click a switch from the device tree or connectivity map.
2. Select **Events** from the list.

   The **All Events** dialog box displays.

SNMP traps

Simple Network Management Protocol (SNMP) provides a means to monitor and control network products and to manage configurations, statistics, performance, and security through authentication and privacy protocols.

The Management application allows you to configure SNMP traps. The SNMP configuration tasks are described in the following sections:

Adding a trap recipient to one or more switches

The **SNMP Trap Recipients** dialog box allows you to register any recipient as a trap recipient on selected products. You can register different recipients for different products.

> NOTE
> You can register and unregister other recipient servers on the Fabric OS switches on a per-switch basis. For IP products, you can perform registration only at the switch level.

To add a trap recipient to one or more switches, complete the following steps.

1. Select **Monitor > SNMP Setup > Product Trap Recipients**.

   The **SNMP Trap Recipients** dialog box, shown in Figure 522, displays.

FIGURE 522 SNMP Trap Recipients dialog box

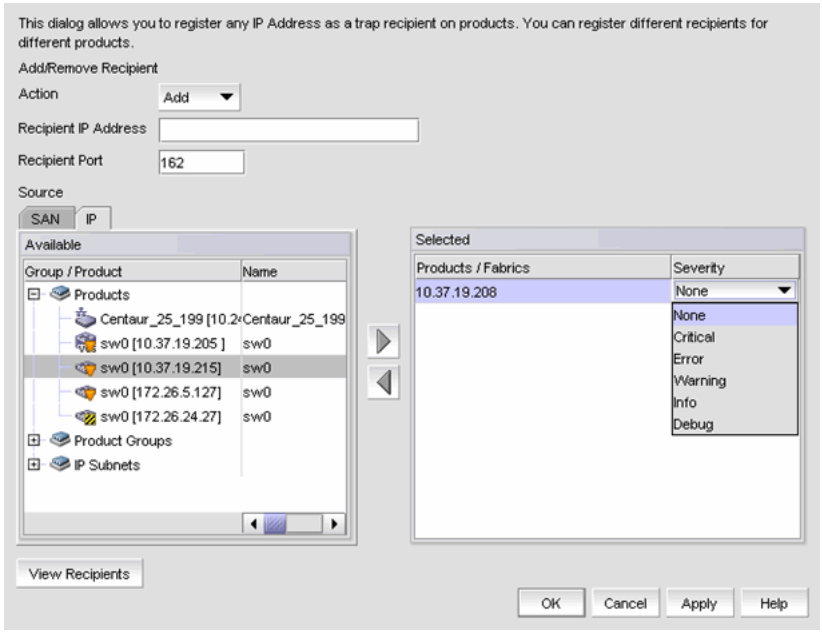

2. Click **Add** from the **Action** list.
3. Enter the IP address of the SNMP trap receiver (the recipient server) in the **Recipient IP Address** field. This is a mandatory field. IPv4 addresses are accepted, but a Domain Name System (DNS) name is not accepted.
4. Enter the SNMP trap port of the recipient in the **Recipient Port** field. This is a mandatory field. Valid numeric values range from 1 through 65535 and 162 is the default.
5. Select the fabric or switches from the **Available** list and click the right arrow button to move it to the **Selected** list. You can select multiple products.

   NOTE
   For IP products and product groups, only switches are available to select.

6. If the selected product is a SAN or Network OS device, select a severity from the **Severity** list. Severity levels can be one of the following: None, Critical, Error, Warning, Info, or Debug. The **Severity** list is disabled for IP products. None is the default.
7. Click the **View Recipients** button to list the recipients that correspond to a selected fabric or product from the **Available** list.

   The **Trap Recipients - Fabric** dialog box or the **Trap Recipients - IP address** dialog box (depending on which product you selected) displays a list of configured recipients.

8. Click **OK**.

   The Management application registers the recipient IP address as an SNMP trap recipient. The SNMP version and credentials from the SNMP profile (for example, SNMP v3) are registered.

Removing a trap recipient from one or more switches

To remove a trap recipient from one or more switches, complete the following steps.

1. Select **Monitor > SNMP Setup > Product Trap Recipients**.

   The **SNMP Trap Recipients** dialog box, shown in Figure 522, displays.

2. Click **Remove** from the **Action** list.
3. Enter the IP address of the SNMP trap port (the recipient server) in the **Recipient IP Address** field.
4. Select the fabric or switches from the **Available** list.

   NOTE
   For IP products, only switches are available to select.

5. Click **OK**.

   The Management application removes the recipient from the managed switches.

SNMP trap forwarding

The **SNMP Trap Forwarding** dialog box allows the Management application to forward received SNMP traps to product trap recipients.

You can use the SNMP Trap Forwarding feature to set up filters to determine which traps will be forwarded. The filters can be one of the following:

- Severity of the trap
- Available products type
- Trap type
- Message types (application messages or pseudo events)

To forward SNMP traps, complete the following steps.

Select **Monitor > SNMP Setup > Trap Forwarding**.

The **SNMP Trap Forwarding** dialog box, shown in Figure 523, displays.

FIGURE 523 SNMP Trap Forwarding dialog box

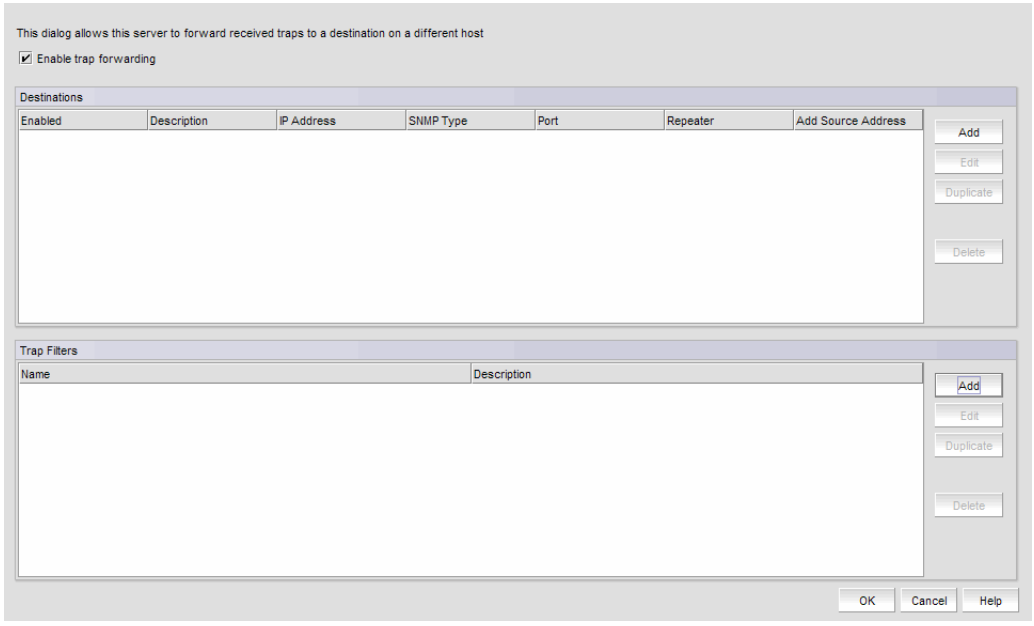

Adding a trap destination

The **Add Trap Destination** dialog box allows you to configure destinations for forwarding SNMP traps.

To add a trap destination, complete the following steps.

1. Select **Monitor > SNMP Setup > Trap Forwarding**.

   The **SNMP Trap Forwarding** dialog box, shown in Figure 523, displays.

2. Select the **Enable trap forwarding** check box.
3. Click **Add** in the **Destinations** area of the **SNMP Trap Forwarding** dialog box.

   The **Add Trap Destination** dialog box, shown in Figure 524, displays.

   FIGURE 524 Add Trap Destination dialog box

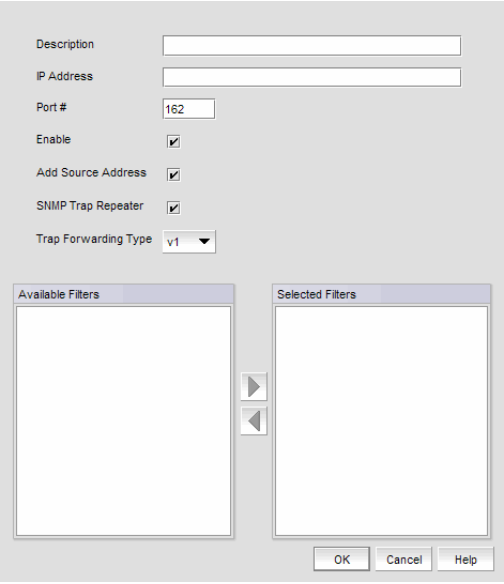

4. Enter a general description of the trap destination in the **Description** field.
5. Enter the IP address of the trap destination in the **IP Address** field. This is a mandatory field. IPv4 and IPv6 addresses are accepted, but a DNS name is not accepted.
6. Enter the SNMP trap listening port of the recipient in the **Port #** field. This is a mandatory field. Valid numeric values range from 1 through 65535.

   The **Enable** check box, **Add Source Address** check box, and **SNMP Trap Repeater** check box are selected by default. When selected, all traps, whether the source is managed or unmanaged, are forwarded. When unselected, only traps from the selected products are forwarded. When selected, an IP Address is added to the variable binding (varbind) value to the trap before forwarding.

7. Select a supported SNMP type from the **Trap Forwarding Type** list. Supported SNMP types are v1, v2c, and v3. The default SNMP type is v1.
8. You can choose not to select a filter (zero), or you can select up to five filters from the **Available Filters** list. Click the right arrow button to move them to the **Selected Filters** list. You must uncheck the **SMNP Trap Repeater** checkbox to assign filters from the **Available Filters** list to the **Selected Filters** list.
9. Click **OK**.

Adding a new trap filter

The **Add Trap Filter** dialog box allows you to configure trap filters for forwarding SNMP traps. You can add trap filters on SAN products, IP products, or hosts. These filters can be on individual switches or the Fabric as a whole.

To add a new trap filter, complete the following steps.

1. Select **Monitor > SNMP Setup > Trap Forwarding**.

   The **SNMP Trap Forwarding** dialog box displays.

2. Click **Add** in the **Trap Filters** area of the **SNMP Trap Forwarding** dialog box.

   The **Add Trap Filter** dialog box, shown in Figure 525, displays.

   FIGURE 525 Add Trap Filter dialog box

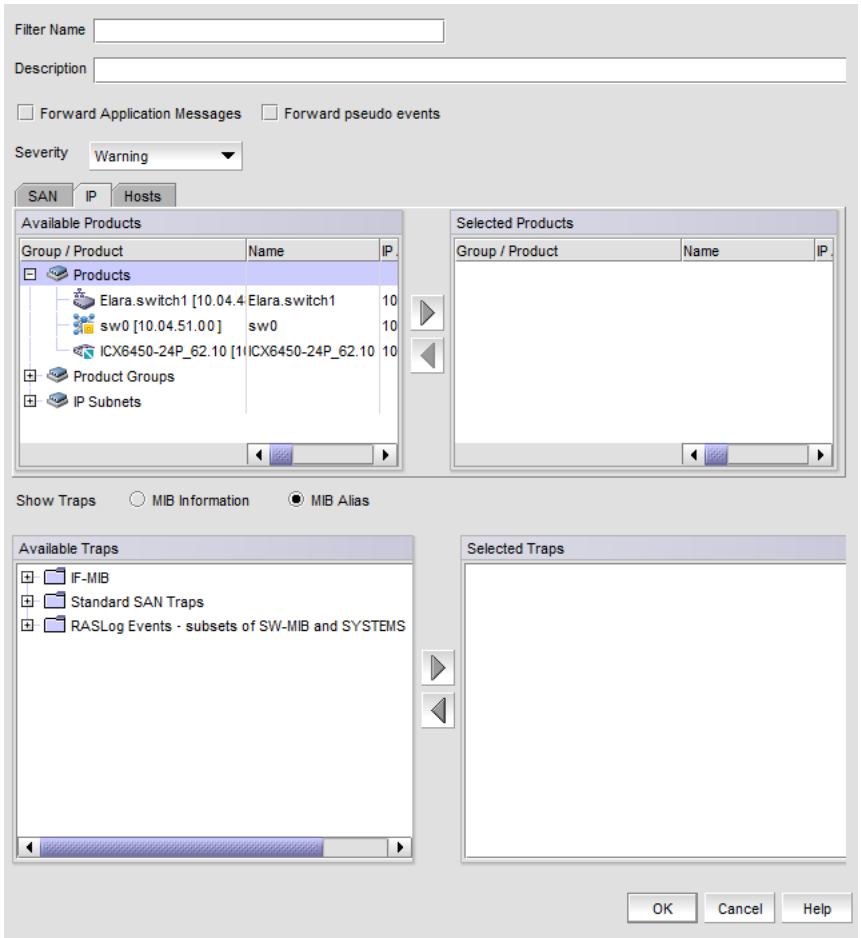

3. Enter a unique name for the trap filter in the **Filter Name** field.
4. Enter a general description of the trap filter in the **Description** field.
5. Select the **Forward Application Messages** check box to forward application events.
6. Select the **Forward pseudo events** check box to forward pseudo events.

NOTE
Select **Forward Application Messages** and/or **Forward pseudo events** check boxes to forward apllication messages or pseudo events of the discovered device from source to the destionation server.

7. Select a severity level from the **Severity** pulldown menu. The severity level can be one of the following, and appear in descending order of severity.
   - Emergency
   - Alert
   - Critical
   - Error
   - Warning
   - Notice
   - Info
   - Debug

   Traps with the selected severity and those with higher severity levels are forwarded. For example, by default, Critical severity is selected. Therefore, traps with Critical, Alert, and Emergency severity levels are forwarded. To have all traps forwarded, select Debug, the lowest severity level.
8. Select the **SAN**, **IP**, or **Hosts** tab. Depending on the tab selected, the products available to which you can add a trap filter display in the **Available Products** list.
9. By default, all traps are listed in the **Available Traps** list, under the folders for the MIB to which they belong. You can limit the list by selecting one of the following MIB types:
   - **MIB Information** - Select this check box if you want the default SNMP name for the traps to be displayed.
   - **MIB Alias** - Select this check box if you want the aliases for traps to be displayed.
10. After limiting the list of available traps, expand the MIB folder to which the trap you want belongs under the **Available Trap Type** list and select that trap. Click the right arrow button to move it to the **Selected Trap Type** list.
11. Click **OK**.

    SNMP Traps and Syslog messages from the selected switches or Fabric will now be forwarded to the configured destination server.

Event reception

The Event Reception feature provides an interface to add the credentials and community strings required to decode traps. You can use the **Event Reception** dialog box to configure the trap message, severity, and alias name that is used by the Event Processor.

The **Event Reception** dialog box contains two tabs:

- The **Trap Credentials** tab allows you to configure the server to accept or drop SNMP traps and add SNMP credentials and community strings for decoding traps.
- The **Trap Configuration** tab allows you to customize the trap description or message, severity, and alias name.

To access the **Event Reception** dialog box, select **Monitor > SNMP > Event Reception**.

The **Event Reception** dialog box, shown in Figure 526, displays.

FIGURE 526 Event Reception dialog box - Trap Credentials tab

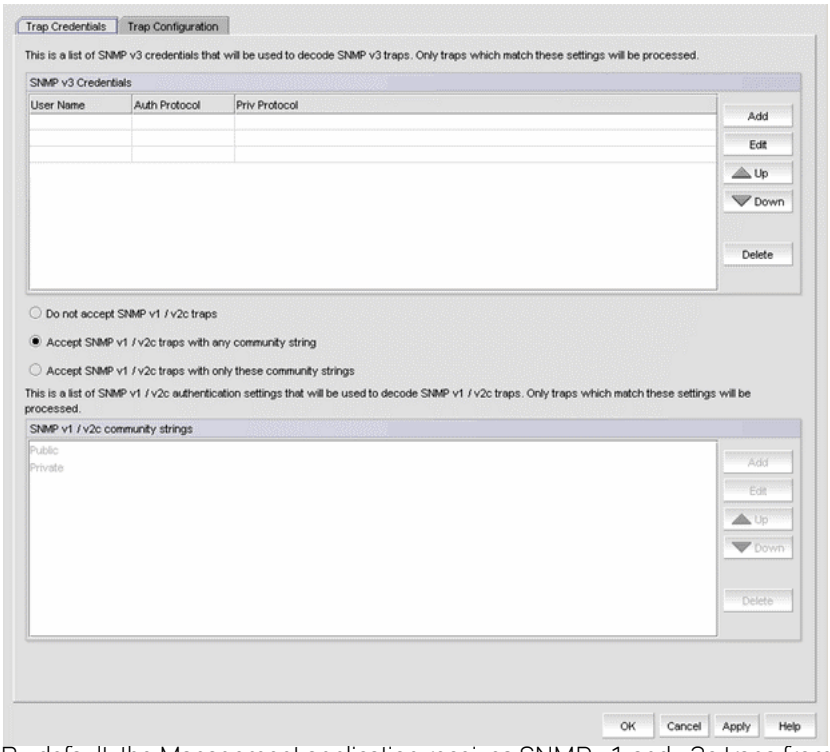

By default, the Management application receives SNMP v1 and v2c traps from IronWare OS and Network OS IP products that have any SNMP community strings. You can accept or restrict SNMP v1 and v2c traps by selecting one of the following check boxes in the **Event Reception** dialog box:

- **Do not accept SNMP v1/v2c traps**

  Use this option to turn off receiving SNMP v1 and v2c traps.If selected, the Management application will not receive any SNMP v1 and v2c traps.

- **Accept SNMP v1/v2c traps with any community string**

  Use this option to turn on receiving SNMP v1 and v2c traps with any community string.

- **Accept SNMP v1/v2c traps with only these community strings**

  Use this option to turn on receiving SNMP v1 and v2c traps with only the specified community strings.

The Management application can receive SNMP v1 traps from Fabric OS SAN switches and directors that have any SNMP community strings. It can receive SNMP v3 traps and informs from these SAN products. The Management application runnning 14.0.0 or later does not support SNMP v3 traps with AES-256.

Table 102 explains the combinations of security and authentication, which will help you when you make your SNMP credentials configuration decisions.

TABLE 102 SNMP security and authentication

SNMP credential type	Privacy protocol	Authentication	Result
v1	No authentication No privacy protocol	Community string	Uses a community string to match for authentication.
v2c	No authentication No privacy protocol	Community string	Uses a community string to match for authentication.
v3	No authentication No privacy protocol	User name	Uses a user name to match for authentication.
v3	Authentication No privacy protocol	MD5 or SHA	Provides authentication based on the HMAC-MD5 (Message Digest Algorithm) or HMAC-SHA algorithms (Secure Hash Algorithm).
v3	Authentication Privacy protocol	MD5 or SHA	Provides authentication based on the HMAC-MD or HMAC-SHA algorithms (Hash-based Message Authentication). Provides privacy based on CBC_DES (Cipher Block Chaining) or CFB_AES_128 (Cipher Feedback).

For information about how to configure SNMP credentials, refer to "Adding an SNMP v3 credential" on page 1144 or "Adding an SNMP v1 or v2c community string" on page 1145.

Adding an SNMP v3 credential

The **SNMP v3 Credentials** dialog box allows you to add the SNMP v3 credentials.

To add an SNMP v3 credential, complete the following steps.

1. Select **Monitor > SNMP Setup > Event Reception**.

   The **Event Reception** dialog box displays.

2. Select an SNMP v3 credential from the **SNMP v3 Credentials** list on the **Event Reception** dialog box.

3. Click **Add**.

   The **Add SNMP v3 Credentials** dialog box, shown in Figure 527, displays.

   FIGURE 527 Add SNMP v3 Credentials dialog box

   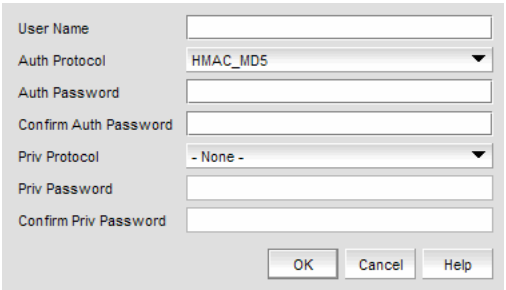

4. Type the user name in the **User Name** field.

   For configurations that do not have authentication or privacy, the Management application uses the user name to match for authentication.

5. Select an authentication protocol from the **Auth Protocol** list. You can select -None-, HMAC-MD5, or HMAC_SHA. HMAC_MD5 is the default.

   If you select no authentication, the Management application uses the user name to match for authentication.

6. Type a password in the **Auth Password** field and re-type the password in the **Auth Confirm Password** field.
7. Select a privacy protocol from the **Priv Protocol** list. You can select -None-, CBC_DES, or CFB_AES_128.

   If you select no privacy, the Management application uses the user name to match for authentication.

8. Type a password in the **Priv Password** field and re-type the password in the **Confirm Priv Password** field.
9. Click **OK**.

Adding an SNMP v1 or v2c community string

The **SNMP v1/2 Community String** dialog box allows you to add the SNMP v1 or v2c credentials.

To add an SNMP v1 or v2c community string credential, complete the following steps.

1. Select **Monitor > SNMP Setup > Event Reception**.

   The **Event Reception** dialog box displays.

2. Click the **Accept SNMPv1/v2c traps with only these community strings** button.
3. Click **Add**.

   The **SNMP v1/v2c Community String** dialog box, shown in Figure 528, displays.

   FIGURE 528 SNMP v1/v2 Community String dialog box

   

4. Enter a unique community string in the **Community String** field, which will be used to match for authentication in SNMP v1 and v2c configurations. This field is case-sensitive.
5. Re-enter the string in the **Confirm Community String** field.
6. Click **OK**.

Importing a new MIB into the Management application

The SNMP traps that the Management application receives must be registered in the Management application in order for these traps to be available. To register a trap, you must first identify the MIB file that contains the trap information in the mibs_to_compile.txt file. Then, you must register the traps using the **Event Reception** dialog box.

To add the MIB file that contains the trap you want to register to mibs_to_compile.txt, complete the following steps.

1. Go to <*install-dir*>\conf\mibs\ (Windows) or <*install-dir*>/conf/mibs/ (UNIX) directory and copy the MIB file into that directory. You may want to copy the MIB into a subdirectory of that directory.
2. In the <*install-dir*>\conf\mibs\ (Windows) or <*install-dir*>/conf/mibs/ (UNIX) directory, search for the mibs_to_compile.txt file.

3. Using a text editor, open the mibs_to_compile.txt file and add the MIB information to the document.

   When adding the MIB information, be aware of the following rules:

   - MIBs are compiled in the order that they are listed in the mibs_to_compile.txt file.
   - You can add composite MIB files (more than one MIB in a single file).
   - MIB file names in the mibs_to_compile.txt file are case-sensitive. Make sure the case of the file name you enter matches the case of the actual MIB file. Also, be sure to enter the complete path of the MIB file, or the portion relative to the mibs directory.

   The following is an example of how to add the two Cisco MIB files.

   ```
   #
   # Cisco Mibs
   #
   CISCO-SMI.mib
   CISCO-CONFIG-COPY-MIB.mib
   #
   # End Cisco Mibs
   #
   ```

4. Save the file.

   The Management application recompiles all the MIB files. If compilation is successful, the traps can now be registered in the **Event Reception** dialog box.

   NOTE
   If there are compilation errors, you can view the errors in the server log:
   <*install dir*>\logs\server\server.log (Windows) or <*install dir*>/logs/server/server.log (UNIX).

5. If you make changes to the MIB file, open the mibs_to_compile.txt file and save the file.

   The Management application recompiles the MIB files and reloads the changes.

Trap customization

The **Trap Configuration** tab of the **Event Reception** dialog box enables you to configure the following settings:

- Register and unregister various Management Information Bases (MIBs)
- Customize trap description messages based on varbinds and severity and specify alias names

Registering traps

Traps must be registered in the **Event Reception** dialog box to make them available.

To register traps, complete the following steps.

1. Select **Monitor > SNMP Setup > Event Reception**.
2. Click the **Trap Configuration** tab.

   The **Trap Configuration** tab of the **Event Reception** dialog box, shown in Figure 529, displays.

   The **Registered** and **Not Registered** buttons at the top of the Traps tree serves as a filter for the traps. If there are unregistered traps, they are listed when you select the **Not Registered** button.

Traps appear under each MIB folder. The MIB folders correspond to the MIBs identified in the mibs_to_compile.txt file.

FIGURE 529 Trap Configuration tab of the Event Reception dialog box

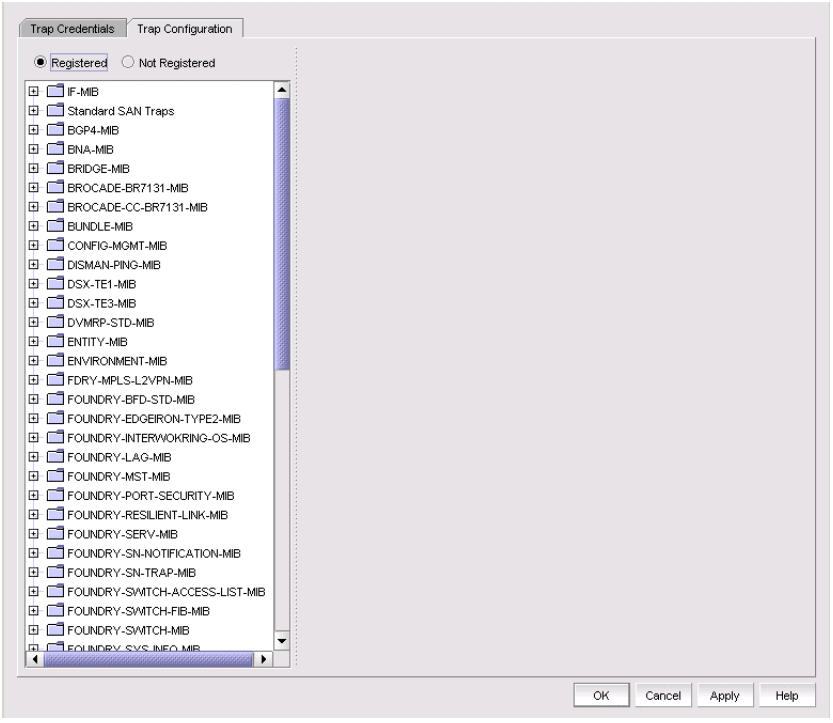

3. Expand a folder for a MIB to display the traps in the MIB. If the list is too long, use the Search tool to find a MIB or trap.
4. Select the trap you want to register.

   The SNMP name and Object Identification (OID) of the trap appear at the top line of the configuration pane. Also, the status of the trap shows **Not Registered**, which is the default definition of the trap.

   Details about the trap appear in the fields beneath the **MIB Name** field.

   Trap details supply the following information:

   - The name of the MIB to which the trap belongs
   - Information about the trap
   - Any variable bindings (varbinds) that the trap uses. Information about the varbind, its name, OID, and type, is displayed
   - Recommended action specified by the user
5. Enter the following information:

   a. Select the severity level you want to assign to the trap from the **Severity** list. If you do not select a severity, it defaults to Emergency.

   b. Enter the message you want to display for this trap in the **Message** field. If the trap has varbinds, use $#, where # represents the varbind number, to indicate the varbind. You must enter a message.

   c. Enter an alias string that serves as a second name for the trap in the **MIB Alias** field. This string might be more understandable to users. This parameter is optional. The Event Processor uses this alias, and this alias is displayed in the Event Action.

   d. Configure the recommended action for the trap.

6. When you have finished, click **OK** to accept your entries.

   The status of the trap changes to **Registered - Customized** and the trap appears in the Event Log.

Unregistering a registered trap

You can unregister only the traps that you have registered. You cannot unregister traps that come with the Management application by default.

To unregister a trap that you have registered, complete the following steps.

1. Select **Monitor > SNMP Setup > Event Reception**.
2. Click the **Trap Configuration** tab.
3. Click the **Registered** button.

   The Trap tree displays the MIBs that contain the registered traps.
4. Expand a MIB folder to display the traps that have been registered for that MIB.
5. Select a trap to display its current definition.
6. Click the **Not Registered** button.
7. Click **OK**.

   Once unregistered, the status of the trap changes to **Not Registered**.

Customizing a registered trap definition

To modify the definitions of registered traps, complete the following steps.

1. Click the **Trap Configuration** tab.
2. Click the **Registered** button.

   The Trap tree displays the MIBs that contain the registered traps.
3. Expand a MIB folder to display the traps that have been registered for that MIB.
4. Select a trap to display its current definition. You can change the severity, message, or alias of the trap.
5. When you have finished, click **OK** or **Apply** to accept your entries.

   If you modified a default trap, its status changes from **Registered - Default** to **Registered - Customized**.

Reverting the customization of a registered trap to default

To revert to the default definitions of registered-customized traps, complete the following steps.

1. Click the **Trap Configuration** tab.
2. Click the **Registered** button.

   The Trap tree displays the MIBs that contain the registered traps.
3. Expand a MIB folder to display the traps that have been registered for that MIB.
4. Select a trap to display its current definition.
5. If the trap has been customized, a button labeled **Default** is available. Click **Default** to revert the previous changes to its default.

SNMP informs

The **SNMP Informs** dialog box allows you to enable or disable informs on informs-capable products. SNMP traps are unreliable because the receiver does not send any acknowledgment when it receives a trap. The sender cannot determine if the trap was received. However, an SNMP manager that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If the manager does not receive an inform request, it does not send a response. If the sender (switch) never receives a response, the inform request can be sent again. For this reason, informs are more likely to reach their intended destination.

When using informs, the engine ID must be set to correspond to the management engine IP address to authenticate the inform request. When informs are enabled, the sender sends initial informs request for engine ID discovery from any of its ephemeral ports (ranging from 32768 to 65535) to port 161 on the Management server. The sender receives the acknowledgment of the informs requests on these ephemeral ports. If there is a firewall between the Management application and the switches, the ephemeral ports must be open for SNMP informs to work.

Enabling or disabling SNMP informs

To enable or disable SNMP informs, complete the following steps.

1. Select **Monitor > SNMP Setup > Informs**.

   The **SNMP Informs** dialog box displays.

2. Select a product group from the **Fabric / Products** list.

   The products display in the **SNMP Informs Capable Products** list, where you can determine if the product's status is enabled or disabled.

3. Select a product in the **SNMP Informs Capable Products** list and click the appropriate **Action** button, depending on whether you want to enable or disable SNMP informs for that product.

4. Click **OK**.

Syslogs

Use the **Options** dialog box to automatically register the Management application server as the syslog recipient on all managed SAN and IP products. The syslog listening port number is 514 by default. If you change the port number from 514, auto-registration is disabled.

> NOTE
> Ethernet Access Switches are not listed in the **Syslog Recipient** dialog box.

Adding a syslog recipient

The **Syslog Recipients** dialog box allows you to register any recipient as a syslog recipient on selected products. You can register different recipients for different products.

You can register and unregister other recipient servers on the Fabric OS switches on a per-fabric basis. For IP products, you can perform registration only at the switch level.

> NOTE
> IPv6 Syslog registration is not supported for IronView OS products.

To add a syslog recipient, complete the following steps.

1. Select **Monitor > Syslog Configuration > Product Syslog Recipients**.

   The **Syslog Recipients** dialog box, shown in Figure 530, displays.

   FIGURE 530 Syslog Recipients dialog box

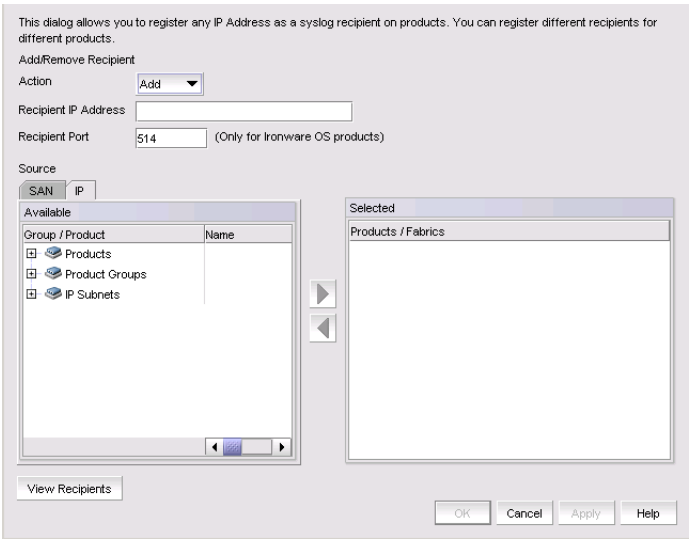

2. Select **Add** from the **Action** list.
3. Enter the IP address of the syslog port (the recipient server) in the **Recipient IP Address** field. This is a mandatory field. IPv4 addresses are accepted, but a DNS name is not accepted.
4. Enter the syslog port of the recipient in the **Recipient Port** field. Valid numeric values range from 1 through 65535. The default value is 514.

   NOTE
   For Network OS and Fabric OS products, non-default ports cannot be registered.

5. Select the fabric or switches from the **Available** list and click the right arrow button to move it to the **Selected** list. You can select multiple products.
6. Click **OK**.

   The Management application registers the recipient IP address as a syslog recipient.

Removing a syslog recipient

To remove a syslog recipient, complete the following steps.

1. Select **Monitor > Syslog Configuration > Syslog Forwarding.**

   The **Syslog Recipients** dialog box displays.

2. Select **Remove** from the **Action** list.
3. Enter the IP address of the syslog port (the recipient server) in the **Recipient IP Address** field.
4. Select the fabric or switches from the **Available** list.

5. Click **OK**.

   The Management application removes the recipient from the managed switches.

Syslog forwarding

The **Syslog Forwarding** dialog box enables the Management application to forward syslog events to a destination on another host. You can use the Syslog Forwarding feature to set up filters to determine which syslog events will be forwarded.

Adding a syslog forwarding destination

The **Add Syslog Destination** dialog box allows you to configure destinations for forwarding syslog events.

To add a syslog destination, complete the following steps.

1. Select **Monitor > Syslog Configuration > Syslog Forwarding**.

   The **Syslog Forwarding** dialog box, shown in Figure 531, displays.

   FIGURE 531 Syslog Forwarding dialog box

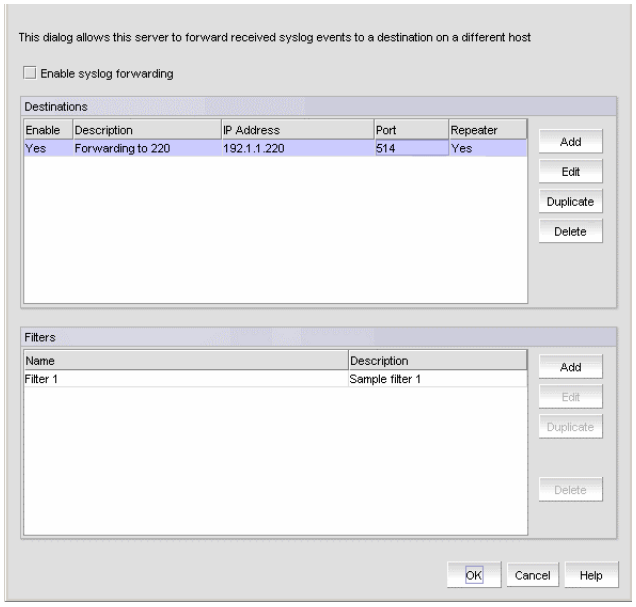

2. Select the **Enable syslog forwarding** check box.
3. Click **Add**.

   The **Add Syslog Destination** dialog box, shown in Figure 532, displays. The **Enable** and **Syslog Repeater** check boxes are selected by default.

FIGURE 532 Add Syslog Destination dialog box

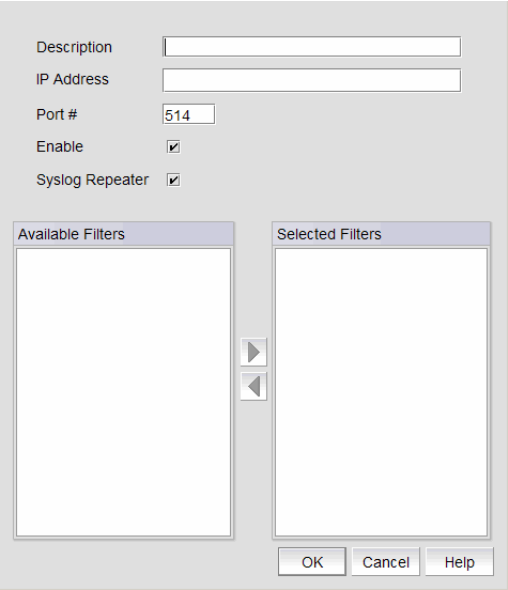

4. Enter a general description of the syslog destination in the **Description** field.
5. Enter the IP address of the syslog destination in the **IP Address** field. This is a mandatory field. IPv4 and IPv6 addresses are accepted, but a DNS name is not accepted.
6. Enter the syslog listening port of the recipient in the **Port #** field. This is a mandatory field. Valid numeric values range from 1 through 65535. The default is 514.
7. Select the **Enable** check box to enable syslog forwarding to this recipient.
8. Select the **Syslog Repeater** check box if you want to forward all syslogs, whether the source is managed or unmanaged. If the Syslog Repeater check box is unselected, syslogs from the managed products are sent to the server. If no filter is selected, then syslogs from all products are sent.
9. You can choose not to select a filter (zero) or you can select up to five filters from the **Available Filters** list. Click the right arrow button to move them to the **Selected Filters** list. This is enabled only when **Syslog Repeater** is not selected.
10. Click **OK**.

Adding a syslog filter

You can add a syslog filter on SAN products, IP products, or hosts.

To add a syslog filter, complete the following steps.

1. Select **Monitor > Syslog Configuration > Syslog Forwarding**.

   The **Syslog Forwarding** dialog box displays.
2. Select the **Enable syslog forwarding** check box.
3. Select **Add in the Filters area.**

   The **Add Syslog Filter** dialog box, shown in Figure 533, displays.

FIGURE 533 Add Syslog Filter dialog box

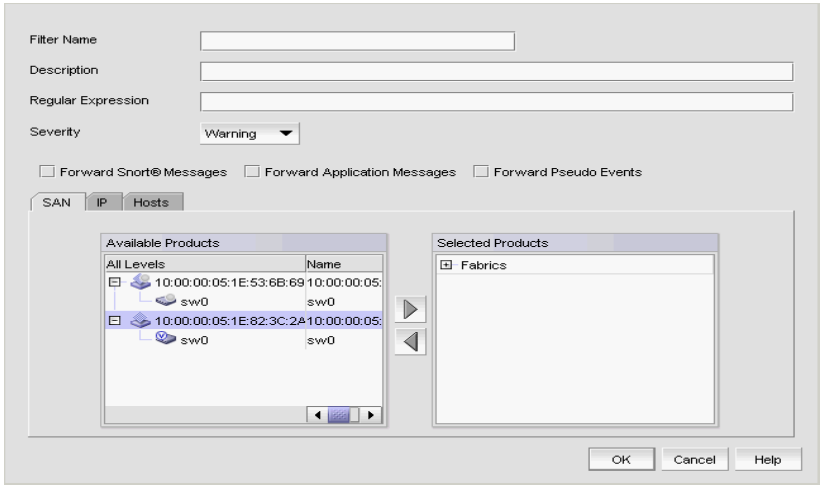

4. Enter a unique name for the syslog filter in the **Filter Name** field.
5. Enter a general description of the syslog filter in the **Description** field.
6. (Optional) For additional filtering, enter a text string using from 1 through 512 characters or wild card symbols in the **Regular Expression** field. The regular expression is used to describe a pattern in text. You can use an asterisk (*) to indicate a wildcard, as in the following examples:
   - *cdef: Matches a message ending with cdef
   - abc*: Matches a message beginning with abc
   - *abc*: Matches a message that contains abc
7. Select a severity level from the **Severity** pulldown menu. The severity level can be one of the following, and appear in descending order of severity.
   - Emergency
   - Alert
   - Critical
   - Error
   - Warning (Default)
   - Notice
   - Info
   - Debug

   Events with the selected severity and those with higher severity levels are forwarded. For example, by default, Critical severity is selected. Therefore, events with Critical, Alert, and Emergency severity levels are forwarded.
   To have all traps forwarded, select Debug, the lowest severity level.
8. Select the **Forward Snort® Messages** check box to turn on Snort message forwarding. Refer to "Snort message forwarding" on page 1154 for more information about Snort messages.
9. Select the **Forward Application Messages** check box to turn on application message forwarding.
10. Select the **Forward Pseudo Events** check box to turn on pseudo event forwarding. Refer to "Pseudo events" on page 1168 for more information about pseudo events.

11. Select the **SAN**, **IP**, or **Hosts** tab. Depending on the tab selected, the products available to which you can add a syslog filter display in the **Available Products** list.
12. Select the product from the **Available Products** list and click the right arrow button to move it to the **Selected Products** list.
13. Click **OK**.

Snort message forwarding

Snort is a third-party tool that monitors network traffic in real time. When Snort detects dangerous payloads or other abnormal behavior, it sends an alert to the syslog in real time. You can turn Snort messages on or off using the **Add Syslog Filter** dialog box

By default, the Forward Snort© Messages feature is not enabled. You must enable it to have Snort messages forwarded to the configured syslog destinations.

You can forward Snort messages, by selecting the **Forward Snort® Messages** check box in the **Add Syslog Filter** dialog box (refer to step 8 in "Adding a syslog filter" on page 1152).

Event action definitions

To reduce the amount of events being logged in the Management application database, the **Event Actions** dialog box allows you to control what events the Management application monitors, on which products they are to be monitored, how often they are to be monitored, and what to do when the monitored events are generated. This information can be defined by creating an event action definition.

For example, you can create an event action definition if you want the Management application to monitor link up and link down traps only, and only on products that belong to Product Group 1. Furthermore, you may want these traps to be logged in the Management application database only if they occur 10 times within a 5-minute interval. You may also want an e-mail message sent to a network administrator when these traps are generated.

In another case, you may not want to log any occurrence of Topology Change traps from Product Group 2. You may also want to disable a port on a product if an event that resembles an attack on the network occurs at a certain frequency.

Creating an event action definition

You can configure event policies for events you want to monitor. Use the **Event Actions** dialog box, shown in Figure 534, to customize the event management policy using triggers and actions.

To customize the event management policy, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box, shown in Figure 534, displays.

FIGURE 534 Event Actions dialog box

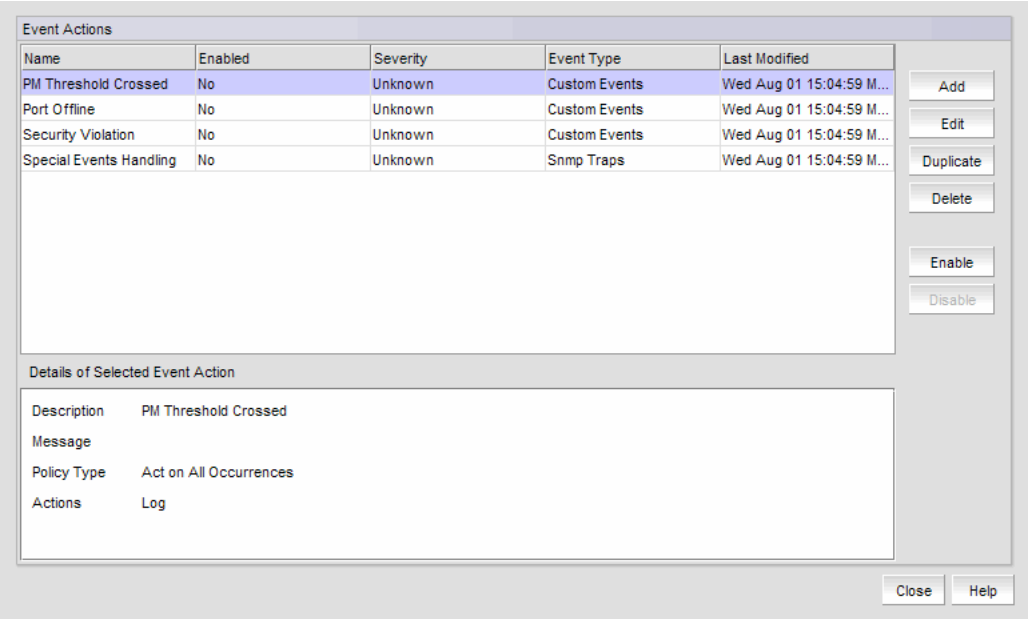

2. Click **Add** to display the **Identification** pane of the **Add Event Action** dialog box.
3. Enter a name and description for the event action and select the **Enabled** check box.
4. Click **Next** to advance to the **Events** pane.

Selecting an event for an event action

To select an event for an event action, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box displays.

2. Click **Add** to display the **Identification** pane of the **Add Event Action** dialog box.

   Enter a name and description for the event action and select the **Enabled** check box.

3. Click **Next** to advance to the **Events** pane.

   By default, the **Events** pane of the **Add Event Action** dialog box displays, shown in Figure 535.

FIGURE 535 Add Event Action dialog box - Events pane

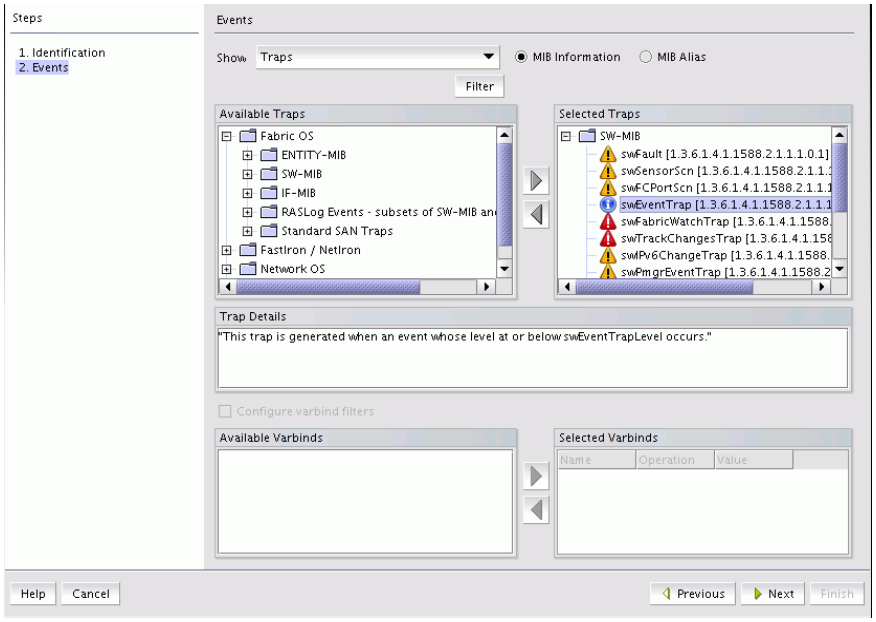

4. Select one of the following event types from the **Show** list:

   - Traps (default)
   - Application Events
   - Pseudo Events
   - Custom Events
   - Snort® Message

   Depending on what the event type you select, a box listing the available events or pseudo events displays.

5. By default, the traps are grouped under the **Fabric OS**, root nodes and listed in the **Available Traps** list, under the folders for the MIB to which they belong. You can limit the list by doing any of the following:

   - Select one of the following options:
     - **MIB Information**, if you want the default SNMP name for the traps to be displayed.
     - **MIB Alias**, if you want the aliases for the traps to be displayed.
   - Use the Trap Filter tool to limit the trap list to the trap severities you want. To use this tool, click the **Filter** button to display the **Trap Filters** dialog box.

     NOTE
     RASLog events specific to Network OS and Fabric OS are listed in the respective Network OS and Fabric OS root nodes. You must select the respective RASLogs for triggering the event action.

6. After limiting the list of available traps, expand the MIB folder to which the trap you want belongs under the **Available Traps** list and select that trap. Click the right arrow button to move it to the **Selected Traps** list. You can select a trap in the **Selected Traps** list and view the trap details in the **Trap Details** section.

7. If you selected **Application Events** in step 4, select the application events in the left list and use the arrow button to move them to the right list.

8. If you selected **Pseudo Events** in step 4, select one or more of the pseudo events you created that you want to include in the definition, then click the right arrow button to move it to the **Selected Pseudo Events** list.

9. If you selected **Custom Events** in step 4, click **Next** to accept the defaults; otherwise, select the Event Category, Severity, Message ID, and Description Contains, as required.
10. If you selected **Snort® Message** in step 4, select the Snort® messages in the left table and use the arrow button to move them to the right.

    To import Snort® rules, click the **Import Snort® Rules** button.
11. Select **Configure varbind filters** to configure filters on varbind values (refer to "Configuring varbind filters" on page 1157 for more information). If you do not want to configure varbind filters, click **Next**.

    The **Sources** pane of the **Add Event Action** dialog box is displayed. You can use the Search tool to search for sources.

Configuring varbind filters

If actions must be confirmed based on a trap variable binding value (varbinds), select the **Configure varbind filters** check box on the **Events** pane of the **Add Event Action** dialog box. This enables you to configure filters on varbind values for this event action.

NOTE
Varbind filter configuration is only available if you selected Traps in step 4 of "Creating an event action definition" on page 1154.

The varbinds for the selected trap are listed in the **Available Varbinds** list, shown in Figure 536.

To configure varbind filters for an event action, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box displays.
2. Click **Next** to advance to the **Events** pane.

   FIGURE 536 Available Varbinds and Selected Varbinds lists

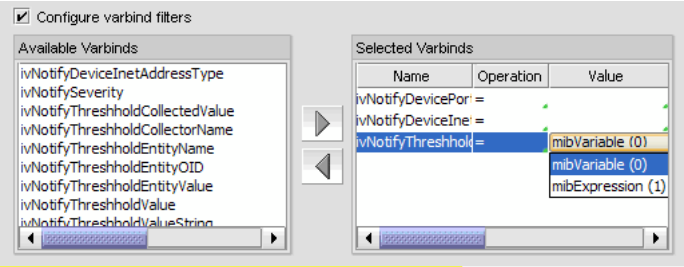

3. Select the varbind you want to include in the configuration and click the right arrow button to move it to the **Selected Varbinds** list.

   If you selected more than one trap and those traps have the same varbinds, then their varbinds are listed in the **Available Varbinds** list. However, if the traps you selected have different varbinds, the **Available Varbinds** list is empty.
4. For each varbind in the **Selected Varbinds** list, select one of the following operations for the condition you want to filter:
   - = – Equal to
   - != – Not equal
   - < – Less than
   - > – Greater than
   - >= – Greater than or equal to

- <= – Less than or equal to
- In – Matches collection
- Not_in – Does not match collection
- ~ – Arbitrary Unicode regular expression

5. Enter the value of the varbind. The value you enter must conform to the data type required by the varbind. For example, if the varbind expects an integer and you enter a text string, your entry will be rejected. Alternatively, you can select values from drop-down lists, shown in Figure 536.
6. Click **Next**.

   The **Sources** pane of the **Add Event Action** dialog box displays. Proceed to "Selecting source address products and ports".

Selecting source address products and ports

The **Sources** pane of the **Add Event Action** dialog box, shown in Figure 537, allows you to enter the IP address, the world wide name, or the name of the source to use as event senders. Alternatively, you can select source address products to use as event senders from the available list of sources. You can select from the available list of SAN products, IP products, or hosts by selecting the appropriate tab.

FIGURE 537 Sources pane of the Add Event Action dialog box

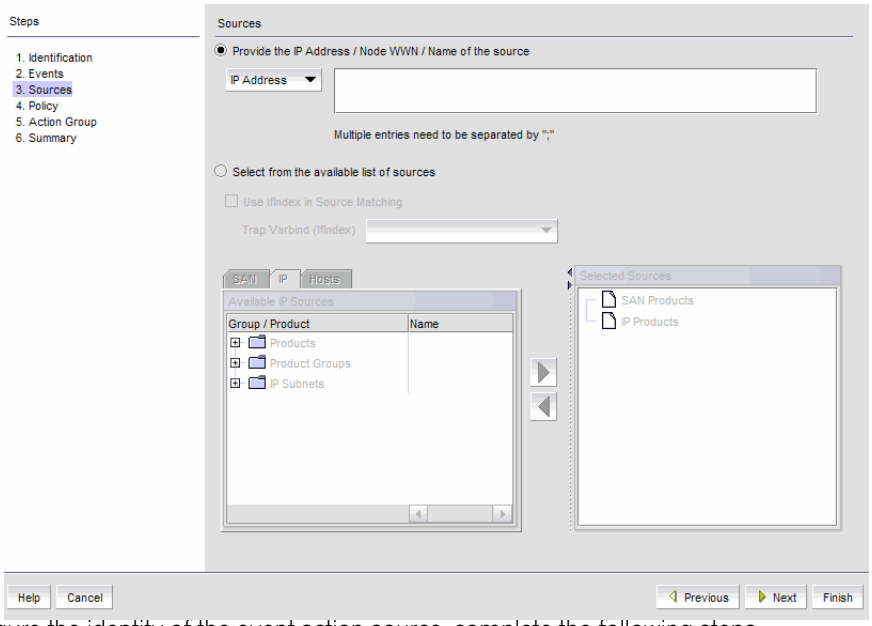

To configure the identity of the event action source, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box displays.
2. Click **Next** to advance to the **Sources** pane. The **Select from the available list of sources** option is selected by default. You can select source products or ports to use as event senders from the available list of sources.
3. Click the **Provide the IP Address / Node WWN / Name of the source** button if you want to manually enter the IP address, the world wide name (WWN), or the name of the source in the **IP Address** field.
4. Select the **Use Ifindex in Source Matching** check box if you want to use ifIndex to filter traps on a specific port of a product; otherwise, the filter is applied globally on a product.

5. If the **Use IfIndex in Source Matching** check box is selected, select the varbind to be used from the **Trap Varbind (IfIndex)** list.
6. Select the event senders you want from the **Available Sources** list, then click the right arrow button to move them in the **Selected Sources** list.
   - If you selected a non-Fabric OSproduct as the source, that product can send e-mail alerts only.
   - If you selected Pseudo Events from the **Events** pane of the **Add Event Action** dialog box, and there is only one pseudo event available, double-click the pseudo event in the **Available Sources** list.
   - If you selected a product group or port group as event senders, select a group from the list.

     NOTE
     The selected source count cannot exceed 100.
7. Click **Next**.

   The **Policy** pane of the **Add Event Action** dialog box displays. Proceed to "Configuring event action policies".

Configuring event action policies

The **Policy** pane of the **Add Event Action** dialog box, shown in Figure 538, allows you to define the frequency of the event, enter a message for an event that will be displayed in the event log, and specify the event severity.

FIGURE 538 Policy pane of the Add Event Action dialog box

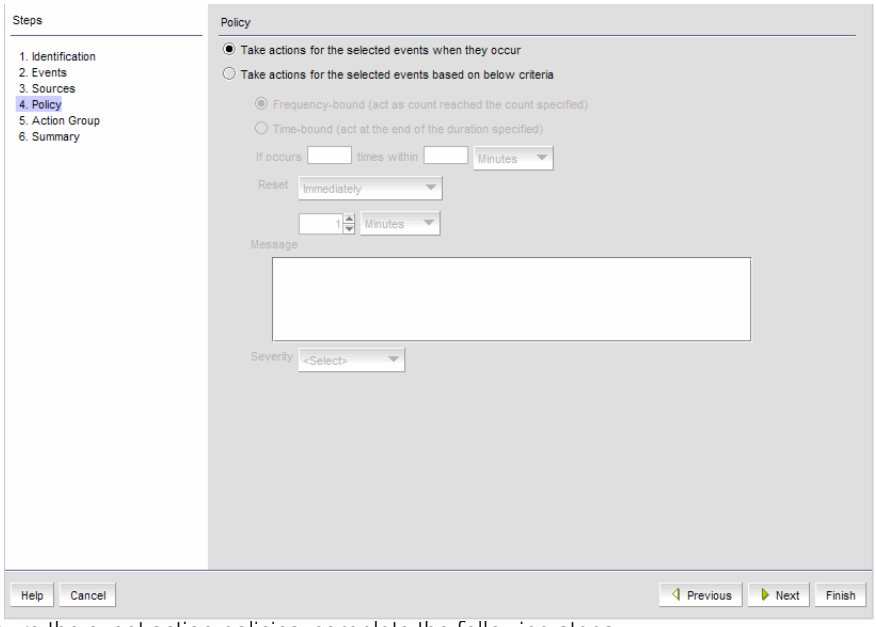

To configure the event action policies, complete the following steps.

1. Click **Take actions for the selected events when they occur (default)** if you want the action to be triggered each time the selected events occur.
2. Click **Take actions for the selected events based on below criteria** if you want the action to be triggered only when the occurrence of the event meets the specified criteria.
   - Click **Frequency bound (act as count reaches the count specified)** if you want the Management application to perform the specified action once the specified number of occurrences has occurred *during* the specified duration. For example, if you want the action to be applied when 10 link down traps occur during a one-minute interval, then the specified action will be applied as soon as 10 link down traps occur, even though the one- minute duration has not elapsed.

- Click **Time bound (act at the end of the duration specified)** if you want the Management application to perform the specified action once the specified number of occurrences has occurred *and* the specified duration has elapsed. For example, if you want the action to be applied when 10 link down traps occur during a one-minute duration, the Management application waits until 10 link down traps occur and one minute has elapsed before the defined action is applied. There is a one-second delay for the action to be applied.

For either option, if the number of occurrences has not been met and the time duration has elapsed, the observation window is advanced to the next occurrence after the first occurrence on the current window.

3. Enter values in the **If occurs __ times within __ fields and select a value from the Minutes list** if you want the action to be applied only if the event occurs at a certain frequency.
4. Indicate how often the policy is to be reset. You can choose one of the following options:
   - **Reset immediately** - Repeats the policy as soon as the specified action has been applied.
   - **Wait until ____ seconds or minutes or hours** - If this parameter is selected, the policy will not be applied to the product for the specified duration of time. Enter the duration in minutes (max 120 minutes), hours (max 2 hours), or one day (1 Day). You can suppress the policy just for the events specified in the policy or for any event that occurs on the product. Once the duration expires, the policy can be repeated.
5. In the **Message** field, enter the message that will be displayed in the Event Log for the generated event. This entry replaces the default message that is displayed for a trap. Also, this message is used as the Event Action message and is displayed in single quotes on the Event Log report.
6. From the **Severity** list, select the severity you want to assign to the generated event.
7. Click **Next.**
8. The **Actions Group - Actions** pane of the **Add Event Action** dialog box displays. Proceed to "Editing event actions". To directly launch the **Actions Group - Actions** pane refer to"Creating an event action from the Master Log" on page 1160

Creating an event action from the Master Log

The Master Log lists the events and alerts that have occurred on the network. You can create an event action from the Master Log**.**

1. Right-click an event in the Master Log**.**
2. Select the event action from the submenu.

   The **Actions Group - Actions** pane of the **Add Event Action** dialog box displays.

   The field values of **Identification**, **Events**, **Sources**, and **Policy** panes will be automatically populated based on the selected event in the Master Log. If required, you can edit the populated fields and create an event action with or without configuring the varbind filter. For more information about adding an event action, refer to "Creating an event action definition" on page 1154.

   NOTE
   If multiple events are selected in the Master Log, the **Event Action** menu is disabled.

   NOTE
   The **Event Action** menu is disabled for unsupported SNMP traps and for the Syslog, Other, and Unknown origins.

Editing event actions

The **Action Group - Actions pane of the Edit Event Action** dialog box, shown in Figure 539, defines what action the Management application takes when the criteria are met.

FIGURE 539 Action Group - Actions pane of the Edit Event Action dialog box

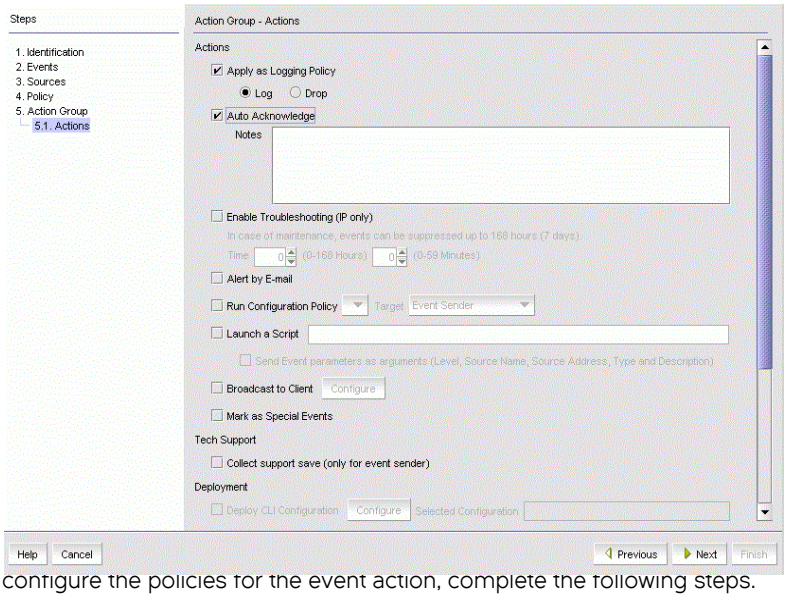

To configure the policies for the event action, complete the following steps.

1. Select **Apply as a Logging Policy** to indicate whether or not you want the event occurrence to be logged in the Management application database:
   - Select **Log** to log the occurrence in the Management application database and Master Log.
   - Select **Drop** to not log the occurrence in the Management application database or Master Log.

     NOTE
     If the policy specifies **Act as specified** on the **Policy** pane of the **Add Event Action** dialog box, and you select **Log** for this parameter, only events that meet the criteria defined in the **Act as specified** area are logged. For example, if the event is logged when 10 link down traps occur during a one-minute interval, then one record will be logged after 10 link down traps occur. If you want all 10 link down traps to be logged, then create a policy where **Act on all occurrences** is selected on the **Policy** pane of the **Add Event Action** dialog box.

2. Select the **Auto Acknowledge** check box to suppress events without being in troubleshooting mode. Activating **Auto Acknowledge** also helps to avoid cluttering the Master Log with unwanted messages without modifying filters. You can enter notes in the **Notes** field. The notes will be added to the matching events.

   NOTE
   Auto Acknowledge is enabled only when **Take actions for the selected events when they occur** is selected in the Policy step of the Event Actions Wizard. If you edit an Event Action that has Auto Acknowledge selected and change this option in the Policy step to **Time-bound** or **Frequency-bound**, you will be required to confirm your choice. The “Acknowledged by” value is displayed as **System** for auto-acknowledged events.

3. **Select** the **Alert by E-mail** check box if you want an e-mail message to be sent to an administrator if the policy criteria have been met.

4. Select the **Run Policy Monitor** check box to execute a policy monitor as an action based on a selected event, and then select the target for the policy monitor from the list. Target options include **Event Sender** and **Specified in Config**.
5. Select the **Launch a Script** check box if you want to execute to an external script file when the matching criteria have been met, and then enter the script in the accompanying field.
6. Select the **Broadcast to Client** check box, and click **Configure** to broadcast a message to all the clients when the matching criteria have been met.

   NOTE
   The remaining parameters are not available if a non-Fabric OS product is selected as an event sender.

   The **Broadcast Message** dialog box displays.

   a. Select a severity level from the list.
   b. Type a message in the **Message Content** field.
   c. Click **OK**.
7. The **Mark as Special Events** check box is unselected by default. Leave it this way if you want the event action to be added to the Special Event Handling event action category. Refer to "Special events handling" for more complete information.
8. Click the **Collect support save** check box to enable SupportSave on the event. The check box is unselected by default.
9. Click **Next** to display the **Action Group - E-mail Settings** pane of the **Add Event Action** dialog box if you selected **Alert by E-mail**. If you did not select **Alert by E-mail**, you will advance to the **Summary** pane.

Special events handling

The following special error conditions are examples of events that are categorized as Special Events Handling events, a separate category that appears in the **Name** list of the **Event Actions** dialog box. All pre-selected events are SNMP traps.

- Invalid T1 zone configuration event
- 48-blade inserted into a non-Virtual Fabric chassis
- Port fencing Fabric Watch trap, when a port is fenced
- Blade Processor FPGA version is incompatible with the Fabric OS firmware version

Though these error conditions are automatically considered "special events handling" events, you can add or edit any event action and mark the action as a special event for special events handling using the **Actions** pane of the **Edit Event Action** dialog box.

Refer to step 7 of "Editing event actions" on page 1161 for information on enabling special events handling for an event using the **Actions** pane of the **Edit Event Action** dialog box.

Acknowledging special events

When the Management application receives and processes events selected as special events, the special icon on the status bar displays (shown in Figure 540).

FIGURE 540 Status bar with highlighted special events icon

To configure special event acknowledgements, complete the following steps.

1. Click the special events icon to launch the **Special Events** dialog box, shown in Figure 541.

   The **Special Events** dialog box lists the 1,000 most-recent events that have been identified as special events. You can add notes while acknowledging or unacknowledging a special event.

   FIGURE 541 Special Events dialog box

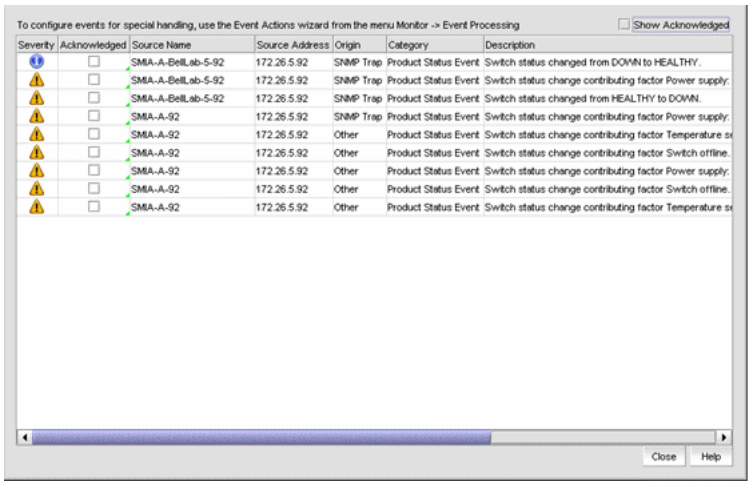

2. Select the **Acknowledged** check box that corresponds to the special event you want to acknowledge.

   The **Add / Edit Note** dialog box displays.

3. Acknowledge or unacknowledge the notes in the **Add / Edit Note** dialog box and click **OK**.

   If an event is marked as acknowledged either in the **Special Events** dialog box or the Master Log, the event is acknowledged in both places.

4. To view all acknowledged special events, select the **Show Acknowledged** check box in the upper right corner of the dialog box. This check box is unselected by default.

   The acknowledged special events display, sorted by the last event server time.

Configuring event action e-mail settings

The **Action Group - E-Mail Settings** pane of the **Add Event Action** dialog box, shown in Figure 542, allows you to select e-mail recipients from a list, add new e-mail recipients, and compose e-mail messages.

FIGURE 542 Action Group - E-Mail Settings pane of the Add Event Action dialog box

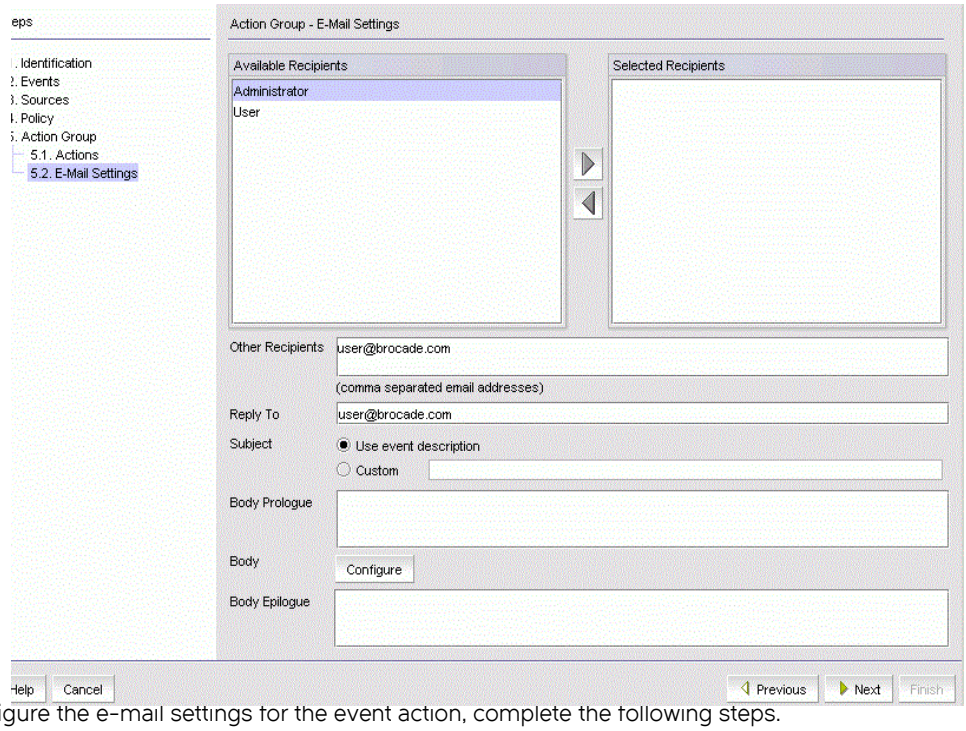

To configure the e-mail settings for the event action, complete the following steps.

1. Select the Management application user to whom the e-mail message will be sent from the **Available Recipients** list, and click the right arrow button to move the recipient to the **Selected Recipients** list.

   NOTE
   Make sure the user you select has an e-mail address defined in a user account.

2. (*Optional*) Add additional e-mail recipient addresses in the **Other Recipients** field. Separate multiple e-mail addresses with a comma. At least one e-mail address must be specified by either selecting an available recipient from the list (step 1) or entering an e-mail recipient.

3. If you want the e-mail message for the alert to display a description on the subject line, perform one of the following actions:

   - Select **Use event description** to use the existing event description.
   - Select **Custom** to enter a new event description in the **Subject** field.

   NOTE
   You can create a prefix that is included in the subject line of every e-mail alert that the Management application sends. The prefix is defined in the configuration.properties file. The prefix plus the text entered in the subject line field cannot exceed 255 characters.

4. If you want a prologue to be inserted at the beginning of the e-mail message, enter up to 255 characters in the **Body Prologue** field. The event action message follows the prologue.

5. If you want to customize and include dynamic content in the body of the e-mail message, click **Configure**.

The **E-mail Content** dialog box displays the following list of parameters selected by default:

- **Event Initiator** - The products or devices or an application from which an e-mail message is sent.
- **Action Event Type** - The event type (Trap, Application Event, or Pseudo Event or Custom Event) selected when configuring an event action.
- **Triggering Event** - The event description with its severity.
- **Event Occurence** - The number of times an event has occurred in the specified time range.
- **Source/Product** - The product or device details.
- **Port** - The port details (such as unit, slot, port number and port name).

You can select or unselect the desired parameters of the content to be displayed in the body of the e-mail message.

6. If you want an epilogue to be placed at the end of the e-mail message, enter up to 255 characters in the **Body Epilogue** field.

   NOTE
   The prologue, the event action message, and the epilogue form the body of the e-mail alert.

7. Click **Finish**.

   The **Summary** pane of the **Edit Event Action** dialog box displays an overview of the e-mail configuration you are creating.

8. Review your entries and take one of the following actions:
   - Click **Finish** to approve the configuration.
   - Click **Previous** to return to the **Action Group - E-Mail Settings** pane of the **Add Event Action** dialog box.
   - Click **Cancel** to cancel the operation.

Creating a new event action definition by copying an existing definition

You can create a new event action definition by copying one that is in the **Event Actions** list.

To create a new event action definition by copying an existing definition, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box displays.

2. Select the definition that you want to copy from the **Event Actions** list.
3. Click the **Duplicate** button to display the **Duplicate Event Actions** dialog box.

   The name of the event action is the name of the selected action with the word “copy” appended. For example, Action1 becomes Action1 copy.

4. Enter a new name for the definition.
5. Change the description of the definition, if needed. You can perform this action in any of the panes of the **Add Event Action** dialog box.
6. Click **Finish** to save the new definition.

Modifying an event action definition

Use caution when you modify an event action. Saving changes to an event action definition resets the runtime information for the events in the definition.

To modify an existing event action definition, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box displays.

2. Select the definition that you want to edit from the **Event Actions** list.
3. Click **Edit** to display the **Edit Event Action** dialog box.
4. Make the changes you want to make to the definition. You can perform this action in any of the panes of the **Add Event Action** dialog box.
5. Click **Finish** to save your definition.

Deleting an event action definition

To delete an event action definition, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box displays.

2. Select the definition that you want to delete from the **Event Actions** list.
3. Click **Delete**.

   A message displays asking you to confirm the deletion request.

4. Click **Yes** to delete the definition, or **No** to cancel the request.

Configuring event actions for Snort messages

To configure an event action for Snort messages, complete the following steps.

1. From the **Identification** pane of the **Add Event Action** dialog box, click **Next** to advance to the **Events** pane. Refer to "Creating an event action definition" on page 1154 for complete instructions on event actions.

   The **Events** pane of the **Add Event Action** dialog box displays, shown in Figure 543. Snort® Message is the default in the **Show** list.

FIGURE 543 Events pane of the Add Event Action dialog box

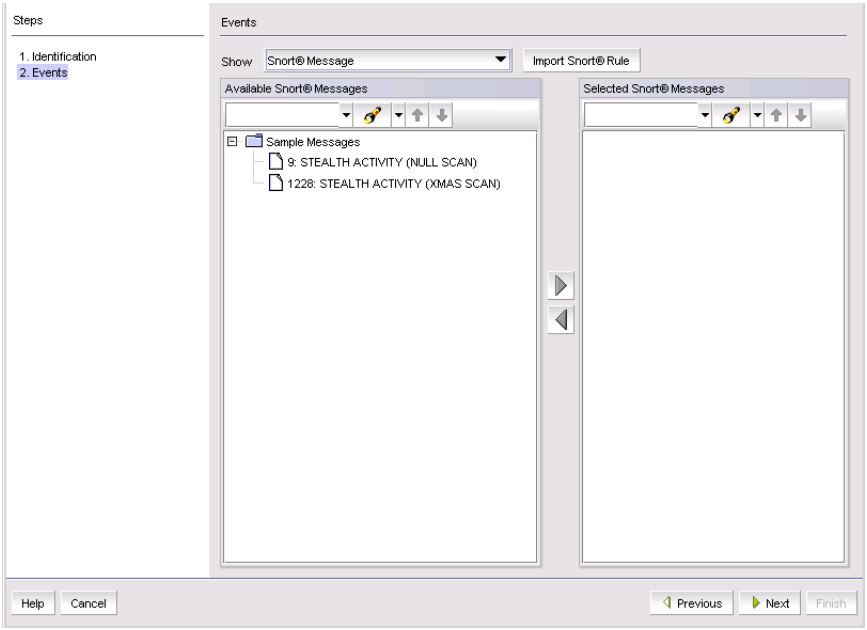

2. Click the **Import Snort® Rule** button.

   The **Import Snort® Rule File** dialog box displays, shown in Figure 544.

FIGURE 544 Import Snort® Rule File dialog box

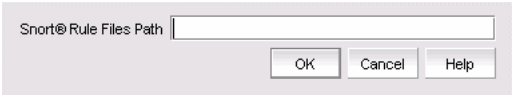

3. Enter the complete path of the Snort rule file located on the Syslog server.
4. Click **OK** to import the Snort rules.
5. While still in the **Add Event Action** dialog box, continue to click **Next** until you advance to the **Action Group - Actions** pane.
6. Select the **Deploy CLI Configuration** check box and click **Configure** if you want to deploy a configuration from CLI Configuration Manager to products if the policy criteria have been met. You can only deploy a CLI configuration for IP products.

   NOTE
   If the CLI configuration you chose from CLI Configuration Manager contains a non-Fabric OS product as a target, the configuration will not be deployed to the non-Fabric OS product.

7. Select one of the following existing CLI configuration parameter sources from the **Parameter** list:
   - **Source IP** — The source IP address of the attack.
   - **Source Port —** The source port of the attack.
   - **Destination IP** — The destination IP address of the attack.
   - **Destination Port** — The destination port of the attack.
8. Continue to advance through the **Add Event Action** dialog box. The **Summary** pane of the **Edit Event Action** dialog box displays an overview of the e-mail configuration you are creating.
9. Review your entries and take one of the following actions:

- Click **Finish** to approve the configuration.
- Click **Previous** to return to the **Action Group - E-Mail Settings** pane of the dialog box.
- Click **Cancel** to cancel the operation.

Pseudo events

A pseudo event is a combination of different SNMP traps that you decide would constitute a single event. For example, there are two separate SNMP traps for link up and link down occurrences. You might decide that these two occurrences should be just one event.

Displaying pseudo event definitions

To display the properties of a pseudo event definition, complete the following steps.

1. Select **Monitor > Event Processing > Pseudo Events**.

   The **Pseudo Events** dialog box, shown in Figure 545, displays.

   FIGURE 545 Pseudo Events dialog box

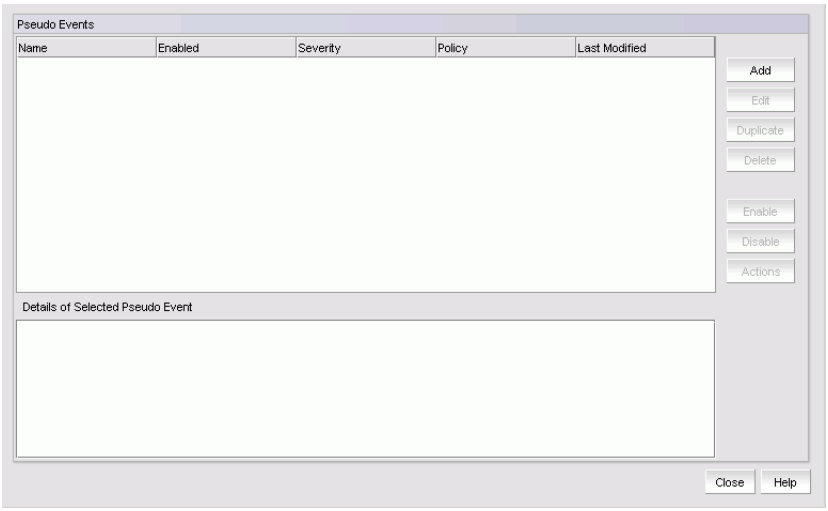

2. To view additional information for a definition, select a definition from the list. Additional information displays in the **Details of Selected Pseudo Event** list at the bottom of the dialog box.

Creating pseudo event definitions

To create a pseudo event definition, complete the following steps.

1. Select **Monitor > Event Processing > Pseudo Events**.

   The **Pseudo Events** dialog box, shown in Figure 545, displays.

2. Click **Add**.
3. The **Identification** pane of the **Add Pseudo Event** dialog box displays.
4. Type a unique name for the pseudo event. Duplicate names are not allowed.
5. Select the **Enabled** check box to enable the pseudo event or clear the check box to disable the pseudo event.

6. Click **Next**.

   The **Policy** pane of the **Add Pseudo Event** dialog box, shown in Figure 546, displays.

Setting pseudo event policies

The **Policy** pane of the **Add Pseudo Event** dialog box, shown in Figure 546, allows you to create escalation, resolve, and flapping policies for the pseudo event, and then specify the time duration for each of these policies in minutes or seconds.

FIGURE 546 Policy pane of the Add Pseudo Event dialog box

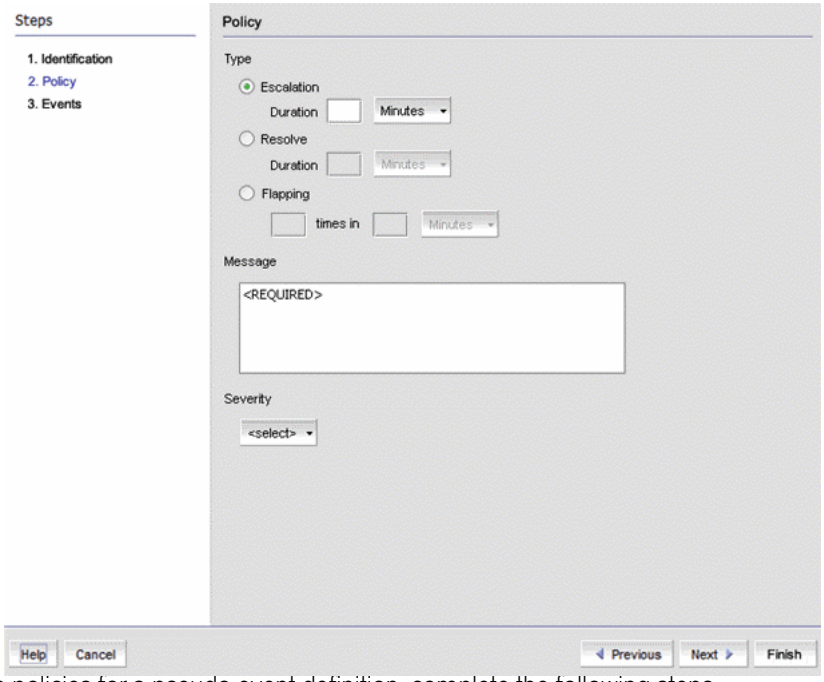

To create policies for a pseudo event definition, complete the following steps.

1. Click the **Escalation** button to create an escalation policy, and then enter the duration of time that the Management application waits before performing the specified action. Specify the escalation time in minutes or seconds.

   When an event occurs, an escalation policy waits for a duration of time to see if the event remains in that state. If it does, then the specified action in the definition is performed.

   Refer to "Adding a pseudo event on the escalation policy" on page 1173 for complete instructions.

2. Click the **Resolve** button to create a resolve policy, and then enter the duration of time the Event Processor waits before generating the pseudo event. Specify the resolve time in minutes or seconds.

   When a down event occurs, a resolving policy waits for a specified duration to see if the event remains in that state by checking if an up event occurs. If an up event occurs, a resolving pseudo event is generated by the Event Processor.

   Refer to "Creating an event action with a pseudo event on the resolving policy" on page 1176 for complete instructions.

3. Click the **Flapping** button to create a flapping policy, and then enter the number of occurrences and the duration of time before the Management application performs the action specified in an event action. Specify the number of flapping times in minutes or seconds.

The flapping policy checks to see if the event consistently transitions between two opposite states during a specified length of time. If it does, then the specified action in the definition is performed.

Refer to "Creating an event action with a pseudo event on the flapping policy" on page 1177 for complete instructions.

4. Enter a description in the **Message** field. This description is displayed in the event log for this pseudo event. The event log displays the exact text you enter in this field; therefore, this message should describe the events in the event action policy.
5. Select a severity from the **Severity** list. You must assign a severity to the pseudo event.
6. Click **Next**.

   The **Events** pane of the **Add Pseudo Event** dialog box, shown in Figure 547, displays.

Refer to the following topics for specific procedures using this dialog box:

- "Creating pseudo event definitions" on page 1168
- "Editing a pseudo event definition" on page 1172

Filtering pseudo event traps

The **Events** pane contains a **Selected Down Trap** list and a **Selected Up Trap** list. The **Selected Down Trap** list defines the traps for the down state of a product or an interface. The **Selected Up Trap** list defines the traps for the up state of the product or an interface.

NOTE
By default in a SAN+IP configuration, all traps known to the Management application are included in the **Available Traps** list, under the folders for the MIB to which they belong.

To filter pseudo event traps, complete the following steps.

1. Select **Monitor > Event Processing > Pseudo Events**.

   The **Pseudo Events** dialog box, shown in Figure 545, displays.
2. Click **Add**.

   The **Events** pane of the **Add Pseudo Event** dialog box, shown in Figure 547, displays.

FIGURE 547 Events pane of the Add Pseudo Event dialog box

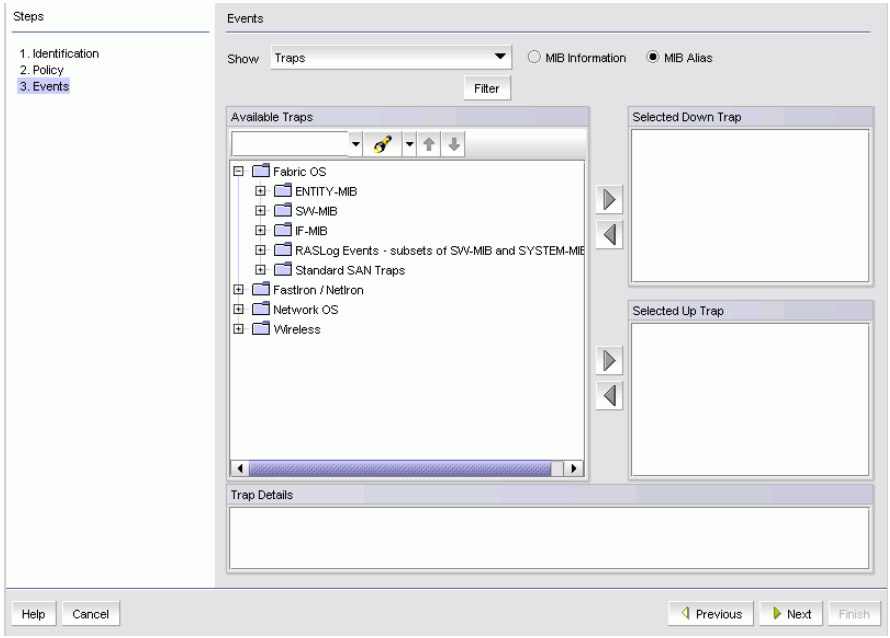

3. From the **Available Traps** list, select the trap for the down state of a product or interface.

   By default, all the traps known to the Management application are grouped under the **Fabric OS**, root nodes and are included in the **Available Traps** list, which is a list of all traps that are available based on the MIB and filter criteria.

4. Select a trap for the **Selected Down Trap** list and a trap for the **Selected Up Trap** list.

   You cannot select the same trap for up and down conditions. Move the traps from the **Available Traps** list to the **Selected Down Trap** and **Selected Up Trap** lists using the right arrow button.

5. You can change the text associated with the selected trap by doing either of the following:

   - Select one of the following options:
     - **MIB Information**, if you want the default SNMP name for the traps to be displayed.
     - **MIB Alias**, if you want the aliases for the traps to be displayed.
   - Use the Trap Filter tool to limit the trap severity. To use this tool, click the **Filter** button to display the **Trap Filters** dialog box.

     NOTE
     RASLog events specific to Network OS and Fabric OS are listed in the respective Network OS and Fabric OS root nodes. You must select the respective RASLogs for triggering the event action.

6. Click **Next** to advance to the **Summary** pane.
7. Click **Finish** to save your definition. The new pseudo event appears on the **Pseudo Event** list on the **Pseudo Event** dialog box.

Creating a pseudo event definition by copying an existing definition

To create a pseudo event definition by copying an existing definition, complete the following steps.

1. Select **Monitor > Event Processing > Pseudo Events**.
2. Select the pseudo event definition that you want to copy from the **Pseudo Events** list.

3. Click the **Duplicate** button.

   The **Pseudo Events** dialog box, shown in Figure 545, displays.

   The name of the event action is the name of the selected action with the word “copy” appended. For example, “Event1” becomes “Event1 copy”.

4. Enter a new name for the pseudo event definition.
5. Make the changes you want to make to the definition. Refer to “Creating pseudo event definitions” on page 1168 for details.
6. Click **Finish** to save your definition.

Editing a pseudo event definition

Use caution when you modify pseudo event definitions. Saving changes to a pseudo event definition resets the run-time information for that pseudo event.

To edit a pseudo event definition, complete the following steps.

1. Select **Monitor > Event Processing > Pseudo Events**.

   The **Pseudo Events** dialog box, shown in Figure 545, displays.

2. Select the pseudo event definition that you want to edit from the **Pseudo Events** list.
3. Click the **Edit** button to display the **Edit Pseudo Event** dialog box.
4. Make the changes you want to make to the definition. Refer to “Creating pseudo event definitions” on page 1168 for details.
5. Click **Finish** to save your definition.

Deleting a pseudo event definition

Use caution when you delete pseudo event definitions. Deleting a pseudo event definition discards the run-time information for that pseudo event.

To delete a pseudo event definition, complete the following steps.

1. Select **Monitor > Event Processing > Pseudo Events**.

   The **Pseudo Events** dialog box, shown in Figure 545, displays.

2. Select the pseudo event definition that you want to delete from the **Pseudo Events** list.
3. Click **Delete**.

   A message displays, prompting you to confirm the deletion request.

4. Click **Yes** to delete the selected definition.

   The definition is removed from the **Pseudo Events** list.

Creating an event action from a pseudo event

You can create an event action from a pseudo event. An event action can be created only for the selected pseudo event.

To create an event action from a pseudo event, complete the following steps.

1. Select **Monitor** > **Event Processing** > **Pseudo Events**.

   The **Pseudo Events** dialog box displays.

2. Select one or more pseudo events and click **Actions.**

   The **Source** pane of the **Add Event Action** dialog box displays.

   The field values of the **Identification** pane and the **Events** pane will be automatically populated based on the selected pseudo events. For more information about adding an event action, refer to "Creating an event action definition" on page 1154.

Adding a pseudo event on the escalation policy

Use the escalation policy to be notified if a critical event occurs on a product, port, or system. When the event occurs, the escalation policy waits for a duration of time to see if the event remains in that state. If it does, then the specified action in the definition is performed.

The following two-part procedure uses both the **Identification** pane of the **Add Pseudo Events** dialog box and the **Add Event Actions** dialog box to create an event action with the escalation policy.

To add a pseudo event definition to the escalation policy, complete the following steps.

1. Select **Monitor > Event Processing > Pseudo Events**.

   The **Pseudo Events** dialog box, shown in Figure 545, displays.

2. Click **Add**.

   The **Identification** pane of the **Add Pseudo Event** dialog box displays.

3. Enter a name for the pseudo event.

4. Select the **Enabled** check box to enable the event, and click **Next**.

   The **Policy** pane of the **Add Pseudo Event** dialog box displays.

5. Click the **Escalation** button, and then enter the duration of time the Event Processor will wait before generating the pseudo event. Specify the escalation time in minutes or seconds.

6. Click **Next**.

   The **Events** pane of the **Add Pseudo Event** dialog box displays.

7. Select a critical event, such as LinkDown, and click the right arrow button to move it to the **Selected Down Trap** list.

8. Select a remediation event, such as LinkUp, and click the right arrow button to move it to the **Selected Up Trap** list.

9. Click **Next** to advance to the **Summary** pane.

10. Click **Finish** to complete the pseudo event configuration.

    Now, you must create a new event action definition using the **Add Event Actions** dialog box. Refer to the following sections for instructions on performing this task:

- "Creating an event action definition" on page 1154
- "Creating a new event action definition by copying an existing definition" on page 1165
- "Creating an event action with a pseudo event on the escalation policy" on page 1174
- "Creating an event action with a pseudo event on the resolving policy" on page 1176
- "Creating an event action with a pseudo event on the flapping policy" on page 1177

Creating an event action with a pseudo event on the escalation policy

To create an event action with a pseudo event on the escalation policy, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box displays.

2. Click **Add** to display the **Identification** pane of the **Add Event Action** dialog box.
3. Enter a name and description for the event action and select the **Enabled** check box to enable the event.
4. Click **Next** to display the **Events** pane.

   By default, the **Events** pane of the **Add Event Action** dialog box displays.

5. Select the **Pseudo Events** event type from the **Show** list.

   The available pseudo events display.

6. Select the pseudo event you created and click **Next**.

   The **Sources** pane of the **Add Event Action** dialog box displays.

7. Select the source that you will use to monitor this event from the **Selected Sources** list.
8. Click **Next** to advance to the **Policy** pane of the **Add Event Action** dialog box.

   The **Policy** pane of the **Add Event Action** dialog box displays.

9. Click the **Take actions for the selected events when they occur** button if you want to take action for the selected events when they occur.
10. Click **Next** to advance to the **Action Group-Actions** pane of the **Add Event Action** dialog box.

    The **Action Group-Actions** pane of the **Add Event Action** dialog box displays.

11. **Select** the **Alert by E-mail** check box. An e-mail notification will be sent to the designated e-mail recipient if the policy criteria have been met.
12. Click **Next** to display the **Action Group - E-mail Settings** pane of the **Add Event Action** dialog box.

    The **Action Group - E-mail Settings** pane of the **Add Event Action** dialog box allows you to select e-mail recipients from a list, add new e-mail recipients, and compose e-mail messages.

13. Select the Management application user to whom the e-mail message will be sent from the **Available Recipients** list, and click the right arrow button to move the recipient to the **Selected Recipients** list.

    NOTE
    Make sure the user you select has an e-mail address defined in a user account.

14. Add additional e-mail recipient addresses in the **Other Recipients** field. Separate multiple e-mail addresses with a semicolon.

15. If you want the e-mail message for the alert to display a description on the subject line, enter the text in the **Subject Line** field.
16. If you want a prologue to be inserted at the beginning of the e-mail message, enter up to 255 characters in the **Body Prologue** field. The event action message follows the prologue.
17. If you want an epilogue to be placed at the end of the e-mail message, enter up to 255 characters in the **Body Epilogue** field.

    NOTE
    The prologue, the event action message, and the epilogue form the body of the e-mail alert.

18. Click **Next** to advance to the **Summary** pane.
19. Click **Finish**.

    The **Summary** pane of the **Add Event Action** dialog box displays an overview of the e-mail configuration you are creating.

For more information about adding an event action, refer to "Event action definitions" on page 1154.

Adding a pseudo event on the resolving policy

When a down event occurs, a resolving policy waits for a specified duration to see if the event remains in that state by checking if an up event occurs. If an up event occurs, a resolving pseudo event is generated by the Event Processor.

The following two-part procedure uses both the **Add Pseudo Events** dialog box and the **Add Event Actions** dialog box to create an event action with the resolving policy.

To add a pseudo event definition to the resolving policy, complete the following steps.

1. Select **Monitor > Event Processing > Pseudo Events**.

   The **Pseudo Events** dialog box displays.

2. Click **Add**.

   The **Identification** pane of the **Add Pseudo Event** dialog box displays.

3. Enter a name for the pseudo event, and select the **Enabled** check box to enable the event.
4. Click **Next**.

   The **Policy** pane of the **Add Pseudo Event** dialog box displays.

5. Click the **Resolve** button, and then enter the duration of time the Event Processor will wait before generating the pseudo event. Specify the resolve time in minutes or seconds.
6. Click **Next**.

   The **Events** pane of the **Add Pseudo Event Events** dialog box displays.

7. Select a critical event, such as LinkDown, and click the right arrow button to move it to the **Selected Down Trap** list.
8. Select a remediation event, such as LinkUp, and click the right arrow button to move it to the **Selected Up Trap** list.
9. Click **Finish** to complete the pseudo event configuration.

   Now, you must create a new event action definition using the **Add Event Actions** dialog box.

Creating an event action with a pseudo event on the resolving policy

To create an event action with a pseudo event on the resolving policy, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box displays.

2. Click **Add** to display the **Identification** pane of the **Add Event Action** dialog box.
3. Enter a name and description for the event action and select the **Enabled** check box to enable the event.
4. Click **Next** to display the **Events** pane.

   By default, the **Events** pane of the **Add Event Action** dialog box displays.

5. Select the **Pseudo Events** event type from the **Show** list.

   The available pseudo events display.

6. Select the pseudo event you created and click **Next**.

   The **Sources** pane of the **Add Event Action** dialog box displays.

7. Select the source that you will use to monitor this event from the **Selected Sources** list.
8. Click **Next** to advance to the **Policy** pane of the **Add Event Action** dialog box.

   The **Policy** pane of the **Add Event Action** dialog box displays.

9. Define the frequency of the event's occurrence that would trigger the action.
   - Click the **Take actions for the selected event when they occur** button if you want to take action for the selected events when they occur.
   - Click the **Take actions for the selected events based on below criteria** button if you want to take action for the selected events based on specified criteria.
10. Click **Next** to advance to the **Action Group-Actions** pane of the **Add Event Action** dialog box.

    The **Action Group-Actions** pane of the **Add Event Action** dialog box displays.

11. Select **Apply as a Logging Policy** to indicate whether or not you want the event occurrence to be logged in the Management application database:
    - Select **Log** to log the occurrence in the Management application database.
    - Select **Drop** to not log the occurrence in the Management application database.
12. Click **Next** to advance to the **Summary** pane.
13. Click **Finish**.

For more information about adding an event action, refer to "Event action definitions" on page 1154.

Adding a pseudo event on the flapping policy

The flapping policy checks to see if the event consistently transitions between two opposite states during a specified length of time. If it does, then the specified action in the definition is performed.

The following two-part procedure uses both the **Add Pseudo Events** dialog box and the **Add Event Actions** dialog box to create an event action with the flapping policy.

To add a pseudo event on the flapping policy, complete the following steps.

1. Select **Monitor > Event Processing > Pseudo Events**.

   The **Pseudo Events** dialog box displays.

2. Click **Add**.

   The **Identification** pane of the **Add Pseudo Event** dialog box displays.

3. Enter a name for the pseudo event, and select the **Enabled** check box to enable the event.
4. Click **Next**.

   The **Policy** pane of the **Add Pseudo Event** dialog box displays.

5. Click the **Flapping** button, and then enter the duration of time the Event Processor will wait before generating the pseudo event. Specify the number of flapping times in minutes or seconds.
6. Click **Next**.

   The **Events** pane of the **Add Pseudo Event** dialog box displays.

7. Select a critical event, such as LinkDown, and click the right arrow button to move it to the **Selected Down Trap** list.
8. Select a remediation event, such as LinkUp, and click the right arrow button to move it to the **Selected Up Trap** list.
9. Click **Next** to advance to the **Summary** pane.
10. Click **Finish** to complete the pseudo event configuration.

    Now, you must create a new event action definition using the **Add Event Actions** dialog box.

Creating an event action with a pseudo event on the flapping policy

To create an event action with a pseudo event on the flapping policy, complete the following steps.

1. Select **Monitor > Event Processing > Event Actions**.

   The **Event Actions** dialog box displays.

2. Click **Add** to display the **Identification** pane of the **Add Event Action** dialog box.
3. Enter a name and description for the event action and select the **Enabled** check box to enable the event.
4. Click **Next** to display the **Events** pane.

   By default, the **Events** pane of the **Add Event Action** dialog box displays.

5. Select the **Pseudo Events** event type from the **Show** list.

   The available pseudo events display.

6. Select the pseudo event you created and click **Next**.

The **Sources** pane of the **Add Event Action** dialog box displays.

7. Select the source that you will use to monitor this event from the **Selected Sources** list.
8. Click **Next** to advance to the **Policy** pane of the **Add Event Action** dialog box.

   The **Policy** pane of the **Add Event Action** dialog box displays.

9. Click the **Take actions for the selected events when they occur** button if you want to take action for the selected events when they occur.
10. Click **Next** to advance to the **Action Group-Actions** pane of the **Add Event Action** dialog box.

    The **Action Group-Actions** pane of the **Add Event Action** dialog box displays.

11. Select the **Deploy CLI Configuration** check box and click the **Configure** button if you want to deploy a configuration from CLI Configuration Manager to products if the policy criteria have been met.

    NOTE
    If the CLI configuration you chose from CLI Configuration Manager contains a non-Fabric OS product as a target, the configuration will not be deployed to the non-Fabric OS product.

12. You can either select an existing CLI configuration or create a new one and select that configuration. After selecting a CLI configuration, the name of the CLI configuration is displayed in the **Selected Configuration** field.

    - **Has Parameters** - Displays **Yes** if the CLI configuration has parameters that require values to be entered before it can be deployed, and displays **No** if no parameter needs to be defined.
    - The **Parameters** list lists the parameters that need to be defined in the configuration.
      - The **Parameter** column displays the parameter and its variables in the CLI configuration.
      - The **Source** column lists the appropriate SNMP attributes for the parameters. Each attribute contains a specific parameter value, such as an IP address. Select the attribute you want from the list.
      - The **Transformation** column uses the product IP addresses and MAC addresses listed in the Address Finder. If the Address Finder list is empty, the product or port will not be found. From this column, specify what you want Event Processor to do with the value in the attribute:
        - **Find Device**: Find the product with the IP address in the attribute and deploy the CLI configuration to that product.
        - **Find Port**: Find the port on a product with the IP address in the attribute and deploy the CLI configuration to that port.
        - **Find Intruder MAC**: Find the product with the IP address in the attribute that matches the intruder MAC address and deploy the CLI configuration to that product.
        - **None**: The Event Processor only reports occurrence of the products.

13. Select the **Deploy Product Configuration** check box if you want to deploy a payload to the products if the policy criteria have been met.
14. Select the **Apply as a Logging Policy** check box to indicate whether or not you want the event occurrence to be logged in the Management application database:

    - Select **Log** to log the occurrence in the Management application database.
    - Select **Drop** to not log the occurrence in the Management application database.

15. Click **Next** to advance to the **Summary** pane.
16. Click **Finish**.

For more information about adding an event action, refer to "Event action definitions" on page 1154.

Event custom reports

The **Event Custom Reports** dialog box allows you to manage customized event filter definitions and schedule when the definitions are run.

To access the dialog box, select **Reports > Event Custom Reports**.

The **Event Custom Reports** dialog box, shown in Figure 548, displays.

FIGURE 548 Event Custom Reports dialog box - Report Definitions tab

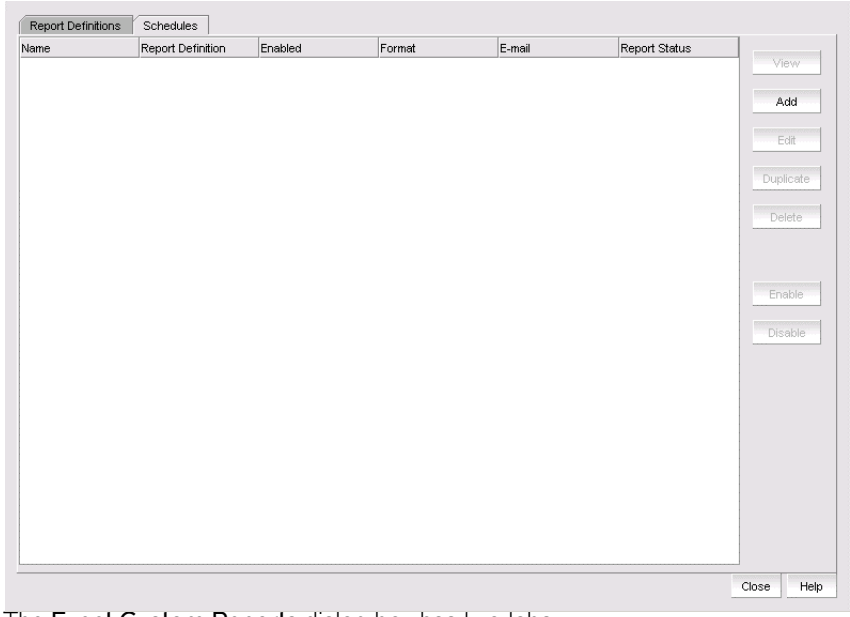

The **Event Custom Reports** dialog box has two tabs:

- The **Report Definitions** tab lists all the previously created report definition objects. This tab enables you to add a new definition or modify, delete, or duplicate existing report definitions.
- The **Schedules** tab lists all the previously created schedules on the report definition. This tab enables you to add a new schedule or modify, delete, or duplicate existing schedules. Users cannot view, edit, or share a schedule that was created by another user.

Defining report settings

You can configure report settings so that you see only a restricted set of information in a report.

NOTE
You can change the number of displayed event custom report records by following the procedure in "Configuring custom report preferences" on page 261. By default, 1000 records display, even if the event count is greater than 1000.

NOTE
You must first enter a name and title on the **Identification** tab before you can run the result settings.

To configure report settings, complete the following steps.

1. Select **Reports > Event Custom Reports**.

   The **Event Custom Reports** dialog box displays.

2. Click the **Add** button.
3. The **Add/Edit Report Definition** dialog box - **Product** tab, shown in Figure 549, displays.
4. Click the **Result Settings** tab.

   The **Add/Edit Report Definition dialog box - Result Settings tab** displays.

   FIGURE 549 Add/Edit Report Definition dialog box - Result Settings tab

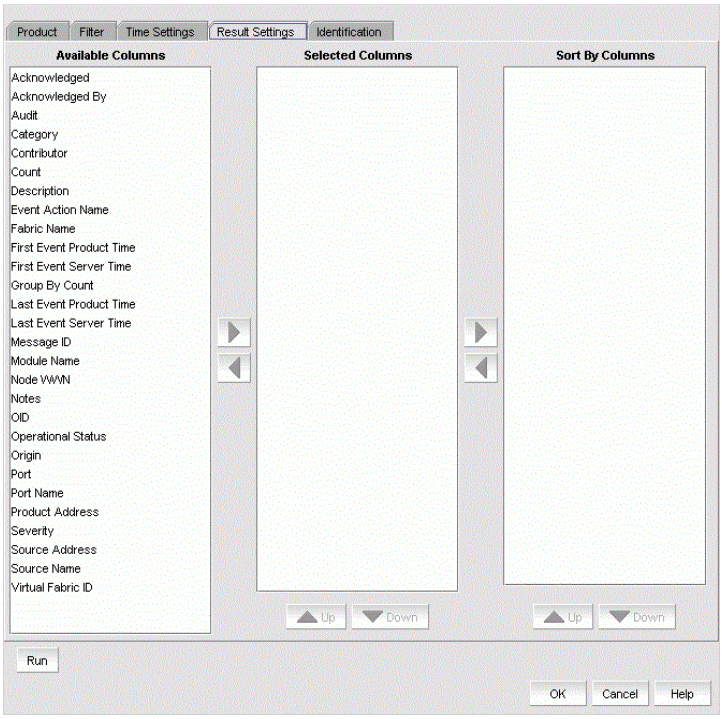

   NOTE
   The **Available Column** list lists the attributes you can include in the report. Each attribute represents a column on the report.

5. Select the attribute you want, then click the right arrow to move your selection to the **Selected Columns** list. To remove an attribute from the **Selected Columns** list, select the attribute that you want to remove, then click the left arrow button.
   - If you selected the **Count** column, the Management application adds the **First Seen** and **Last Seen** columns to a report.
   - For products that support stacking, the **Port** column shows the port.
6. Data for all attributes is sorted in ascending order and is sorted in the sequence that the attributes appear in the **Sort By Columns** list. In the **Selected Columns** list, select which attribute will be used to sort the generated report. Then click the right arrow button to move your selection to the **Sort by Columns** list. To remove an entry from the **Sort by Columns** list, select the entry, then click the left arrow button.
7. Click **OK** to save the definition, **Run** to launch the report, or click the **Identification** tab to display the parameters that you use to identify the definition.

Defining the report identity

The **Identification** tab in the **Event Custom Reports** dialog box allows you to enter the identity information of the report information.

To define the report identity, complete the following steps.

1. Select **Reports > Event Custom Reports**.

   The **Event Custom Reports** dialog box displays.

2. Click the **Add** button.
3. The **Add/Edit Report Definition** dialog box - **Product** tab displays.
4. Click the **Identification** tab.

   The **Add/Edit Report Definition dialog box - Identification tab, shown in** Figure 550, displays.

   FIGURE 550 Add/Edit Report Definition dialog box - Identification tab

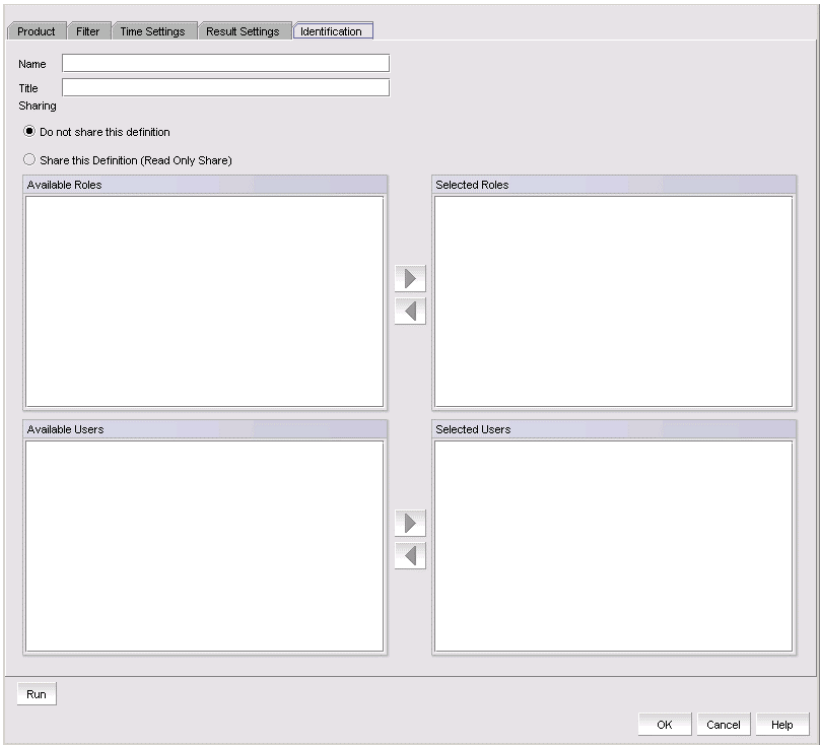

5. In the **Name** field, enter a name for the definition.

   This name appears under the **Name** column on the **Report Definitions** tab of the **Event Custom Reports** dialog box. This name must be unique for each report group. This is a required parameter.

6. In the **Title** field, enter a title for the definition, which will be used as the title of a generated report. This is a required parameter.
7. Click the **Do not share this definition** button if you do not want to share this definition with other Management application users.

   If you select this button, no Management application users will see this definition on the **Report Definitions** tab of the **Event Custom Reports** dialog box when they log in.

8. Click the **Share this definition (Read only)** button if you want other Management application users to have Read Only permission for this definition.

If you selected the **Share this definition (Read only)** button, a list of Management application roles appears in the **Available Roles** list.

9. Select the roles that will have view and run access to this definition, then press the right arrow button to move the role in the **Selected Roles** list.

   NOTE
   All Management application users who have the selected roles will be able to view, copy, and run the definition.

10. Select the roles that will have view and run access to this definition, then press the right arrow button to move the role in the **Selected Roles** list.

    All Management application users who have the selected roles will be able to view, copy, and run the definition.

    NOTE
    You can share the available users definition with specific Management application users. If you click the **Share this definition (Read only)** button, a list of Management application user accounts appears in the **Available Users** list.

11. Select the user account that will be able to view and run this definition, then press the right arrow button to move that user account in the **Selected Users** list.
12. Click **OK** to save the definition, or click **Run** to launch the report.

Filtering a report definition

You can filter a report definition. To do so, you must first enter a name and title on the **Identification** tab and select at least one column in the **Results Setting** tab to run or save a filter. You can select from the available list of SAN products, IP products, or hosts by selecting the appropriate tab.

NOTE
The swDeviceStatusTrap (OID 1.3.6.1.4.1.1588.2.1.1.0.15) trap is sent from the switch whenever there is a device login or logout. This trap is part of the SW-MIB and is listed under the SW-MIB of the **SNMP Trap Recipients** dialog box, the **Event Actions** dialog box, and the **SNMP Trap Forwarding** dialog box. For a complete list of event categories, refer to "Event Categories" on page 1327.

To filter a report definition, complete the following steps.

1. Select **Reports > Event Custom Reports**.

   The **Event Custom Reports** dialog box displays.

2. Click the **Add** button.
3. The **Add/Edit Report Definition** dialog box - **Product** tab, shown in Figure 551, displays.

FIGURE 551 Add/Edit Report Definition dialog box - Product tab

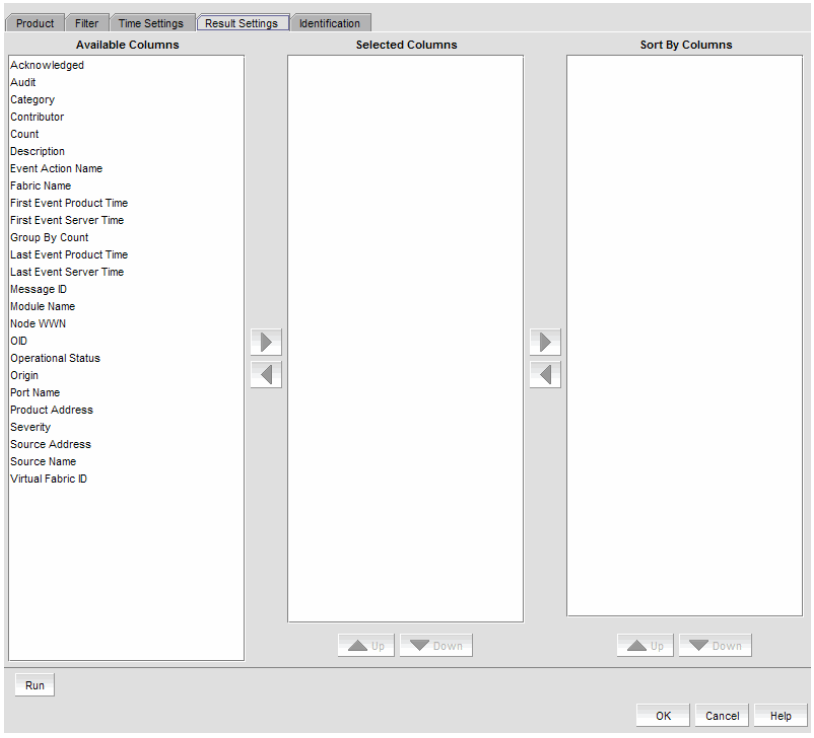

4. Click the **Filter** tab.

   The **Add/Edit Report Definition** dialog box - **Filter** tab, shown in Figure 552, displays.

FIGURE 552 Add/Edit Report Definition dialog box - Filter tab

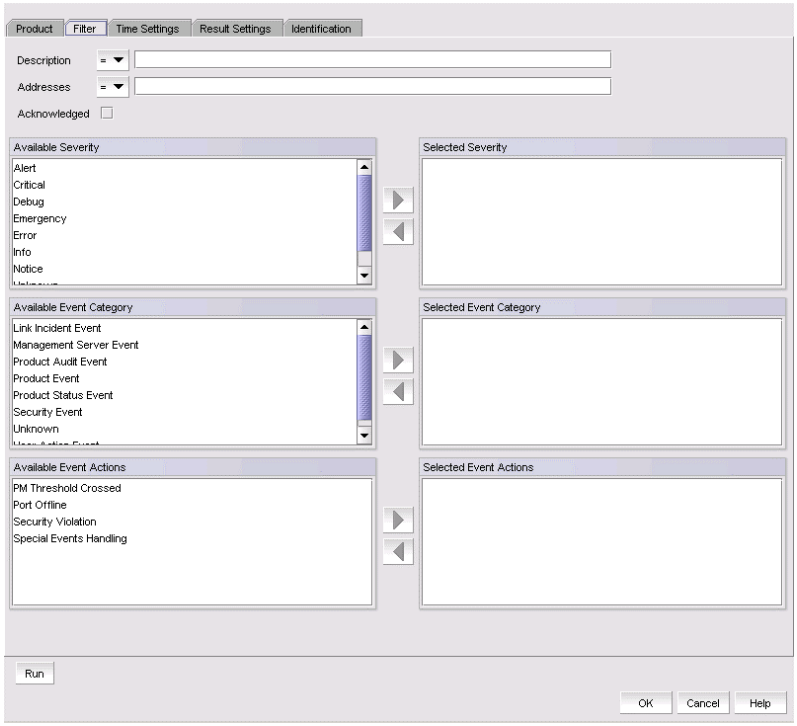

5. To limit the search results to traps, syslog, and pseudo event messages with a specific text string, enter the text string in the **Description** field.

   You can use an asterisk (*) to indicate a wildcard, as in the following examples:

   - *cdef: Matches a message ending with cdef
   - abc*: Matches a message beginning with abc
   - *abc*: Matches a message that contains abc

   For example, if you want to find the events that have the text "Auth" in the message, enter "*Auth*".

   NOTE
   You can view all port history events for a switch by creating an event custom report and entering a description of **Port Login/Logout History** for that particular switch. The Port Login/Logout history trap will be listed under the **Available traps** list of the **Add Trap Filter** dialog box and the **Add Event Action** dialog box — **Events** pane.

   For information about event categories, refer to "Event Categories" on page 1327.

6. To limit the search results to traps, syslog, and pseudo event messages from a specific IP address, enter the IP address or the AP MAC address in the **Address** field. You can enter multiple addresses. Separate each address with a comma.
7. Select the **Acknowledge** check box if you want messages that have been acknowledged to be included in the report.
8. Select the severity from the **Available Severity** list, and click the right arrow button to move your selection to the **Selected Severity** list. Events with the selected severity are included in the report.
9. Select the event type you want to include in the report from the **Available Event Category** list. Click the right arrow button to move your selection to the **Selected Event Category** list.
10. Select the event action you want to include in the report from the **Available Event Actions** list. Click the right arrow button to move your selection to the **Selected Event Actions** list.
11. Click **OK** to save the definition, **Run** to launch the report, or click the **Time Settings** tab on the **Add/Edit Report Definition** dialog box if you want to filter the events by date and time.

Filtering report events by date and time

The **Event Custom Reports dialog box — Time Settings** tab allows you to specify the time range of the events to be reported.

To filter report events by date and time, complete the following steps.

1. Select **Reports > Event Custom Reports**.

   The **Event Custom Reports** dialog box displays.

2. Click the **Add** button.

   The **Add/Edit Report Definition** dialog box - **Product** tab displays.

3. Click the **Time Settings** tab.

The Add/Edit Report Definition dialog box - Time Settings tab, shown in Figure 553, displays.

FIGURE 553 Add/Edit Report Definition dialog box - Time Settings tab

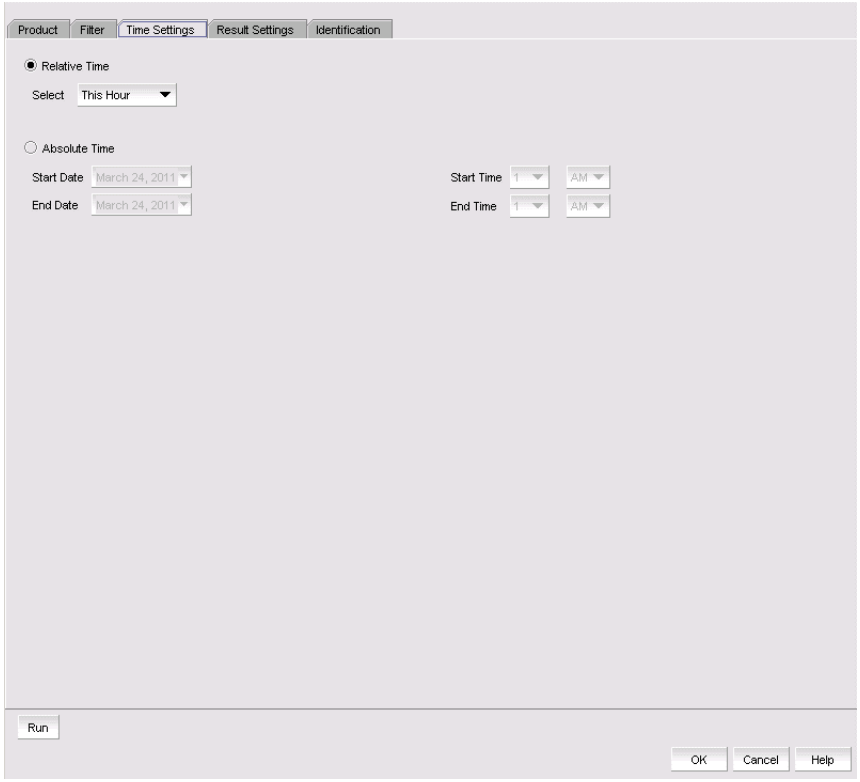

4. Choose between relative time (the default) and absolute time.
   - Click **Relative Time** if you want to filter traffic based on when the report is generated, and then select a relative time from the **Range** list. Relative time is calculated based on the date and time the report is generated.
   - Click **Absolute Time** if you want to filter traffic sent at a specific date and time.
     a. Select the specific start date from the **Start Date** list.
     b. Select the specific hour time for the start time from the **Start Time** list, and select AM or PM.
     c. Select the specific end date from the **End Date** list.
     d. Select the specific hour for the end time from the **End Time** list, and select AM or PM.
5. Click **OK** to save the definition, or click **Run** to launch the report.

Creating a new report definition by copying an existing definition

The simplest way to create a new report definition is by copying an existing definition.

To create a new report definition is by copying an existing definition, complete the following steps.

1. Select the definition you want to copy from the **Report Definitions** tab of the **Event Custom Reports** dialog box.
2. Click **Duplicate**.

   The name of the definition is the name of the selected definition with the word “copy” appended. For example, “SelectedPortName” becomes “SelectedPortName copy”.

3. Click the **Identification** tab to enter a new name and description for the new definition.
4. Make changes to the report as required.
5. Perform one of the following tasks when you are finished modifying the definition:
    - Click **OK** to save the report.
    - Click **Cancel** to discard your changes and exit from the **Report Definitions** tab of the **Event Custom Reports** dialog box.
    - Click **Reset** to discard your changes without exiting from the **Report Definitions** tab of the **Event Custom Reports** dialog box.
    - Click **Run** to launch the report.

The new definition is added to the **Report Definitions** tab of the **Event Custom Reports** dialog box.

Editing a report definition

For your definitions, you can modify a definition and save the changes you have made. For a shared definition from another user, you can modify the definition, then run that definition to obtain the desired report; however, you will not be able to save your changes.

To edit a report definition, complete the following steps.

1. Click the **Report Definitions** tab of the **Event Custom Reports** dialog box and select the definition you want to modify.
2. Click **Edit**.
3. When the **Add/Edit Report Definition** dialog box displays, modify the definition. (Refer to "Filtering a report definition" on page 1182.)
4. When you have finished, perform one of the following tasks:
    - If you own this definition, the **OK** button is available. Click **OK** to save your changes.
    - Click **Run** to generate the report.
    - Click **Cancel** to discard your changes and exit the **Report Definitions** tab of the **Event Custom Reports** dialog box.

Deleting a report definition

You can delete a report definition, but only if it belongs to you.

To delete a report definition, complete the following steps.

1. To access the dialog box, select **Reports > Event Custom Reports**.

    The **Event Custom Reports** dialog box displays.
2. Click the **Report Definitions** tab of the **Event Custom Reports** dialog box and select the definition you want to delete.
3. Click the **Delete** button.

    A message displays, prompting you to confirm the deletion.
4. Click **Yes** to delete the definition or **No** to cancel your request.

Event custom report schedules

Click the **Schedules** tab, shown in Figure 554, to display its contents. The **Schedules** list shows the definitions that have been scheduled to automatically run at a specified date and time.

FIGURE 554 Schedules tab of the Event Custom Reports dialog box

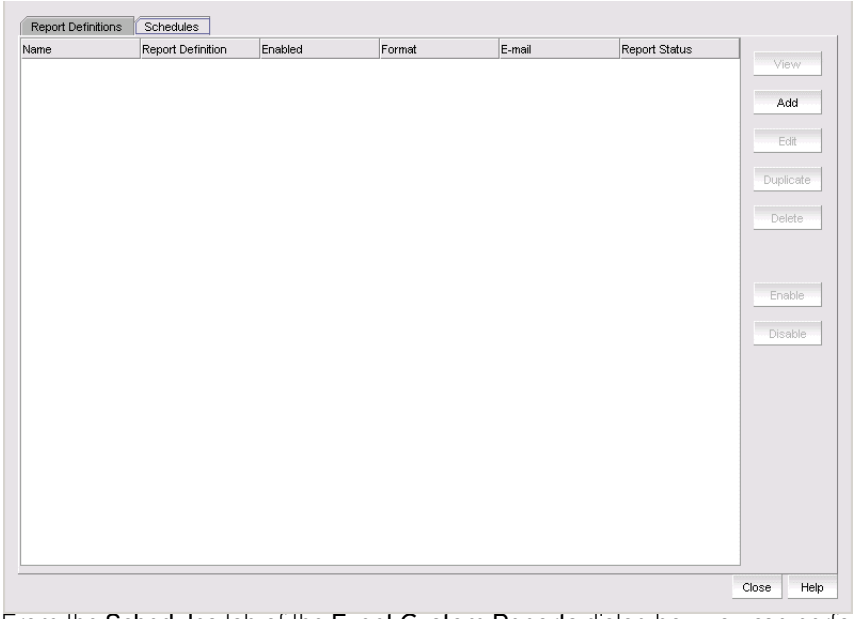

From the **Schedules** tab of the **Event Custom Reports** dialog box, you can perform the following tasks:

- **View —** Displays the report data of the scheduled report definition. The **View** button is not enabled for a report that is listed as Not Available.
- **Add** — Launches the **Add Schedule** dialog box.
- **Edit** — Launches the **Edit Schedule** dialog box with the selected schedule information pre-populated.
- **Duplicate** — Creates a copy of the selected report schedule.
- **Delete —** Deletes the selected schedule from the **Schedules** list.
- **Enable** — Enables the selected schedule.
- **Disable** — Disables the selected schedule.

Adding or editing an event report schedule

The **Add Schedule** dialog box, shown in Figure 555, allows you to select an existing report definition and configure the parameters, such as the schedule's format, frequency, recipients, and message content, for when the report is run and to whom the report is sent.

To add or edit an event report schedule, complete the following steps.

1. Select **Reports > Event Custom Reports**.

   The **Event Custom Reports** dialog box displays.

2. Click the **Schedules** tab.
3. Click the **Add** button.

   The **Add Schedule** dialog box displays.

FIGURE 555 Add Schedule dialog box

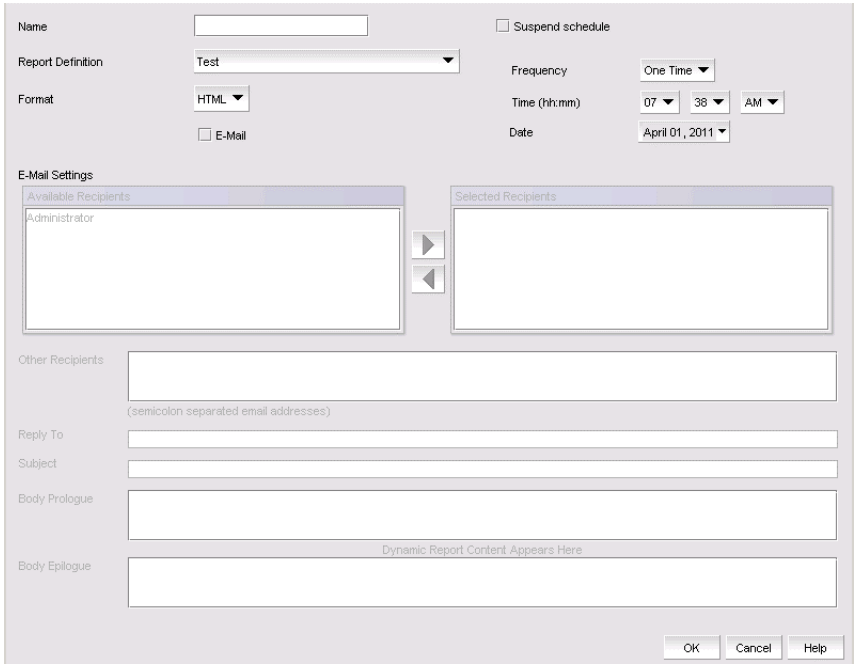

4. Enter the name of the new schedule in the **Name** field. You must enter a unique name for the schedule. The name can be up to 64 characters in length and it is case-sensitive.
5. Select the **Suspend schedule** check box if you want to disable the schedule. For example, you may want to temporarily prevent a report from being generated until further notice. You can clear the check mark to resume the automatic generation of the report.
6. Select the report definition you want to schedule from the **Report Definition** list. If a report is deleted, the corresponding schedule will be deleted.
7. Select one of the following periods from the **Frequency** list:
   - **One Time**
   - **Hourly** — If you selected **Hourly** as the schedule type, **Minutes past the hour** appears. Select the minutes after the hour when the report will be generated.
   - **Daily** — If you selected **Daily** as the schedule type, **Time (hh:mm)** appears.
   - **Weekly** — If you selected **Weekly** as the schedule type, **Day of the week** appears. Select the day of the week when the report will be generated.
   - **Monthly** — If you selected **Monthly** as the schedule type, **Day of the month** appears. Select the day of the month when the report will be generated.
   - **Yearly** — If you selected **Yearly** as the schedule type, **Day of the year** appears. Select the day of the year when the report will be generated.
8. Select a report format from the **Format** list: HTML or CSV.
9. Select the time when the report will be generated. Indicate the hour, minute, and whether it is AM or PM. This parameter appears if you selected any schedule type except **Hourly**.
10. Select the **E-mail** check box if you want the report to be sent to e-mail recipients. The server limits the displayed or sent report to 1000 records.

11. Change the value of the customReports.MaxRecordsToDisplay parameter in the configuration.properties file to the number of records you want displayed or sent.
12. Indicate the date when the report is generated. Open the calendar and select the date. This parameter appears if you selected **One Time** or **Yearly** as the schedule type.
13. Enter an e-mail address to which the e-mail recipient can send a response. The e-mail address is a mandatory field.
14. From the **Available Recipients** list, select the user to whom the report will be sent. Click the right arrow button to move that user name to the **Selected Recipients** list. Click the left arrow button to remove the name from the **Selected Recipients** list and return it to the **Available Recipients** list.

    NOTE
    Make sure an e-mail address is configured in the user's account for the selected user.

15. Enter other e-mail addresses to which the report should be sent in the **Other Recipients** field, separating multiple addresses with a semicolon. At least one e-mail address from the **Application Recipients** or **Other Recipients** must be entered.
16. In the **Reply To** field, enter an e-mail address to which the e-mail recipient can send a response. This is a mandatory field.
17. In the **Subject Line** field, enter the text that you want to appear in the subject line of the e-mail message. You can leave this field empty.
18. If you want introductory text to be included at the beginning of the e-mail message, enter the text in the **Body Prologue** field. The maximum number of characters supported by the **Body Prologue** field is 256.
19. If you want specific text to be included at the end of the e-mail message, enter that text in the **Body Epilogue** field. The maximum number of characters supported by the **Body Epilogue field** is 256.

Event logs

You can view all events that take place through the Master Log at the bottom of the main window. You can also view a specific log by selecting an option from the **Logs** submenu of the **Monitor** menu. The logs are described in the following list:

- **Audit Log —** Displays all Application Events raised by the application modules and all Audit Syslog messages from the switches and Brocade HBAs.
- **Product Event Log —** Displays all Product Event type events from all discovered switches and Brocade HBAs.
- **Fabric Log. (SAN only) —**Displays 'Product Events', 'Device Status', and 'Product Audit' type events for all discovered fabrics.
- **FICON Log —** Displays all the 'RLIR' and 'LRIR' type events, for example, 'link incident' type events.
- **Product Status Log** — (SAN only) Displays events which indicate a change in Switch Status for all discovered switches and Brocade HBAs.
- **Security Log —** Displays all security events for the discovered switches.
- **Syslog Log —** Displays syslog messages from switches and HBAs.

The Management application also has an event notification feature. By configuring event notification, you can specify when the application should alert you of an event. For details, refer to "Configuring e-mail notification" on page 1132.

For information about the Master Log interface, fields, and icons, refer to "Master Log" on page 300.

Viewing event logs

You can view log data through the Master Log on the main window. If you want to see only certain types of events; for example only security events, open a specific log through the **Logs** dialog box.

NOTE
You can also launch the Fabric logs and the Product Status logs from the status bar.

To view an event log, complete the following steps.

1. Select **Monitor > Logs >** *<Log_Type>*.

   The *<Log_Type>* **Logs** dialog box displays the type of log you selected.

2. Review the information in the log.
3. Click **Close**.

Viewing event logs with background color

You can view the event logs in the Master Log with different colors based on the event severity. For more information about enabling the background color, refer to "Highlighting events in the Master Log" on page 18.

Table 103 lists the color legend based on the event severity.

TABLE 103 Event severity color

Color	Severity
Red	Emergency, Alert, Critical, and Error
Yellow	Warning
No Color	Notice, Info, and Debug

Th color change will be reflected in the Events dialog launched from device context. Color change will be reflected in the MAPS drill down widgets launched from Dashboard. The color change will reflect on the following event logs.

- Audit Log
- Fabric Log
- FICON Log
- Master Log
- Special Events
- Product Event Log
- Product Status Log
- Security Log
- Syslog Log

NOTE
The event logs background color setting is applicable only for the current user.

NOTE
Changes do not take effect until after you restart the client.

Copying part of a log entry

You can copy data from logs to other applications. Use this method to analyze or store the data using another tool.

To copy part of an event log, complete the following steps.

1. Select **Monitor > Logs >** <*Log_Type*>.

   The <*Log_Type*> **Logs** dialog box displays the type of log you selected.

2. Select the rows you want to copy:

   - To select contiguous rows, select the first row you want to copy, press **Shift**, and click the contiguous row or rows you want to copy.
   - To select non-contiguous rows, select the first row you want to copy, press **CTRL**, and click the additional row or rows you want to copy.

3. Right-click one of the selected rows and select **Copy Rows**.
4. Open the application to which you want to paste the data.
5. Click where you want to paste the data.
6. Press **CTRL+V** (or select **Edit > Paste** from the other application).

   All data and column headings are pasted.

7. Click **Close** to close the dialog box.

Copying an entire log entry

You can copy data from logs to other applications. Use this method to analyze or store the data using another tool.

To copy an event log, complete the following steps.

1. Select **Monitor > Logs >** <*Log_Type*>.

   The <*Log_Type*> **Logs** dialog box displays the type of log you selected.

2. Right-click a row and select **Copy Table**.
3. Open the application to which you want to paste the data.
4. Click where you want to paste the data.
5. Press **CTRL+V** (or select **Edit > Paste** from the other application).

   All data and column headings are pasted.

6. Click **Close** to close the dialog box.

Exporting the entire log

You can export the log data to a tab-delimited text file.

To export an event log, complete the following steps.

1. Select **Monitor > Logs >** <*Log_Type*>.

   The <*Log_Type*> **Log** dialog box displays the type of log you selected.

2. Right-click a row and select **Export Table**.

   The **Save table to a tab delimited file** dialog box displays.

3. Browse to the location where you want to export the data.
4. Enter a name for the file in the **File Name** field.
5. Click **Save**.

   All data and column headings are exported to the text file.

6. Click **Close** to close the dialog box.

E-mailing all event details from the Master Log

NOTE
You must configure e-mail notification before you can e-mail event details from the Master Log. To configure e-mail notification, refer to "Configuring e-mail notification" on page 1132.

To e-mail all event details from the Master Log, complete the following steps.

1. Right-click an entry in the Master Log.
2. Select **E-mail > All.**

   The **E-mail** dialog box displays.

3. Enter the e-mail address of the person to receive the e-mail notifications in the **To** field.
4. Enter your e-mail address in the **From** field.
5. Click **OK**.

E-mailing selected event details from the Master Log

NOTE
You must configure e-mail notification before you can e-mail event details from the Master Log. To configure e-mail notification, refer to "Configuring e-mail notification" on page 1132.

To e-mail selected event details from the Master Log, complete the following steps.

1. Right-click the selected events in the Master Log.
2. Select the events that you want to e-mail.
3. Select **E-mail > Selection.**

   The **E-mail** dialog box displays.

4. Enter the e-mail address of the person to receive the e-mail notification in the **To** field.
5. Enter your e-mail address in the **Reply From** field.
6. Click **OK**.

Displaying event properties from the Master Log

You can view detailed information for an event.

To display event details from the Master Log, complete the following steps.

1. Right-click an entry in the Master Log.
2. Select **Properties**.

   The **Event Properties** dialog box displays.
3. Review the information.

TABLE 104 Event Properties

Event field	Description
Probable Cause	The most likely reason the event occurred.
Description	A description of the event.
Count	Number of times this event occurred on the host.
Origin	The event's origin, for example, SNMP trap.
Message ID	The message associated with the event.
Port Name	The port name associated with the event.
First Event Server Time	The time the event occurred.
Fabric Name	The VCS fabric name.
Product Address	The IP address of the product on which the event occurred.
Audit	Information regarding the audit.
Category	One of the following event categories, which are detailed in "Event Categories" on page 1327: • Product Event • Link Incident Event • Product Audit Event • Product Status Event • Security Event • User Action Event • Management Server Event
Last Event Product Time	The day, date, and time the last event occurred on the product.
Last Event Server Time	The day, date, and time the last event occurred on the server.
Severity	The event severity.
Source Name	The source of the event.
Virtual Fabric ID	The virtual fabric identifier.
Contributor	The contributor to this event.
Recommended Action	The recommended action to take to remedy the event.
First Event Product Time	The day, date, and time the first event occurred on the product.
Operational Status	The product's operational status.
Module Name	The module associated with the event.
Source Address	The IP address of the source.
Acknowledged	Indicates whether the event has been acknowledged.

TABLE 104 Event Properties

Event field	Description
Acknowledged by	Displays the information about the user who acknowledged the event.
Notes	An optional entry of information.

4. Click **Close** to close the **Event Properties** dialog box.

Finding the device associated with an event

You can locate a device on which an event was triggered by right-clicking the event and selecting **Locate**.

The device displays highlighted in the Product List and Topology Map.

NOTE
Locate does not locate devices in automatically collapsed fabrics. You must expand the fabric (right-click and select **Expand**) and repeat the locate command.

Copying part of the Master Log

You can copy data from logs to other applications. Use this method to analyze or store the data using another tool.

To copy part of the Master Log, complete the following steps.

1. Select the rows you want to copy in the Master Log:
   - To select contiguous rows, select the first row you want to copy, press **Shift**, and click the contiguous row or rows you want to copy.
   - To select non-contiguous rows, select the first row you want to copy, press **CTRL**, and click the additional row or rows you want to copy.
2. Right-click one of the selected rows and select **Table > Copy Rows**.
3. Open the application to which you want to paste the data.
4. Click where you want to paste the data.
5. Press **CTRL+V** (or select **Edit > Paste** from the other application). All data and column headings are pasted.

Copying the entire Master Log

You can copy the entire Master Log to other applications. Use this method to analyze or store the data using another tool.

To copy the entire Master Log, complete the following steps.

1. Right-click an entry in the Master Log.
2. Select **Table > Copy Table**.
3. Open the application to which you want to paste the data.
4. Click where you want to paste the data.
5. Press **CTRL+V** (or select **Edit > Paste** from the other application). All data and column headings are pasted.

Exporting the Master Log

You can export the Master Log to a tab-delimited text file. Use this method to analyze or store the data using another tool.

To export the Master Log, complete the following steps.

1. Right-click an entry in the Master Log.
2. Select **Table > Export Table**.

   The **Save table to a tab delimited file** dialog box displays.
3. Browse to the location where you want to export the data.
4. Enter a name for the file in the **File Name** field.
5. Click **Save**. All data and column headings are exported to the text file.
6. Click **Close** to close the dialog box.

Filtering events in the Master Log

You can filter the events that display in the Master Log on the main window. By default, all event types display in the **Selected Events** list.

When you select a filter from the **Filter** drop-down menu, the Master Log refreshes to display the events associated with that filter. This filter setting is kept when you exit the client.

For more information about the Master Log, refer to "Master Log" on page 300.

To filter events in the Master Log, complete the following steps.

1. Select the filter you want from the **Filter** drop-down menu at the top of the Master Log panel.

   FIGURE 556 Master Log Filter menu

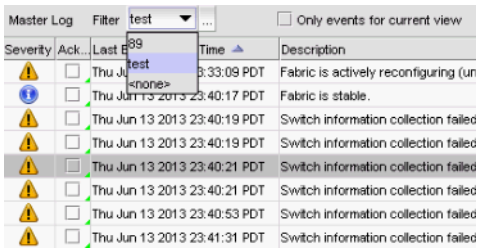

2. If you do not see the filter you want, click the **...** button immediately to the left of the menu.

   The **Define Filters** dialog box displays.

FIGURE 557 Define Filters dialog box - Basic tab, IP tab selected

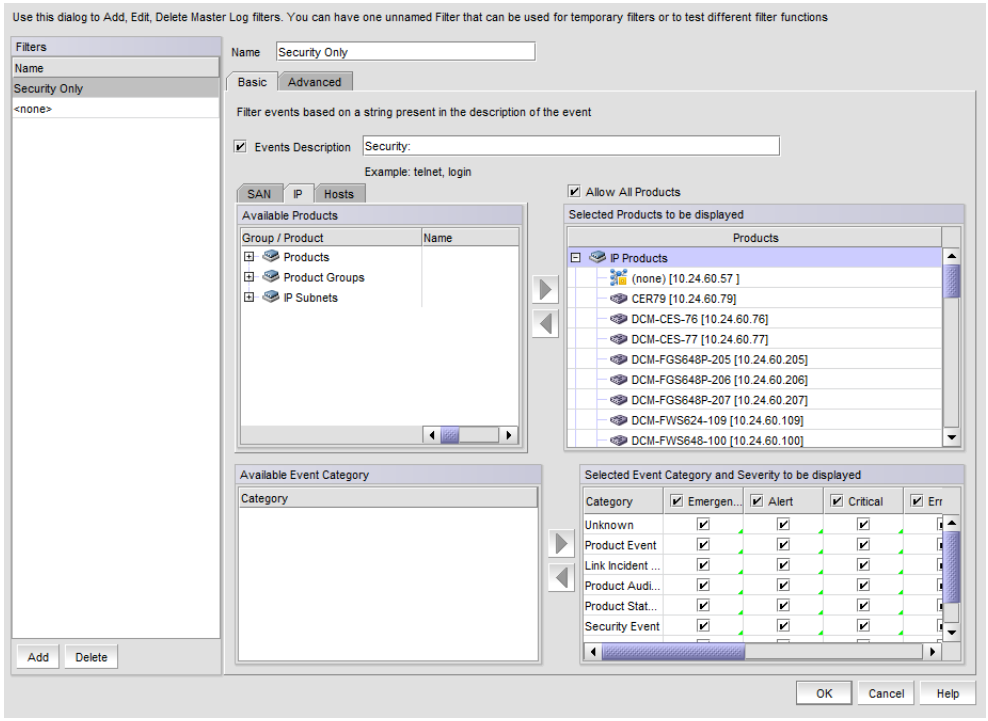

3. Use the following to include or exclude products.
   - To include an event type in the filter, select the event from the **Available Products list** and click the right arrow.
   - To exclude an event type from the filter, select the event from the **Selected Products to be displayed** list and click the left arrow.
   - To include all products, select the **Allow All Products** check box.
4. Select from the following to include or exclude event types.
   - To include an event type in the filter, select the event category from the **Available Event Category list** and click the right arrow.
   - To exclude an event type from the filter, select the event from the **Selected Event Category and Severity to be displayed** list and click the left arrow.
5. From the **Selected Event Category and Severity to be displayed** list, select one of the following severity levels to assigned to the selected event action:
   - Emergency
   - Alert
   - Critical
   - Errors
   - Warning
   - Notice
   - Info
   - Debug
   - Unknown

Clear the severity level check boxes to turn off the filter for the selected events.

6. (*Optional*) To filter events based on a string (such as telnet or login) that appears in the event description, select the **Events Description** check box and enter the string that the filter is to use in the associated text box.
7. Enter a name for the filter in the **Name** field. The Filter name length is limited to 128 alphanumeric characters. You cannot use other characters in this text box.
8. If you want to create multiple filters, click **Add** after you define the filter. This adds the defined filter to the Filters list, but does not close the dialog box.
9. When you have created all the filters you want, click **OK**.

   The **Define Filters** dialog closes and you are returned to the main window.

The “unnamed” filter

If a filter is migrated from a previous release, it is saved with the name **unnamed**. If a filter was not present in the release you are migrating from, then there will be no **unnamed** filter. If the **unnamed** filter was the default filter for you in the previous release, it will be set as the default filter for you in the current release.

The “none” filter

The filter named **none** is the default configuration filter. You cannot to edit or delete this filter. Selecting this filter lets you view Master Log events with no filtering applied. This is the default filter selected when no other filter is applied by the user or when there is no migrated filter.

Editing a filter

To edit a filter, select the filter you want to edit in the **Filters** panel and make the desired changes to the filter configuration. Any changes you make will be reflected in the **Filters** panel when navigating to another filter, but changes are not made permanent until you click **OK**.

Duplicating a filter

To duplicate a filter, select the filter you want to duplicate **Filters** panel and click **Add**. The content of the selected filter will be loaded, but with the **Name** field left blank. Enter a name for the new filter and click **OK**.

Deleting a filter

To delete a filter, select the filter and click **Delete**. Deleting a filter removes the filter name from the **Filters** panel of the **Define Filters** dialog box. A filter is not permanently deleted until you click **OK**.

Notes on filters

- Changing the filter in one client session does not alter the filter selection on other clients. However, if the currently selected filter is updated, once the filter is saved, the Master Log is reloaded to reflect the changes to that filter. This affects all your client sessions.
- If the currently selected filter is deleted, the Master Log is reloaded, and changes the selected filter to **none** for all your client sessions.
- Copying user preferences includes all user-created filters.
- For default filters, you are allowed to select a host from the **Available Products** list and include it in the **Selected Products to be displayed** list only if the host is managed by the Management application.

Adding notes while acknowledging or unacknowledging events in the Master Log

You can add a note and acknowledge an event in the Master Log. The notes entered by one user can be edited by another user assigned to the same AOR.

To add a note for an event in the Master Log, complete the following steps.

1. Select the **Notes** and **Acknowledges by** column in the Master Log. To add columns, refer to "Displaying columns" on page 309.
2. Right-click an entry in the Master Log and select the **Notes** from the menu.

   The **Add** / **Edit Note** dialog box (shown in Figure 558) displays.

   FIGURE 558 Add / Edit Note dialog box

   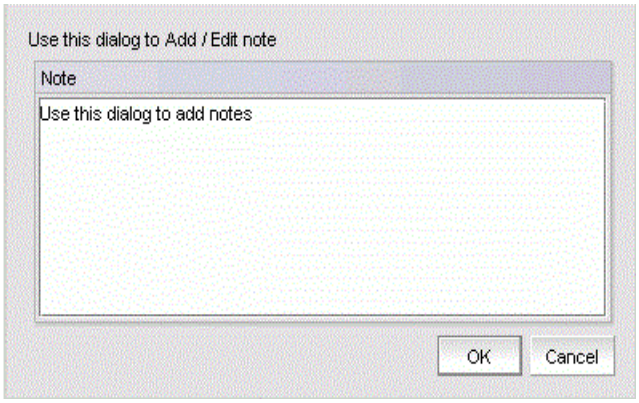

3. Enter the notes in the **Add** / **Edit Note** dialog box and click **OK**.

   The notes entered can be viewed in the **Notes** column. If you place your pointer over the **Notes** icon, a tool tip displays the recently added notes.

To add a note while acknowledging or unacknowledging an event in the Master Log, complete the following steps.

1. Select the **Acknowledge with notes** check box in the Master Log.
2. Select the **Acknowledge** check box to acknowledge an event or right-click an entry in the Master Log and select **Acknowledge/ Unacknowledge** from the menu.

   The **Add** / **Edit Note** dialog box displays.

3. Enter the notes in the **Add** / **Edit Note** dialog box and click **OK**.

   The notes entered are displayed in the **Notes** column and the "acknowledged-by" user information is displayed in the **Acknowledged by** column. If you place your pointer over the **Notes** icon, a tool tip displays the recently added notes. The full set of notes for an event can be viewed in the **Events Properties** dialog box from the Master Log. For more information, refer to "Displaying event properties from the Master Log" on page 1193. The notes added will be displayed when a user selects the same event to add a new note or append the existing note.

   NOTE
   If a single event is selected, you can add a note or append an existing note. If multiple events are selected, the existing note will not be displayed.

NOTE

You can select the **Notes** and **Acknowledged by** columns from the **Available columns** list in the **Add/Edit Report Definition** dialog box - **Result Settings** tab and generate an event custom report.

Monitoring and Alerting Policy Suite

In this chapter

Monitoring and Alerting Policy Suite overview

The Monitoring and Alerting Policy Suite (MAPS) is an optional storage area network (SAN) health monitor that allows you to enable each switch to constantly monitor its fabric for potential faults and automatically alerts you to problems long before they become costly failures.

MAPS tracks a variety of SAN fabric measures and events. Monitoring fabric-wide events, ports, and environmental parameters enables early fault detection and isolation, as well as performance measurement. You can configure fabric measures and alert thresholds on an individual port and group basis.

For Fabric OS devices, MAPS provides customizable monitoring thresholds. You can configure MAPS to provide notifications before problems arise, such as reporting when network traffic through a port is approaching the bandwidth limit. This information enables you to perform preemptive network maintenance, such as trunking or zoning, and avoid potential network failures.

For Fabric OS devices, MAPS enables you to define how often to check each switch and fabric measure and specify notification thresholds. Whenever fabric measures exceed these thresholds, MAPS automatically provides notification using several methods, including e-mail messages, SNMP traps, and log entries.

Supported hardware

MAPS is only supported on Fabric OS devices running Fabric OS 7.2.0 and later.

NOTE
MAPS is not supported on DCB devices.

MAPS license requirements

MAPS is supported on all versions of the Management application with SAN management.

MAPS is supported on Fabric OS devices running Fabric OS 7.2.0 or later with the Fabric Vision license or the combination of Fabric Watch and APM license. MAPS must be enabled on the device (refer to "Enabling MAPS on a device" on page 1203).

The Table 105 lists the MAPS supported and unsupported features for Fabric OS devices.

TABLE 105 Fabric OS Supported and Unsupported Features

Features	Fabric OS devices	Unlicensed MAPS (supported from Fabric OS 7.4.0 and later)
Enable MAPS	Yes	NA
Disable MAPS	No	No
Default Policies	Yes	Yes
Activate Default Policy	Yes	No
View Policy	Yes	Yes
Create Custom Policy	Yes	No
Activate Custom Policy	Yes	No
Edit Custom Policy	Yes	No
Delete Custom Policy	Yes	No
Distribute custom Policy	Yes	No
Import Custom Policies	Yes	No
Compare Policies	Yes	Yes
Manage Groups	Yes	No
View Violations	Yes	Yes
FPI Option	Yes	No
Export Policies	Yes	Yes
E-mail Setup	Yes	No
Actions	Yes	Yes
Receive MAPS violations occurred during the last 1 hour on switch discovery	Yes	Yes

MAPS role-based access control

NOTE
MAPS configuration requires read and write permissions to the MAPS Management privilege.

The Management application user accounts contain the identification of the Management application user, as well as privileges, roles, and areas of responsibility (AORs) assigned to the user. Privileges provide access to the features in the Management application. A role is a group of selected privileges. An AOR contains selected Ethernet fabrics, devices, and groups that a Management application user is allowed to manage.

By default, the SAN System Administrator and Network Administrator roles have read and write permissions to the MAPS Management privilege. The Operator role has read only permissions.

MAPS Management read permissions enable you to perform the following actions:

- View the **Out of Range Violations** and **Port Health Violations** widgets on the Dashboard.
- View MAPS violations.
- Access additional data from the MAPS-specific widgets.

- View existing MAPS policies.

For Fabric OS devices, in addition to the read actions, MAPS Management read and write permissions enable you to perform the following actions:

- Configure, edit, and delete user-defined MAPS policies.
- Activate policies on a device.
- Distribute MAPS policies from one device to another.
- Configure and edit user-defined MAPS rules.
- Configure, edit, and delete custom groups.

Enabling MAPS on a device

Beginning with Fabric OS 7.4.0, Fabric Watch is not supported for monitoring the switches. MAPS is explicitly enabled during the firmware upgrade to Fabric OS 7.4.0.

You can enable MAPS on one or more Fabric OS devices at the same time.

1. Select the **SAN** tab and select **Monitor > Fabric Vision > MAPS > Enable**.

   The **Enable MAPS** dialog box displays (Figure 559).

   FIGURE 559 Enable MAPS dialog box for SAN

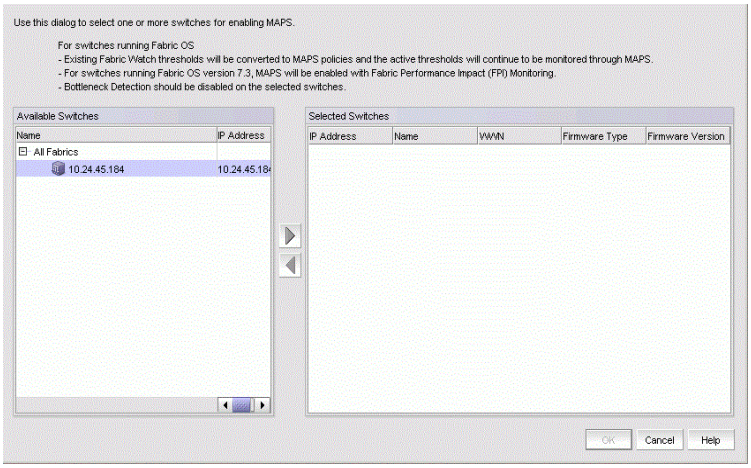

2. Select one or more devices under **All Fabrics** on which you want to enable MAPS in the **Available Switches** list.
3. Click the right arrow button to move the selected devices to the **Selected Switches** list.

   Remove switches from the **Selected Switches** list by selecting them and clicking the left arrow button.

4. Click **OK** on the **Enable MAPS** dialog box.

   A 'warning' confirmation message displays.

5. Click **Yes** on the confirmation message to commit the changes to the selected devices.

MAPS interoperability with other features

Virtual Fabrics

MAPS is a logical switch-specific feature. Different logical switches can have different MAPS configurations for the needs of the specific logical switch.

When you enable MAPS on the Virtual Fabric-enabled switch, MAPS is enabled, with the same active policy, on all Fabric Identifiers (FIDs).

Configuration upload and download

MAPS configuration is stored in separate configuration files. The default MAPS configuration is stored in one configuration file.

The user-created configuration is stored in another configuration file. One user configuration file exists for each logical switch.

You cannot upload and download the default MAPS configuration file. A configuration upload or download affects only the user-created configuration files.

High availability

MAPS configuration is maintained across a failover or reboot; however, MAPS statistics collected are lost.

Admin Domains

MAPS is not supported with Admin Domains. If MAPS is enabled, you cannot create Admin Domains. If user-created Admin Domains are present on the switch, migration to MAPS fails.

If MAPS is enabled, make sure you do not download configuration files that have Admin Domains defined.

Fabric Watch

MAPS cannot coexist with Fabric Watch. For Fabric OS switches running 7.4.0 or later, Fabric Watch is disabled and MAPS support is enabled by default.

Fabric Watch behavior

The Management application provides a launch point to Fabric Watch configuration (**Monitor > Fabric Watch > Configure**). In addition to launching Fabric Watch, the Management application allows certain Fabric Watch configurations through Port Fencing, Frame Monitoring, and performance thresholds and allows replication of Fabric Watch configurations. Also, note that some features require the Fabric Watch license to work (such as port optics and Call Home).

Once you upgrade a switch to Fabric OS 7.2.0 or later and enable MAPS, Fabric Watch configuration and any Fabric Watch-related features are no longer supported.

- Launch Fabric Watch — A “None of the Fabric Watch specific operations can be performed on this switch because the MAPS (Monitoring and Alerting Policy Suite) are enabled.” error message displays.
- Replicate Fabric Watch configuration — If you select a Fabric Watch configuration to replicate, the Management application filters the MAPS enabled switches from **Source Configuration** and **Destination Switches** steps of the replicate configuration wizard.

- Port Fencing — Only displays switches that do not have MAPS enabled. Depending on your discovered devices, displays one of the following error messages:
  - "None of the Fabric Watch specific operations can be performed on this fabric because the MAPS (Monitoring and Alerting Policy Suite) is enabled on all the switches."
  - "Port Fencing cannot be configured on one or more switches in this fabric because MAPS is enabled on them. Do you want to configure Port Fencing on the remaining switches?"
- Frame Monitor — Only displays switches that do not have MAPS enabled. A "None of the Fabric Watch specific operations can be performed on this switch because the MAPS (Monitoring and Alerting Policy Suite) are enabled." error message displays.
- Performance Thresholds — Only displays switches that do not have MAPS enabled. A "None of the Fabric Watch specific operations can be performed on this switch because the MAPS (Monitoring and Alerting Policy Suite) are enabled." error message displays.
- Port Optics — Displays port optics details. For Fabric OS products running 7.2.0 or later, displays combined status and allows threshold-based monitoring.
- Call Home — Does not generate Fabric Watch events from MAPS enabled switches. Therefore, does not generate Call Home notification for existing Fabric Watch events. You must configure the new MAPS Call Home event (MAPS-1021) to receive the Call Home message. the Management application running 14.0.0 or later does not list failed FRU components for MAPS Call Home event (MAPS-1021).
- Fabric Watch dashboard support — The MAPS dashboard widgets display the number of MAPS threshold violations for all network objects (such as ports, trunks, switches, and circuits) for all MAPS-capable devices. In addition, the MAPS dashboard widgets include the Fabric Watch threshold violations for devices with the Fabric Watch license or FC devices running Fabric OS 7.2.0 or later with the Fabric Vision license but not migrated to MAPS.

  The Fabric Watch support only applies to the following widgets and dialog box:

  - **Out of Range Violations** widget (refer to "Out of Range Violations widget" on page 229)
  - **Port Health Violations** widget (refer to "Port Health Violations widget" on page 231)
  - **Violations** dialog box (refer to "Viewing MAPS violations" on page 1255)

  Table 106 details the supported RAS log event identifiers that display in the MAPS dashboard widgets and the **Violations** dialog box.

TABLE 106 Fabric Watch supported RAS event IDs

Category	Measure	Unit label	RAS event ID
Port Health	CRC — CRC errors	Count	1182
	ITW — Invalid transmit words	Count	1178
	LOSS_SYNC — Loss of synchronization	Count	1166
	LF — Link failure	Count	1162
	LOSS_SIGNAL — Signal loss	Count	1170
	PE — Protocol errors	Count	1174
	LR — Link reset	Count	1198
	C3TXTO — Class 3 timeout	Count	1202
	STATE_CHG — State changes	Count	1194
	CURRENT — SFP transceiver current	mA	1046 1047
	RXP — SFP transceiver receive power	microWatts	1038 1039
	TXP — SFP transceiver transmit power	microWatts	1042 1043
	VOLTAGE — SFP transceiver voltage	mV	1050 1051
	SFP_TEMP — SFP transceiver temperature	Degrees celsius	1034 1035
	PWR_HRS[1] — SFP transceiver power on hours	Count	1053 1054
Switch Policy Status	BAD_PWR — Absent or faulty power supply	Count	1426 1427 1428 1429
	BAD_TEMP — Temperature sensors outside range	Count	1430
	BAD_FAN — Absent or faulty fans	Count	1431
	FLASH_USAGE — Flash usage	%	1435
	MARG_PORTS — Percentage of marginal ports	%	1436
	FAULTY_PORTS — Percentage of faulty ports	%	1437
	MISSING_SFP — Percentage of missing SFP transceivers	%	1438
	ERR_PORTS — Percentage of error ports	%	1448
	WWN_DOWN — WWN card faulty or down	Count	1432
	DOWN_Core — Core blade monitoring	Count	1447
	FAULTY_BLADE — Faulty blades	Count	1434
	HA_SYNC — HA monitoring	Count	1433
	FAN_AIR_FLOW_MISMATCH — Fan airflow mismatch	N/A	N/A

TABLE 106 Fabric Watch supported RAS event IDs (Continued)

Category	Measure	Unit label	RAS event ID
Fabric State Changes	DID_CHG — Domain ID change	Count	1123
	FLOGI — Fabric login	Count	1135
	FAB_CFG — Fabric reconfigurations	Count	1119
	EPORT_DOWN — E_Ports down	Count	1115
	FAB_SEG — Fabric segmentation	Count	1127
	ZONE_CHG — Zone changes	Count	1131
	L2_DEVCNT_PER — Layer 2 device count	Count	N/A
	LSAN_DEVCNT_PER — LSAN device count	Count	N/A
	ZONE_CFGSZ_PER — Zone configuration size	Count	N/A
	BB_FCR_CNT — FCR count	Count	N/A
FRU Health	PS_STATE — Power supply state changes	N/A	N/A
	FAN_STATE — Fan state changes	N/A	1440 1441 1442 1443 1444
	BLADE_STATE — Blade state changes	N/A	1440 1441 1442 1443 1444
	SFP_STATE — SFP transceiver state changes	N/A	1337
	WWN — WWN card state changes	N/A	1440 1441 1442 1443 1444

TABLE 106 Fabric Watch supported RAS event IDs (Continued)

Category	Measure	Unit label	RAS event ID
Security Health	SEC_DCC — Device connection control violations	Count	1338
	SEC_HTTP — HTTP violations	Count	1302
	SEC_CMD — Illegal command violations	Count	1378
	SEC_IDB — Incompatible security DB	Count	1374
	SEC_LV — Login violations	Count	1342
	SEC_CERT — Invalid certifications	Count	1354
	SEC_FCS — Primary Fabric Configuration Server (FCS) contact losses	Count	1370
	SEC_SCC — Switch connection control violations	Count	1334
	SEC_AUTH_FAIL: — Packet authentication failures	Count	1358
	SEC_TELNET — Telnet violations	Count	1298
	SEC_TS — Time server out of synchronization	Count	1366
	DAYS_TO_EXPIRE — Number of days before certificate expiry	N/A	N/A
	EXPIRED_CERTS — Number of expired certificate greater than threshold	N/A	N/A
Switch Resources	TEMP — Temperature sensor	N/A	1002 1003 1004
	FLASH_USAGE — Flash usage	%	1402
	CPU — CPU usage	%	1404
	MEMORY_USAGE — Memory usage	%	1404 1406
	ETH_MGMT_PORT_STATE — Ethernet management port state	%	N/A
FCIP	CIR_STATE — FCIP circuit state changes	N/A	3020
	CIR_UTIL — FCIP circuit utilization	%	3012
	CIR_PKTLOSS — FCIP packet loss	%	3016
	IP_UTIL — IP circuit utilization	N/A	N/A
	IP_PKTLOSS — IP packet loss	N/A	N/A
	IP_RTT — IP round-trip time of the circuit	N/A	N/A
	IP_JITTER — IP variance in round-trip time of the circuit	N/A	N/A

TABLE 106 Fabric Watch supported RAS event IDs (Continued)

Category	Measure	Unit label	RAS event ID
Traffic Performance	RX — Receive bandwidth usage percentage	%	1186
	TX — Transmit bandwidth usage percentage	%	1190
	UTIL — Trunk utilization	%	1206
	TX_FCNT — Tx Frame Count	Count	N/A
	RX_FCNT — Rx Frame Count	Count	N/A
	TX_THPUT — Tx Throughput	Count	N/A
	RX_THPUT — Rx Throughput	Count	N/A
	IO_RD — IO Read Command Count	Count	N/A
	IO_WR — IO Write Command Count	Count	N/A
	IO_RD_BYTES — IO Read Data	mbps	N/A
	IO_WR__BYTES — IO Write Data	mbps	N/A
	RD_PENDING_IO — IO Pending Read Data	mbps	N/A
	WR_PENDING_IO — IO Pending Write Data	mbps	N/A
	RD_STATUS_TIME — Read Completion Time	mbps	N/A
	WR_STATUS_TIME — Write Completion Time	mbps	N/A
	RD_1stDATA_TIME — First Read Response Time	mbps	N/A
	WR_1stXFER_RDY — First Write Response Time	mbps	N/A

1. Only valid for 16 Gbps, 10 Gbps, and QSFP transceivers.

MAPS category, object, and measure hierarchy

Fabric measures and events are organized in a hierarchy by category, object, and measure. There is a category, object, and measure associated with every monitored behavior. Categories are the highest level in the system, subdivided into one or more objects. Objects contain one or more measures.

For a list of all categories, objects, and measures, refer to "MAPS categories, measures, and actions" on page 1210.

MAPS structural elements

MAPS has the following structural elements: categories, groups, rules, and policies. Table 107 provides a brief description of each structural element.

TABLE 107 MAPS structural elements

Element	Description
Action	What MAPS is to do if a condition defined in a rule evaluates to true. For more information, refer to "MAPS actions" on page 1224.
Category	A grouping of similar elements that can be monitored (for example, "Security Violations"). For more information, refer to "MAPS monitoring categories" on page 1213.
Condition	A true/false trigger created by the combination of a time base and a threshold value. For more information, refer to "MAPS conditions" on page 1223.
Measure	A value that can be monitored. This includes switch conditions, data traffic levels, error messages, and many other values. For more information, refer to the specific category in "MAPS monitoring categories" on page 1213.

TABLE 107 MAPS structural elements

Element	Description
Object	An object that can be monitored. This includes FC ports, SFP transceivers, local switch, flow, and other values. For more information, refer to "MAPS categories, measures, and actions" on page 1210.
Group	A collection of similar objects that you can monitor as a single entity. For example, you can create a group of E_Ports from different devices to be monitored by the same policies. A flow can be imported as a group. For more information, refer to "MAPS groups" on page 1245.
Rule	A rule associates a condition with one or more actions that need to occur when the specified condition is triggered. For more information, refer to "MAPS rules" on page 1223.
Policy	A set of rules defining thresholds for triggering the actions MAPS takes when that threshold is triggered. When a policy is enabled, all of the rules in the policy are in effect. For more information, refer to "MAPS policies" on page 1221.

MAPS categories, measures, and actions

Table 108 details the Fabric OS object types for each category, the threshold measures for each object type, and the action you can configure when a threshold is crossed.

TABLE 108 Monitors a nd actions by category for Fabric OS

Category	Objects	Measures	Possible actions
Port Health	FC Port	• CRC — CRC errors • ITW — Invalid transmit words • LOSS_SYNC — Loss of synchronization • LF — Link failure • LOSS_SIGNAL — Signal loss • PE — Protocol errors • LR — Link reset • C3TXTO — Class 3 timeout • STATE_CHG — State changes • DEV_NPIV_LOGINS — Device NPIV Logins	• RAS Log Event • Port Decommission • Fence • SNMP Trap • E-mail • FMS
	SFP transceiver	• CURRENT — SFP transceiver current • RXP — SFP transceiver receive power • TXP — SFP transceiver transmit power • VOLTAGE — SFP transceiver voltage • SFP_TEMP — SFP transceiver temperature • PWR_HRS[1] — SFP transceiver power on hours	• RAS Log Event • SNMP Trap • E-mail • SFP Marginal • FMS
Switch Policy Status	Chassis	• BAD_PWR — Absent or faulty power supply • BAD_TEMP — Temperature sensors outside range • BAD_FAN — Absent or faulty fans • FLASH_USAGE[2] — Flash usage • WWN_DOWN — WWN faulty or down • DOWN_Core — Core blade monitoring • FAULTY_BLADE — Faulty blades • HA_SYNC — HA monitoring	• Status Critical • Status Marginal • FMS
	Local Switch	• MARG_PORTS — Percentage of marginal ports • FAULTY_PORTS — Percentage of faulty ports • MISSING_SFP — Percentage of missing SFP transceivers • ERR_PORTS — Percentage of error ports	

TABLE 108 Monitors a nd actions by category for Fabric OS (Continued)

Category	Objects	Measures	Possible actions
Fabric State Changes	Local Switch	• DID_CHG — Domain ID change • FLOGI — Fabric login • FAB_CFG — Fabric reconfigurations • EPORT_DOWN — E_Ports down • FAB_SEG — Fabric segmentation • ZONE_CHG — Zone changes • L2_DEVCNT_PER — Layer 2 device count • LSAN_DEVCNT_PER — LSAN device count • ZONE_CFGSZ_PER — Zone configuration size • BB_FCR_CNT — FCR count	• RAS Log Event • SNMP Trap • E-mail • FMS
FRU Health	Power Supply	• PS_STATE — Power supply state	• RAS Log Event • SNMP Trap • E-mail • FMS
	Fan	• FAN_STATE — Fan state	
	Blade	• BLADE_STATE — Blade state	
	SFP transceiver	• SFP_STATE — SFP transceiver state	
	WWN	• WWN — World Wide Name state	
Security Health	Local Switch	• SEC_DCC — Device connection control violations • SEC_HTTP — HTTP violations • SEC_CMD — Illegal command • SEC_IDB — Incompatible security DB • SEC_LV — Login violations • SEC_CERT — Invalid certifications • SEC_FCS — No Fabric Configuration Server (FCS) switch • SEC_SCC — Switch Connection Control violations • SEC_AUTH_FAIL: — Authentication failures • SEC_TELNET — Telnet violations • SEC_TS — Time server out of synchronization • DAYS_TO_EXPIRE — Number of days before certificate expiry • EXPIRED_CERTS — Number of expired certificate greater than threshold	• RAS Log Event • SNMP Trap • E-mail • FMS
Switch Resources	Temperature sensor	• TEMP — Temperature	• RAS Log Event • SNMP Trap • E-mail • FMS •
	Local Chassis	• FLASH_USAGE[2] — Flash usage • CPU — CPU usage • MEMORY_USAGE — Memory usage • ETH_MGMT_PORT_STATE — Ethernet management port state	• RAS Log Event • SNMP Trap • E-mail • FMS •
	Fan	• FAN_AIRFLOW_MISMATCH — Fan airflow direction mismatch.	• RAS Log Event • SNMP Trap • E-mail • FMS • SW_MARGINAL • SW_CRITICAL •

TABLE 108 Monitors a nd actions by category for Fabric OS (Continued)

Category	Objects	Measures	Possible actions
FCIP	Circuit	• CIR_STATE — FCIP circuit state changes • CIR_UTIL — FCIP circuit utilization • CIR_PKTLOSS — FCIP packet loss • PKTLOSS — FCIP tunnel packet loss[3] • RTT — Round-trip time of the circuit • Jitter — The variance in round-trip time of the circuit • TUNNEL_UTIL — Tunnel utilization • TUNNEL_STATE — FCIP tunnel status • IP_UTIL — IP circuit utilization • IP_PKTLOSS — IP packet loss • IP_RTT — IP round-trip time of the circuit • IP_JITTER — IP variance in round-trip	• RAS Log Event • Fence (CIR_STATE) • SNMP Trap • E-mail • FMS
Traffic /Flows Performance	Flow	• TX_FCNT — Tx Frame Count • RX_FCNT — Rx Frame Count • TX_THPUT — Tx Throughput • RX_THPUT — Rx Throughput • IO_RD — IO Read Command Count • IO_WR — IO Write Command Count • IO_RD_BYTES — IO Read Data • IO_WR__BYTES — IO Write Data • RD_PENDING_IO — IO Pending Read Data • WR_PENDING_IO — IO Pending Write Data • RD_STATUS_TIME — Read Completion Time • WR_STATUS_TIME — Write Completion Time • RD_1stDATA_TIME — First Read Response Time • WR_1stXFER_RDY — First Write Response Time	• RAS Log Event • SNMP Trap • E-mail • FMS
FPI FC Port		• DEV_LATENCY_IMPACT — Device latency impact • RX — Receive bandwidth usage percentage • TX — Transmit bandwidth usage percentage • UTIL — Trunk utilization	• RAS Log Event • SNMP Trap • E-mail • FMS • SDDQ • Un-Quarantine • Toggle
		• ALL_LOCAL_PIDS - Initiator to Target ratio	• RAS Log Event • SNMP Trap • E-mail • FMS
		• IT_FLOWS - IT Flow ratio	• RAS Log Event • SNMP Trap • E-mail
GigE Port[4]	GigE Port	• GE_CRC — Frames received with CRC error. • GE_LOS_OF_SIG — Frames aborted because of carrier sense error, no carrier, or loss of carrier. • GE_INV_LEN — Frames received with length error when Length_Type field does not match frame size.	• RAS Log Event • SNMP Trap • E-mail
Backend Port Monitoring	Backend Port	• LR — Link reset • BAD_OS — Invalid ordered set • CRC — CRC errors • ITW — Invalid transmit words • FRM_TRUNC — Frames shorter than minimum • FRM_LONG — Frames longer than maximum	• RAS Log Event • SNMP Trap • E-mail

1. Only valid for 16 Gbps, 10 Gbps, and QSFP transceivers.
2. For FLASH_USAGE, you can configure RAS Log Event, Fence, SNMP Trap, E-mail, Switch Status Critical, or Switch Status Marginal.

3. Supported only on the Fabric OS 16 Gbps 24-FC port, 18 GbE port switch .

4. GigeE Port is supported only on FX8-24, 32 Gbps, Router Extension Blade, and 16 Gbps 24-FC port, 18 GbE port switch.

MAPS monitoring categories

MAPS enables you to monitor the independent components that are listed in this section by creating policies. Policies are a series of rules that define thresholds for measures and actions to take when a threshold is triggered.

In version 14.2.0, MAPS monitors the front end ports on switches running Fabric OS 8.1.0 and above. MAPS alerts when a port counter error occur in the encryption port that are reported from the ASIC.

Port monitoring category

The Port category monitors port statistics and takes action based on the configured thresholds and actions. You can configure thresholds per port type and apply the configuration to all ports of the specified type. Configurable ports include physical ports, E_Ports, optical F_Ports (FOP_Ports), copper F_Ports (FCU_Ports), and Virtual E_Ports (VE_Ports).

The Port category also monitors the physical aspects of an SFP transceiver, such as voltage, current, RXP, TXP, and state changes in physical ports, E_Ports, FOP_Ports, and FCU_Ports.

In Fabric OS 8.1.0, the following QSFP rule and groups are not supported in the **Add** and **Edit** policy dialogs and the default policies will not display default rules with QSFP groups:

- All_QSFP rules.
- ALL_32GSWL_QSFP
- ALL_QSFP
- ALL_100M_16GSWL_QSFP

Table 109 lists measures in the Port category and describes each measure.

TABLE 109 Port measures

Measure	Description
Cyclic redundancy check (CRC)	The number of times an invalid cyclic redundancy check error occurs on a port or a frame that computes to an invalid CRC. Invalid CRCs can represent noise on the network. Such frames are recoverable by retransmission. Invalid CRCs can indicate a potential hardware problem.
Invalid transmission words (ITW)	The number of times an invalid transmission word error occurs on a port. A word did not transmit successfully, resulting in encoding errors. Invalid word messages usually indicate a hardware problem.
Sync loss (LOSS_SYNC)	The number of times a synchronization error occurs on the port. Two devices failed to communicate at the same speed. Synchronization errors are always accompanied by a link failure. Loss of synchronization errors frequently occur due to a faulty SFP transceiver or cable.
Link failure (LF)	The number of times a link failure occurs on a port or the port sends or receives a Not Operational Sequence (NOS) state. Both physical and hardware problems can cause link failures. Link failures also frequently occur due to a loss of synchronization or a loss of signal.
Signal loss (LOSS_SIGNAL)	The number of times that a signal loss occurs in a port. Signal loss indicates that no data is moving through the port. A loss of signal usually indicates a hardware problem.
Protocol error (PE)	The number of times a protocol error occurs on a port. Invalid state due to link reset response sequence (LRR) on an online link. Occasionally these errors occur due to software glitches. Persistent errors occur due to hardware problems.
Link reset (LR)	The ports on which the number of link resets exceed the specified threshold value.

TABLE 109 Port measures (Continued)

Measure	Description
Class 3 timeouts (C3TXTO)	The number of class 3 frames discarded due of timeouts.
State changes (STATE_CHG)	The state of the port has changed for one of the following reasons: • The port has gone offline. • The port has come online. • The port is faulty.
SFP Current	The amount of supplied current to the SFP transceiver. Current area events indicate hardware failures.
SFP Receive power (RXP)	The amount of incoming laser, in microwatts (?w). This is used to help determine if the SFP transceiver is in good working condition. If the counter often exceeds the threshold, the SFP transceiver is deteriorating.
SFP Transmit power (TXP)	The amount of outgoing laser, in microwatts (?w). This is used to help determine if the SFP transceiver is in good working condition. If the counter often exceeds the threshold, the SFP transceiver is deteriorating.
SFP Voltage	The amount of voltage supplied to the SFP transceiver. If this value exceeds the threshold, the SFP transceiver is deteriorating.
SFP Temperature (SFP_TEMP)	The physical temperature of the SFP transceiver, in degrees Celsius. A high temperature indicates that the SFP transceiver might be in danger of damage.
RX_ABN_FRAME	Missing termination character.
SFP Power on hours (PWR_HRS)	The number of hours the 16 Gbps SFP transceiver is powered on.
Device NPIV Logins	Monitors the number of logins on F_Ports in a switch.

Switch Status monitoring category

The Switch Status category enables you to monitor the health of the switch by defining the number of types of errors that transitions the overall switch state into a state that is not healthy. For example, you can specify a switch policy so that if a switch has two port failures, it is considered to be in a marginal state; if it has four failures, it is in a down state.

In Fabric OS 8.1.0, the FAN_ARIFLOW_MISMATCH measure is not supported. Default rules are supported to monitor ERR_PORTS measure.

Table 110 lists the current overall Switch Status category parameters in a switch and identifies the factors that affect their health. Note that not all switches use the listed monitors.

TABLE 110 Switch status measures

Measure	Description
Power Supplies (BAD_PWR)	Power supply thresholds detect absent or failed power supplies, and power supplies that are not in the correct slot for redundancy.
Temperatures (BAD_TEMP)	Temperature thresholds, faulty temperature sensors.
Fans (BAD_FAN)	Fan thresholds, faulty fans.
WWN (WWN_DOWN)	Faulty WWN card (applies to modular switches).
FAN_AIRFLOW_MISMATCH	Fan airflow direction.
HA out of sync (HA_SYNC)	High availability (HA) state of the active CP is out of synchronization with the HA state of the standby CP.
Blades (FAULTY_BLADE)	Faulty blades (applies to modular switches).
Core Blade (DOWN_CORE)	Faulty core blades.
Flash (FLASH_USAGE)	Flash thresholds.
Marginal Ports[1] (MARG_PORTS)	Physical port, E_Port, FOP_port (optical), and FCU_Port (copper) port thresholds. Whenever these thresholds are persistently high, the port is Marginal.

TABLE 110 Switch status measures (Continued)

Measure	Description
Faulty Ports[1] (FAULTY_PORTS)	Hardware-related port faults.
Missing SFPs[1] (MISSING_SFP)	Ports that are missing SFP transceiver.
Error Ports[1] (ERR_PORTS)	Ports with errors.

1. Marginal ports, faulty ports, error ports, and missing SFP transceivers are calculated as a percentage of the physical ports (excluding FCoE and VE_Ports).

Fabric monitoring category

The Fabric category groups areas of potential problems arising between devices, such as zone changes, fabric segmentation, downed E_Ports, fabric reconfiguration, domain ID changes, and fabric logins.

Table 111 lists measures in the Fabric category and describes each measure.

TABLE 111 Fabric measures

Measure	Description
Domain ID changes (DID_CHG)	Monitors forced domain ID changes. Forced domain ID changes occur when there is a conflict of domain IDs in a single fabric and the principal switch must assign another domain ID to a switch.
Fabric logins (FLOGI)	Activates when ports and devices initialize with the fabric.
Fabric reconfigurations (FAB_CFG)	Tracks the number of reconfigurations of the fabric. Fabric reconfiguration occurs when: • Two fabrics with the same domain ID are connected. • Two fabrics are joined. • An E_Port or VE_Port goes offline. • A principal link segments from the fabric.
Down E_Port (EPORT_DOWN)	Tracks the number of times that an E_Port or VE_Port goes down. E_Ports and VE_Ports go down each time you remove a cable or an SFP transceiver (where there are SFP transceiver failures or transient errors).
Segmentation changes (FAB_SEG)	Tracks the cumulative number of segmentation changes. Segmentation changes occur because of one of the following: • Zone conflicts. • Incompatible link parameters. During E_Port and VE_Port initialization, ports exchange link parameters, and incompatible parameters result in segmentation. This is a rare event. • Domain conflicts. • Segmentation of the principal link between two switches.
Zone changes (ZONE_CHG)	Tracks the number of zone changes. Because zoning is a security provision, frequent zone changes might indicate a security breach or weakness. Zone change messages occur whenever there is a change in zone configurations.

TABLE 111 Fabric measures (Continued)

Measure	Description
L2_DEVCNT_PER	Tracks the number of device connections in Layer 2 fabrics. The maximum supported limits in Layer 2 fabrics are the following: • 4096 for all platforms supported by Fabric OS v7.0.x • 6000 for fabrics with the following platforms: - 8 Gbps 8-slot Backbone Chassis - 8 Gbps 4-slot Backbone Chassis - 16 Gbps 4-slot Backbone Chassis - 16 Gbps 8-slot Backbone Chassis - **96-port, 16 Gbps switch** - **48-port, 16 Gbps switch** - 24-port, 16 **Gbps** Edge switch - 80-port, 8 Gbps FC Switch - 40-port, 8 Gbps FC Switch - 24-port, 8 Gbps FC Switch - 8 Gbps Extension Switch - Encryption switch
LSAN_DEVCNT_PER	Tracks the maximum number of imported LSAN devices per the total number of devices imported from all the edge fabrics. The maximum supported limits are the following: • 1000 with the following devices: - 8 Gbps 8-slot Backbone Chassis - 8 Gbps 4-slot Backbone Chassis - 16 Gbps 4-slot Backbone Chassis - 16 Gbps 8-slot Backbone Chassis • 5000 with the following devices: - 80-port, 8 Gbps FC Switch - 40-port, 8 Gbps FC Switch - 24-port, 8 Gbps FC Switch - 8 Gbps Extension Switch - **96-port, 16 Gbps switch** - **48-port, 16 Gbps switch**
BB_FCR_CNT	Tracks the number of Fibre Channel routers present on the backbone fabric.
ZONE_CFGSZ_PER	Tracks the zone configuration size limit per switch.

FRU monitoring category

The FRU category enables you to define rules for field replaceable units (FRU), including ports, power supplies, and flash memory.

In Fabric OS 8.1.0, the following WWN rules and group are not supported in the **Add** and **Edit** policy dialogs and the default policies will not display default rules with ALL_WWN groups:

- defALL_WWNWWN_FAULTY
- defALL_WWNWWN_ON
- defALL_WWNWWN_OUT
- ALL_WWN group

Table 112 lists measures in the FRU category and describes each measure. Possible states for all FRU measures are faulty, inserted, on, off, ready, and up.

TABLE 112 FRU measures

Measure	Description
Power Supplies (PS_STATE)	State of a power supply has changed.
Fans (FAN_STATE)	State of a fan has changed.
Blades (BLADE_STATE)	State of a blade has changed.
SFPs (SFP_STATE)	State of the SFP has changed.
SFMs (SFM_STATE)	State of the SFM has changed.
WWN	State of a world wide name (WWN) card has changed.

Security monitoring category

The Security category monitors different security violations on the switch and takes action based on the configured thresholds and their actions.

Table 113 lists measures in the Security category and describes what each measure indicates.

TABLE 113 Security measures

Measure	Description
DCC violations (SEC_DCC)	An unauthorized device attempts to log in to a secure fabric.
HTTP violations (SEC_HTTP)	A browser access request reaches a secure switch from an unauthorized IP address.
Illegal command (SEC_CMD)	Commands permitted only to the primary fabric configuration server (FCS) are executed on another switch.
Incompatible security DB (SEC_IDB)	Secure switches with different version stamps have been detected.
Login violations (SEC_LV)	Login violations that occur when a secure fabric detects a login failure.
Invalid certifications (SEC_CERT)	There is a missing or invalid certificate file.
No-FCS (SEC_FCS)	The switch has lost contact with the primary FCS.
SCC violations (SEC_SCC)	SCC violations that occur when an unauthorized switch tries to join a secure fabric. The WWN of the unauthorized switch appears in the ERRLOG.
Security authentication failures (SEC_AUTH_FAIL)	Authentication failures that occur when packets try to pass from a nonsecure switch to a secure fabric.
Telnet violations (SEC_TELNET)	Telnet violations that occur when a Telnet connection request reaches a secure switch from an unauthorized IP address.
TS out of sync (SEC_TS)	Time Server (TS) violations that occur when an out-of-synchronization error has been detected.
DAYS_TO_EXPIRE	Notifies the user about the number of days before which a certificate will expire is less than the threshold specified in the rule.
EXPIRED_CERTS	Monitors if the number of expired certificates is greater than the configured threshold.
ALL_CRTS	Monitors security certificates.

Resource monitoring category

The Resource category monitors the system RAM, flash, CPU, and memory.

The Resource category uses monitors to perform the following tasks:

- Configure thresholds for MAPS event monitoring and reporting for the environment and resource classes. Environment thresholds enable temperature monitoring, and resource thresholds enable monitoring of flash memory.
- Configure memory or CPU usage parameters on the switch or display memory or CPU usage. Configuration options include setting usage thresholds which, if exceeded, trigger a set of specified MAPS alerts. You can set up the system monitor to poll at certain intervals and specify the number of retries required before MAPS takes action.

Table 114 lists measures in the Resource category and describes what each measure indicates.

TABLE 114 Resource measures

Measure	Description
Temperature (TEMP)	Refers to the ambient temperature inside the switch, in degrees Celsius. Temperature sensors monitor the switch in case the temperature rises to levels at which damage to the switch might occur.
Flash (FLASH_USAGE)	Monitors the compact flash space available by calculating the percentage of flash space consumed and comparing it with the configured high threshold value.
CPU	Monitors the CPU available by calculating the percentage of CPU consumed and comparing it with the configured threshold value.
Memory (MEMORY_USAGE)	Monitors the available memory by calculating the percentage of memory consumed and comparing it with the configured threshold value.
VTAP IOPS(VTAP_IOPS)	Monitors mirrored traffic IO per second per ASIC chip.

FCIP monitoring category

The FCIP category enables you to define rules for Fibre Channel over IP (FCIP) health, including circuit state changes and utilization as well as packet loss.

Table 115 lists measures in the FCIP category and describes what each measure indicates.

TABLE 115 FCIP measures

Measure	Description
FCIP circuit state changes (CIR_STATE)	The state of the circuit has changed for one of the following reasons: • The circuit has gone offline. • The circuit has come online. • The circuit is faulty.
FCIP circuit utilization (CIR_UTIL)	The percentage of utilization for the circuit at the time of the last poll.
FCIP circuit packet loss (CIR_PKTLOSS)	The number of packets routed through a port exceeds the port bandwidth.
FCIP tunnel packet loss (PKTLOSS)[1]	Monitors the number of retransmitted packets in the tunnel.
Round trip time (RTT)[1]	Monitors the round-trip time of the circuit.
Jitter[1]	Monitors the variance in round-trip time of the circuit.
Tunnel utilization (TUNNEL_UTIL)[1]	Monitors throughput in the channel.
FCIP Tunnel Status (TUNNEL_STATE)[1]	Monitors the tunnel state.

1. Supports only on 16 Gbps 24-FC port, 18 GbE port switch .

Traffic/Flows monitoring category

The Traffic/Flows category groups areas that track the source and destination of traffic and flows. Use traffic and flow thresholds and alarms to determine traffic load and flow and to reallocate resources appropriately.

Table 116 lists measures in the Traffic/Flows category and describes each measure.

TABLE 116 Traffic measures

Measure	Description
Tx Frame Count (TX_FCNT)	The number of transmitted frames for the flow that exceeds the configured thresholds.
Rx Frame Count (RX_FCNT)	The number of received frames for the flow that exceeds the configured thresholds.
Tx Throughput (TX_THPUT)	The number of transmitted words for the flow that exceeds the configured thresholds.
Rx Throughput (RX_THPUT)	The number of received words for the flow that exceeds the configured thresholds.
IO Read Command Count (IO_RD)	The number of SCSI read commands for the flow that exceeds the configured thresholds.
IO Write Command Count (IO_WR)	The number of SCSI write commands for the flow that exceeds the configured thresholds.
IO Read Data Rate (IO_RD_BYTES)	The SCSI read data rate (mbps) for the flow that exceeds the configured thresholds.
IO Write Data Rate (IO_WR_BYTES)	The SCSI write data rate (mbps) for the flow that exceeds the configured thresholds.

Beginning with Fabric OS 8.0.1, MAPS supports FCP IO latency monitoring. User can create rules with the following measures on user–defined static flows only, and not on learning flows.

MAPS monitors following measure matrix (Table 117):

- Pending IO stats – Indicate the number of pending IO requests.
- Completion time stats – Indicates total request completion time.
- First Read or First Write time stats – Indicate how quick the target respond to the command.

NOTE
FCP IO latency monitoring is supported only on 32G platforms.

TABLE 117 Measure matrix

Matrix	Bucket size	Measure
Pending IO	Less than or equal to 8K	RD_PENDING_IO_LT_8K
		WR_PENDING_IO_LT_8K
	Greater than 8K but less than or equal to 64K	RD_PENDING_IO_8_64K
		WR_PENDING_IO_8_64K
	Greater than 64K but less than or equal to 512K	RD_PENDING_IO_64_512K
		WR_PENDING_IO_64_512K
	Greater than or equal to 512K	RD_PENDING_IO_GE_512K
		WR_PENDING_IO_GE_512K
Completion Time	Less than or equal to 8K	RD_STATUS_TIME_LT_8K
		WR_STATUS_TIME_LT_8K
	Greater than 8K but less than or equal to 64K	RD_STATUS_TIME_8_64K
		WR_STATUS_TIME_8_64K

TABLE 117 Measure matrix

Matrix	Bucket size	Measure
	Greater than 64K but less than or equal to 512K	RD_STATUS_TIME_64_512K
		WR_STATUS_TIME_64_512K
	Greater than or equal to 512K	RD_STATUS_TIME_GE_512K
		WR_STATUS_TIME_GE_512K
First read or write response time	Less than or equal to 8K	RD_1stDATA_TIME_LT_8K
		WR_1stXFER_RDY_LT_8K
	Greater than 8K but less than or equal to 64K	RD_1stDATA_TIME_8_64K
		WR_1stXFER_RDY_8_64K
	Greater than 64K but less than or equal to 512K	RD_1stDATA_TIME_64K_512K
		WR_1stXFER_RDY_64K_512K
	Greater than or equal to 512K	RD_1stDATA_TIME_GE_512K
		WR_1stXFER_RDY_GE_512K

FPI monitoring category

The FPI category groups areas that measures the thresholds on the performance of the fabric.

Table 118 lists measures in the FPI category and describes the measure.

TABLE 118 FPI measures

Measure	Description
Device Latency Impact (DEV_LATENCY_IMPACT)	The latency impact of the device.
Rx Bandwidth Usage (RX)	The Rx bandwidth of the device.
Tx Bandwidth Usage (TX)	The Tx bandwidth of the device.
Trunk Utilization (UTIL)	The Tx bandwidth of the device.
Backend Port Latency Impact (BE_LATENCY_IMPACT)	The backend port latency impact of the device.
IT Flow ratio (IT_FLOW)	Initiator to Target flow ratio.

GigE Port monitoring category

The GigE Port category groups areas that measures thresholds to monitor the GigE port health.

Table 119 lists measures in the GigE category and describes the measure.

TABLE 119 GigE measures

Measure	Description
CRC Errors (GE_CRC)	The number of times an invalid cyclic redundancy check error occurs on a GigE port or a frame that computes to an invalid CRC.
Invalid Length (GE_INV_LEN)	The number frames received with an invalid length.
Frames Aborted (GE_LOS_OF_SIG)	The number of frames aborted.

Backend Port monitoring category

The Backend Port monitoring category groups areas that measures thresholds to monitor the back-end port health.

Table 120 lists measures in the Backend Port category and describes the measure.

TABLE 120 Backend Port measures

Measure	Description
Link reset (LR)	The ports on which the number of link resets exceed the specified threshold.
BAD_OS	The number of invalid ordered set (platform and port-specific)
CRC errors (CRC)	The number of times an invalid cyclic redundancy check error occurs on a port or a frame that computes to an invalid CRC.
Invalid transmit words (ITW)	The number of times an invalid transmission word error occurs on a port.
FRM_TRUNC	Frames shorter than minimum
FRM_LONG	Frames longer than minimum

MAPS policies

A MAPS policy is a set of rules that define thresholds for measures and action to take when a threshold is triggered. When you enable a policy, all of the rules in the policy are in effect.

A device can have multiple policies. For example, you can have a policy for everyday use and you can have another policy for when you are running backups or performing switch maintenance. However, only one policy can be active at a time. When you enable a policy, it becomes the active policy and the rules in the active policy take effect.

At any time, one policy must always be active on the switch. You can have an active policy with no rules, but you must have an active policy. You cannot disable the active policy. You can only change the active policy by enabling a different policy.

Preconfigured policies

MAPS provides three preconfigured policies. You cannot modify or delete these policies. The preconfigured policies include the following:

- dflt_aggressive_policy

  Contains rules with very strict thresholds. Use this policy if you need a pristine SAN fabric (for example, FICON fabrics).

- dflt_conservative_policy

  Contains rules with more lenient thresholds that allow a buffer and do not immediately trigger actions. Use this policy in environments where the elements are resilient and can accommodate errors.

  This is the default policy unless you specify another policy when you enable MAPS on the device and migrate from Fabric Watch to MAPS using the **mapsconfig --migrate –enablepolicy** *policy_name* command.

- dflt_moderate_policy

  Contains rules with threshold values between the aggressive and conservative policies.

- dflt_base_policy

  For Fabric OS 7.4.0 devices, when MAPS is unlicensed (without the Fabric Vision license), dflt_base_policy will be the default policy activated. The dflt_base_policy policy provides limited monitoring support.

Although you cannot modify the preconfigured policies, you can create a policy based on these policies. For more information, refer to "Cloning a MAPS policy" on page 1236.

User-defined policies

MAPS allows you to define your own policies. You can create a policy and add rules to it, or you can clone one of the default policies and modify the cloned policy. For information on configuring a user-defined policies, refer to "Configuring a MAPS policy" on page 1232.

The following lists the user-defined limitation for switches running Fabric OS 8.1.0 or later:

- Maximum number of user-defined rules per logical switch is 500.
- Maximum number of rules in a user-defined policy per logical switch is 350.
- Maximum number of user-defined policies per logical switch is 20.
- Maximum number RoR rules in a user-defined policy is 50.

Fabric Watch legacy policies

You cannot return to Fabric Watch once you activate MAPS (or migrate to MAPS).

When you migrate from Fabric Watch to MAPS, three policies are automatically created:

- fw_custom_policy

  Contains all of the monitoring rules based on the custom thresholds configured in Fabric Watch, even if the rules have the same parameters as the default rules.

- fw_default_policy

  Contains all of the monitoring rules based on the default thresholds configured in Fabric Watch.

- fw_active_policy

  Contains all of the monitoring rules based on the active thresholds in Fabric Watch at the time of the migration.

The following factors also apply to Fabric Watch conversions:

- Converted active Fabric Watch policies reference either custom or default Fabric Watch rules.
- Converted custom Fabric Watch policies reference either custom or default Fabric Watch rules.
- Fabric Watch rule conversions use the following rule name formats:
  - "_LBC" suffixes are changed to "AL" (indicating that it uses an "Above Low" threshold)
  - Converted Fabric Watch rule names will have the threshold number as the suffix. This is the same pattern as MAPS rules use. For example, defALL_10GLWL_SFPSFP_TEMP_AH will be changed to defALL_10GLWL_SFPSFP_TEMP_90.
  - Converted rules are prefixed with fw_def_xxx or fw_cust_xxx.

These policies are treated as user-defined policies. You can modify them by adding and deleting rules, and you can delete the policy.

MAPS rules

A rule associates a condition with actions that need to be triggered when the specified condition is evaluated to be true.

Each rule specifies the following items:

- A group of objects to be evaluated.

  Refer to "MAPS groups" on page 1245 for additional information.

- The measure to be monitored.

  Refer to "MAPS monitoring categories" on page 1213 for additional information.

- The condition.

  Each rule specifies a single condition. A condition includes a time base and a threshold. Refer to "MAPS conditions" on page 1223 for additional information.

- The actions to take if the condition is evaluated to be true.

  Refer to "MAPS actions" on page 1224 for additional information.

- The severity to be monitored.

  Refer to "MAPS severity" on page 1224 for additional information.

The combination of actions, conditions, and measures allow you to create a rule for almost any scenario required for your environment.

MAPS conditions

A condition includes a time base and a threshold. If the condition is evaluated as true, the rule is triggered. The condition depends on the element that is to be monitored.

Consider the following rule:

> For all F_Ports, if the change in the CRC counter in the last minute is greater than 10, then fence the port and issue a RASLog message.

In this rule, the condition is whether the change in the CRC counter in the last minute is greater than 10.

Thresholds

Thresholds are the values at which potential problems might occur. For example, in configuring a port threshold, you can select a specific value at which an action is triggered because of too many threshold violations.

Time base

Time bases specify the time interval between two samples to be compared. You can set the time base to **Day** (samples are compared once a day), **Hour** (samples are compared once an hour), or **Minute** (samples are compared every minute).

The time base affects the comparison of sensor-based data with user-defined threshold values.

For measures where the time base is not applicable, the time base is automatically set to **None**.

MAPS severity

MAPS severity allows you to configure the severity during a policy rule creation. You can select Default, Critical, Error, Warning, or Info values from the list. The default severity value is “Default”. You can import and distribute severity only for switches running Fabric OS 8.1.0 or later.

MAPS actions

MAPS provides actions (event notifications) in several different formats to ensure that event details are accessible from all platforms and operating systems. In response to an event, MAPS can record event data as any (or all) of the following alarm options.

To enable MAPS actions, refer to “Enabling or disabling policy actions for all policies” on page 1227. For example, if you define a rule in the Management application with an SNMP action and a violation of that rule occurs, the switch with the violation only sends the SNMP trap if you configured SNMP on that switch.

Status

NOTE
Status is available for all measures in the Switch Status category and the flash usage measure in the Resource category.

When a threshold is triggered, the switch status changes to a marginal or critical status icon on the Dashboard and SAN tabs. Marginal status displays with a yellow icon. Critical status displays with a red icon.

SFP Marginal

NOTE
SFP Marginal is available for SFP transceiver measures in the Port category.

When a threshold is triggered, the SFP transceiver status changes to a marginal status icon on the Dashboard and SAN tabs. Port SFP transceiver measures include Current, Receive Power, Transmit Power, Voltage, Temperature, and Power On Hours. Marginal status displays with a yellow icon in the **Port Optics** dialog box.

RAS log events

Following an event, MAPS adds an entry to the internal event log for an individual switch. The RAS log stores event information but does not actively send alerts.

Port decommission

Port decommission automatically takes ports offline when the configured thresholds in a rule are exceeded. Port fencing is auto-enabled if port decommission is enabled for a MAPS rule or action. You can enable port decommissioning for E_Ports and F_Ports.

Fence

Fence the port, if port fencing is enabled. Port fencing takes the ports offline if the user-defined thresholds are exceeded. Supported port types include physical ports, E_Ports, optical F_Ports (FOP_Ports), copper F_Ports (FCU_Ports), and Virtual E_Ports (VE_Ports).

E-mail

An e-mail alert sends information about a switch event to a specified e-mail address. An e-mail alert can send information about any error from any element, area, and class (only one e-mail recipient can be configured per class). The e-mail alert specifies the threshold and describes the event, much like an error message. To configure multiple e-mail recipients, refer to "Configuring e-mail notification" on page 1229. You must separate the e-mail addresses with a semi-colon and include the complete e-mail address. For example, abc@12.com is a valid e-mail address; abc@12 is not.

SNMP traps

In environments in which you have a high number of messages coming from a variety of switches, you may want to receive them in a single location and view them using a graphical user interface (GUI). In this type of scenario, the Simple Network Management Protocol (SNMP) notifications may be the most efficient notification method. You can avoid logging in to each switch individually as you would have to do for error log notifications.

SNMP performs an operation called a ***trap*** that notifies a management station using SNMP when events occur. Log entries can also trigger SNMP traps if the SNMP agent is configured. When the SNMP agent is configured to a specific error message level, error messages at that level trigger SNMP traps.

An SNMP trap forwards the following information to an SNMP management station:

- Name of the element with a counter that registered an event
- Class, area, and index number of the threshold that the counter crossed
- Event type
- Value of the counter that exceeded the threshold
- State of the element that triggered the alarm
- Source of the trap

NOTE
The SNMP trap stores event information but does not actively send alerts.

You must configure the software to receive trap information from the network device. You must also configure the SNMP agent on the switch using the **snmpConfig** (Fabric OS devices) command to send the trap to the management station. You can configure SNMP notifications using the Management application (refer to "Event notification" on page 1132).

For information on configuring the SNMP agent using the **snmpConfig** command, refer to the ***Fabric OS Command Reference***.

SNMP MIB support

MAPS requires SNMP management information base (MIB) support on the device for management information collection. For a list of required MIBs, refer to Table 197.

TABLE 121 Required MIB support for Fabric OS devices

MIB name	Required MIB object	Data collected
Brocade MAPS MIB	mapsTrapAM	mapsConfigRuleName
		mapsConfigObjectGroupType
		mapsConfigObjectKeyType
		mapsConfigObjectKeyValue
		mapsConfigNumOfMS
		mapsConfigMsList
		mapsConfigSeverityLevel

FMS (FICON Management Server)

FMS is used to notify the FICON Host management service of MAPS violation. Rules with FICON notification action will be a part of all three default policies (dflt_aggressive_policy, dflt_moderate_policy, and dflt_conservative_policy). In an active policy, if FICON notification is configured for any triggered events, then MAPS sends the notification to FMS with events information.

SDDQ (Slow Drain Device Quarantine)

SSDQ to used to isolate the traffic targeted to slow-drain devices and reduce the impact of traffic targeted on other devices. Due to this automatic isolation from the regular flows, the effects of the slow-drain flows on the fabric are reduced. If the quarantined ports go offline or disabled, the ports remain in Slow Drain Quarantined state. Once the ports come online, the flows destined to the port are quarantined.

Un-Quarantine

Un-Quarantine is used to automatically quarantine a port when slow drain device quarantine does not take place for a given timeout period. Beginning with Fabric OS 8.1.0 or later, un-quarantine action is supported on FPI Monitoring. All dflt_conservative_policy and dflt_moderate_policy policies with IO_LATENCY_CLEAR state will support un-quarantine action and default un-quarantine timeout configuration. You can set the timeout period for days, hours, minutes, or seconds.

Toggle

Use to toggle ports and recover from bottleneck conditions caused by the target device. MAPS toggle action helps to recover from the bottleneck or altogether forces the Fibre Channel traffic to switch over on the redundant path. You can configure the Toggle action for a shorter or longer duration. The minimum timeout duration is 2 seconds and the maximum is 3600 seconds (24 hours).

Quiet time

Beginning with Fabric OS 7.4.0, Quiet time is used as a rule parameter to avoid sending alerts for the quiet time duration (Days, Hours, Minutes, Second) set, after alerting you for the first time. You can configure quiet time for all rules except "RAS Log" and "E-mail" actions. You can also configure quite time for “SNMP Trap” action for devices running Fabric OS 8.0.1 and later. By default the **Quiet Time** check box is in disabled state.

Enabling or disabling policy actions for all policies

You can define what actions are allowable on the device, regardless of the actions specified in the individual rules in a policy.

Enabling and disabling actions at a global level allows you to configure rules with stricter actions, such as port fencing, but disable the action globally until you can test the configured thresholds. After validating the thresholds, you can enable port fencing globally without having to change all of the rules.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select an object in the MAPS policies list and click **Actions**.

   Select **All Fabrics** to configure actions for all policy rules on all MAPS-enabled devices in all fabrics.

   Select one or more fabrics to configure actions for all policy rules on all MAPS-enabled devices in the selected fabrics.

   Select one or more devices to configure actions for all policy rules on the selected devices.

   The **MAPS Policy Actions** dialog box displays.

3. Select the associated check box for each action you want to enable.

   Enable all actions by clicking the **Enable All** button. Disable all actions by clicking the **Disable All** button.

   Not all actions are available for all objects. Actions availability is based on the Fabric OS devices selected. Options include:

   - **RAS Log Event** — Use to log a RAS event.
   - **SNMP Trap** — Use to send an SNMP trap event.
   - **Port Decommission** — Use to decommission the offending port.
   - **Fence** — Use to fence the offending port.
   - **E-mail** — Use to send an e-mail notification.
   - **SFP Status Marginal** — Use to set the SFP status to marginal.
   - **Switch Status Critical** — Use to set the switch status to critical
   - **Switch Status Marginal** — Use to set the switch status to marginal.
   - **FMS** — Use to notify the FICON Host management service of MAPS violations.
   - **SDDQ** (Slow Drain Device Quarantine) — Use to isolate the traffic targeted to slow-drain devices and reduce the impact of traffic targeted on other devices.
   - Un-Quarantine — Use to automatically quarantine a port when the slow drain device quarantine does not take place for a given timeout period.
   - **Toggle** — Use to toggle ports and recover from bottleneck conditions. The minimum timeout duration is 2 seconds and the maximum is 3600 seconds (24 hours).

     NOTE
     For switches running Fabric OS 8.0.0 or later , Switch Status Critical and Switch Status Marginal are enabled by default and cannot be disabled.

   For switches running Fabric OS 8.0.0 or later , the following actions are available:

- **Switch Status Actions** - Enabled by default on a single-switch fabric and disabled when multiple switches are selected. In case multiple switch selection, the user is allowed to configure Switch Status Actions on supported switches only.
  - Switch Status Critical - Checked and disabled by default.

- Switch Status Marginal - Checked and disabled by default.

For a complete list of categories and the associated measures and actions, refer to "MAPS categories, measures, and actions" on page 1210.

4. Click **OK** on the **MAPS Policy Actions** dialog box.
5. Click **Close** on the **MAPS Configuration** dialog box.

Enabling FPI monitoring

Beginning with Fabric OS 8.0.0, **FPI (Fabric Performance Impact) Monitoring** is enabled by default for switches running Fabric OS 8.0.0 or later with or without a Fabric Vision license.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays (Figure 560).

   FIGURE 560 MAPS Configuration dialog box

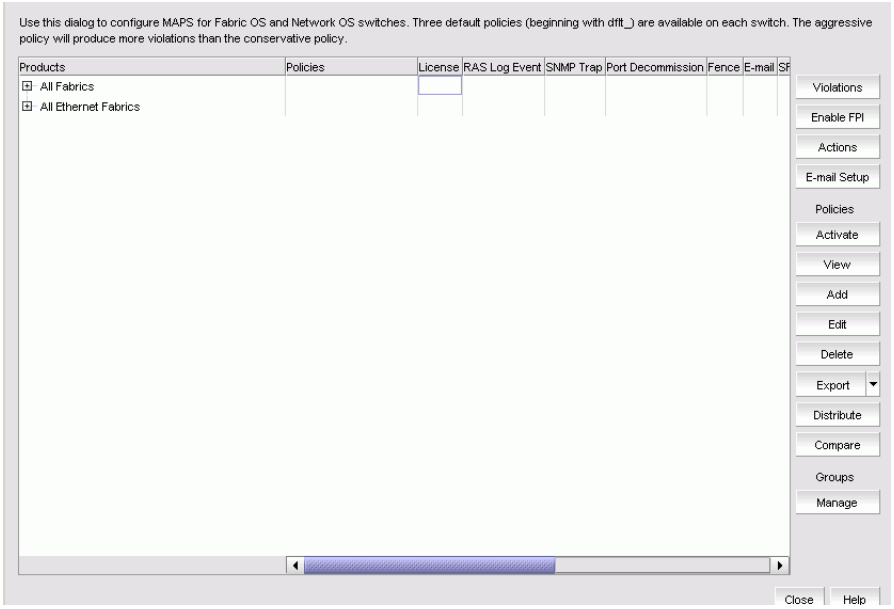

2. Select one or more devices under **All Fabrics** on which you want to enable FPI monitoring and click **Enable FPI**.

   A message "FPI monitoring is enabled by default on switches running Fabric OS v8.0 and above" displays. Click **OK**.

   The **Enable FPI Monitoring** dialog box displays (Figure 561).

FIGURE 561 Enable FPI Monitoring dialog box

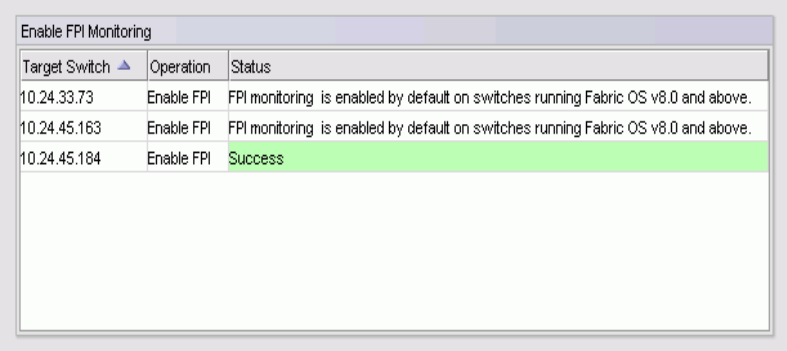

The monitoring status is displayed with different colors as follows:

- **Green -** FPI monitoring is successful.
- **Red -** FPI monitoring failed.
- **No color -** FPI monitoring is not applicable or already enabled.

If you select a combination of switches running Fabric 7.4.0 and Fabric 8.0.0 and click **Enable FPI**, the **Fabric Performance Impact (FPI)** dialog box displays (Figure 562). Click **OK** to enable the FPI monitoring for the selected switches.

FIGURE 562 Fabric Performance Impact (FPI) dialog box.

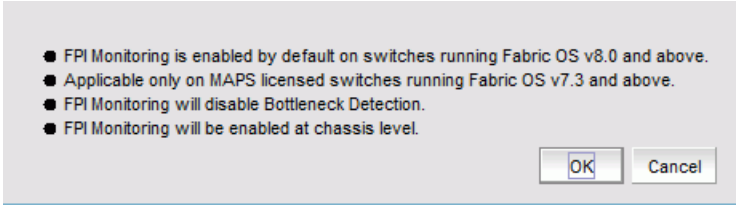

Configuring e-mail notification

In environments where it is critical that you are notified about errors quickly, you can use e-mail notifications. With e-mail notifications, you can be notified of serious errors by e-mail, text message, or pager, so you can react quickly.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select an object in the MAPS policies list and click **E-mail Setup**.

   Select All Fabrics to configure e-mail notification for all policy rules for all MAPS-enabled devices in all fabrics.

   Select one or more fabrics to configure e-mail notification for all policy rules on all MAPS-enabled devices in the selected fabrics.

   Select one or more devices to configure e-mail notification for all policy rules on the selected devices.

   The **MAPS E-Mail Setup** dialog box displays. This dialog box enables you to configure e-mail notification for all policy rules on a device.

3. Enter one or more addresses in the **E-mail address** text box.

   You can enter up to 5 addresses, separated by semi-colons. E-mail addresses are logical switch specific, not physical chassis specific. You can configure different e-mail addresses for different logical switches in the same physical chassis.

   To send a text message or page via e-mail, use the following format *number@carrier.com*, where *number* is your phone number and *carrier.com* is the SMS server. For example, 3035551212@txt.att.net (text message) or 3035551212@page.att.net (page).

   NOTE
   Check with your carrier for the exact e-mail address.

4. Enter the IP address of the relay host in the **Relay Host** field.

   Relay host is a physical chassis setting. This setting affects all logical switches in the physical chassis.

5. Enter the domain name in the **Domain Name** field.

   Domain name is a physical chassis setting. This setting affects all logical switches in the physical chassis.

6. Click **Apply** to configure an e-mail in the switch.
7. Click **Test** to send a test e-mail to all the e-mail address in **E-mail address** text box.
8. Click **OK** on the **MAPS E-Mail Setup** dialog box.
9. Click **Close** on the **MAPS Configuration** dialog box.

NOTE
You can clear the E-mail address by removing the configured addresses from the **E-mail address** text box.

Viewing MAPS policy data

You can view the MAPS-capable devices and the associated MAPS policies and actions.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays (Figure 563).

FIGURE 563 MAPS Configuration dialog box for SAN

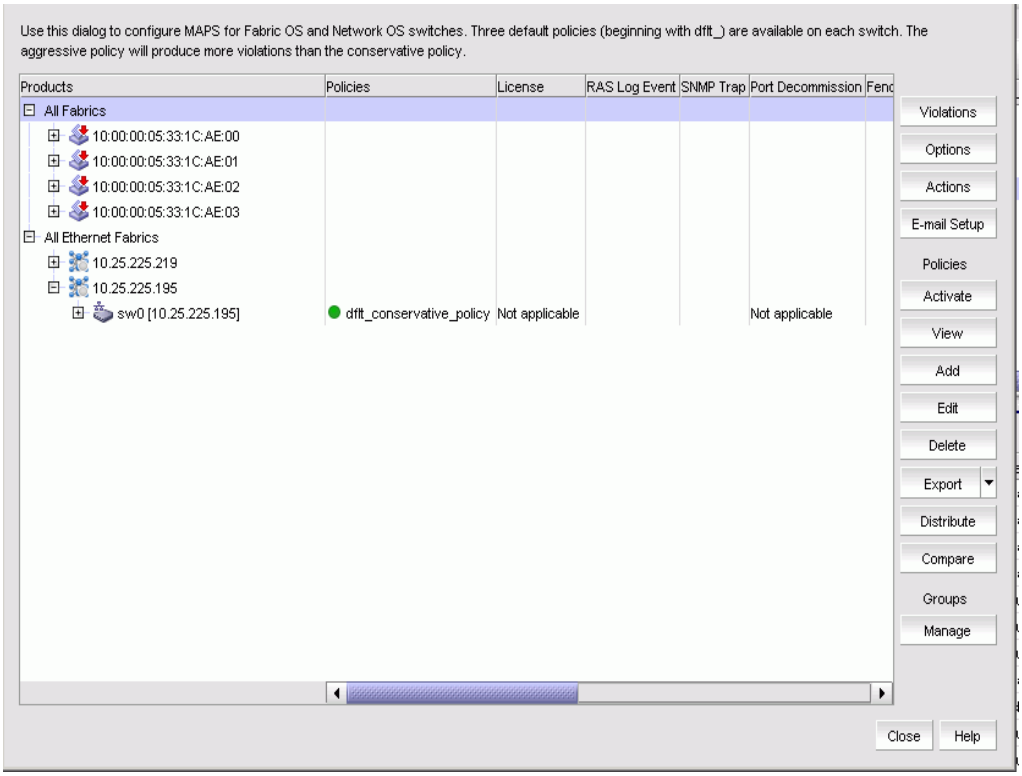

Sort the contents by clicking the column header. Click the same column header again to reverse the sort order.

2. Review the MAPS data:

- MAPS policies list — Lists the MAPS-capable devices and associated policies and actions.
  - **All Fabrics/Fabric/Switch** — All fabrics that contain MAPS-capable switches. You can expand the fabric node to view switches under each fabric. MAPS policies deployed to the switch display under the switch node.
  - **Policies** — Policies available on the associated switch. The active policy on the switch displays in the cell adjacent to the associated switch.
  - **License** — Beginning with Fabric OS 7.4.0, Fabric Watch is not supported for monitoring the performance and status of switches. MAPS is enabled implicitly for monitoring the unlicensed features. A check mark displayed in the license column indicates whether the switch is monitored with or without the MAPS license.

    NOTE
    The active policy indicates an active status icon in the **Policies** column of the **Switch** row and **Active Policy** row.

  - **RAS Log Event** — If check mark displays, logs a RAS event when triggered.
  - **SNMP Trap** — If check mark displays, sends an SNMP trap event when triggered.
  - **Fence** — If check mark displays, fences the offending port when triggered.
  - **E-mail** — If check mark displays, sends an e-mail notification when triggered.
  - **SFP Status Marginal** — If check mark displays, sets the SFP status to marginal when triggered.
  - **Port Decommission** — If check mark displays, decommissions the port, when triggered.
  - **Switch Status Marginal**. — If check mark displays, sets the switch status to marginal when triggered.
  - **Switch Status Critical** — If check mark displays, sets the switch status to critical when triggered.
  - **FMS** — Use to notify the FICON Host management service of MAPS violations.

  - **SDDQ** (Slow Drain Device Quarantine) — Use to isolate the traffic targeted to slow-drain devices and reduce the impact of traffic targeted on other devices.
  - **Toggle** — Use to toggle ports and recover from bottleneck conditions. The minimum timeout duration is 2 seconds and the maximum is 3600 seconds (24 hours).
- **Violations** button — Select an object (switch or fabric) and click to open the **Violations** dialog box for the selected object. For more information, refer to "Viewing MAPS violations" on page 1255.
- **Enable FPI** button — Select an object (switch or fabric) and click the **Enable FPI** button in the **MAPS Configuration** dialog box to enable FPI monitoring on an object. For switches running Fabric OS 8.0.0 or later, the FPI (Fabric Performance Impact) Monitoring will be enabled by default. For more information refer to "Enabling FPI monitoring" on page 1228.
- **Actions** button — Select an object (switch, fabric, or all fabrics) and click to enable or disable actions for all policy rules on an object.
- **E-mail Setup** button — Select to configure e-mail notification. For more information, refer to "Configuring e-mail notification" on page 1229.
- **Activate** button — Select an inactive policy and click to activate the policy during deployment. Only one policy can be active on a switch at a time. You can activate policies for multiple switches at once by selecting the policy you want to activate for each switch and clicking **Activate**. For more information, refer to "Activating a MAPS policy" on page 1239.
- **View** button — Select a policy and click to open the **View Policy** dialog box and view the rules defined for the policy. For more information, refer to "Viewing MAPS policy rules" on page 1242.
- **Add** button — Click to create a new policy or select a policy in the **Policies** list and click to clone a policy. For more information, refer to "Configuring a MAPS policy" on page 1232 and "Cloning a MAPS policy" on page 1236.
- **Edit** button — Select a policy and click to open the **Edit Policy** dialog box. You cannot edit a default policy. For more information, refer to "Editing a MAPS policy" on page 1236.
- **Delete** button — Select one or more policies and click to delete. For more information, refer to "Deleting a MAPS policy" on page 1241.
- **Export** button — Click to export a policy definition to an XML file. For more information, refer to "Exporting a MAPS policy" on page 1241.
- **Import** option (on the **Export** button) — Click to import a policy definition. For more information, refer to "Importing a MAPS policy" on page 1241.
- **Distribute** button — Select a policy and click to replicate the policy to all devices in a fabric or SAN. For more information, refer to "Replicating a policy to other devices" on page 1239.
- **Manage** button in the **Groups** area — Select the SAN fabric or switch for which you want to edit groups and click to open the **Manage Groups - MAPS** dialog box. For more information, refer to "Managing MAPS groups" on page 1252.
- **Compare** button — Click to compare two policies across the fabric. For more information, refer to "Comparing MAPS policies" on page 1244.

3. Click **Close**.

Configuring a MAPS policy

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Click **Add**.

   The **Add Policy** dialog box displays(Figure 564).

FIGURE 564 Add Policy dialog box for SAN

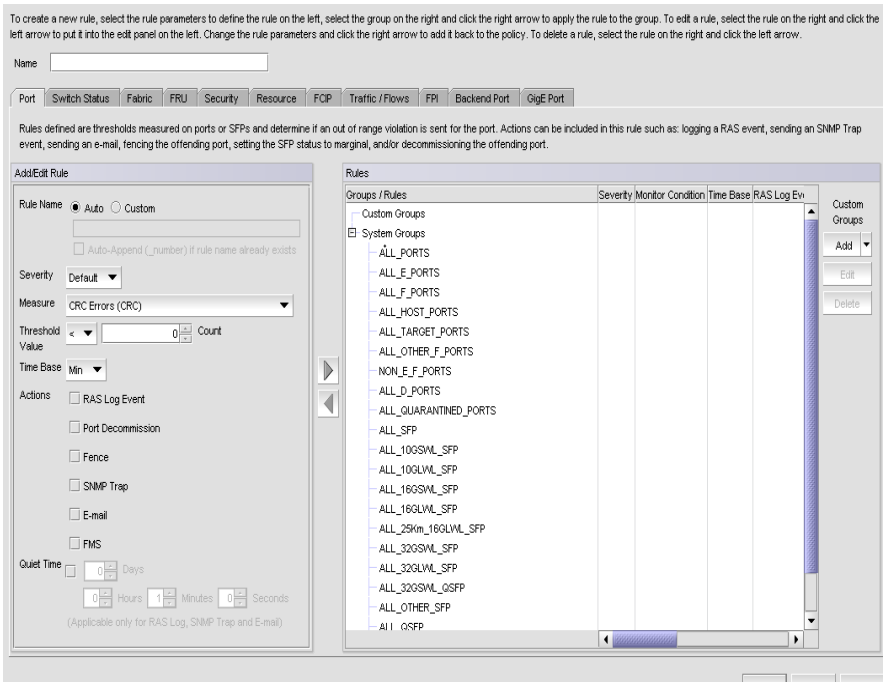

3. Enter a name for the policy in the **Name** field.

   The policy name can be up to 31 characters and can only contain of alphanumeric and underscore characters.

4. Select one of the following category tabs to configure the policy measures.

   For a complete list of categories and the associated measures and actions, refer to "MAPS categories, measures, and actions" on page 1210. Options include:

   Beginning with Fabric OS 8.1.0 or later, the FCIP tab, Backend Port tab, and GigE Port tab will be displayed in th ADD, Edit, or View policy dialog boxes based on the platforms supported.

   - **Port** tab — Rules defined on this tab measure thresholds on ports or SFPs to determine if an out of range violation is sent for the port.
   - **Switch Status** tab — Rules defined on this tab measure thresholds at the switch or chassis level to determine the switch operational status.
   - **Fabric** tab — Rules defined on this tab measure thresholds at the switch level to detect out of range fabric-wide changes.
   - **FRU** tab — Rules defined on this tab measure thresholds on fans, power supplies, SFPs, blades, or WWN cards to detect out of range FRU changes.
   - **Security** tab — Rules defined on this tab measure thresholds at the switch level to detect out of range security changes.
   - **Resource** tab — Rules defined on this tab measure thresholds on temperature sensors or at the chassis level to detect out-of-range resource usage.
   - **FCIP** tab — Rules defined on this tab measure thresholds on FCIP circuits to detect out of range state, utilization, or packet loss.
   - **Traffic / Flows** tab — Rules defined on this tab measure thresholds on ports in the configured group to detect out-of-range link utilization or on flows to detect out-of-range flow changes.
   - **FPI** tab — Rules defined on this tab measure thresholds on the performance of the fabric. In Fabric OS 8.1.0, new default policy with Rx, Tx, and UTIL measures can be monitored with Time base in minute and Quiet Time for 1 hour.

- **GigE Port** tab — Rules defined on this tab measure thresholds to monitor Gigabit Ethernet ports (GE) for ALL_EXT_GE_PORTS system group with default policy rules. GigE Port tab supports RAS Log Event, SNMP Trap, and E-mail actions.

  NOTE
  Beginning with 8.0.1 or later, GigE Port monitoring is supported on 16 Gbps 24-port, 18 GbE port switch, FX8-24, and 32 Gbps, Router Extension Blade.

- **Backend Port** tab — Rules on this tab measure thresholds to monitor the backend port health for ALL_BE_PORTS system group with default policy rules. Backend Port rules are predefined and obtained directly from the switch. You cannot create custom policy rules for the ALL_BE_PORTS system group; you can only select predefined rules for the ALL_BE_PORTS system group.

5. Select one of the following option to set the rule type in the **Rule Type** area:

   - **Base Rule** — Default rule type. By default Base Rule radio-button is selected.
   - **Rule on Rule (RoR)** — RoR allows you to define rules on an existing rule to monitor the frequency of rule execution on the switches running Fabric OS v 8.1.0a or later.You can create a RoR rule on an existing Base Rule. A base rule must exist in order to create a RoR. RoR time base can be created greater than the Base Rule time base. RoR quiet time must be > = (greater than or equal to) the time base selected.You can create a maximum of 50 RoR per policy. You can import and distribute RoR for switches running Fabric OS 8.1.0a or later, and compare policies with or without RoR.

     For CHASSIS measure, RoR rules can be created for default switches only, which means the ROR rule is not supported for FRU health and Switch Resource categories on non-default Fabric Identifiers (FIDs).

     The RoR rule is not supported for the following categories, measures, and actions:

     - Categories and Measures
     - Security Health (DAYS_TO_EXPIRE and EXPIRED_CERTS are not supported)
     - Fabric State Changes (All measures are not supported)
     - Switch Policy Status (All measures are not supported)
     - FPI Monitoring (IT_FLOW is not supported)
     - Traffic /Flows Performance (All measures not supported)
     - Switch Resources (All measures are not supported except CPU and MEMORY_USAGE)
     - Action
     - FMS (FICON Management Server) action is not supported.

6. Select one of the following options to name the rule in the **Rule Name** area:

   - Select the **Auto** option (default) to auto-generate the rule name.

     The Management application auto-generates a name for the rule based on the rule parameters (measure, threshold, time base, actions, and group). Auto-generated rule names use the following naming convention: *<group_name_abbreviation>_<measure_abbreviation><logical_operator><value>_<timebase>_<actions>*.

     For example, AP_CRCL5_M_RxTxxxx, where “AP” is an abbreviation of the group name, “CRC” is an abbreviation of the selected measure, “L” is the selected logical operator, “5” is the selected count value, “M” is the selected timebase, and “RxTxxxx” defines the selected actions.

     Logical operators are represented using the following abbreviations: L represents < (less than), LE represents <= (less than or equal to), G represents > (greater than), GE represents >= (greater than or equal to), and EQ represents = (equal to).

     Timebase durations are represented using the following abbreviations: M represents minute, H represents hour, and D represents day.

Actions are represented in a bit-wise format, where each "X" represents a possible action you can configure in a rule. This format uses the following order: RASlog (R), fence (F), SNMP trap (T), e-mail (E), switch critical (C), switch marginal (M), SFP marginal (S), Port Decommission (D)**,** FMS (N), SDDQ (Q), and Un-Quarantine (U), and Toggle (D). For example, if you configure a rule with SNMP trap and RASlog actions, the actions portion of the rule name would be "RxTxxxx".

If you are editing an existing rule, you can change the auto-generated rule name by selecting the **Custom** option and editing the name in the **Custom** field.

If you select **Auto_Append**, the custom rule name will be appended with an integer on selecting multiple groups.

- Select the **Custom** option to provide a user-defined name and enter a name in the **Custom** field.

  The rule name can be up to 72 characters and can only contain of alphanumeric and underscore characters.

7. Select a severity from the Severity list.

   Valid values include: Default, Critical, Error, Warning, and Info

8. Select a measure from the **Measure** list.

   Available measures depend on the selected category. For a complete list of categories and the associated measures and actions, refer to "MAPS categories, measures, and actions" on page 1210.

9. Select a logical operator from the **Threshold** list.

   Valid values include: < (less than), <= (less than or equal to), > (greater than), >= (greater than or equal to), or = (equal to).

10. Enter a threshold value in the **Threshold** field.

    Valid values include 1 through 1,000 for numerical values and 0.00 through 100.00 for percentage measures. For the SFP_TEMP measure in Port category, valid values are -40 through 100. For FRUs, valid values include: IN, READY, UP, ON, OFF, and FAULTY. For the TEMP measure in the Resource category, valid values are IN_RANGE and OUT_OF_RANGE.

11. Select one of the following durations to monitor the counter from the **Time Base** list.

    Valid durations include: **Min** (default), **Hour**, or **Day**. If a duration is not applicable for the selected measure (such as MEMORY_USAGE), the list displays **None**. The RoR rule does not support **None** timebase but, supports any timebase greater than the Base Rule timebase.

12. From the **Actions** check boxes, select the check box for each action you want to occur when a threshold is crossed.

    Not all actions are available for all objects. Options include: **Status Critical**, **Status Marginal**, **RAS Log Event**, **Port Decommission**, **Fence**, **SNMP Trap**, **E-mail**, **SFP Marginal**, **FMS**, **SDDQ**, Un-Quarantine, and **Toggle**. For a complete list of categories and the associated measures and actions, refer to "MAPS categories, measures, and actions" on page 1210.

13. Add the rule to a group or multiple groups by selecting the group in the **Rules** area and clicking the right arrow button to move the new rule to the selected group (or imported flow).

    The **Rules** area displays the default groups (under the **System Groups** node) and user-defined groups (under the **Custom Groups** node) for the selected switch. Even though all groups display available in the Rules area, you can only add the rule to an appropriate group. For example, if you selected an SFP measure, you can only add the SFP measure to an SFP group. If you try to add a measure to an inappropriate group, an error message displays.

    You can only configure a user-defined group on the **Port tab** and **FCIP** tabs. For more information, refer to "Configuring a group" on page 1248.

    Rules display below the appropriate group node based on rule targets.

14. (**FCIP** tab only) Add a group to the **Rules** area by clicking **Add** in the **Custom Groups** area.

The **Add Group** dialog box displays. For more information, refer to "Configuring a group" on page 1248.

15. Click **OK** to add the policy to the **MAPS Configuration** dialog box.
16. Click **Close** on the **MAPS Configuration** dialog box.

Editing a MAPS policy

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select any non-default policy in the list and select **Edit**.

   You can also select the switch in the list and select **Edit** to edit the active policy. When you edit the active policy on the switch, updated rules activate on the switch automatically.

   NOTE
   You cannot edit a default policy.

   The **Edit Policy** dialog box displays.

3. To edit an existing rule, select the rule in the **Rules** area and click the left arrow button.

   The rule displays in the **Add/Edit Threshold** area. Note that this removes the rule from the associated group, once you edit the rule, you must add it back to the group (refer to "Configuring a MAPS policy" on page 1232 and complete step 5 through step 13). To edit groups, refer to "Editing a group" on page 1251.

4. Click **OK** on the **Edit Policy** dialog box to update the policy and return to the **MAPS Configuration** dialog box.
5. Click **Close** on the **MAPS Configuration** dialog box.

Cloning a MAPS policy

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select a policy in the list and click **Add**.

   The **Add Policy** dialog box displays.

3. Enter a name for the policy in the **Name** field.

   The policy name can be up to 31 characters and can only contain of alphanumeric and underscore characters.

4. To create a new policy, refer to "Configuring a MAPS policy" on page 1232 and complete step 4 through step 14.
5. Click **OK** on the **Add Policy** dialog box.
6. Click **Close** on the **MAPS Configuration** dialog box.

Importing Flow definitions

You can import a flow definition into MAPS for threshold monitoring.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select a policy in the list and click **Add**.

   The **Add Policy** dialog box displays.

3. Click the **Traffic / Flows** tab.

4. Click **Import**.

   The **Import Flow Definitions** dialog box displays. The **Available Flow** list displays all flows defined on the switch except learning flows for switches running Fabric OS v7.2.0. For Fabric OS v7.3.0 and later, import and learning flows are supported. Learning flows use an asterisk (*) for both the source and destination devices which enables the flow monitor to learn all the source device and destination device pairs passing through the devices using a particular source port or destination port. Each subflow that is a part of the imported learning flow is monitored.

   FIGURE 565 Import Flow Definitions dialog box

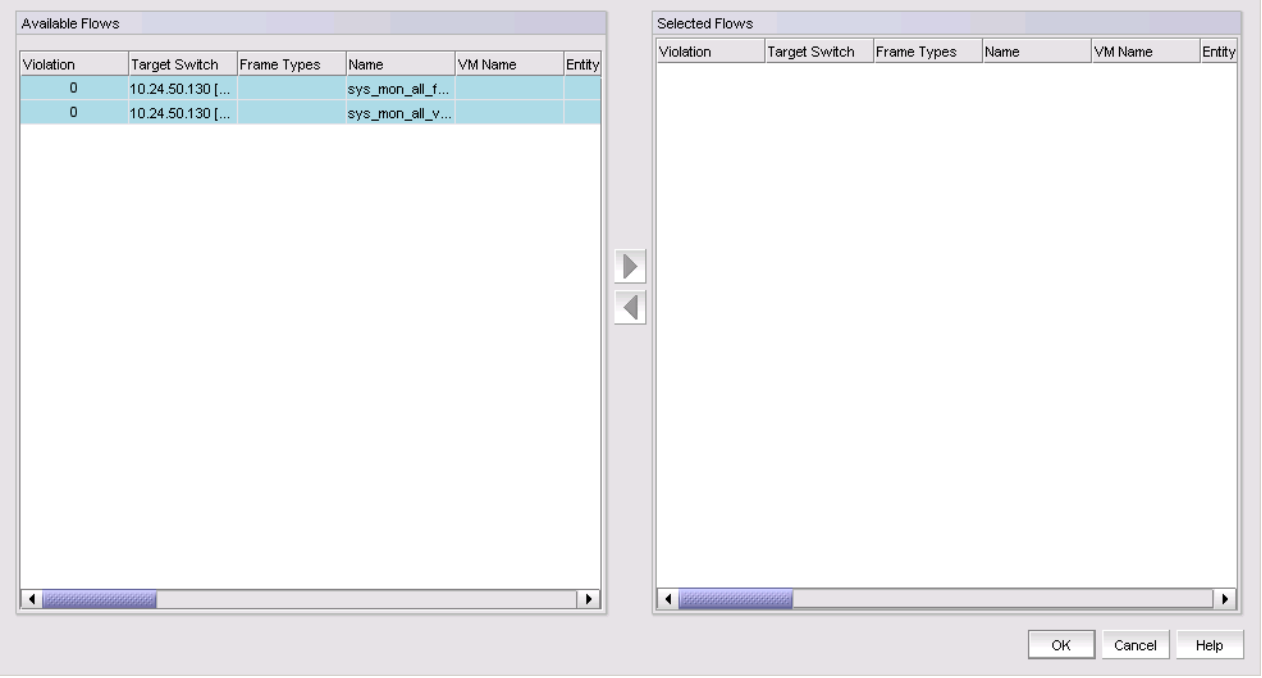

5. Select the flow definition you want to import in the **Available Flow** list and click the right arrow button.

   The selected flow moves from the **Available Flow** list to the **Selected Flow** list. The **Available Flow** and **Selected Flow** lists contain the following data:

   - Violation — The violation triggered for the flow.
   - Target Switch — The device on which the flow definition was created.
   - Name — The name of the flow.
   - VM Name — The VM name of the flow.

- Monitor — Whether or not the Monitor feature is active or not.
- Mirror — Whether or not the Mirror feature is active or not.
- Generator — Whether or not the Generator feature is active or not.
- Source — The source device identifier.
- Source Info — Icon and name for the source device. The device name is a hyper link to the device's properties.
- Destination — The destination device identifier
- Destination Info — Icon and name for the destination device. The device name is a hyper link to the device's properties.
- Source Port — The port number where the flow originates.
- Destination Port — The port number where the flow ends.
- Source Domain — The domain where the flow originates.
- Destination Domain — The domain where the flow ends.
- Source Fabric ID — The fabric identifier where the flow originates.
- Destination Fabric ID — The fabric identifier where the flow ends.
- Rx Port — The receive port.
- Tx Port — The transmit port.
- LUN — The LUN values defined in the flow.
- Bi-direction — Whether or not the flow is bi-directional.
- Zone Check — The zone checks defined in the flow
- Flow Definition Persistence — Whether or not to persist flow definition over device reboot.
- Data Type — The data type defined for the flow.
- Routing Control — The routing control defined in the flow.
- QOS — The Quality of Service (QOS) defined for the flow.
- Offset — The offset value defined in the flow.
- Originator — The FC originator defined for the flow.
- SCSI Commands — The SCSI command defined for the flow.
- Protocol — The protocol type defined in the flow.

NOTE
You can import flows on which the Monitor feature is enabled.

6. Click **OK** on the **Import Flow Definitions** dialog box.

   The imported flow displays in the Imported Flows group in the **Rules** area. You can now configure a rule and add it to the imported flow (refer to "Configuring a MAPS policy" on page 1232 and complete step 5 through step 13).

7. Click **OK** on the **Add Policy** dialog box.
8. Click **Close** on the **MAPS Configuration** dialog box.

Removing imported Flows

You can remove a flow from MAPS threshold monitoring.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select a policy in the list and click **Add**.

   The **Add Policy** dialog box displays.

3. Click the **Traffic Flows** tab.
4. Select the imported flow you want to remove and click **Remove**.
5. Click **OK** on the **Add Policy** dialog box.
6. Click **Close** on the **MAPS Configuration** dialog box.

Activating a MAPS policy

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select an inactive policy in the list and click **Activate**.

   Only one policy can be active on a switch at a time. You can activate policies for multiple switches at once by selecting the policy you want to activate for each switch and clicking **Activate**.

   When you edit the active policy on the switch, updated rules activate on the switch automatically.

   NOTE
   The active policy indicates an active status icon in the **Policies** column of the **Switch** row and **Active Policy** row.

3. Click **Close** on the **MAPS Configuration** dialog box.

Replicating a policy to other devices

You can replicate a non-default policy on a device to all MAPS-capable devices in a fabric or SAN .

NOTE
Copying a policy from one device to another overwrites any policy with a matching name on the target devices.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select a non-default policy on a device (source) you want to replicate in the list and click **Distribute**.

   The **Policy Distribution** dialog box displays (Figure 566).

FIGURE 566 Policy Distribution dialog box for SAN

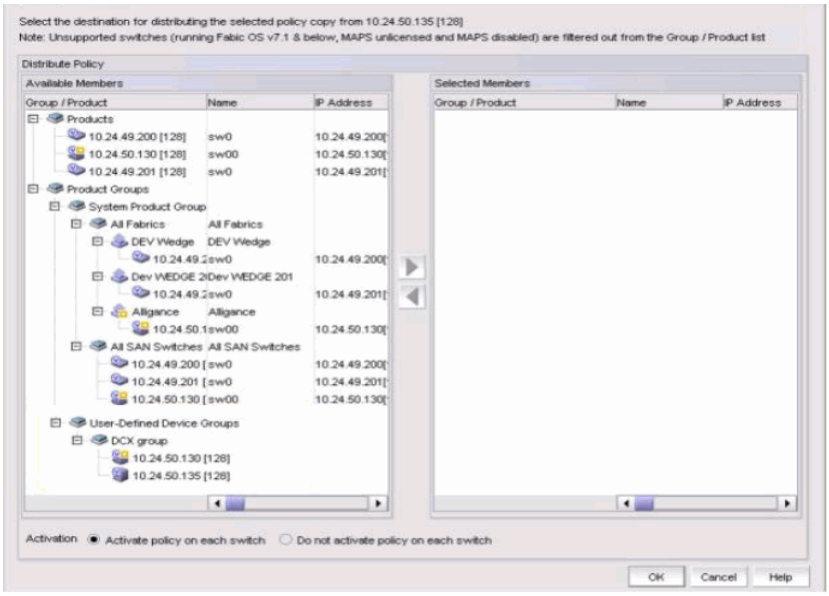

3. Select a group, product group, or an user-defined device group created in COMPASS, from the Available Members area and click the right arrow button to move the available members to the Selected Members area. You can replicate the policy as follows:

   - Select a specific switch in any fabric to replicate the policy on a specific switch in any fabric.
   - Select all switches in a single fabric to replicate the policy on all switches in the selected fabric.
   - Select all switches in all fabrics to replicate the policy on all switches in all fabrics.

4. Set the activation parameters by choosing one of the following options:

   - **Activate policy on each switch** — Select to immediately activate the policy on the target devices after distribution.

     If the selected policy is not an active policy, **Activate after distribution** activates the policy on the source device as well as the target devices.

   - **Do not activate policy on each switch** — Select to not activate the policy on the target devices after distribution.

5. Click **OK** on the **Policy Distribution** dialog box.

   The selected policy is replicated on all MAPS-capable devices in the selected fabric or SAN. If you chose to activate the policy after distribution, the selected policy is activated on the target devices and the source device, if necessary.

   NOTE
   If the fabric contains a switch running an earlier version of Fabric OS, the rules supporting the earlier version are discarded. Beginning with Fabric OS 8.1.0a or later, when you distribute any policy with RoR rule from Fabric OS 8.1.0a to earlier version, only base rule is distributed.

   NOTE
   A warning message is displayed if you try to replicate a policy name that already exists in the switch.

   NOTE
   The “Select any switch to distribute” information message displays if you click **OK** without selecting any switch for distribution.

6. Click **Close** on the **MAPS Configuration** dialog box.

Exporting a MAPS policy

You can export a policy to an xml file format.

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select the policy you want to export in the list and click **Export**.
3. Browse to the location you want to save the policy and click **Save**.
4. Click **Close** on the **MAPS Configuration** dialog box.

Importing a MAPS policy

You can import a policy with an xml file format to a device.

NOTE
You cannot import policies at a SAN or fabric level.

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select the device to which you want to import the policy and select the **Import** option (on the **Export** button list).
3. Browse to the location of the policy you want to import and click **Open**.

   You cannot import a policy with the same name as a default policy. The policy is imported to the selected device.

4. Click **Close** on the **MAPS Configuration** dialog box.

Deleting a MAPS policy

NOTE
You cannot delete default or active policies.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select the policies you want to delete in the list and click **Delete**.

   You can delete one or more policies from the same switch or multiple switches.

   A confirmation message displays.

3. Click **Yes** on the confirmation message.
4. Click **Close** on the **MAPS Configuration** dialog box.

Deleting MAPS rules for a custom group or imported flows

Beginning with the 12.4.0 release, automatic deletion is allowed for policy rules assigned to a custom group and flows imported as a custom group.

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select the custom group policies you want to delete from the list and click **Delete**.

   A confirmation dialog box with a list of policies and associated rules or flows displays (Figure 567).

   FIGURE 567 Policy Group Delete window

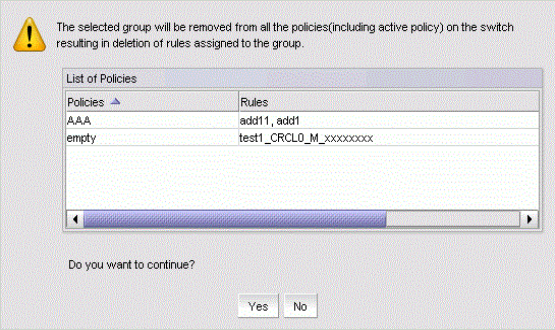

3. Click **Yes** to delete all the rules from the assigned policies on the switch or click **No** to abort the operation.

   NOTE
   Beginning with version 14.2.0 or later, when you de-associate the rules from a policy, the rules are deleted from the switches also.

Viewing MAPS policy rules

You can open more than one **View Policy** dialog box at the same time.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select a policy and click **View**.

   You can also select the switch in the list and select **View** to view the active policy. The **View Policy** dialog box displays (Figure 568).

FIGURE 568 View Policy dialog box for SAN

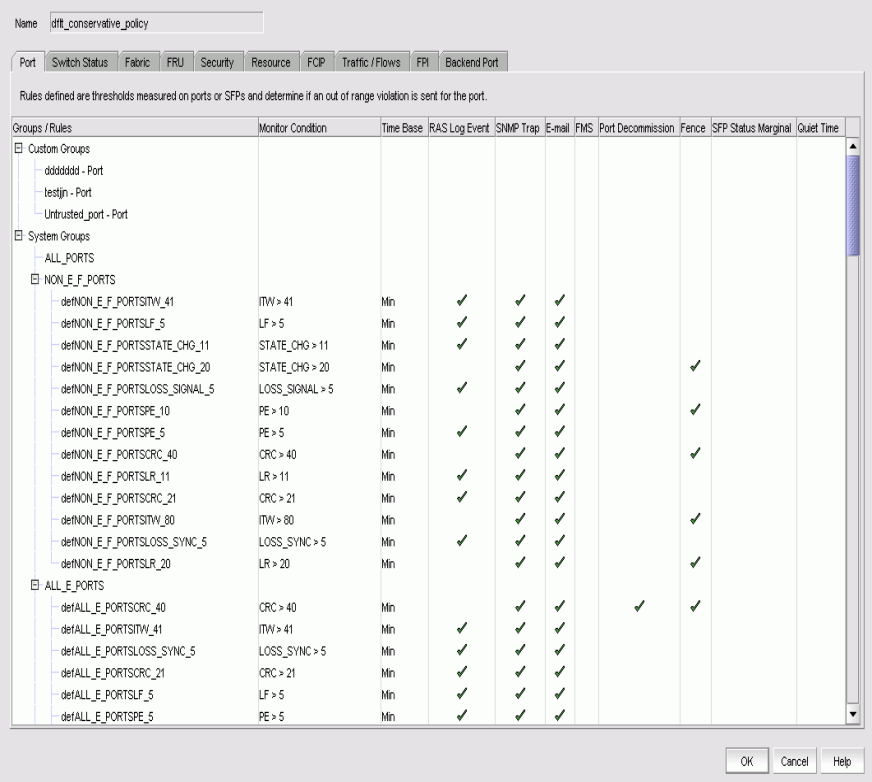

3. Select one of the category tabs to view the rules defined for the policy.

   For a complete list of categories and the associated measures and actions, refer to "MAPS categories, measures, and actions" on page 1210. Options include:

   - **Port** tab — Rules defined on this tab measure thresholds on ports or SFPs to determine if an out of range violation is sent for the port.
   - **Switch Status** tab — Rules defined on this tab measure thresholds at the switch or chassis level to determine the switch operational status.
   - **Fabric** tab — Rules defined on this tab measure thresholds at the switch level to detect out of range fabric-wide changes.
   - **FRU** tab — Rules defined on this tab measure thresholds on fans, power supplies, SFPs, blades, or WWN cards to detect out of range FRU changes.
   - **Security** tab — Rules defined on this tab measure thresholds at the switch level to detect out of range security changes.
   - **Resource** tab — Rules defined on this tab measure thresholds on temperature sensors or at the chassis level to detect out of range resource usage.
   - **FCIP** tab — Rules defined on this tab measure thresholds on FCIP circuits to detect out of range state, utilization, or packet loss.
   - **Traffic / Flows** tab — Rules defined on this tab measure thresholds on ports in the configured group to detect out of range link utilization or on flows to detect out of range flow changes.
   - **FPI** tab — Rules defined on this tab measures thresholds on the performances of the fabric.
   - **GigE Port** tab — Rules defined on this tab measure thresholds to monitor Gigabit Ethernet ports (GE) for ALL_EXT_GE_PORTS system group with default policy rules. GigE Port tab supports RAS Log Event, SNMP Trap, and E-mail actions.

- **Backend Port** tab — Rules on this tab measure thresholds to monitor the backend port health for ALL_BE_PORTS system group with default policy rules. Backend Port rules are predefined and obtained directly from the switch. You cannot create custom policy rules for the ALL_BE_PORTS system group; you can only select predefined rules for the ALL_BE_PORTS system group.

Each tab contains the following fields and components:

- **Rules** list — Lists the rules defined for the selected policy.
- **Groups/Rules** — Displays the default groups (under the **System Groups** node) and user-defined groups (under the **Custom Groups** node) for the selected switch. The available groups in the **Rules** table depend on the measure you selected in the **Add/Edit Threshold** area. For example, if you selected an SFP measure, only SFP groups become available. You can only configure a user-defined group on the **Port** and **FCIP** tabs. For more information, refer to "Configuring a group" on page 1248.
  Rules display below the appropriate group node based on rule targets.
- **Severity — Severity configured for the rule.**
- **Monitor Condition** — A combination of the measure, time base, and threshold expression for the rule.
- **Time Base** — The duration by which to monitor the measure.
- **Actions** check boxes — Actions configured for the rule. Each action has a column and each action selected for a rule has a check mark. Supported actions include:
  - **Status Marginal** (switch status)
  - **Status Critical** (switch status)
  - **RAS Log Event**
  - **Port Decommission**
  - **Fence (Port Fencing)**
  - **SNMP Trap**
  - **E-mail**
  - **SFP Marginal** (Port SFP status)
  - **FMS**
  - **SDDQ**
  - **Un-Quarantine**
  - **Toggle**

4. Click **OK** on the **View Policy** dialog box.
5. Click **Close** on the **MAPS Configuration** dialog box.

Comparing MAPS policies

You can compare any two policies across a fabric at a time. To compare MAPS policies, complete the following steps.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select two policies across the fabric that you want to compare and click **Compare**.

   The **Compare** dialog box displays (Figure 569).

FIGURE 569 Compare MAPS dialog box for SAN

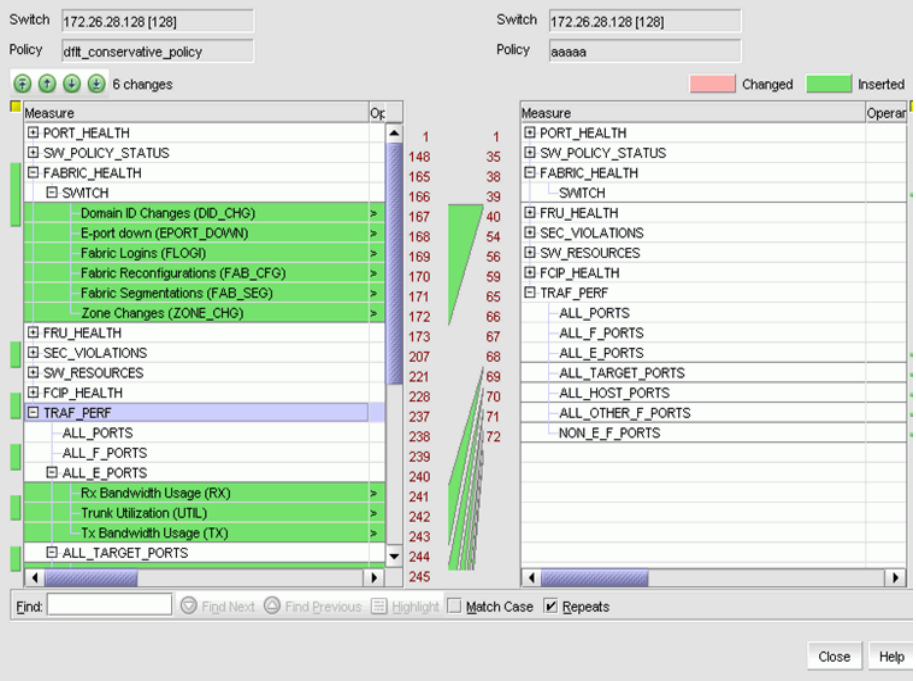

The **Compare** dialog box displays the following information:

- **Switch** - Displays the switch containing the selected policy.
- **Policy** -Displays the name of the selected policy.
- **Change Navigator** buttons - These buttons help to navigate between the differences in the policy contents.
- **Differences Legend** - Displays the color legend for differences in policy contents.
  - **Changed** status displays in peach - When the rule is present in the other policy with differences in a few values.
  - **Inserted** status displays in green - When the rule is not present in the other policy.
- **Find** tool bar - Displays the text entered for the selected policy.

3. Click **Close**.

MAPS groups

A MAPS group is a collection of similar objects that you can monitor as a single entity.

You can create a group of objects and then use that group in rules, thus simplifying rule configuration and management. For example, you can create a group of UNIX ports, and then create specific rules for monitoring this group.

Preconfigured groups

MAPS provides several preconfigured groups. You cannot edit or delete a preconfigured group. You can add user-defined rules to the preconfigured groups. For more information, refer to "Configuring a MAPS policy" on page 1232.

For Network OS devices you cannot edit or delete the pre-dfined groups except ALL_iSCSI_PORTS, ALL_NAS_PORTS and ALL_DAS_PORTS. For more information, refer to "Configuring a MAPS policy" on page 1232.

Table 122 lists the Fabric OS preconfigured groups and their descriptions.

TABLE 122 Preconfigured groups for Fabric OS

Preconfigured group name	Element type	Description
ALL_PORTS	FC Port	All FC ports physically present in the logical switch.
ALL_F_PORTS	FC Port	All F_Ports present in the logical switch, including all ports in F_Port trunks.
ALL_E_PORTS	FC Port	All E_Ports and EX_Ports present in the logical switch, including all ports in E_Port and EX_Port trunks.
ALL_D_PORTS	FC Port	All D_Ports (Diagnostic ports) present in the logical switch.
ALL_TARGET_PORTS	FC Port	All logical switch ports connected to targets. MAPS automatically detects if a device connected on this port is a target port and adds it to this set.
ALL_HOST_PORTS	FC Port	All ports in the logical switch connected to hosts. MAPS automatically detects if a device connected on this port is a server port and adds it to this set.
ALL_SFP	FC Port	All gigabit interface converters (GBIC) and SFP transceivers present in the logical switch.
ALL_QSFP	FC Port	All QSFP transceivers present in the logical switch.
ALL_32GSWL_SFP	FC Port	All 32 Gbps Short Wavelength (SWL) SFP transceivers present in logical switch.
ALL_25Km_16GLWL_SFP	FC Port	All 25Km_16 Gbps Long Wavelength (LWL) SFP transceivers present in logical switch.
ALL_16GLWL_SFP	FC Port	All 16 Gbps Long Wavelength (LWL) SFP transceivers present in logical switch.
ALL_16GSWL_SFP	FC Port	All 16 Gbps Short Wavelength (SWL) SFP transceivers in logical switch.
ALL_10GLWL_SFP	FC Port	All 10 Gbps LWL SFP transceivers on FC ports in logical switch.
ALL_10GSWL_SFP	FC Port	All 10 Gbps SWL SFP transceivers on FC ports in logical switch.
ALL_EXT_GE_PORTS	GigE Port	All GigE Port physically present in the logical switch.
ALL_SLOTS **NOTE:** Beginning with Fabric OS 8.1.0 or later, ALL_SLOTS group is removed from non-chassis system.	Slot	All slots present in the chassis.
ALL_SW_BLADES **NOTE:** Beginning with Fabric OS 8.1.0 or later, ALL_SW_BLADES group is removed from non-chassis system.	Blade	All port and application blades in the chassis.

TABLE 122 Preconfigured groups for Fabric OS (Continued)

Preconfigured group name	Element type	Description
ALL_CORE_BLADES **NOTE:** Beginning with Fabric OS 8.1.0 or later, ALL_CORE_BLADES group is removed from non-chassis system.	Blade	All core blades in the chassis.
ALL_PS	Power Supply	All power supplies present in the chassis.
ALL_TS	Temperature Sensor	All temperature sensors present in the chassis.
ALL_FAN	Fan	All fans present in the chassis.
ALL_CIRCUITS **NOTE:** Beginning with Fabric OS 8.1.0 or later, ALL_CIRCUITS group is removed from non-chassis system.	Circuit	All FCIP circuits present in the logical switch.
ALL_TUNNELS[1]	Tunnel	All FCIP tunnels present in the logical switch.
ALL_TUNNEL_HIGH_QOS	Tunnel	All tunnel high QoS high monitors 50% of the available bandwidth.
ALL_TUNNEL_MED_QOS	Tunnel	All tunnel high QoS high monitors 30% of the available bandwidth.
ALL_TUNNEL_LOW_QOS	Tunnel	All tunnel high QoS high monitors 20% of the available bandwidth.
ALL_TUNNEL_F_QOS	Tunnel	All tunnel F QoS monitors bandwidth at the expense of the lowest priority.
SWITCH	Switch	Default group used to define rules on global parameters for the entire switch; for example, security violations or fabric health.
CHASSIS	Chassis	Default group used to define rules on global parameters for the entire chassis; for example, CPU, Flash, and so on.
ALL_FLASH	Flash	All monitored flash.
ALL_WWN	WWN	All monitored WWN cards.
ALL_2K_QSFP	SFP	All 2K QSFP transceivers present in the logical switch.
ALL_100M_16GSWL_QSFP	SFP	All 100M 16 Gbps SWL QSFP transceivers in the logical switch.
ALL_QUARANTINED_PORTS	FC Port	All ports in the logical switch which have been quarantined for slow-drain performance.
ALL_CIRCUIT_F_QOS	Circuit	All circuit F QoS monitors bandwidth at the expense of the lowest priority.
ALL_CIRCUIT_HIGH_QOS	Circuit	All circuit high QoS high monitors 50% of the available bandwidth.
ALL_CIRCUIT_MED_QOS	Circuit	All circuit high QoS high monitors 30% of the available bandwidth.

TABLE 122 Preconfigured groups for Fabric OS (Continued)

Preconfigured group name	Element type	Description
ALL_CIRCUIT_LOW_QOS	Circuit	All circuit high QoS high monitors 20% of the available bandwidth.
ALL_BE_PORTS **NOTE:** Beginning with Fabric OS 8.1.0 or later, ALL_BE_PORTS group is removed from non-chassis system.	Ports	All back-end ports in the physical switch.
ALL_ASICS	Asic	All ASIC chip in the physical switch.
ALL_LOCAL_PIDS	Ports	All local PIDS in the physical switch.
ALL_TUNNEL_IP_HIGH_QOS[1]	Tunnel	All tunnel IP high QoS high monitors 50% of the available bandwidth.
ALL_TUNNEL_IP_MED_QOS[1]	Tunnel	All tunnel IP high QoS high monitors 30% of the available bandwidth.
ALL_TUNNEL_IP_LOW_QOS[1]	Tunnel	All tunnel IP high QoS high monitors 20% of the available bandwidth.
ALL_CIRCUIT_IP_HIGH_QOS[1]	Circuit	All circuit IP high QoS high monitors 50% of the available bandwidth.
ALL_CIRCUIT_IP_MED_QOS[1]	Circuit	All circuit IP high QoS high monitors 30% of the available bandwidth.
ALL_CIRCUIT_IP_LOW_QOS[1]	Circuit	All circuit IP high QoS high monitors 20% of the available bandwidth.
ALL_CERTS	Security	All security certificates

1. The feature is supported only on 16 Gbps 24-port, 18 GbE port switch and FX8-24.

User-defined groups

NOTE
You can only create user-defined custom groups for ports, SFPs,and FCIP circuits.

You can create a group of ports, circuits, or SFPs to which you can assign thresholds. This enables you to configure different monitoring conditions for each group. For more information, refer to "Configuring a group" on page 1248.

Configuring a group

Often on a device there are sets of ports that behave in a similar manner and have a different behavior from other sets of ports. For example, the behavior of ports connected to UNIX hosts and servers is different from the behavior of ports connected to Windows hosts and servers.

MAPS allows you to group ports, SFP transceivers, or FCIP circuits together across network devices You can create groups of ports that behave in a similar manner and monitor these ports using the same rules and thresholds.

NOTE
You can create up to 64 groups for each logical switch.

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Click **Add**.

   The **Add Policy** dialog box displays.

3. Choose one of the following options:

   a. (**Port** tab) Create a port group by clicking **Add > Port** in the **Custom Groups** area.

      The **Add Port Group** dialog box displays (Figure 570).

FIGURE 570 Add Port Group dialog box for SAN

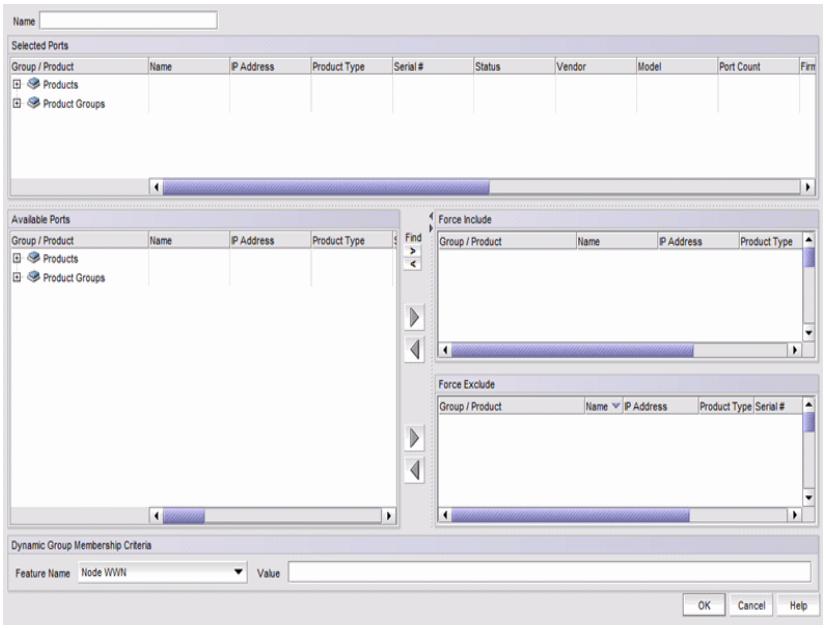

FIGURE 571

i. Enter a unique name for the group in the **Name** field.
   The name can be up to 32 characters and can only contain alphanumeric and underscore characters.

ii. Add objects to the group by selecting the object (port or circuit) in the **Available Ports/Circuits** list and clicking the right arrow button.
   The selected objects move from the **Available Ports/Circuits** list to the **Force Include** list.

iii. Remove objects from the group by selecting the object (port or circuit) in the **Available Ports/ Circuits** list and clicking the right arrow button.
   The selected objects move from the **Available Ports/Circuits** list to the **Force Exclude** list.

iv. To remove ports from **Force Include** list, select the ports in the **Force Include** list and click the left arrow button. The ports move to the **Available Ports/Circuits** list. Select the port that is moved from **Force Include** list and click the right arrow button of **Force Exclude** list.

v. Enter a value in the **Dynamic Group Membership Criteria** area, to select the feature of the group:

   - Select **Node WWN** to create a port group based on the connected device node WWN.
   - Select **Port Name** to create a port group based on the name of the port.

b. (**Port** tab) Create an SFP group by clicking **Add > SFP** in the **Custom Groups** list.

c. (**FCIP** tab) Create an FCIP circuit group by clicking **Add** in the **Custom Groups** list.

The **Add Group** dialog box displays (Figure 572).

FIGURE 572 Add Group dialog box for SAN

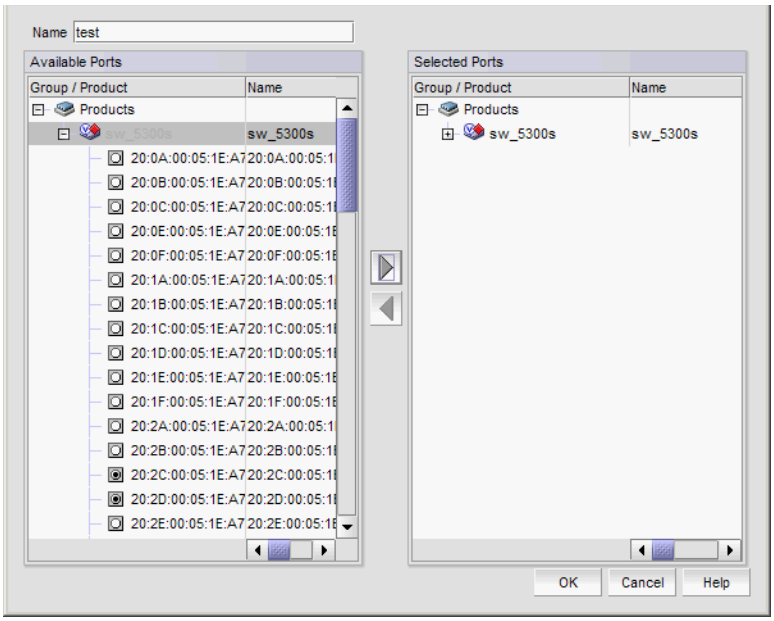

i. Enter a unique name for the group in the **Name** field.
The name can be up to 32 characters and can only contain alphanumeric and underscore characters.

ii. Add objects to the group by selecting the object (port or SFPor circuit) in the **Available Ports/Circuits** list and clicking the right arrow button.
The selected objects move from the **Available Ports/Circuits** list to the **Selected Ports/Circuits** list.

iii. Remove objects from the group by selecting the object (port or SFP or circuit) in the **Selected Ports/Circuits** list and clicking the left arrow button.
The selected objects move from the **Selected Ports/Circuits** list to the **Available Ports /Circuits** list.

4. Click **OK** on the **Add Group** dialog box.

The new group displays in the **Custom Groups** folder of the **Rules** area.

5. Configure policies and rules for the group. For more information, refer to "Configuring a MAPS policy" on page 1232.

6. Click **OK** on the **Add Policy** dialog box.

7. Click **Close** on the **MAPS Configuration** dialog box.

Editing a group

If a new object, such as host, target, or SFP transceiver is added to a fabric, you can monitor the object using existing rules for similar objects.

The group must be the same type as the new object you want to monitor (port, circuit, or SFP).

The object is automatically monitored using the existing rules that have been set up for the group, as long as the rules are in the active policy. You do not need to re-enable the active policy.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select the policy associated with the group you want to edit and click **Edit**.

   The **Edit Policy** dialog box displays.

3. (**Port or FCIP**, or **FPI** tab only) Select the group you want to edit in the **Rules** area and click **Edit** in the **Custom Groups** area.

   The **Edit Group** dialog box displays (Figure 573) .

   FIGURE 573 Edit Group dialog box for SAN

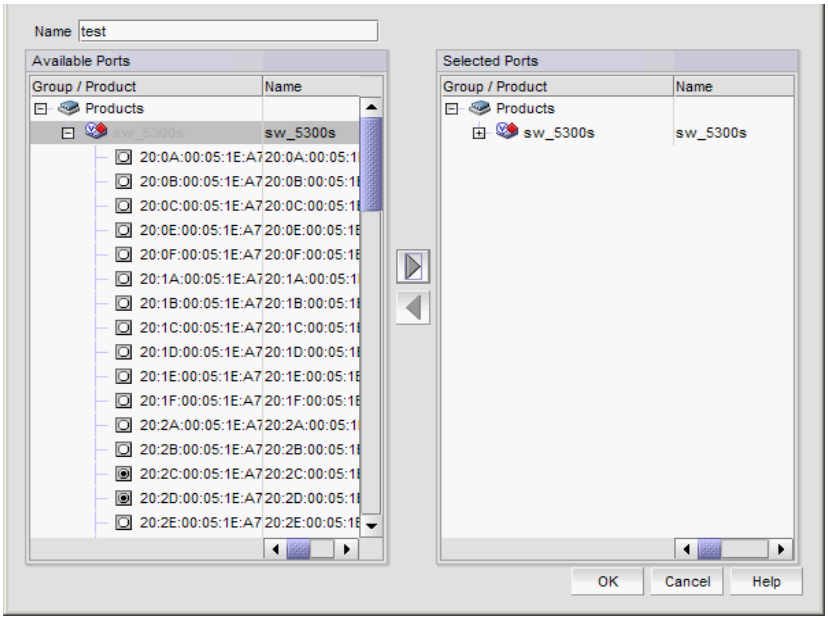

4. Add objects to the group by selecting the object (port or SFP or circuit) in the **Available Ports/ Circuits** area and clicking the right arrow button.

   The selected objects move from the **Available Ports/Circuits** area to the **Selected Ports/ Circuits** area.

5. Remove objects from the group by selecting the object (port or SFP or circuit) in the **Selected Ports/ Circuits** area and clicking the left arrow button.

   The selected objects move from the **Selected Ports/Circuits** area to the **Available Ports/Circuits** area.

6. Click **OK** on the **Edit Group** dialog box.

7. Configure policies and rules for the group. For more information, refer to "Configuring a MAPS policy" on page 1232.

8. Click **OK** on the **Edit Policy** dialog box.

9. Click **Close** on the **MAPS Configuration** dialog box.

Deleting a group

NOTE
You cannot delete a default group or any group that contains a rule.

1. Select the **SAN** tab or right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select the policy associated with the group you want to delete and click **Edit**.

   The **Edit Policy** dialog box displays.

3. Select the **Port** tab (depending on the type of group you want to delete).
4. Select the custom group you want to delete in the **Rules** area and click **Delete** in the **Custom Groups** area.
5. Click **Yes** on the confirmation message.
6. Click **OK** on the **Edit Policy** dialog box.
7. Click **Close** on the **MAPS Configuration** dialog box.

Managing MAPS groups

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select a fabric or device in the **All Fabrics /Fabric/Switch/** list and click **Manage**.

   The *Fabric/Device _Name* - **Manage MAPS Groups** dialog box (for Fabric OS devices only) displays (Figure 574) with a list of all configured SFP, or FCIP groups on the selected fabric or device in the **Groups** area.

FIGURE 574 Fabric/Device_Name - Manage MAPS Groups dialog box for SAN

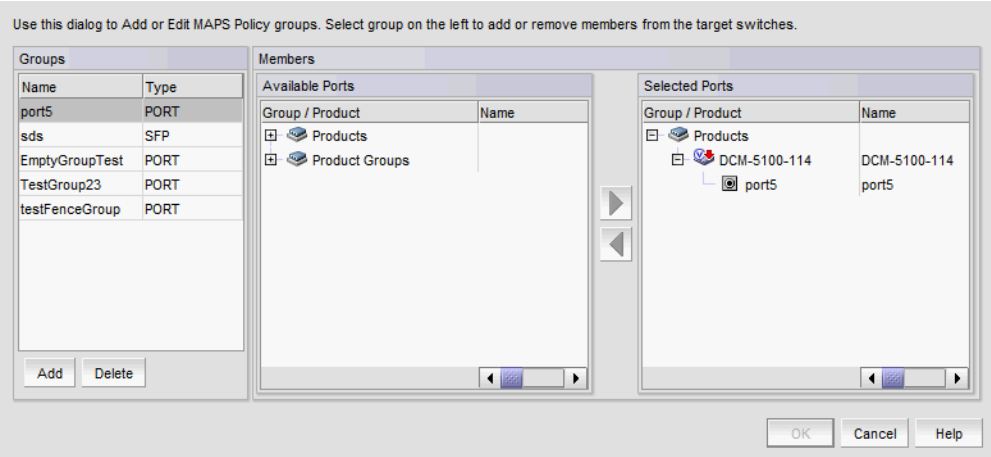

3. Review the group details:

   Sort the contents by clicking the column header (**Name** or **Type**). Click the same column header again to reverse the sort order.

4. Review the group details:

   Sort the contents by clicking the column header (**Name** or **Type**). Click the same column header again to reverse the sort order.

   - **Groups** list — List of groups available on the selected fabric or device.
     - **Name** — Group name
     - **Type** — Group type (Port, SFP, or Circuit)
   - **Available Ports/Circuits** list — List of available ports, SFPs, or circuits and the associated products for the selected group.
     - **Group/Product** — Available devices and ports
     - **Name** — Device name, port name, or circuit name
     - **IP Address** — IP address of the device
     - **Product Type** — Product type (such as switch or blade)
   - Left and right arrow buttons — Click to move ports, SFPs, or circuits between the **Selected Ports/Circuits** list and **Selected Ports/Circuits** list.
   - **Selected Ports/Circuits** list — List of selected ports, SFPs, or circuits and the associated products for the selected group.
     - **Group/Product** — Selected devices and ports
     - **Name** — Device name, port name, or circuit name
     - **IP Address** — IP address of the device
     - **Product Type** — Product type (such as switch or blade)
   - **Add** button — Click to add a group to the **Groups** list. To create one or more groups on the selected device or fabric, refer to "Creating multiple groups" on page 1253.
   - **Delete** button — Click to delete the selected group from the **Groups** list ("Deleting a group" on page 1254).

5. Click **OK** on the *Fabric Device _Name* - **Manage MAPS Groups** dialog box.
6. Click **Close** on the **MAPS Configuration** dialog box.

Creating multiple groups

You can create groups that are in the same fabric or device.

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select a fabric or device in the **All Fabrics /Fabric/Switch/** list and click **Manage**.

   The *Fabric/Device _Name* - **Manage MAPS Groups** dialog box (for Fabric OS devices only) displays with a list of all configured SFP, or FCIP groups on the selected fabric or device in the **Groups** area.

3. Click **Add**.
4. The **Add Group** dialog box displays Enter a unique name (maximum 32 characters) for the group in the **Name** field.
5. Select the type of group you want to create from the **Type** list.

   Options include port or SFP, or Circuit.

6. Add objects to the group by selecting the object (port or SFP or circuit) in the **Available Ports / Circuits** list and clicking the right arrow button.

   The selected objects move from the **Available Ports / Circuits** list to the **Selected Ports/Circuits** list.

7. Repeat step 3 through step 6 for each group you want to add.
8. Click **OK** on the *Fabric Device _Name* - **Manage MAPS Groups** dialog box.

9. Click **Close** on the **MAPS Configuration** dialog box.

Editing multiple groups

You can edit one or more groups that are in the same fabric or device.

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select a fabric or device in the **All Fabrics /Fabric/Switch/**list and click **Manage**.

   The *Fabric/Device _Name* - **Manage MAPS Groups** dialog box displays with a list of all configured SFP, or FCIP groups on the selected fabric or device in the **Groups** area.

3. Select the group you want to edit from the **Groups** area.

   Sort the contents by clicking the column header (**Name** or **Type**) to find the group you want to edit. Click the same column header again to reverse the sort order.

   The available and selected ports, SFPs, or circuits display in the **Available Ports/Circuits** list and the **Selected /Circuits** list.

4. Add objects to the group by selecting the object (port or SFP, or Circuit) in the **Available Ports/ Circuits** list and clicking the right arrow button.

   The selected objects move from the **Available Ports/Circuits** list to the **Selected Ports/ Circuits** list.

5. Remove objects from the group by selecting the object (port, SFP, or circuit) in the **Selected Ports/ Circuits** list and clicking the left arrow button.

   The selected objects move from the **Selected Port/ Circuits** list to the **Available Ports/ Circuits** list.

6. Repeat step 2 through step 5 for each group you want to edit.
7. Click **OK** on the *Fabric Device _Name* - **Manage MAPS Groups** dialog box.
8. Click **Close** on the **MAPS Configuration** dialog box.

Deleting a group

You can delete a group from the fabric or device.

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **MAPS Configuration** dialog box displays.

2. Select a fabric or device in the **All Fabrics /Fabric/Switch/**list and click **Manage**.

   The *Fabric/Device _Name* - **Manage MAPS Groups** dialog box displays with a list of all configured SFP, or FCIP groups on the selected fabric or device in the **Groups** area.

3. Select the group you want to delete in the **Groups** list.

   A confirmation message displays.

4. Click **Yes** on the confirmation message.

   The selected group is deleted from **Groups** list.

5. Repeat step 3 and step 4 for each group you want to delete.

6. Click **OK** on the *Fabric Device _Name* - **Manage MAPS Groups** dialog box.
7. Click **Close** on the **MAPS Configuration** dialog box.

MAPS violations

MAPS violation data is stored in the database for 30 days. The system purges old data (over 30 days) every night at 12:00 AM. The system also purges violations from deleted or unmonitored devices.

Viewing MAPS violations

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Violations**.

   The **Violations** dialog box displays (Figure 575).

   FIGURE 575 Violations dialog box

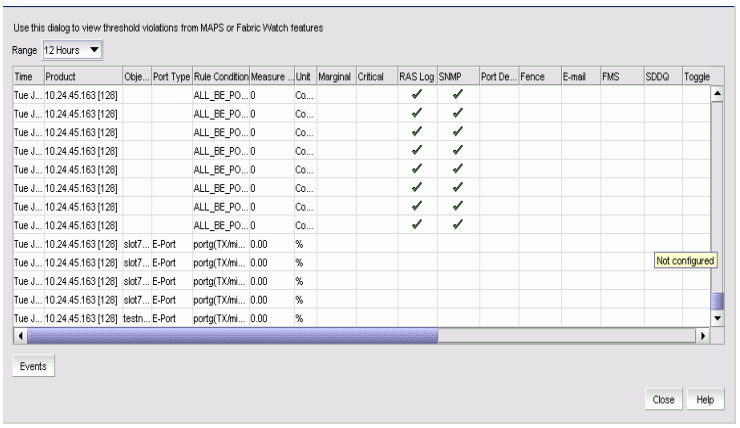

2. Display data for a specific duration by selecting one of the following options from the **Range** list:
   - **30 Minutes** (default) — Displays data for the previous half hour beginning when the **Violations** dialog box is displayed.
   - **1 Hour** — Displays data for the previous hour beginning when the **Violations** dialog box is displayed.
   - **6 Hours** — Displays data for the previous 6 hours beginning when the **Violations** dialog box is displayed.
   - **12 Hours** — Displays data for the previous 12 hours beginning when the **Violations** dialog box is displayed.
   - **1 Day** — Displays data for the previous day beginning when the **Violations** dialog box is displayed.
   - **3 Days** — Displays data for the previous 3 days beginning when the **Violations** dialog box is displayed.
   - **1 Week** — Displays data for the previous week beginning when the **Violations** dialog box is displayed.
   - **1 Month** — Displays data for the previous month beginning when the **Violations** dialog box is displayed**.**
3. Review the detailed data.

   You can sort the contents by clicking the column header. Click the same column header again to reverse the sort order.
   - **Time** (MAPS and Fabric Watch support) — The time on the server when the violation was reported.
   - **Fabric Name** (MAPS and Fabric Watch support) — The Fabric name to which the object belongs.
   - **Product** (MAPS and Fabric Watch support) — The device name.

- **Object Name** (MAPS and Fabric Watch support) — The object name (such as switch name, port name, FRU name, and so on).
- **Category** (MAPS and Fabric Watch support) — The category of the measure violated corresponding to the **Dashboard** tab and configuration dialog boxes.
- **Rule Name** (MAPS only support) — The rule name.
- **Rule Condition** (MAPS and Fabric Watch support) — Associates the condition with the action triggered when the condition occurs.
- Action cells (MAPS only support) — Actions taken as a result of rule violation. Each action has a column and actions triggered for a rule have a check mark. Possible values for the Action cells include:
  - Green check mark — The action is configured and enabled on the device. Tool tip displays as "Configured and triggered".
  - Disabled — The action is configured, but disabled, on the device. Tool tip displays as "Configured in rule, but disabled at switch".
  - Empty — The action is not configured on the device. Tool tip displays as "Not configured".
  - Greyed-out — Data cannot be determined because event collection occurred during discovery (not MAPS violation). Tooltip displays as "Unknown".

  Supported actions include:

  - **Marginal** (switch status)
  - **Critical (switch status)**
  - **RAS Log**
  - **SNMP**
  - **Port Decommission**
  - **Fence** (Port Fencing)
  - **E-mail**
  - **SFP Marginal** (Port SFP status)
  - **FMS**
  - **SDDQ**
  - Un-Quarantine
  - **Toggle**
- **Measure Value** (MAPS and Fabric Watch support) — Value of the measure when the violation occurred.
- **Units** (MAPS and Fabric Watch support) — The units description of the measure value.
- **Recommended Action** (MAPS and Fabric Watch support) — Fabric OS recommended action for the violation. You can wrap text in this column by right-clicking the column header and selecting the **Wrap** check box.

4. Select one or more violations and click **Events** to launch the **MAPS Violation Master Log Events** dialog box (refer to "Viewing MAPS events" on page 1257).
5. Click **Close**.

MAPS events

Once you configure MAPS rule violations to trigger RASLOG messages, the Management application starts receiving SNMP traps for the MAPS rule violations. The Management application processes the RASLOG messages by an event processor and displays them in the Master Log and the historical graphs and monitors the same as any other events.

The event processor also receives dashboard change events in the form of traps from the switch for all rule violations whether you configure the rule to trigger a RASLOG action or not. The Management application uses these notifications internally to process and persist MAPS violations information. These notifications also display in the Master Log and historical graphs and monitors although they may not be seen in the switch RASLOG.

You can determine event severity by the event icons that display on the historical graphs and monitors and Master Log. The following table lists the event icons that display on the historical graphs and monitors and Master Log. For more information about events, refer to the "Fault Management" on page 1131.

TABLE 123 Event severity icon

Event Icon	Description
	Critical
	Error
	Warning
	Informational

Viewing MAPS events

MAPS events allows you to view the events that occur before and after a violation. You can display Master Log events for all MAPS violations at the product level. Port violations display Master Log events at the product level only.

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

   The **Violations** dialog box displays.

2. Display data for a specific duration by selecting one of the following options from the **Range** list:
   - **30 Minutes** (default) — Displays data for the previous half hour beginning when the **Violations** dialog box is displayed.
   - **1 Hour** — Displays data for the previous hour beginning when the **Violations** dialog box is displayed.
   - **6 Hours** — Displays data for the previous 6 hours beginning when the **Violations** dialog box is displayed.
   - **12 Hours** — Displays data for the previous 12 hours beginning when the **Violations** dialog box is displayed.
   - **1 Day** — Displays data for the previous day beginning when the **Violations** dialog box is displayed.
   - **3 Days** — Displays data for the previous 3 days beginning when the **Violations** dialog box is displayed.
   - **1 Week** — Displays data for the previous week beginning when the **Violations** dialog box is displayed.
   - **1 Month** — Displays data for the previous month beginning when the **Violations** dialog box is displayed.
3. Select one or more rows in the **Violations** dialog box and click **Events**.

   The **MAPS Violation Master Log Events** dialog box displays (Figure 576) .

FIGURE 576 MAPS Violation Master Log Events dialog box for SAN

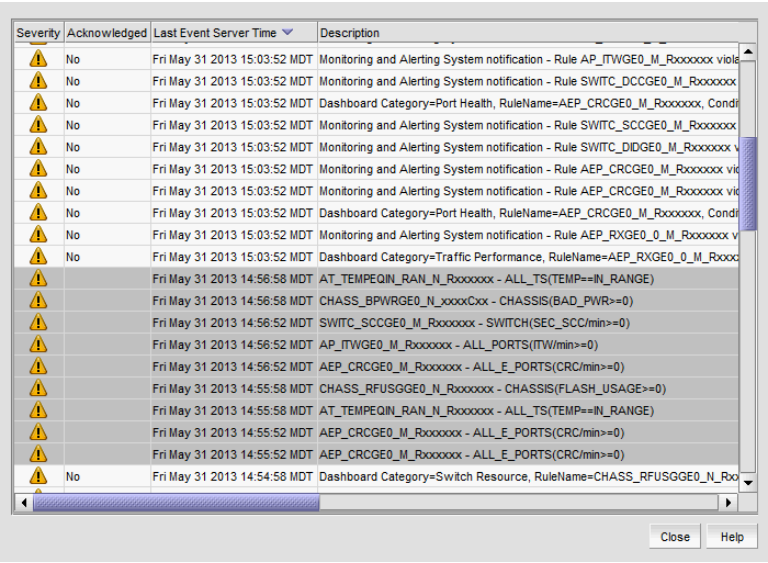

The events display for the selected time range (50% of the events before the selected violations and 50% after) up to a maximum of 200 event rows. For example, if you select 1 MAPS violation and set the time range to 1 hour, events display for 30 minutes before and after the selected violations.

If the number of events within the selected the time range exceeds the maximum number of events (200), the time range changes for the maximum number of events. For example, if you selected 1 hour as the time range but the maximum number of events occurred within 30 minutes, then events display for 15 minutes before and after the selected violations.

4. Review the detailed data.

   The **MAPS Violation Master Log Events** dialog box contains the same fields as the Master Log; however, the MAPS violations only displays content in the MAPS related fields.

   You can sort the contents by clicking the column header. Click the same column header again to reverse the sort order.

   TABLE 124 MAPS violation master log event fields

Event field	Description
Severity	The MAPS event severity is Warning.
Acknowledged	Whether the event is acknowledged or not.
Last Event Server Time	The time range selected in the **MAPS Violations** dialog box.
Description	The rule name and condition.
Origin	The event source type
Source Name	The source name of the event.
Source Address	The IP address (IPv4 or IPv6 format) of the product on which the event occurred.
Category	The event category is MAPS.
Count	The number of times the violation occurred on the device.
Module Name	The name of the module on which the event occurred.
Message ID	The message ID of the event.
Product Address	The IP address of the product on which the event originated.

Event field	Description
Contributor	The name of the contributor on which the event occurred.
Node WWN	The world wide name of the node on which the event occurred.
Fabric Name	The SAN fabric name.
Port Name	The port name on which the violation occurred.
Operational Status	The operational status (such as, unknown, healthy, marginal, or down) of the product on which the event occurred.
First Event Product Time	The time and date the event first occurred on the product.
Last Event Product Time	The time and date the event last occurred on the product.
First Event Server Time	The time and date the event first occurred on the server.
Audit	The audit of the event.
Virtual Fabric ID	The virtual fabric identifier.

5. Click **Close** on the **MAPS Violation Master Log Events** dialog box.

MAPS integration with other features

Dashboard MAPS widgets

The MAPS widgets display on the main **Dashboard** tab (refer to "Monitoring and Alerting Policy Suite widgets" on page 229). The Management application provides the following preconfigured MAPS widgets:

- Out of Range Violations widget — Table view of all out-of-range threshold violations reported in your SAN (refer to "Out of Range Violations widget" on page 229).
- Port Health Violations widget — Table view of out-of-range port health violations (refer to "Port Health Violations widget" on page 231). There are four port health violation widgets: All, ISL, Initiator, and Target.

Master Log

The Master Log displays MAPS events the same as any other events. MAPS events display in the following format:

severity="warning" message="Monitoring and Alerting System notification - Rule ***rule_name*** violated. Obj: object_number/name-from_trap"

To view detailed information for an event, refer to "Displaying event properties from the Master Log" on page 1193.

Performance graphs and monitors

You can enable events on historical graphs and monitors. For instructions, refer to "", as shown in Figure 707 on page 1710Configuring the performance graph display" on page 1002.

Once enabled, if the Management application receives any MAPS violation events during the time range specified on the historical graph or monitor, event icons (indicating the severity) display on the historical graphs and monitors. Place the cursor on an event icon to view the event details. MAPS event details include the following information:

- Time base

- Switch or port information
- Name of the rule with a violation
- Condition of the rule that caused a violation

Technical Support

Server and client support save

You can use Technical Support to collect SupportSave data for the Management server and clients.

Server Support save data includes:-

- Engineering logs
- Events
- Configuration files
- Operating system-specific information
- Environment information
- Vital CPU, memory, network resources
- Agent and driver logs
- Install logs
- Core files
- Database (partial or full)
- Web Tools data

Client Support save data includes:-

- Client Log Files
- Client data model log

Capturing Server and Client support save data

To capture both server and client support save files, complete the following steps.

1. Select **Monitor > Technical Support > SupportSave**.

   The **SupportSave** dialog box displays.

2. Select the **Server SupportSave** check box to run supportsave on the server.
3. Enter a file name for the server support save file in the **File Name** field.

   The default file name is DCM-SS-*Time_Stamp*.

4. Select the **Include Database** check box to include the database in the support save and choose one of the following options.

- Select the **Partial** (Excludes historical performance data and events) option to exclude historical performance data and events from the database capture.
- Select the **Full** option to capture the entire database.

Clear the **Include Database** check box to exclude the database in the support save.

5. Select the **Client SupportSave** check box to run supportsave on the client.
6. Enter a file name for the client support save file in the **File Name** field.

   The default file name is DCM-Client-SS-*Time_Stamp*.
7. Click **OK** on the **SupportSave** dialog box.
8. Click **OK** on the message.

   A progress message displays with a list of the steps to be performed:

   - Capturing client support save
   - Capturing logs and server data
   - Capturing partial/full database
   - Capturing data from the products

   You cannot close the progress message and you cannot perform any other actions until the SupportSave is complete.

   The application generates separate master logs to show the status of the Server and Client Support save collection.

   You cannot change the destination directory for Server and Client support save. Here are the default directories:

   - Server Support save location: *Install_Home*/support
   - Client Support save locations:
   - (Local client) *User_Home*/*Management_Application_Name*/localhost/support
   - (Remote client) *User_Home*/*Management_Application_Name*/Server IP/support

   NOTE
   Server support save initiated from the remote client is only available from a client installed on the server. However, you can copy the server support save from the **View Repository** dialog box (using the **Save** button) to the remote client location.

Capturing Server support save data

To capture server support save files, complete the following steps.

1. Select **Monitor > Technical Support > SupportSave**.

   The **SupportSave** dialog box displays.
2. Select the **Server SupportSave** check box to run supportsave on the server.
3. Make sure the **Client SupportSave** check box is clear.
4. Enter a file name for the server support save file in the **File Name** field.

   The default file name is DCM-SS-*Time_Stamp*.
5. Select the **Include Database** check box to include the database in the support save and choose one of the following options.

- Select the **Partial** (Excludes historical performance data and events) option to exclude historical performance data and events from the database capture.
- Select the **Full** option to capture the entire database.

  NOTE
  Selecting the **Full** option may increase the time needed for the SupportSave to complete.

Clear the **Include Database** check box to exclude the database in the support save.

6. Click **OK** on the **SupportSave** dialog box.
7. Click **OK** on the message.

   A progress message displays with a list of the steps to be performed:

   - Capturing logs and server data
   - Capturing partial/full database
   - Capturing data from the products

   You cannot close the progress message and you cannot perform any other actions until the SupportSave is complete.

   The application generates separate master logs to show the status of the Server Support save collection.

Capturing Client support save data

To capture client support save files, complete the following steps.

1. Select **Monitor > Technical Support > SupportSave**.

   The **SupportSave** dialog box displays.

2. Select the **Client SupportSave** check box to run supportsave on the client.
3. Make sure the **Server SupportSave** check box is clear.
4. Enter a file name for the client support save file in the **File Name** field.

   The default file name is DCM-Client-SS-*Time_Stamp*.

5. Click **OK** on the **SupportSave** dialog box.
6. Click **OK** on the message.

   A progress message displays with a the step to be performed: Capturing client support save.

   You cannot close the progress message and you cannot perform any other actions until the SupportSave is complete.

   The application generates separate master logs to show the status of the Client Support save collection.

Client support save using a command line interface

Use the following procedures to capture client support save files through the command line interface (CLI).

Capturing client support save using the CLI (Windows)

To capture client support save files through the CLI, complete the following steps.

1. Go to the following location:
   - (Local client) *User_Home*/*Management_Application_Name*/localhost
   - (Remote client) *User_Home*/*Management_Application_Name*/Server IP
2. Run the clientsupportsave.bat file.
3. Define a capture location by typing `clientsupportsave` *<path>* in the CLI. If the path has spaces, enclose it in double quotes.

   By default, the capture location is one of the following:
   - (Local client) *User_Home*/*Management_Application_Name*/localhost
   - (Remote client) *User_Home*/*Management_Application_Name*/Server IP
4. Use an archive tool to create a ZIP file of the support save.

Capture client support save using the CLI (Linux)

To capture client support save files through the CLI, complete the following steps.

1. Go to /root /*Management_Application_Name_Folder*/Server IP.
2. Run the clientsupportsave.sh file.
3. Define a capture location by typing `sh clientsupportsave` *<path>* in the CLI. If the path has spaces, enclose it in double quotes.

   By default, the capture location is /root /*Management_Application_Name_Folder*/Server IP/support.
4. Use an archive tool to create a ZIP file of the support save.

Device technical support

You can use Technical Support to collect SupportSave data (such as RASLOG, TRACE, and so on) and switch events from Fabric OS devices.

To gather technical support information for the Management application server, refer to "Capturing technical support information" on page 389.

Scheduling technical support information collection

You can configure the Management application to capture technical support information for any number of devices at one time; however, the Management application processes the requests to the devices in batches of 50. Technical SupportSave uses the built-in FTP, SCP, or SFTP server configured on the Management server to save data. If the switch is running Fabric OS 7.0 or later, the Management application uses the SCP server to save data, if configured. To make sure the built-in FTP, SCP, or SFTP server is configured correctly, refer to "Configuring an external FTP, SCP, or SFTP server" on page 125.

NOTE
Fabric OS switches must be running Fabric OS 7.0 or later to collect technical support data.

NOTE
Technical SupportSave uses the built-in FTP root if the switch is running Fabric OS 7.0 or later and if the external FTP is configured with the Linux host IP address. You must configure the Linux FTP configuration file by setting chroot_local_user to Yes.

NOTE
The Host must be a managed Brocade HBA.

NOTE
Scheduling technical support data collection is not supported on ESXi Servers.

NOTE
You must have the SupportSave privilege to perform this task. For more information about privileges, refer to "User Privileges" on page 1333.

To capture technical support and event information, complete the following steps.

1. Select **Monitor > Technical Support > Product/Host SupportSave**.

   The **Technical SupportSave** dialog box displays.

2. Click the **Schedule** tab.
3. Select the **Enable scheduled Technical Support Data** check box.
4. Select how often you want the scheduled collection to occur from the **Frequency** list.
5. Select the start date for the scheduled collection from the **Start Date** list.

   This list is only available when you select Weekly or Monthly from the **Frequency** list.

6. Select the time you want the scheduled collection to begin from the **Start Time Hour** and **Minute** lists.
7. Click the **SAN Products** tab, if necessary, and complete the following steps.

   The **Available SAN Products** table displays the following information:

   - **All Levels** — All discovered devices and ports as both text and icons.
   - **Name** — The name of the available switch.
   - **Product Type** — The type of product.
   - **Tag** — The tag number of the device.
   - **Serial #** — The serial number of the device.
   - **WWN** — The switch port's world wide name.
   - **IP Address** — The switch port's IP address.
   - **Domain ID** — The switch port's top-level addressing hierarchy of the domain.
   - **Vendor** — The hardware vendor's name.
   - **Model** — The name and model number of the hardware.
   - **Port Count** — The total number of ports.

- **Firmware** — The firmware version.
- **Location** — The customer site location.
- **Contact** — The primary contact at the customer site.
- **Description** — A description of the customer site.
- **State** — The switch state, for example, online or offline.
- **Status** — The operational status of the switch, for example, unknown or marginal.

a. Right-click in the **Available SAN Products** table and select **Expand All**.

b. Select the switches you want to collect data for in the **Available SAN Products** table and click the right arrow to move them to the **Selected Products and Hosts** table.

Technical SupportSave data for Fabric OS devices is saved to the following directory: *Install_Home*\data\ftproot\technicalsupport\

Technical SupportSave uses the following naming convention for the Fabric OS device support save files: Supportinfo-Day-mm-dd-yyyy-hh-mm-ss*Switch_Type-Switch_IP_Address-Switch_WWN.*

8. Click the **Hosts** tab and complete the following steps.

The Available Hosts table displays the following information:

- **Name** — The name of the available host.
- **IP Address** — The host port's IP address.
- **Network Address** — The network address of the host.
- **Fabrics** — The fabric of the host.

a. Right-click in the **Available Hosts** table and select **Expand All**.

b. Select the products you want to collect data for in the **Available Hosts** table and click the right arrow to move them to the **Selected Products and Hosts** table.

The **Selected Products and Hosts** table displays the following information:

- **IP Address** — The IP address of the selected product or host.
- **Name** — The name of the selected product or host.
- **WWN** — The world wide name of the selected product or host.
- **Firmware Type** — The type of firmware: FOS (Fabric OS).
- **Firmware version** — The firmware version of the selected product or host.
- **Support Save Credentials** — Whether the product or host has supportSave credentials or not.

Technical SupportSave data for SAN devices is saved to the following directory: *FTP_Host*\ftproot\technicalsupport\

9. Select how often you want to purge the support data from the **Purge Support Data** list.
10. Click **OK** on the **Technical SupportSave** dialog box.
11. Click **OK** on the confirmation message.

Data collection may take 20-30 minutes for each selected switch. This estimate may increase depending on the number of switches selected. Check the Master Log for status information.

NOTE
Unreachable switches increase the time needed to collect supportSave data.

Starting immediate technical support information collection

Technical SupportSave uses the built-in FTP, SCP, or SFTP server configured on the Management server to save data. If the switch is running Fabric OS 7.0.X or earlier, the Management application uses the SCP server to save data, if configured. If the switch is running Fabric OS 7.1 or later, the Management application uses the SFTP server to save data, if configured. To make sure the built-in FTP, SCP, or SFTP server is configured correctly, refer to"Configuring an external FTP, SCP, or SFTP server" on page 125.

NOTE
Fabric OS switches must be running Fabric OS 7.0 or later to collect technical support data.

NOTE
The HBA must be a managed Brocade HBA.

NOTE
You must have the SupportSave privilege to perform this task. For more information about privileges, refer to "User Privileges" on page 1333.

To capture technical support and event information for specified devices, complete the following steps.

1. Select **Monitor > Technical Support > Product/Host SupportSave**.

   The **Technical SupportSave** dialog box displays.

2. Click the **Generate Now** tab, if necessary.
3. Click the **SAN Products** tab, if necessary, and complete the following steps.
   a. Right-click in the **Available SAN Products** table and select **Expand All**.
   b. Select the switches you want to collect data for in the **Available SAN Products** table and click the right arrow to move them to the **Selected Products and Hosts** table.

      Technical SupportSave data for Fabric OS devices is saved to the following directory: *Install_Home*\data\ftproot\technicalsupport\

      Technical SupportSave uses the following naming convention for the Fabric OS device support save files: Supportinfo-Day-mm-dd-yyyy-hh-mm-ss*Switch_Type-Switch_IP_Address-Switch_WWN.*
4. Click the **Hosts** tab, if necessary, and complete the following steps.
   a. Right-click in the **Available Hosts** table and select **Expand All**.
   b. Select the hosts you want to collect data for in the **Available Hosts**table and click the right arrow to move them to the **Selected Products and Hosts** table.

      Technical SupportSave data for SAN devices is saved to the following directory: *FTP_Host*\ftproot\technicalsupport\
5. Click **OK** on the **Technical SupportSave** dialog box.

   Data collection may take 20-30 minutes for each selected switch. This estimate my increase depending on the number of switches selected.

   The **Technical SupportSave Status** dialog box displays with the following details.

Field	Description
Name	The name of the product.
IP Address	The product's IP address.
Firmware Type	The type of product.
Progress	The status of the supportsave. On products running Fabric OS 7.0 or later, this field shows the percentage complete and is updated every minute. For Host products, as well as Fabric OS products running 6.4 or earlier, this field cannot display the percentage (only displays whether it is 'in Progress' or' Completed').
Status	The status of the support save, for example, Ceases or Failure.
Remarks	Displays the administrator name and the client IP address of the original initiator who triggered the switch support save.

6. Click **Close** on the **Technical SupportSave Status** dialog box.

Viewing the technical support repository

You can only view technical support save files that are captured in the default location. Table 125 details the default locations for the technical support save files.

TABLE 125 Technical support save defaults

Type	Default location	Default naming convention
Client SupportSave	*User_Home*/ServerIP/Managed Product Name/support	DCM-Client-SS-*Time_Stamp*
Server SupportSave	*Install_Home*\support	DCM-SS-*Time_Stamp*
Host (discovered from the SAN tab)	*Install_Home*\data\ftproot\technicalsupport\host	Supporinfo-HostName-Day-mm-dd-yyyy-hh-mm-ss
SAN Product	*Install_Home*\data\ftproot\technicalsupport\	Supporinfo-HostName-Day-mm-dd-yyyy-hh-mm-ss
Auto Trace Dump	*Install_Home*\data\ftproot\tracedump\	

To view the technical support repository, complete the following steps.

1. Select **Monitor > Technical Support > View Repository**.

   The **Technical Support Repository** dialog box displays.

2. Review the technical support repository details:

Field/Component	Description
Available SupportSave and Upload Failure Data Capture Files table	Select the support data file you want to view. Displays the following information: **File Name** — The name of the SupportSave file. **Size (MB)** — The name of the SupportSave file. **Last Modified** — The date the SupportSave file was generated. **Firmware Type** — The type of file (Client, Server, FOS (Fabric OS), or First Failure Data Capture (FFDC)). Blank for Host support save files.
E-mail button	Click to e-mail the support data file. For the procedure, refer to "E-mailing technical support information" on page 1269.
FTP button	Click to copy the support data file to an external FTP server. For the procedure, refer to "Copying technical support information to an external FTP server" on page 1270.
Save button	Click to save a copy of the support data. For the procedure, refer to "Saving technical support information to another location" on page 1269.
Delete button	Click to delete the support data file. For the procedure, refer to "Deleting technical support files from the repository" on page 1271.

3. Click **OK** on the **Technical Support Repository** dialog box.

Saving technical support information to another location

To save technical support information to a location other than the default, complete the following steps.

1. Select **Monitor > Technical Support > View Repository**.

   The **Technical Support Repository** dialog box displays.

2. Select a device support save file and click **Save.**

   The **Save** dialog box displays.

3. Browse to the location where you want to save the support file.
4. Click **Save** on the **Save** dialog box.
5. Click **OK** on the message.
6. Click **OK** on the **Technical Support Repository** dialog box.

E-mailing technical support information

NOTE
You cannot e-mail technical support information collected from the remote client.

To e-mail technical support information, complete the following steps.

1. Select **Monitor > Technical Support > View Repository**.

   The **Technical Support Repository** dialog box displays.

2. Select the file you want to e-mail in the table.
3. Click **E-mail** to e-mail the event and supportsave files (zip).

NOTE
The **E-mail** button is unavailable from the remote client.

You must configure the Management application e-mail server before you can define the e-mail action. For more information, refer to "Configuring e-mail notification" on page 1132.

The **E-mail** dialog box displays.

4. Enter the e-mail address of the person to receive the e-mail in the **To** field.
5. Enter your e-mail address in the **From** field.
6. Click **OK**.

   The e-mail is sent and the **Technical Support Repository** dialog box closes automatically.

Copying technical support information to an external FTP server

NOTE
You cannot copy technical support information to an external FTP server collected from the remote client.

To copy the SupportSave data located in the built-in FTP server to an external FTP server, complete the following steps.

1. Select **Monitor > Technical Support > View Repository**.

   The **Technical Support Repository** dialog box displays.

2. Select the file you want to copy in the table.
3. Click **FTP** to send the switch event and supportsave files (zip) by FTP.

   NOTE
   The **FTP** button is unavailable from the remote client.

   The **FTP Credentials** dialog box displays.

4. Enter the network address or domain name of the external FTP server in the **Network Address** field.
5. Enter your user name and password.
6. Enter the destination directory where you want to copy the data on the external FTP server in the **Destination Directory** field.

   The destination directory should be the sub directory of the external FTP server's root directory. For example, if you enter "repository" as the destination directory, then the support save file is copied to the "/repository" directory of the external FTP server.

7. Click **OK**.

   The data is copied and the **Technical Support Repository** dialog box closes automatically.

Uploading SupportSave information

To upload the SupportSave data using anonymous FTP user and blind FTP sites, complete the following steps.

1. Select **Monitor > Technical Support > View Repository**.

   The **Technical Support Repository** dialog box displays (Figure 577).

2. Select the file you want to upload in the table.
3. Click **FTP** to upload the supportsave files (zip) by FTP.

   NOTE
   The **FTP** button is unavailable from the remote client.

   The **FTP Credentials** dialog box displays.
4. Enter the network address or domain name of the external FTP server in the **Network Address** field.
5. Enter your user name and password in the **User Name** field and **Password** field.

   Or

   Select the **Anonymous Login** check box to upload SupportSave using anonymous FTP user site or blind FTP site. If you select **Anonymous Login** check box, the **User Name** field displays Anonymous and **Password** field displays blank and grayed out.
6. Enter the destination directory where you want to upload the data on the external FTP server in the **Destination Directory** field.

   If supportsave files for the switch already exists in destination directory, it will be replaced with the new SupportSave files.

   FIGURE 577 FTP Credential dialog box

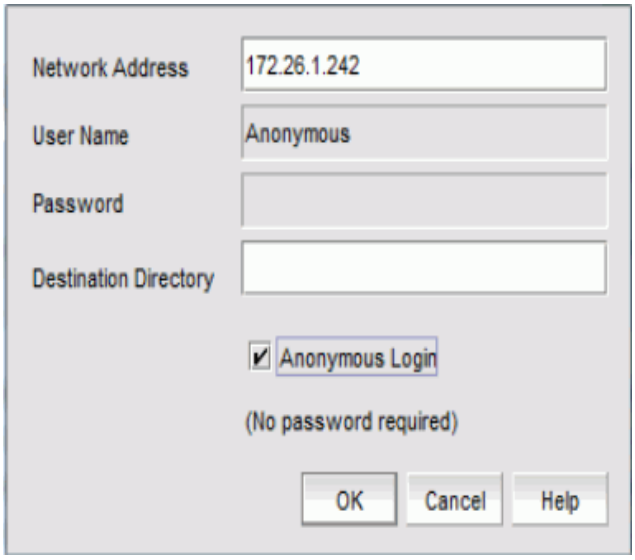

7. Click **OK**.

   The data is uploaded and the **Technical Support Repository** dialog box closes automatically.

Deleting technical support files from the repository

To delete a technical support file from the repository, complete the following steps.

1. Select **Monitor > Technical Support > View Repository**.

   The **Technical Support Repository** dialog box displays.
2. Select the file you want to delete in the table.
3. Click **Delete**.

4. Click **OK** on the **Technical Support Repository** dialog box.

Upload failure data capture

You can use upload failure data capture to enable, disable, and purge failure data capture files as well as configure the FTP Host for the switch.

NOTE
Upload failure data capture is only supported on Fabric OS devices.

Enabling upload failure data capture

1. Select **Monitor > Technical Support > Upload Failure Data Capture**.

   The **Upload Failure Data Capture** dialog box displays.

   FIGURE 578 Upload Failure Data Capture dialog box

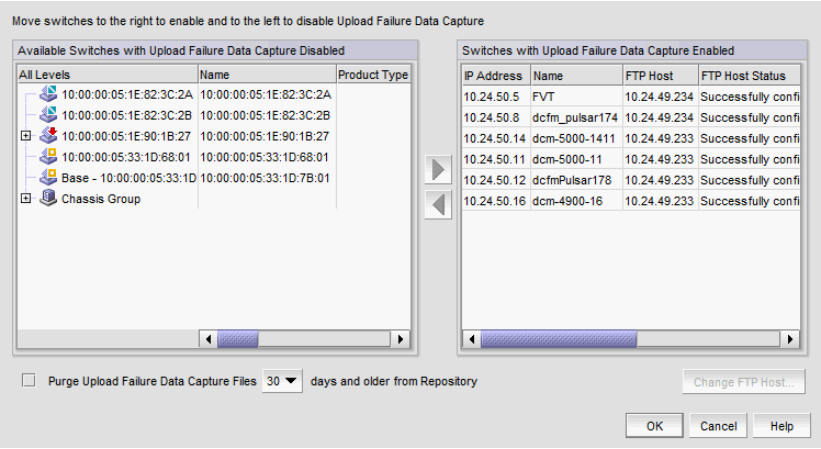

2. Select a one or more devices on which you want to enable automatic trace dump from the **Available Switches with Upload Failure Data Capture Disabled** table.

   The **Available Switches with Upload Failure Data Capture Disabled** table displays the following information:

   - **All Levels —** All discovered devices and ports as both text and icons.
   - **Name —** The name of the available switch.
   - **Product Type —** The type of product.
   - **Tag —** The tag number of the device.
   - **Serial # —** The serial number of the device.
   - **WWN —** The switch port's world wide name.
   - **IP Address —** The switch port's IP address.
   - **Domain ID —** The switch port's top-level addressing hierarchy of the domain.
   - **Vendor —** The hardware vendor's name.
   - **Model —** The name and model number of the hardware.
   - **Port Count** — The total number of ports.

- **Firmware —** The firmware version.
- **Location —** The customer site location.
- **Contact —** The primary contact at the customer site.
- **Description —** A description of the customer site.
- **State —** The switch state, for example, online or offline.
- **Status —** The operational status of the switch, for example, unknown or marginal.

3. Click the right arrow button.

   The selected devices move from the **Available Switches with Upload Failure Data Capture Disabled** table to the **Switches with Upload Failure Data Capture Enabled** table.

   The **Switches with Upload Failure Data Capture Enabled** table displays the following information:

   - **IP Address** — The switch's IP address.
   - **Name** — The name of the switch.
   - **FTP Host** — The current FTP host configured on the switch.
   - **FTP Host Status** — The status of the FTP host configured on the switch.
   - **FTP User** — The user for the current FTP host configured on the switch.
   - **FTP Root** — The root location where failure data capture files are saved.

4. Click **OK** on the **Upload Failure Data Capture** dialog box.
5. Click **OK** on the confirmation message, if necessary.

Disabling upload failure data capture

NOTE
Upload Failure Data Capture is only supported on Fabric OS devices.

1. Select **Monitor > Technical Support > Upload Failure Data Capture**.

   The **Upload Failure Data Capture** dialog box displays.

2. Select one or more devices on which you want to disable automatic trace dump from the **Available Switches with Upload Failure Data Capture Enabled** table.
3. Click the left arrow button.

   The selected devices move from the **Switches with Upload Failure Data Capture Enabled** table to the **Available Switches with Upload Failure Data Capture Disabled** table.

4. Click **OK** on the **Upload Failure Data Capture** dialog box.
5. Click **OK** on the confirmation message, if necessary.

Purging upload failure data capture files

NOTE
Upload Failure Data Capture is only supported on Fabric OS devices.

1. Select **Monitor > Technical Support > Upload Failure Data Capture**.

   The **Upload Failure Data Capture** dialog box displays.

2. Select the **Purge Upload Failure Data Capture Files** check box to enable purging the trace dump files.
3. Select how often (days) you want to purge the trace dump data from the **Purge Upload Failure Data Capture Files** list.
4. Click **OK** on the **Upload Failure Data Capture** dialog box.

Configuring the upload failure data capture FTP server

NOTE
Upload Failure Data Capture is only supported on Fabric OS devices.

NOTE
Some external FTP software (such as, Filezilla and Xlight) are not supported.

1. Select **Monitor > Technical Support > Upload Failure Data Capture**.

   The **Upload Failure Data Capture** dialog box displays.

2. Select a device from the **Available Switches with Upload Failure Data Capture Enabled** table.
3. Click **Change FTP Host.**

   The **Change FTP Server** dialog box displays.

4. Choose one of the following options:
   - Select the **Use** *Management_Application* option to use the Management application FTP server.
   - Select the **Custom** option and complete the following steps to configure a FTP server for the selected device.
     a. Enter the server's IP address in the **Host IP** field.
     b. Enter a user name for the server in the **User Name** field.
     c. Enter a password for the server in the **Password** field.
     d. Enter the path to where the trace dump data is saved in the **Directory Path** field.
5. Click **Test** to test the server credentials.
6. Click **OK** on the **Change FTP Host** dialog box.
7. Click **OK** on the **Upload Failure Data Capture** dialog box.
8. Click **OK** on the confirmation message, if necessary.

Saving the upload failure data capture repository

NOTE
Upload Failure Data Capture is only supported on Fabric OS devices.

1. Select **Monitor > Technical Support > View Repository**.

   The **Repository** dialog box displays.

2. Select the **Switches** tab to view upload failure data capture information.
3. Select the trace dump file you want to save and click **Save**.
4. Browse to the location you want to save the file and click **OK**.
5. Click **OK** on the **Repository** dialog box.

Reports

In this chapter

Reports overview

Reports are available from the **Reports** menu. You must have the Reports privilege to access the reports. For more information about privileges, refer to "User Privileges" on page 1333.

Browser requirements

SAN reports can printed from a web browser. Reports are supported in the following browsers:

- Internet Explorer 11.0.9 or later (Windows 8.1, Windows Server 2008 R2, and Windows Server 2012 R2)
- Firefox 24 or later on Windows or Linux
- Google Chrome 33 or later on Windows

SAN report types

Presenting and archiving data about a SAN is equally as important as gathering the data. Through the Management application, you can generate reports about the SAN. You can send the reports to network administrators, support consultants, and others interested in the SAN's architecture, or archive the reports for future reference.

The following standard report types are available from the **Generate Reports** dialog box:

- **Fabric Ports** — Lists discovered ports including used and unused ports. Port data for each fabric is divided into three parts: Fabric-wide port details, Switch-wide port details, and individual port details.
- **Fabric Summary** — Lists information about discovered fabrics including fabric and switch details, device information, and ISL and trunk summary.
- **Host Adapter Inventory** — Lists all Brocade adapters discovered through Host discovery.
- **Host Adapter Faulty SFP** — Lists all Brocade adapters with unsupported or faulty SFPs.

The following device-specific reports are available through the **Monitor** (**Monitor > Performance > Historical Report**) or **Reports** menus and right-click menus:

- **Performance** — Lists historical performance-related data.

  NOTE
  Performance reports require a SAN Trial or Licensed version.

- **Zone** — Lists zoning objects.

Generating SAN reports

To generate reports, complete the following steps.

1. Select **Reports > Generate**.

   The **Generate Reports** dialog box displays.

2. Select the types of reports you want to generate:
   - Fabric Summary Report
   - Fabric Ports Report
   - Adapters Inventory report
   - Adapters Faulty SFP report
3. Select the fabrics or host for which you want to generate reports.
4. Click **OK**.

   The generated reports display in the **View Reports** dialog box.

   NOTE
   Hyperlinks in reports are active only if the source data is available.

5. Click **Close** to close the **View Reports** dialog box.
6. Click **Yes** on the "are you sure you want to close" message.

Viewing SAN reports

You can view any report generated in the SAN. To view reports, complete the following steps.

1. Select **Reports > View** or click the **View Report** icon.

   The **View Reports** dialog box displays.

2. Select the report you want to view in the **All Reports** list.

   If you do not see the report you want to view, generate it first by following the instructions in "Generating SAN reports" on page 1278.

   You can select reports by Time, Report Type, or User.

3. Use the buttons in the following table to navigate through and resize the report.

Icon	Description
	First — Click to return to the first page in the report. Unavailable when you are on the first page of the report.
	Previous — Click to return to the previous page in the report. Unavailable when you are on the first page of the report.
	Next — Click to move to the next page in the report. Unavailable when you are on the last page of the report.
	Last — Click to move to the last page in the report. Unavailable when you are on the last page of the report.
	Actual Size — Click to display the report in its actual size.
	Fit to Page — Click to resize the report to display entirely in the view.
	Fit to Width — Click to resize the report to fit in the view by width.
	Zoom In — Click to zoom in on the report.
	Zoom Out — Click to zoom out on the report.

4. Click **Show in Browser** to view the selected report in your default browser window.
5. Click **Close** to close the **View Reports** dialog box.
6. Click **Yes** on the “are you sure you want to close” message.

Fabric Summary Report

The Fabric Summary Report (Figure 898) provides a summary of the discovered fabrics as well as Switch and Access Gateway devices associated with the fabric.

FIGURE 579 **Fabric Summary Report**

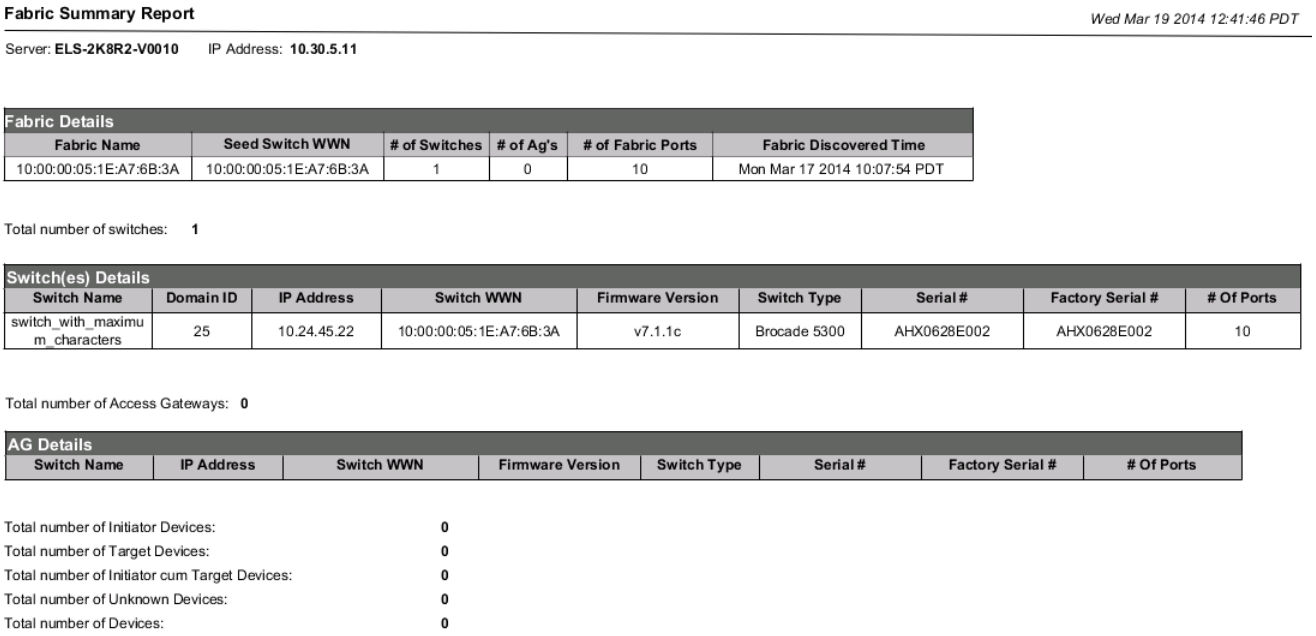

Fabric Summary Report *Wed Mar 19 2014 12:41:46 PDT*

Server: **ELS-2K8R2-V0010** IP Address: **10.30.5.11**

Fabric Details

Fabric Name	Seed Switch WWN	# of Switches	# of Ag's	# of Fabric Ports	Fabric Discovered Time
10:00:00:05:1E:A7:6B:3A	10:00:00:05:1E:A7:6B:3A	1	0	10	Mon Mar 17 2014 10:07:54 PDT

Total number of switches: **1**

Switch(es) Details

Switch Name	Domain ID	IP Address	Switch WWN	Firmware Version	Switch Type	Serial #	Factory Serial #	# Of Ports
switch_with_maximum_characters	25	10.24.45.22	10:00:00:05:1E:A7:6B:3A	v7.1.1c	Brocade 5300	AHX0628E002	AHX0628E002	10

Total number of Access Gateways: **0**

AG Details

Switch Name	IP Address	Switch WWN	Firmware Version	Switch Type	Serial #	Factory Serial #	# Of Ports

Total number of Initiator Devices: **0**
Total number of Target Devices: **0**
Total number of Initiator cum Target Devices: **0**
Total number of Unknown Devices: **0**
Total number of Devices: **0**

Table 225 describes the fields and components of the Fabric Summary Report.

TABLE 126 Fabric Summary Report fields and components

Field/Component	Description
Server	The name of the Management application server.
IP Address	The IP address of the Management application server.
Fabric Details table	
Fabric Name	The name of the fabric.
Seed Switch WWN	The world wide name of the seed switch.
# of Switches	The number of switches in the fabric.
# of AGs	The number of AGs in the fabric.
# of Fabric Ports	The number of ports in the fabric.
Fabric Discovered Time	The date and time the fabric was discovered.
Total number of switches	The number of switches in the fabric.

TABLE 126 Fabric Summary Report fields and components (Continued)

Field/Component	Description
Switch(es) Details table	
Switch Name	The name of the switch.
Domain ID	The Domain ID for the switch.
IP Address	The IP address (IPv4 or IPv6 format) of the device.
Switch WWN	The world wide name of the switch.
Firmware Version	The firmware version of the switch.
Switch Type	The type of the switch. For example, Encryption SAN switch.
Serial #	The serial number of the switch.
Factory Serial #	The factory serial number of the switch.
# of Ports	The number of ports on the switch.
Total number of AGs	The number of AGs in the fabric.
AG Details table	
Switch Name	The name of the device.
IP Address	The IP address (IPv4 or IPv6 format) of the device.
Switch WWN	The world wide name of the device.
Firmware Version	The firmware version of the device.
Switch Type	The type of the device.
Serial #	The serial number of the device.
Factory Serial #	The factory serial number of the device.
# of Ports	The number of ports on the switch.
Total number of Initiator Devices	The number of initiator devices in the fabric.
Total number of Target Devices	The number of target devices in the fabric.
Total number of Initiator cum Target Devices	The number of initiator and target devices in the fabric.
Total number of Unknown Devices	The number of unknown devices in the fabric.
Total number of Devices	The number of devices in the fabric.
Device(s) Information table	
Switch Name	The name of the switch.
Device Name	The manufacturer's name for the device.
Device Port WWN	The world wide name of the switch port.
Status	The status of the device. For example, Online or Offline.
Device Type	The type of device. For example, Physical.
Role	The role of the device. For example, Target or Initiator.
Slot #	The slot number.
Speed	The speed of the port
Port #	The port number.
Port Type	The type of port. For example, L-Port or F-Port.

TABLE 126 Fabric Summary Report fields and components (Continued)

Field/Component	Description
Vendor	The vendor who manufactures the device.
ISL/Trunk Summary table	
From Switch/To Switch	The switch, port, and ISL/trunk data for the devices on either end of a connection.
Switch Name	The name of the switch.
Domain ID	The Domain ID of the switch.
Port #	The port number.
ISL/Trunk	Whether it is an ISL or Trunk.

Fabric Ports Report

The Fabric Ports Report (Figure 898) provides a summary of the discovered ports including used and unused ports. Port data for each fabric is divided into three parts: Fabric-wide port details, Switch-wide port details, and individual port details.

FIGURE 580 Fabric Ports Report

Fabric Ports Report *Wed Mar 19 2014 12:41:43 PDT*

Server : **ELS-2K8R2-V0010** @ **10.30.5.11**

Fabric : **10:00:00:05:1E:A7:6B:3A**

	Director Utilization				Switch Utilization			
Total Fabric Ports	Total Number of Ports	Number of Ports connected	Number of Ports Free	Number of Ports allocated	Total Number of Ports	Number of Ports connected	Number of Ports Free	Number of Ports allocated
10	0	0	0	0	10	0	10	0

Details

IP Address	Switch Name	Domain/ Port #	Zone	Connected Device					Port Name	Port Speed (GBPS)	Port Status	Port State	Port Type	Physical/ Logical
				DeviceName	Vendor	Device Type	Model	PortWwn						
10.24.45.22	switch_with_maximum_characters	D25 /P52	(Online)						new_52	8	No_Module	Offline	U-Port	Physical
10.24.45.22	switch_with_maximum_characters	D25 /P51	(Online)						new_51	8	No_Module	Offline	U-Port	Physical
10.24.45.22	switch_with_maximum_characters	D25 /P56	(Online)						56	8	No_Module	Offline	U-Port	Physical
10.24.45.22	switch_with_maximum_characters	D25 /P55	(Online)						55	8	No_Module	Offline	U-Port	Physical
10.24.45.22	switch_with_maximum_characters	D25 /P53	(Online)						Port-53	8	No_Module	Offline	U-Port	Physical
10.24.45.22	switch_with_maximum_characters	D25 /P50	(Online)						new_50	8	No_Module	Offline	U-Port	Physical
	switch_with_													

Table 225 describes the fields and components of the Fabric Ports Report.

TABLE 127 Fabric Ports Report fields and components

Field/Component	Description
Server	The name of the Management application server.
@	The IP address of the Management application server.

TABLE 127 Fabric Ports Report fields and components (Continued)

Field/Component	Description
Fabric	The name of the fabric.
Fabric table	
Total Fabric Ports	The number of initiator devices in the fabric.
Director Utilization	
Total Number of Ports	The total number of director ports.
Number of Ports connected	The number of connected ports on a director.
Number of Ports Free	The number of free ports on a director.
Number of Ports allocated	The number of allocated ports on a director.
Switch Utilization	
Total Number of Ports	The total number of switch ports.
Number of Ports connected	The number of connected ports on a switch.
Number of Ports Free	The number of free ports on a switch.
Number of Ports allocated	The number of allocated ports on a switch.
Details table	
IP Address	The IP address of the switch.
Switch Name	The name of the switch.
Domain/ Port #	The domain ID and port number of the switch.
Zone	Whether the zone is online or offline.
Connected Device	Details about the connected device.
Device Name	The name of the connected device.
Vendor	The manufacturer of the connected device.
Device Type	The device type of the connected device. For example, Initiator or Target.
Model	The model of the connected device.
PortWwn	The port world wide name of the connected device.
Port Name	The port name.
Port Speed (GBPS)	The port speed in gigabits per second.
Port Status	The port status. For example, Online, Offline, No light and so on.
Port State	The port state. For example, Online or Offline.
Port Type	The port type. For example, F-Port, L-Port, and so on.
Physical/Logical	Whether the switch is physical or logical.

Exporting SAN reports

To export reports, complete the following steps.

1. Select **Reports > View** or click the **View Report** icon.

   The **View Reports** dialog box displays.

2. Select the report you want to export in the **All Reports** list.

   If you do not see the report you want to export, generate it first by following the instructions in "Generating SAN reports" on page 1278.

   You can select reports by Time, Report Type, or User.

3. Select the format (**PDF**, **HTML**, or **XML**) you want to export to from the list to the left of the **Export** button.
4. Click **Export**.

   The **Save** dialog box displays.

5. Browse to the file location where you want to save the report and click **Save**.
6. Click **Close** to close the **View Reports** dialog box.
7. Click **Yes** on the "are you sure you want to close" message.

Printing SAN reports

You can print reports through a web browser.

1. Select **Reports > View** or click the **View Report** icon.

   The **View Reports** dialog box displays.

2. Select the report you want to print in the **All Reports** list.

   If you do not see the report you want to view, generate it first by following the instructions in "Generating SAN reports" on page 1278.

   NOTE
   Hyperlinks in reports are active only if the source data is available.

3. Click **Show in Browser**.

   The selected report displays in your default web browser.

4. Select **File > Print** (in the web browser).

   The **Print** dialog box displays.

5. Select the printer to which you want to print and click **Print**.
6. Close the web browser.
7. Click **Close** in the **View Reports** dialog box.
8. Click **Yes** on the "are you sure you want to close" message.

Deleting SAN reports

To delete reports, complete the following steps.

1. Select **Reports > View** or click the **View Report** icon.

   The **View Reports** dialog box displays.

2. Select the report you want to delete in the **All Reports** list.

   You can select reports by Time, Report Type, or User.

3. Click **Delete Report**.

   ATTENTION
   Once you click **Delete Report**, the report is deleted without confirmation.

4. Click **Close** to close the **View Reports** dialog box.
5. Click **Yes** on the "are you sure you want to close" message.

Generating SAN performance reports

NOTE
Performance reports require a SAN Trial or Licensed version.

To generate a historical performance report for a device, complete the following steps.

1. Select the device for which you want to generate a performance report.
2. Choose one of the following options:
   - Select **Monitor > Performance > Historical Report**.

     OR

   - Right-click the device and select **Performance > Historical Report**.

   The **Historical Performance Table** dialog box displays.

3. Filter the historical data by completing the following steps.
   a. Select the number of results to display from the **Display** list.
   b. Select the ports from which you want to gather performance data from the **From** list.

      NOTE
      Devices with 10GE ports must be running Fabric OS 6.4.1ltd or later to obtain the correct TE port statistics (TX/RX).

      If you select **Custom**, complete the following steps.

   i. Select the type of ports from the **Show** list.
   ii. Right-click a device in the **Available** table and select **Expand All**.
   iii. Select the ports (**Ctrl** or **Shift** + click to select multiple ports) from which you want to gather performance data from the **Available** table and click the right arrow button.
   The selected ports move to the **Select Ports** table.
   iv. Click **OK**.

c. Select the historical period from which you want to gather performance data from the **For** list.

If you select **Custom**, complete the following steps.

i. Select the **Last** option and enter the number of minutes, hours, or days.
OR
Select the **From** option and enter the date and time.

ii. Click **OK**.

d. Select the granularity at which you want to gather performance data from the **Granularity** list:

- 5 minutes for last 8 days.
- 30 minutes granularity for last 30 days
- 2 hour granularity for last 30 days
- 1 day granularity for last 730 days.

e. Select the measure by which you want to gather performance data from the **Measures** list.

To select more than one measure, click the **Additional Measures** expand arrows and select the check box for each additional measure.

f. Save this configuration by selecting **Save**.

The **Save Favorites** dialog box displays. This enables you to save the selected configuration so that you can use it to generate the same type of report at a later date.

g. Enter a name for the configuration in the **Favorites Name** field and click **OK**.

h. Click **Apply**.

The selected report automatically displays in the **View Reports** dialog box.

NOTE
Hyperlinks in reports are active only if the source data is available.

To print the selected report, refer to "Printing SAN reports" on page 1284.

To export the selected report, refer to "Exporting SAN reports" on page 1284.

To delete the selected report, refer to "Deleting SAN reports" on page 1285.

4. Click the close button (X) to close the **View Reports** dialog box.
5. Click the close button (X) to close the **Historical Performance Table** dialog box.

For more information about performance, refer to "Performance Data" on page 959.

Generating SAN zoning reports

The Management application enables you to generate a report for the current zone DB in the fabric. To generate a report for the edited zone DB, you must save it to the fabric first. Make sure no one else is making changes to the same area prior to submitting or your changes may be lost.

To generate zoning reports, complete the following steps.

1. Select **Configure > Zoning** or right-click the device and select **Zoning**.

   The **Zoning** dialog box displays.

2. Click **Report**.
3. Click **OK** on the message.

   The selected report automatically displays in the **View Reports** dialog box.

   NOTE
   Hyperlinks in reports are active only if the source data is available.

   To print the selected report, refer to "Printing SAN reports" on page 1284.

   To export the selected report, refer to "Exporting SAN reports" on page 1284.

   To delete the selected report, refer to "Deleting SAN reports" on page 1285.

4. Click **Close** to close the **View Reports** dialog box.
5. Click **Yes** on the "are you sure you want to close" message.

   For more information about zoning, refer to "Zoning" on page 779.

CLI reports

The Management application provides the following reports via the CLI.

Report Name	Report Alias
Detailed Port Report	101
Fabric Summary Report	102
Host Adapter Inventory Report	103
Host Adapter Unsupported and Faulty SFP Report	104
Port SFP Performance Report	105
Port SFP Time Series Performance Report	106
Port FCIP Time Series Performance Report	107
Product Time Series Performance Report	108
Switch Report	109
Zone Summary Report	110

The report CLI scripts are available in <Install_Home>/bin.

Generating a CLI report

You must specify all mandatory parameter-value pairs in the command line. If you do not provide a mandatory parameter-value pairs, an error message or command usage message displays. If the parameter-value contains a space or special character, it must use double quotes (such as --t "1 Hour").

1. Open a command prompt
2. Navigate to <Install_Home>/bin.
3. Enter one of the following options:
   - On Windows systems, **report-cli --report-type** <Report Name> <parameter-name> <parameter-value> <parameter-name> <parameter-value>... --username <user>
   - On Linux systems, **sh report-cli.sh --report-type** <Report Name> <parameter-name> <parameter-value> <parameter-name> <parameter-value>... --username <user>

   NOTE
   You must use the exact syntax for generating the reports (for example, you must provide the correct parameter options and the values must be entered with double quotes). If you enter incorrect parameters and values, the report generation fails.

CLI parameter	Description
n	The network scope of the report. Select "All" to choose all available network scope at the time you run the report.
t	The time scope of the report. Choose one of the following options: • 30 Minutes • 1 Hour • 6 Hours • 12 Hours • 1 Day • 3 Days • 1 Week • 1 Month
st	The start time when the time scope is applied to the report. The format is mm:dd:yyyy hh:mm.
et	The end time when the time scope is applied to the report. The format is mm:dd:yyyy hh:mm.
ps	The port scope to be considered for the reporting. Choose one of the following options: All Ports, ISL Ports, Initiator Ports, or Target Ports. The default is All Ports.
o	The output directory path where the generated report is saved. The format is <platform specific absolute path>. The default is <Install_Home>\data\reports\<Report Type>-<Date and Time in MM-DD-YYYY-HH:MM:SS> In Linux, the following folder structure report is not created (/root path and sub folders). If this path is required, you must provide 777 access to all folder levels in the specified path. This parameter is applicable to all reports.
fab	The fabric name for summary report. The format is <fabricname;IP>.
s	The identification information of the switch(es) for which the report is generated. The format is <WWN>.

tun	The input tunnel information for the report. The format is <Tunnel ID>;<Local Switch WWN>, <Tunnel ID>;<Local Switch WWN>.
m	The measure(s) for the report. FCIP Port measures include Cumulative Compression Ratio, Latency, Dropped Packets, Link Retransmits, Time Out Retransmits, Fast Retransmits, Duplicate Ack Received, Window Size RTT, TCP Out Of Order Segments, Slow Start Status Errors, or Current Compression Ratio. SFP Port measures include SFP Power, SFP Voltage, SFP Current, and SFP Temperature. Product Performance measures include Memory Utilization Percentage, CPU Utilization Percentage, Temperature, Fan Speed, System Up Time, and Ports Not in Use.
p	The port WWN(s) for the report. The format is <WWN>,<WWN>.

Example Detailed Port Report

(Windows) `report-cli --report-type "Detailed Port Report"  --n "All"  --t "1 Hour"  --et "11:10:2016 00:00" --username "administrator"`
(Linux) `sh report-cli.sh --report-type "Detailed Port Report"  --n "All"  --t "1 Hour"  --et "11:10:2016 00:00" --username "administrator"`

Example Fabric Summary Report

(Windows) `report-cli --report-type "Fabric Summary Report" --fab "fabric1;10.24.42.22" --username "administrator"`
(Linux) `sh report-cli.sh --report-type "Fabric Summary Report" --fab "fabric1;10.24.42.22" --username "administrator"`

The --fab parameter is required.

Example Host Adapter Inventory Report

(Windows) `report-cli --report-type "Host Adapter Inventory Report" --username administrator`
(Linux) `sh report-cli.sh --report-type "Host Adapter Inventory Report" --username administrator`

Example Host Adapter Unsupported and Faulty SFP Report

(Windows) `report-cli --report-type "Host Adapter Unsupported and Faulty SFP Report" --username administrator`
(Linux) `sh report-cli.sh --report-type "Host Adapter Unsupported and Faulty SFP Report" --username administrator`

Example Port FCIP Time Series Performance Report

(Windows) `report-cli --report-type "Port FCIP Time Series Performance Report" --n "All"  --s "10:00:00:27:F8:D8:8B:73"  --tun "tunnel1" --m "Latency" --st "11:01:2016 00:00"  --et "11:10:2016 00:00"`
(Linux) `sh report-cli.sh --report-type "Port FCIP Time Series Performance Report" --n "All"  --s "10:00:00:27:F8:D8:8B:73"  --tun "tunnel1" --m "Latency" --st "11:01:2016 00:00"  --et "11:10:2016 00:00"`

Example Port SFP Time Series Performance Report

(Windows) `report-cli --report-type " Port SFP Time Series Performance Report" --n "All"  --s "10:00:00:27:F8:D8:8B:73" --p "10:00:10:34:F8:D8:8B:73"  --ps "All Ports"  --m "SFP Power"    --username administrator`
(Linux) `sh report-cli.sh --report-type " Port SFP Time Series Performance Report" --n "All"  --s "10:00:00:27:F8:D8:8B:73" --p "10:00:10:34:F8:D8:8B:73"  --ps "All Ports"  --m "SFP Power"    --username administrator`

The --m (SFP port measures), --st, and --et parameters are required.

Example Port SFP Performance Report

(Windows) `report-cli --report-type " Port SFP Performance Report"  --t "1 Hour" --ps "All Ports" --et "11:10:2016 00:00" --username administrator`
(Linux) `sh report-cli.sh --report-type " Port SFP Performance Report"  --t "1 Hour" --ps "All Ports" --et "11:10:2016 00:00" --username administrator`

The --m (FCIP measures), --st, and --et parameters are required.

Example Product Time Series Performance Report

(Windows) `report-cli --report-type " Product Time Series Performance Report"  --n “All”  --s “10:00:00:43:F5:T7:G5:D4” --m “Fan Speed” --st “11:10:2016 00:00”  --et “11:20:2016 00:00” --username administrator`
(Linux) `sh report-cli.sh --report-type " Product Time Series Performance Report"  --n “All”  --s “10:00:00:43:F5:T7:G5:D4” --m “Fan Speed” --st “11:10:2016 00:00”  --et “11:20:2016 00:00” --username administrator`

The --m (Product Performance measures), --st, and --et parameters are required.

Example Switch Report

(Windows) `report-cli --report-type " Switch Report" --s “10:00:00:43:F5:T7:G5:D4” --username administrator`
(Linux) `sh report-cli.sh --report-type " Switch Report" --s “10:00:00:43:F5:T7:G5:D4” --username administrator`

The --s parameter is required.

Example Zone Summary Report

(Windows) `report-cli --report-type " Zone Summary Report" --fab “fabricname1”  --username administrator`
(Linux) `sh report-cli.sh --report-type " Zone Summary Report" --fab “fabricname1”  --username administrator`

The --fab parameter is required.

Host adapter reports

NOTE
Host Adapter reports are only available for Brocade adapters.

NOTE
Host Adapter reports only display hosts and adapters in your AOR.

Viewing host adapter reports

From the Product List, right-click the Host node and select one of the following options:

- **Adapters Inventory Report**
- **Adapters Faulty SFP Report**

  The report for the selected Host displays.

From the SAN tab, complete the following steps.

1. Select **Reports > Generate Reports**.

   The **Generate Reports** dialog box displays.

2. Select the report you want to generate:
   - Host Adapter Inventory
   - Host Adapter Faulty SFP
3. Select the fabrics for which you want to generate reports.
4. Click **OK**.

   The selected report displays in the **View Reports** dialog box. This report includes data for all Hosts discovered through Host Adapter discovery.

Adapters Inventory report

The Adapters Inventory report only displays adapters and ports discovered through Host Adapter discovery. This report is only available for Brocade adapters.

Table 128 describes the fields and components of the Adapters Inventory Report.

TABLE 128 Adapters Inventory report fields and components

Field/Component	Description
Host Name	The name of the host.
Host IP	The IP address of the host.
Server Vendor	The vendor name of the server.
Server Model	The model of the server.
OS	The host operating system; for example, Microsoft Windows or Red Hat Linux.
Server Location	The location of the server.
Server Contact	The name of the person or group you should contact about the server.
Server Description	A description of the server.
HCM Agent/CIM Provider Version	The version of the HCM Agent or CIM Provider.
Adapter Name	The name of the adapter configured in the Management application.
Adapter HCM Name	The name of the adapter configured in the HCM application.
Adapter Node	The world wide node (WWN) or MAC address of the adapter.
Adapter Status	Whether the adapter is enabled or disabled.
Adapter Model	The model of the adapter.
Adapter Serial #	The serial number of the adapter.
Port Count	The physical port count of the adapter.
Firmware Version	The firmware version of the adapter.
BIOS Version	The BIOS version of the adapter.
OEM Info	The OEM information for the adapter.
Port ID	The port number or PCI function index, depending on adapter type and driver version.
Port Name	The port name or Eth device name, depending on port type.
Port HCM Name	The name of the port configured in HCM.
Symbolic Name	The symbolic name of the port
Node WWN/MAC	The port's WWN or MAC address.
Port WWN/MAC	The node's (parent device) WWN or MAC address.
Type	The type of port, for example, IP-Port, N, or NL.
QoS Enabled	Whether QoS is enabled (True) or disabled (False).
Frame Field Size (Bytes)	The frame field size in bytes.
Media	The type of media; for example, 8G-sw (8 Gbps software).
FC Address	The port's Fibre Channel address.
Fabric Name	The name of the Fabric.
Zone Alias	The alternate name of the zone.

TABLE 128 Adapters Inventory report fields and components (Continued)

Field/Component	Description
Fabric Assigned Address	The state (enabled, disabled, or N/A) of the fabric assigned address for the adapter.
WWN Source	The source of the world wide name. Options include: Fabric — The WWN is assigned from the fabric. The fabric assigned address must be enabled. Factory — The WWN is assigned at the factory. N/A — Not applicable.
Boot Over SAN	Whether Boot over SAN is enabled or disabled.
Max Supported Speed	The maximum speed that is supported on the port. For the FC port, the maximum speed is 8 Gbps.
Status	The status of the Ethernet port; for example, Linkup or Linkdown.

Adapters Faulty SFP report

The Adapters Faulty SFP report lists all Brocade adapters with unsupported or faulty small form-factor pluggable (SFP) transceivers. This report is only available for Brocade adapters.

For adapters with unsupported SFPs, this report is the same as the Adapters Inventory report filtered to show only those hosts with adapter ports that have unsupported SFPs. For a complete list of fields and components of the Adapters Inventory report, refer to the "Adapters Inventory report" on page 1291.

For adapters with degrading SFPs, there are additional columns in the table that show the SFP/SFP+ and pluggable optical module (POM) properties for the degrading SFP.

Table 129 describes the SFP and POM fields of the Adapters Faulty SFP Report.

TABLE 129 Adapters Faulty SFP report fields and components

Field/Component	Description
SFP Supported	Whether the SFP is supported or not.
Connector Type	The type of port connector; for example, LC. SC, or Cu (copper cable).
Transceiver	The type of transceiver; for example, SFP or SFP+.
Media	The type of media for the transceiver; for example, single mode.
Speed	The port speed.
Extended ID	The identifier for the extended link.
Encoding	Displays how the extended link is encoded, for example, 8B10B.
Baud Rate	The transmission rate, roughly equivalent to the number of bits per second.
Length 9U1	The length of the single-mode fiber-optic cable, used in situations where gigabit performance is not required (for distances greater than 100 meters).
Length 9U2	The length of the single-mode fiber-optic cable, used in situations where gigabit performance is not required (for distances greater than 100 meters).
Length 50u	The length of the 50u fiber-optic cable (for distances greater than 10 meters).
Length 62.5u	The length of the 62.5u fiber-optic cable (for distances greater than 10 meters).
Length Cu	The length of the copper cable (for distances greater than 1 meter, where optimum performance is required).
Vendor Name	The vendor of the extended link.
Vendor OUI	The vendor's organizational unique identifier (OUI).

TABLE 129 Adapters Faulty SFP report fields and components (Continued)

Field/Component	Description
Vendor Part	The part number of the extended link.
Revision	The revision level of the extended link.
Wavelength	The wavelength translation, which enables longer reach through lower attenuation.
Options	Displays details about the transceiver; for example, the type of port connector, type of transceiver, and enable/disable status.
BR Max	The upper bit rate limit at which the SFP transceiver meets its specifications.
BR Min	The lower bit rate limit at which the SFP transceiver meets it specifications.
Serial #	The serial number of the SFP transceiver.
Date Code	The date the SFP was manufactured transceiver.
Temperature (c)	The port temperature, measured in Celsius.
Bias Current (mA)	The low-level DC current (the Bias Current), measured in mA.
Tx Power (mW)	The transmitted power, measured in mW.
Rx Power (mW)	The received power, measured in mW.
Voltage (V)	The voltage; for example, 1.8V, 3.3V, or 5.0V.
Alarm/Warning	Indicates whether an alarm has been triggered.

Application Menus

Dashboard main menus

The menu bar is located at the top of the main window. The following table outlines the many functions available on each menu.

Menu	Command	Command Options
Server Menu		
	Users — Select to configure users and user groups.	
	User Profile — Select to configure user profiles.	
	Active Sessions — Select to display the active Management application sessions.	
	Server Properties — Select to display the Server properties.	
	Options — Select to configure the Management application options.	
	Exit — Select to close the Management Client.	
View Menu		
	Show Main Tab — Select to choose which tab to display.	
		Dashboard — Select to show the dashboard.
		SAN — Select to show the SAN tab.
		IP — Select to show the IP tab.
	Show Panels — Select to choose which widgets to display.	
		All Panels — Select to show the Dashboard and Master Log.
		Dashboard — Select to only show the Dashboard.
		Master Log — Select to only show the Master Log.
Help Menu		
	Contents — Select to open the Online Help.	
	Find — Select to search the Online Help.	
	License — Select to view or change your License information.	
	About *Management_Application_Name* — Select to view the application information, such as the company information, link to patent web page, and release number.	

SAN main menus

The menu bar is located at the top of the main window. The following table outlines the many functions available on each menu.

Menu	Command	Command Options
Server Menu		
	Users — Select to configure users and user groups.	
	User Profile — Select to configure user profiles.	
	Active Sessions — Select to display the active Management application sessions.	
	Server Properties — Select to display the Server properties.	
	Options — Select to configure the Management application options.	
	Exit — Select to close the Management Client.	
Edit Menu		
	Copy — Select to copy information and move it to another location.	
	Show Connections — Select to show connections in a group.	
	Select All — Select to select all objects in the Product List.	
	Properties — Select to display the selected objects properties.	
View Menu		
	Show Main Tab — Select to choose which tab to display.	
		Dashboard — Select to show the dashboard.
		SAN — Select to show the SAN tab.
		IP — Select to show the IP tab.
	Show Panels — Select to select which panels to display.	
		All Panels — Select to show all panels.
		Topology Map — Select to only show the topology map.
		Product List — Select to only show the Product List.
		Master Log — Select to only show the Master Log.
	Manage View — Select to set up the Management application view.	
		Create View — Select to create a new view.
		Display View — Select to display by View All or by a view you create.
		Levels — Select to display by All Levels, Products and Ports, Product Only, or Ports Only.
		Copy View — Select to copy a view.
		Delete View — Select to delete a view.
		Edit View — Select to edit a view.
	Zoom — Select to configure the zoom percentage.	
	Show — Select to determine what products display.	
		Fabrics Only — Select to display only fabrics.
		Groups Only — Select to display only groups.
		All Products — Select to display all products.

Menu	Command	Command Options
		All Ports — Select to display all ports.
	Enable Flyover Display check box — Select to enable flyover display.	
	Show Ports check box — Select to show utilized ports on the selected device.	
	Connected End Devices — Select to show or hide all connected end devices.	
	Include Virtual Devices check box — Select to include virtual devices.	
		Hide All — Select to hide all connected end devices.
		Show All — Select to show all connected end devices.
		Custom — Select to set a custom display for all connected end devices.
		MyCustomList — Lists all custom views.
	Map Display — Select to customize a group's layout to make it easier to view the SAN and manage its devices.	
	Domain ID/Port # — Select to set the display domain IDs and port numbers in decimal or hex format.	
		Decimal — Select to display all domain IDs and port numbers in decimal format.
		Hex — Select to display all domain IDs in hex format.
	Product Label — Select to configure which product labels display.	
		Name — Select to display the product name as the product label.
		Node WWN — Select to display the node name as the product label.
		IP Address — Select to display the IP Address (IPv4 or IPv6 format) as the product label.
		Domain ID — Select to display the domain ID as the product label.
		Zone Alias — Select to display the zone alias as the product label.
	Port Label — Select to configure which port labels display.	
		Name — Select to display the name as the port label.
		Port — Select to display the port number as the port label.
		Port Address — Select to display the port address as the port label.
		Port WWN — Select to display the port world wide name as the port label.
		User Port # — Select to display the user port number as the port label.
		Zone Alias — Select to display the zone alias as the port label.
	Port Display — Select to configure how ports display.	

Menu	Command	Command Options
		Occupied Product Ports — Select to display the ports of the devices in the fabrics (present in the Connectivity Map) that are connected to other devices.
		UnOccupied Product Ports — Select to display the ports of the devices (shown in the Connectivity Map) that are not connected to any other device.
		Attached Ports — Select to display the attached ports of the target devices.
		Switch to Switch Connections — Select to display the switch-to-switch connections.
Discover Menu		
	Fabrics — Select to discover fabrics.	
	Host Adapters — Select to discover hosts.	
	VM Managers — Select to discover VM managers.	
	Host Port Mapping — (Trial and Licensed version Only) Select to manually map HBA ports to a host.	
	Storage Port Mapping — (Trial and Licensed version Only) Select to manually map Storage Ports to a Storage Device or other Storage Ports.	
Configure Menu		
	Element Manager — Select to configure the selected device.	
		Hardware — Select to launch the Element Manager or Web Tools application for the selected device.
		Ports — Select to launch Web Tools - Port Administration for the selected device.
		Admin — Select to launch Web Tools - Switch Administration for the selected device.
		Router Admin — Select to launch Web Tools - FCR Administration for the selected device.
		Name Server — Select to launch Web Tools - Name Server for the selected device.
		HCM — (HBA or CNA only) Select to launch the HCM Agent for the selected device.
	Enable/Disable — Select to enable or disable Virtual Fabrics, switches, and ports.	
		Enable — Select to enable Virtual Fabrics, switches, and ports.
		Disable — Select to disable Virtual Fabrics, switches, and ports.
		Persistent Enable — Select to persistently enable ports.
		Persistent Disable — Select to persistently disable ports.
	Allow/Prohibit Matrix — (Enterprise Licensed version Only) Select to allow FICON users to configure an Allow/Prohibit Matrix table. You can select any matrix tables and compare them either vertically or horizontally.	
	COMPASS — Select to to monitor configuration drifts between an configuration setting and the switch configuration.	

Menu	Command	Command Options
	Configuration — Select to manage the selected device.	
		Save — Select to save device configurations to the repository.
		Save Running to Startup — Select to save the DCB running configuration to the startup configuration on selected switches. Requires at least one discovered DCB switch.
		Restore — Select to restore device configurations from the repository.
		Configuration File Manager — (Trial and Licensed version Only) Select to manage device configurations from the repository.
		Schedule Backup — (Trial and Licensed version Only) Select to schedule configuration backup.
		Replicate — (Trial and Licensed version Only) Select to replicate the switch Configuration or Security.
	Task Scheduler — Select to manage deployment.	
	DCB — Select to manage a DCB switch, port, or link aggregation group (LAG).	
	Encryption — Select to configure encryption for your SAN.	
	Fabric Assigned WWN — Select to configure fabric assigned world wide names to a switch port or AG port.	
	Fabric Binding — (Trial and Licensed version Only) Select to configure whether switches can merge with a selected fabric, which provides security from accidental fabric merges and potential fabric disruption when fabrics become segmented because they cannot merge.	
	FCIP Tunnels — Select to configure tunnels and circuits on FCIP-capable devices.	
	FCoE — Select to manage an FCoE port.	
	FICON — (Enterprise Licensed version Only) Select to configure FICON.	
		Configure Fabric — Select to configure cascaded FICON from the selected fabric.
		Merge Fabrics — Select to merge the selected fabrics.
	Firmware Management — Select to download firmware to devices.	
	High Integrity Fabric — (Trial and Licensed version Only) Select to activate the SCC policy; sets Insistent Domain ID and sets Fabric Wide Consistency Policy for SCC in tolerant mode.	
	Host — Select to manage a selected host.	
		Adapter Software — Select to launch HCM.
		Adapter Ports — Select to configure Host Adapter ports.
	Names — Select to provide familiar simple names to fabrics, products, and ports in your SAN.	
	Port Groups — (Trial and Licensed version Only) Select to configure a group of ports from one or more switches within the same fabric.	

Menu	Command	Command Options
	Port Commissioning — Select to manage port commissioning.	
		Setup — Select to configure port commissioning.
		Decommission — Select to decommission an individual port or all ports on a blade or switch.
		Recommission — Select to recommission an individual port or all ports on a blade or switch.
	Routing — (Trial and Licensed version Only) Select to manage a selected router.	
		Configuration — (Trial and Licensed version Only) Select to view the R_Ports on a router.
		Domain IDs — (Trial and Licensed version Only) Select to configure the router domain IDs.
	Security — Select to manage security.	
		L2 ACL — Select to configure Layer 2 Access Control Lists on products and ports.
	Switch Password Management — Select to manage the password.	
	Swap Blades — (Trial and Licensed version Only) Select to swap blades.	
	Virtual Fabric — (Trial and Licensed version Only) Select to configure logical switches for your SAN.	
		Enable — (Trial and Licensed version Only) Select to enable virtual fabrics for your SAN.
		Disable — (Trial and Licensed version Only) Select to disable virtual fabrics for your SAN.
		Logical Switches — (Trial and Licensed version Only) Select to configure logical switches for your SAN.
		Locate Logical Switches — (Trial and Licensed version Only) Select to locate logical switches.
		Locate Chassis — (Trial and Licensed version Only) Select to locate the chassis.
	VLANs — Select to launch the VLAN Manager.	
	Zoning — Select to configure zones.	
		Fabric — Select to configure fabric zones.
		LSAN Zoning (Device Sharing) — (Trial and Licensed version Only) Select to configure LSAN zones.
		Set Change Limits — Select to set zone limits for zone activation.
		List Zone Members — (Trial and Licensed version Only) Select to display all members in a zone.
Monitor Menu		
	Fabric Vision — Select to configure MAPS or Flow Vision.	

Menu	Command	Command Options
		Flow Vision — Select to define or monitor network traffic by choosing one of the following options: • **Monitor** — Select to monitor network traffic and provide statistics for the defined flows. • **Performance Graph** — Select to monitor performance through a graph, which displays transmit and receive data. The graphs show historical data. • **Add** — Select to define a traffic flow. • **SIM Mode** — Select to enable or disable port mode.
		MAPS — Select to configure or monitor Monitoring and Alerting Policy Suite policies by choosing one of the following options: • **Violations** — Select to view MAPS violations. • **Configure** — Select to configure MAPS policies. • **Enable** — Select to enable MAPS.
	Fabric Watch — Select to manage Fabric Watch.	
		Configure — Select to launch Fabric Watch.
		Port Fencing — (Trial and Licensed version Only) Select to configure port fencing to protect your SAN from repeated operational or security problems experienced by ports.
		Frame Monitor — Select to configure frame monitors.
		Performance Thresholds — (Trial and Licensed version Only) Select to monitor thresholds.
	Policy Monitor — Select to manage best practice policies.	
	Performance — Select to monitor SAN devices.	
		View Utilization — (Trial and Licensed version Only) Select to display connection utilization.
		View Bottlenecks — (Trial and Licensed version Only) Select to display bottlenecks.
		Historical Data Collection — (Trial and Licensed version Only) Select how to monitor historical data by choosing one of the following options: • Enable SAN Wide • Enable Selected • Disable
		End-to-End Monitors — (Trial and Licensed version Only) Select to monitor end-to-end connections.
		Bottlenecks — Select to monitor bottlenecks.
		Clear Counters — Select to clear all port statistics counters at switch level and fabric level.
		Favorites — Select a custom favorite.
		Top Talkers — (Trial and Licensed version Only) Select to monitor performance through a real-time list of top conversations for a switch or port along with related information.
		Real-Time Graph — Select to monitor performance through a graph, which displays transmit and receive data. The graphs show real-time data.

Menu	Command	Command Options
		Historical Graph — (Trial and Licensed version Only) Select to monitor performance through a graph, which displays transmit and receive data. The graphs show historical data.
		Historical Report — (Trial and Licensed version Only) Select to monitor a performance through a table, which displays transmit and receive data. The table shows historical data.
		Bottleneck Graph — (Trial and Licensed version Only) Select to monitor a bottleneck through a graph.
	Discarded Frames — Select to monitor discarded frames.	
	Events — Select to display all events triggered on the selected device.	
	Event Notification — Select to configure the Management application to send event notifications at specified time intervals.	
		E-mail — Select to configure the Management application to send event notifications through e-mail.
		Call Home — (Trial and Licensed version Only) Select to configure the Management Server to automatically dial in to or send an e-mail to a support center to report system problems.
	Event Processing — Select to configure event processing.	
		Pseudo Events — Select to configure pseudo events.
		Event Actions — Select to configure events actions.
	Fabric Tracking — Select to track fabrics.	
		Track Fabric Changes — Select to track fabric changes on the selected fabric.
		Accept Change(s) — (Trial and Licensed version Only) Select to accept changes to the selected fabric.
		Accept All Changes — (Trial and Licensed version Only) Select to accept all changes to all available fabrics in the current view.
	Logs — Select to display logs.	
		Audit — Select to display a history of user actions performed through the application (except login or logout).
		Fabric — Select to display the events related to the selected fabric.
		FICON — Select to display the FICON events related to the selected device or fabric.
		Product Event — Select to display errors related to SNMP traps and Client-Server communications.
		Product Status — Select to display operational status changes of managed products.
		Security — Select to display security information.
		Syslog — Select to display Syslog events related to the selected device or fabric.
	Port Auto Disable — Select to configure the port auto disable flag on individual FC ports or all ports on a selected device, as well as unblock currently blocked ports.	

Menu	Command	Command Options
	Port Connectivity — Select to view port connectivity on the selected device.	
	Port Optics (SFP) — Select to display the properties associated with a selected small form-factor pluggable (SFP) transceiver on the selected device.	
	SNMP Setup — Select to configure SNMP traps.	
		Trap Forwarding — Select to configure trap forwarding.
		Product Trap Recipients — Select to register a host as a trap recipient.
		Event Reception — Select to configure the server to accept or drop traps and specify SNMP credentials and community strings, which are required to decode traps on receiving them.
		Informs — Select to enable or disable SNMP informs on the device.
	Syslog Configuration — Select to configure Syslog for the Management server.	
		Syslog Forwarding — Select to configure Syslog forwarding.
		Product Syslog Recipients — Select to register a host as a syslog recipient.
	Technical Support — Select to configure technical support data.	
		SupportSave — Select to capture server and client support data.
		Product/Host SupportSave — (Fabric OS devices only) Select to configure technical support data collection.
		Upload Failure Data Capture — Select to configure capture failure data for Fabric OS devices.
		View Repository — Select to view repository data.
	Troubleshooting — Select to troubleshoot your SAN.	
		FC — Select how to troubleshoot FC by choosing one of the following options: • **FC Trace Route** — Select to view the route information between two device ports. • **Device Connectivity** — Select to view the connectivity information for two devices. • **Fabric Device Sharing** — (Trial and Licensed version Only) Select to determine if the selected fabrics are configured to share devices. • **Diagnostic Port Test** — Select to run a diagnostic port test.
		FCIP — Select how to troubleshoot FCIP by choosing one of the following options: • **Ping** — Select to perform a zoning check between the selected device port WWNs. • **Trace Route** — (Trial and Licensed version Only) Select to view the route information from a source port on the local device to a destination port on another device. • **Performance** — (Trial and Licensed version Only) Select to view IP performance between two devices.
Reports Menu		

Menu	Command	Command Options
	Event Custom Reports — Select to generate custom event reports.	
	Generate — Select to determine which reports to run.	
	View — Select to view reports through the application or through an Internet browser.	
Tools Menu		
	Setup — (Trial and Licensed version Only) Select to set up the applications that display on the **Tools** menu.	
	Product Menu — Select to access the tools available on a device's shortcut menu.	
	Plug-in for SCOM — Select to configure a SCOM server.	
	Tools List (determined by user settings) — Select to open a software application. You can configure the **Tools** menu to display different software applications. Recommended tools to include in this menu include an Internet browser, the command prompt application, and Notepad.	
Help Menu		
	Contents — Select to open the Online Help.	
	Find — Select to search the Online Help.	
	License — Select to view or change your license information.	
	About *Management_Application_Name* — Select to view the application information, such as the company information, link to patent web page, and release number.	

SAN shortcut menus

You can use the Management application interface main menu to configure, monitor, and troubleshoot your SAN components. The instructions for using these features are documented in the associated chapters of this manual.

For each SAN component, you can optionally right-click the component and a shortcut menu displays. The table below details the command options available for each component.

Component	Menu/Submenu Commands	Comments
FC Fabric / Backbone Fabric		
	View > Connected End Devices > Include Virtual Devices check box Hide All Show All Custom MyCustomList Create View Automatically Port List Node List	
	Fabric Binding	Trial and Licensed version Only

Component	Menu/Submenu Commands	Comments
	FCIP Tunnels	Only launches the wizard when FCIP-capable switches are in the selected fabric.
	FICON >	Trial and Licensed version Only
	Configure Fabric	
	Merge Fabrics	
	High Integrity Fabric	Trial and Licensed version Only
	Routing >	Trial and Licensed version Only
	Configuration	
	Domain IDs	
	Zoning >	
	Fabric	
	LSAN Zoning (Device Sharing) Trial and Licensed version Only Only enabled for Backbone fabrics.	
	Performance >	
	End-to-End Monitors (Trial and Licensed version Only)	
	Real-Time Graph	
	Historical Graph (Trial and Licensed version Only)	
	Historical Report (Trial and Licensed version Only)	
	Bottleneck Graph (Trial and Licensed version Only)	
	Events	
	Track Fabric Changes check box	Trial and Licensed version Only
	Accept Changes	Trial and Licensed version Only
	Port Connectivity	
	Technical Support >	
	SupportSave	
	Product/Host SupportSave	
	Upload Failure Data Capture	
	View Repository	
	FC Trace Route	
	Create Meta SAN View	Only available for Backbone fabrics. Automatically creates a view with the selected fabric. View name is same as the current label.
	Map Display	
	Port Display >	Only available from Product List.
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	

Component	Menu/Submenu Commands	Comments
	Table >	Only available from Product List.
	Copy '*Fabric_Name*'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
	Collapse or Expand	Only available from Connectivity Map
	Properties	
Device Group		
	Host Port Mapping	Only available for hosts or host group.
	Zoning	Only available for switch group.
	Storage Port Mapping	Trial and Licensed version Only Only available for storage group.
	Map Display	
	Port Display >	Only available from Product List.
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	
	Table >	Only available from Product List.
	Copy '*Device_Name* Group'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
	Collapse or Expand	Only available from Connectivity Map
Fabric OS Switch/Chassis/Access Gateway		

Component	Menu/Submenu Commands	Comments
	Element Manager > Hardware Ports Admin Router Admin Name Server	
	Enable / Disable > Enable Disable Persistent Enable Persistent Disable	
	Allow/Prohibit Matrix	Enterprise Edition Only Only available for Fabric OS devices. Only enabled when the Fabric OS device is FICON-capable and has the Enhanced Group Management license.
	Configuration > Save Save Running to Startup (DCB-capable switch) Restore Configuration Repository Schedule Backup (Trial and Licensed version Only) Replicate > Configuration (Trial and Licensed version Only) Security (Trial and Licensed version Only)	
	Fabric Assigned WWN	
	FICON > Configure Fabric Merge Fabrics	Trial and Licensed version Only
	Firmware Management	
	Decommission > All Ports on the Switch All Ports on the Blade	
	Recommission > All Ports on the Switch All Ports on the Blade	
	Swap Blades	
	Virtual Fabric > Disable Logical Switches Locate Logical Switches > *List_of_Logical_Switches* (Fabric OS only) (Virtual Fabric-capable switches only)	

Component	Menu/Submenu Commands	Comments
	Zoning > Fabric	Does not display when switch is in a Core Switch group, Chassis group or Isolated device group, or when it is in Access Gateway mode.
	DCB (DCB-capable switch)	
	FCoE (DCB-capable switch)	
	Performance > Clear Counters Top Talkers (Trial and Licensed version Only) Real-Time Graph Historical Graph (Trial and Licensed version Only) Historical Report (Trial and Licensed version Only) Bottleneck Graph (Trial and Licensed version Only)	
	Events	
	Accept Change	Trial and Licensed version Only Only enabled in tracked FC Fabrics. Only enabled when a plus or minus icon is present.
	Fabric Watch > Configure Port Fencing (Trial and Licensed version Only) Frame Monitor Performance Thresholds	
	Port Connectivity	
	Port Optics (SFP)	
	Technical Support > SupportSave Product/Host SupportSave Upload Failure Data Capture View Repository	
	FC Trace Route	
	Telnet	
	Telnet through Server	
	<User-defined menu item>	Trial and Licensed version Only Configured in Setup Tools. May be more than one item.
	Setup Tools	Trial and Licensed version Only
	Product	Only enabled when the fabric is tracked, and the product is removed and joins another fabric.
	Other Ports > <Fabric Name 1> <Fabric Name 2>	Does not display when an Access Gateway mode device is attached to multiple fabrics.
	Show Ports check box	

Component	Menu/Submenu Commands	Comments
	Show Connections	
	Port Display >	Only available from Product List.
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	
	Table >	Only available from Product List.
	Copy '*Device_Name* Group'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
	Properties	
Core Switch		
	Element Manager	Only available from Product List.
	Enable/Disable	
	Enable	
	Disable	
	Configuration > (Fabric OS only)	
	Save	
	Restore	
	Schedule Backup (Trial and Licensed version Only)	
	Configuration Repository	
	Replicate >	
	Configuration (Trial and Licensed version Only) Security (Trial and Licensed version Only)	
	Swap Blades	
	Firmware Management (Fabric OS only)	
	Virtual Fabric >	Only available from Product List.
	Enable	
	Disable	
	Logical Switches	
	Locate Chassis	
	Events	

Component	Menu/Submenu Commands	Comments
	Technical Support > (Fabric OS only)	
	Product/Host SupportSave	
	Upload Failure Data Capture	
	View Repository	
	Port Display >	Only available from Product List.
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	
	Table >	Only available from Product List.
	Copy '*Device_Name* Group'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
	Properties	
DCB		
	Element Manager >	Launches Web Tools.
	Hardware	
	Ports	
	Admin	
	Router Admin	
	Name Server	
	Configuration >	
	Save	
	Save Running to Startup	
	Restore	
	Configuration Repository	
	Schedule Backup	
	Replicate >	
	Configuration	
	Security	

Component	Menu/Submenu Commands	Comments
	Enable / Disable > Enable Disable Persistent Enable Persistent Disable	
	Firmware Management	
	Swap Blades	Only available from chassis.
	Zoning	
	DCB	
	FCoE	
	VLAN	
	Allow / Prohibit Matrix	
	Security > L2 ACL	
	Performance > Clear Counters Top Talkers Real-Time Graph Historical Graph Historical Report Bottleneck Graph	
	Fabric Watch > Configure Port Fencing Frame Monitor Performance Thresholds	
	Technical Support > Product / Host SupportSave Upload Failure Data Capture** View Repository	
	Events	
	Port Connectivity	
	Port Optics (SFP)	
	Telnet	
	Telnet through Server	
	<User-defined menu item>	
	Setup Tools	
	Product	Only enabled when the fabric is tracked, and the product is removed and joins another fabric.

Component	Menu/Submenu Commands	Comments
	<Other Ports >	Visible only for AGs that are attached to multiple fabrics.
	<Fabric Name 1>	
	<Fabric Name 2>	
	Show Ports	
	Accept Changes	
	Show Connections	
	Port Display >	Only available from Product List.
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	
	Properties	
	Table >	Only available from Product List.
	Copy '*Device_Name* Group'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
HBA, iSCSI Host, and HBA Enclosure		
	Element Manager	Launches Element Manager for Fabric OS HBAs discovered using JSON agent. Launches blank window for unmanaged Fabric OS HBAs.
	Host Port Mapping	Only available for Brocade, Emulex, and Qlogic HBAs and HBA enclosures.
	Performance >	Disabled when all ports are offline.
	Real Time Graphs	Does not display for Node Origin and Routed instance in a routed fabric.
	LightPulse Utility/NT	Only available for Emulex devices. Launches with Origin in context for routed device.
	Emulex Configuration Tool	Only available for Emulex devices. Launches with Origin in context for routed device.
	SANSurfer	Only available for Qlogic HBAs.
	<User-defined menu item>	Configured in Setup Tools. May be more than one item.
	Host	Only available in Fabric view for managed HBAs.
	Setup Tools	Trial and Licensed version Only

Component	Menu/Submenu Commands	Comments
	Show Ports	
	Show Connections	
	Fabric > Fabric1 Fabric2	Only available for HBAs under the Host node.
	Origin	Only available for HBAs under the Host node or devices routed in. Not available for enclosures.
	Destination	Only available for devices routed out. Not available for enclosures.
	Port Display > Occupied Product Ports UnOccupied Product Ports Attached Ports Switch to Switch Connections	Only available from Product List.
	Expand All	Only available from Product List.
	Collapse All	Only available from Product List.
	Properties	
Storage, iSCSI Storage, and Storage Enclosure		
	Storage Port Mapping	Trial and Licensed version Only Disabled for routed device.
	<User defined menu item>	
	Setup Tools	Trial and Licensed version Only
	Show Ports	
	Show Connections	
	Origin	Only available for devices routed in. Not available for enclosures.
	Destination	Only available for devices routed out. Not available for enclosures.
	Port Display > Occupied Product Ports UnOccupied Product Ports Attached Ports Switch to Switch Connections	Only available from Product List.

Component	Menu/Submenu Commands	Comments
	Table >	Only available from Product List.
	Copy '*Device_Name* Group'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
	Properties	
Router Phantom Domains		
	Accept Change	Trial and Licensed version Only Only available for tracked FC Fabrics. Only enabled when a plus or minus icon is present.
	Show Connections	Displays as disabled because this component does not display in the Connectivity Map.
	Origin	
	Port Display >	Only available from Product List.
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	
	Table >	Only available from Product List.
	Copy '*Device_Name* Group'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
	Properties	
Switch Port FC		

Component	Menu/Submenu Commands	Comments
	Enable / Disable > Enable Disable Persistent Enable Persistent Disable	
	Decommission > Port	
	Recommission > Port	
	Zoning > List Zone Members	
	Performance > Real-Time Graph Historical Graph (Trial and Licensed version Only) Historical Report (Trial and Licensed version Only) Bottleneck Graph (Trial and Licensed version Only)	
	MAPS Violations	
	Discarded Frames	
	Locate Connected Port	
	Port Display > Occupied Product Ports UnOccupied Product Ports Attached Ports Switch to Switch Connections	Only available from Product List.
	Table > Copy '*Device_Name* Group' Copy Row Copy Table Export Row Export Table Search Select All Size All Columns To Fit Expand All Collapse All Customize	Only available from Product List.
	Collapse All	Only available from Product List.
	Properties	
HBA and iSCSI Initiator		
	Host Port Mapping	Only available for Brocade, Emulex, and Qlogic HBAs and HBA enclosures.
	Performance > Real Time Graphs	Disabled when all ports are offline.

Component	Menu/Submenu Commands	Comments
	FC Security Protocol	Only available for Managed JSON HBA Ports. Only available when you have the Security Privilege.
	Zoning	
	List Zone Members	Trial and Licensed version Only
	Connected Port	
	Port Display > Occupied Product Ports UnOccupied Product Ports Attached Ports Switch to Switch Connections	Only available from Product List.
	Table > Copy '*Device_Name* Group' Copy Row Copy Table Export Row Export Table Search Select All Size All Columns To Fit Expand All Collapse All Customize	Only available from Product List.
	Properties	
HBA Port		
	Host Port Mapping	Does not display for routed devices.
	Performance > Real Time Graphs	Only available for occupied, managed ports. Disabled when all ports are offline.
	FC Security Protocol	Only available for Managed JSON HBA Ports. Only available when you have the Security Privilege.
	Zoning	
	List Zone Members	Trial and Licensed version Only
	Connected Port	
	Port Display > Occupied Product Ports UnOccupied Product Ports Attached Ports Switch to Switch Connections	Only available from Product List.
	Expand All	Only available from Product List.
	Collapse All	Only available from Product List.
	Properties	
Storage Node		

Component	Menu/Submenu Commands	Comments
	Storage Port Mapping	Trial and Licensed version Only
	Show Ports	Does not display for routed devices.
	Show Connections	
Storage FC and iSCSI Storage port		
	Storage Port Mapping	Trial and Licensed version Only
	Zoning	
	List Zone Members	Trial and Licensed version Only
	Connected Port	
	Port Display > Occupied Product Ports UnOccupied Product Ports Attached Ports Switch to Switch Connections	Only available from Product List.
	Table > Copy '*Device_Name* Group' Copy Row Copy Table Export Row Export Table Search Select All Size All Columns To Fit Expand All Collapse All Customize	Only available from Product List.
	Properties	
Giga-Bit Ethernet Port		
	Performance > Real-Time Graph	
	Enable / Disable > Enable Disable Persistent Enable Persistent Disable	
	Modify	Launches Element Manager.
	IP Troubleshooting > Ping Trace Route Performance (Trial and Licensed version Only)	

Component	Menu/Submenu Commands	Comments
	Port Display >	Only available from Product List.
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	
	Table >	Only available from Product List.
	Copy '*Device_Name* Group'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
	Properties	
Connection		
	Properties	
FCIP Tunnel		
	Properties	
Trunk		
	Port Display >	Only available from Product List.
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	
	Table >	Only available from Product List.
	Copy '*Device_Name* Group'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
	Properties	

Component	Menu/Submenu Commands	Comments
VE Port (physical)		
	Enable / Disable >	
	Enable	
	Disable	
	Persistent Enable	
	Persistent Disable	
	Decommission > Port	
	Recommission > Port	
	Zoning >	
	List Zone Members	
	Performance >	
	Real-Time Graph	
	Historical Graph (Trial and Licensed version Only)	
	Historical Report (Trial and Licensed version Only)	
	Bottleneck Graph (Trial and Licensed version Only)	
	MAPS Violations	
	Locate	
	Locate Connected Port	
	Port Display >	Only available from Product List.
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	
	Table >	Only available from Product List.
	Copy '*Device_Name* Group'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
	Properties	
White Area of the Connectivity Map		
	Accept All Changes	
	Zoom	
	Zoom In	
	Zoom Out	

Component	Menu/Submenu Commands	Comments
	Map Display	
	Expand	
	Collapse	
	Export	
White Area of the Product List		
	Port Display >	
	Occupied Product Ports	
	UnOccupied Product Ports	
	Attached Ports	
	Switch to Switch Connections	
	Table >	
	Copy '*Component*'	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
Product List		
	Table >	Some form of this shortcut menu is available for all tables in the Management interface.
	Copy '*Component*'	
	Copy Table	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	

Right-click option to enable or disable multiple ports

You can enable or disable multiple switch ports by selecting enable or disable option in the right-click shortcut menu. This multiple ports selection option is hidden, if a combination of switch ports and end device ports are selected.

You can disable multiple ports if the port selections are already enabled, and vice versa. If the multiple port selection is a combination of enabled and disabled ports, both the options are displayed. The following warning messages display when you enable or disable multiple ports.

“You are trying to enable the selected ports. If the ports are connected to another switch, the fabric might reconfigure. Do you want to continue?”.

“You are trying to disable the selected ports. If the ports are connected to another switch, the fabric might reconfigure. Devices connected to the ports can no longer communicate with the fabric. Do you want to continue?”.

Click Yes to continue. Port enable or disable status dialog displaying if the operation is a success or failure for each port.

FIGURE 581 Port enable or disable status dialog

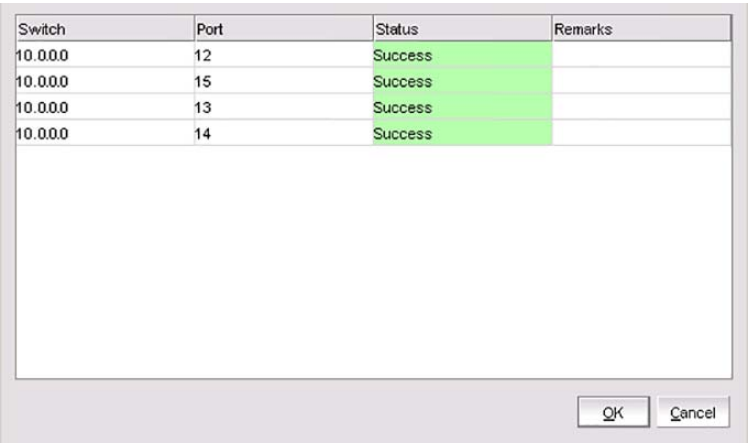

Call Home Event Tables

This appendix provides information about the specific events that display when using Call Home. This information is shown in the following Event Tables.

TABLE 130 # CONSRV Events

Event reason code	FRU code/Event type	Description	Severity
504	DVP/LIM/HW	Port module failure.	3
506	DVP/PORT	Fibre Channel port failure.	3
509	DVP/PORT	Fibre Channel path failure.	0
511	LIM/DVP	LIM SPP failure.	3
514	DVP/ LIM/PORT	SFP/XFP optics failure.	3
517	LIM	LIM SPP offline.	3
530	LIM/DVP	LIM Power-up diagnostic failure.	3
536	LIM/DVP	Internal Frame Error port anomaly - threshold exceeded.	2
604	SBAR/SWM/HW	SBAR module failure.	3
607	SBAR/SWM/HW	Switch contains no operational SBAR cards.	4
610	SWM/INFO	SWM BMAC Link Down.	0
622	SBAR/INFO	SWM powered off.	0
625	SBAR/INFO	SWM NV RAM failure.	0

TABLE 131 # Thermal Events

Event reason code	FRU code/Event type	Description	Severity
800	DVP/LIM/HW	High temperature warning.	3
801	DVP/LIM/HW	Critically hot temperature warning.	3
802	DVP/LIM/HW	Port card shutdown due to thermal violations.	3
805	SWM/SBAR/HW	High temperature warning.	3
806	SWM/SBAR/HW	Critically hot temperature warning.	3
807	SWM/SBAR/HW	SBAR module shutdown due to thermal violations.	3
810	CTP/HW	High temperature warning.	3

TABLE 131 # Thermal Events (Continued)

Event reason code	FRU code/Event type	Description	Severity
811	CTP/HW	Critically hot temperature warning.	3
812	CTP/HW	CTP shutdown due to thermal violations.	3
850	CTP/HW	System shutdown due to CTP thermal threshold violations.	4

TABLE 132 Fabric OS Events

Event reason code	FRU code/Event type	Description	Severity
N/A	Ethernet	Switch is not reachable.	3
N/A	SW-Missing	Switch is missing from the fabric.	3
1009	MS-1009	Error in registered link incident record (RLIR).	4
1021	MAPS-1021	Core blade redundancy.	3
1021	MAPS-1021	Error ports.	3
1021	MAPS-1021	Faulty CPs.	3
1021	MAPS-1021	Faulty or absent blades.	3
1021	MAPS-1021	Faulty or absent fans.	3
1021	MAPS-1021	Faulty or absent power supplies.	3
1021	MAPS-1021	Faulty ports.	3
1021	MAPS-1021	Faulty temperature sensors.	3
1021	MAPS-1021	Faulty WWN cards.	3
1021	MAPS-1021	Flash usage is out of range.	3
1021	MAPS-1021	Marginal ports.	3
1021	MAPS-1021	Missing SFPs.	3
1034	EM-1034	Faulty FRU.	4
1402	FW-1402	Flash usage is out of range (Fabric OS version 6.0 or earlier).	3
1426	FW-1426	Faulty or missing power supply.	3
1427	FW-1427	Faulty power supply.	3
1428	FW-1428	Missing power supply.	3
1429	FW-1429	Problem in power supply arrangement.	3
1430	FW-1430	Faulty temperature sensors.	3
1431	FW-1431	Faulty fans.	3
1432	FW-1432	Faulty WWN cards.	3
1433	FW-1433	Faulty CPs.	3

TABLE 132 Fabric OS Events (Continued)

Event reason code	FRU code/Event type	Description	Severity
1434	FW-1434	Faulty blades.	3
1435	FW-1435	Flash usage is out of range (Fabric OS version 6.1 or later).	3
1436	FW-1436	Marginal port.	3
1437	FW-1437	Faulty port.	3
1438	FW-1438	Faulty or missing SFPs.	3
1444	FW-1444	Faulty FRU.	3
1447	FW-1447	Core blades/SFM failures.	3

Event Categories

This section provides information about the events that display in each of the following categories:

Link incident events

The following link incident events indicate FICON link status changes:

- Link RNID device registration
- Link RNID device de-registration
- Link listener added RLIR
- Link listener removed
- Link RLIR failure

Traps that begin with OID 1.3.6.1.4.1.1588.2.1..1.50 are categorized as link incident events.

Product status events

Product status events indicate a change in the status of the product; for example, changes in the state of the port, the field replaceable unit (FRU), the sensor, or the CP.

Traps that begin with any of the following OIDs are categorized as product status events.

- 1.3.6.1.3.94.0.1 [connUnitStatusChange]
- 1.3.6.1.3.94.0.5 [connUnitSensorStatusChange]
- 1.3.6.1.3.94.0.6 [connUnitPortStatusChange]
- 1.3.6.1.4.1.1588.2.1.1.1.0.3 [swFCPortScn]
- 1.3.6.1.4.1.1588.2.1.1.1.0.15 [swDeviceStatusTrap]
- 1.3.6.1.4.1.1588.2.1.2.2.0.1 [fruStatusChanged]
- 1.3.6.1.4.1.1588.2.1.2.2.0.2 [cpStatusChanged]

If the event is a RASLOG and if the RASLOG ID matches any of the RASLOGS listed below, then the event is categorized as a product status event.

- FW-1424
- FW-1425
- FW-1426

- FW-1427
- FW-1428
- FW-1429
- FW-1430
- FW-1431
- FW-1432
- FW-1433
- FW-1434
- FW-1435
- FW-1436
- FW-1437
- FW-1438
- FW-1439
- FW-1440
- FW-1441
- FW-1442
- FW-1443
- FW-1444

Product audit events

Events that are used to track audit information are categorized as product audit events. Audit Syslog messages from HBAs and the messages with the IDs listed below are categorized as product audit events.

- TRCK-1001
- TRCK-1002
- TRCK-1003
- TRCK-1004
- TRCK-1005
- TRCK-1006

Security events

Security events are those that indicate authentication success or failure, a security violation, or user login and logout.

Security events for FC devices

For FOS switches, if the event is a RASLOG event and the RASLOG ID contains 'SEC', then the event is categorized as a security event.

User action events

User action events are generated for user actions that are performed through the Management applications, such as:

- User creation
- User deletion
- Event action enable
- Event action disable

These events are usually generated to notify status of configuration or data collection operations initiated by the user from the Management application.

Management server events

Management Server Events are those that are generated by the Management application server, such as:

- Service start and stop
- Memory usage
- Device discovery status
- Asset collection status

These events are usually generated to notify the status of server tasks that are running regularly and periodically.

Product events

All other events originating from the product are categorized as product events.

IP Performance monitoring events

IP performance monitoring events, listed in Table 133, occur when users select the option to forward events to the vCenter during VM Manager discovery.

TABLE 133 Performance monitoring IP threshold events

Trap name	OID	Description
bnaRisingThresholdCrossed	1.3.6.1.4.1.1991.1.13.2.0.1	The value of monitored SNMP variable or expression has exceeded the value specified as the higher threshold.
bnaFallingThresholdCrossed	1.3.6.1.4.1.1991.1.13.2.0.1	The value of the monitored SNMP variable or expression has failed below the value specified as the lower threshold.

RASLog Events

The supported events for Event Triggered backup are listed in Table 134.

TABLE 134 Configuration change events

Event ID	Type	Description
AG-1006	LOG	Access Gateway mode has been enabled or disabled.
AG-1013	LOG	N_port failover.

TABLE 134 Configuration change events (Continued)

Event ID	Type	Description
AG-1014	LOG	N_port failback.
AG-1016	LOG	Failing over F_Ports mapped to N_Port <*port number*> to other N_Ports.
AG-1020	LOG	F_Ports to N_Ports route mapping has been changed.
AG-1022	LOG	F_Port <*f_port*> is failed over to its preferred N_Port <*n_port*>.
AG-1023	LOG	F_Port <*f_port*> mapped to offline N_Port <*n_port*> is failed over to its preferred N_Port <*preferred port*>.
AG-1033	AUDIT | LOG	F_Port to N_Port mapping has been updated for N_Port (<*n_port*>).
AG-1041	LOG	Static F_Ports mapped to N_Port <*port number*> are disabled as Trunking is enabled on the N_Port.
AN-1006	AUDIT	Bottleneck detection configuration is successfully changed.
AN-1006	AUDIT	Bottleneck detection configuration is successfully changed.
BL-1001	LOG	Port initialization completed.
BLS-1002	AUDIT | LOG	An IPsec/IKE policy was added.
BLS-1003	AUDIT | LOG	An IPsec/IKE policy was deleted.
CONF-1000	LOG | AUDIT	configDownload completed successfully.
CONF-1031	LOG	configDefault completed successfully.
CONF-1032	LOG	configRemove completed successfully.
CONF-1042	LOG | AUDIT	Indicates that the fabric configuration parameter value has been changed.
CONF-1043	LOG | AUDIT	Indicates that the fabric configuration parameter value has been changed.
CONF-1044	LOG | AUDIT	Indicates that the fabric configuration parameter value has been changed by a user.
FABR-1016	LOG	FICON mode enabled.
FABR-1017	LOG	FICON mode disabled.
FABR-1030	LOG	Domain ID changed
FCR-1005	LOG	Indicates that a device is removed from the logical storage area network (LSAN) zone in the edge fabric.
FCR-1006	LOG	Indicates that a device is added to a logical storage area network (LSAN) zone in the edge fabric.
FCR-1007	LOG	Indicates that a logical storage area network (LSAN) zone attached to the specified port was deleted in the edge fabric.
FCR-1008	LOG	Indicates that a logical storage area network (LSAN) zone was created at the specified port in the edge fabric.
FCR-1009	LOG	Indicates that a logical storage area network (LSAN) zone was enabled in the edge fabric attached to the specified port. The enabled LSAN zone configuration is listed.
FCR-1010	LOG	Indicates that a logical storage area network (LSAN) zone is disabled in the edge fabric attached to the specified port.
FCR-1011	LOG	Indicates that a logical storage area network (LSAN) zone update was received from another domain.
FCR-1015	LOG	Indicates that an EX_Port was created on the specified port in the specified domain.
FCR-1016	LOG	Indicates that a fabric is no longer accessible through the backbone fabric. This may be caused by a link or switch failure.
FCR-1034	LOG	LSAN zone added in backbone fabric.
FCR-1035	LOG	LSAN zone device <*device WWN*> added in the backbone fabric.
FCR-1036	LOG	LSAN zone <*zone name*> enabled in the backbone fabric.
FCR-1037	LOG	LSAN zone disabled in the backbone fabric.
FCR-1061	LOG	Backbone fabric created on a port

TABLE 134 Configuration change events (Continued)

Event ID	Type	Description
FCR-1068	LOG	FCR disabled.
FCR-1069	LOG	FCR enabled.
FCR-1071	LOG	Port is changed from non-FCR port to FCR port.
FCR-1072	LOG	Port is changed from FCR port to non-FCR port.
FCR-1088	LOG	LSAN <*Enforce/Speed*> tag <*Tag Name*> added.
FCR-1089	LOG	LSAN <*Enforce/Speed*> tag <*Tag Name*> removed.
FCR-1091	LOG	Backbone Fabric ID changed to <*Tag*>.
FCR-1102	LOG	ICL EX_Port <*Port Numbers*> need to be present in base switch to make a recommended topology.
FICU-1008	LOG	FMS mode enabled.
FICU-1012	LOG	FMS mode disabled.
FV-3000	AUDIT	Flow definition created.
FV-3001	AUDIT	Flow definition deleted.
FV-3002	AUDIT	Flow definition activated.
FV-3003	AUDIT	Flow definition de-activated.
FV-3004	AUDIT	Flow definition modified.
FV-3005	AUDIT	Flow reset.
FW-1424	LOG	Indicates that the switch is not in a healthy state. This occurred because of a policy violation.
FW-1425	LOG	Indicates that the switch status has changed to a healthy state. This occurred because a policy is no longer violated.
IPAD-1000	LOG	Switch IP change
IPAD-1002	AUDIT | LOG	Switch name has been successfully changed to <*Switch name*>.
MAPS-1100	LOG | AUDIT	Rule <*Rule name*> is created.
MAPS-1101	LOG | AUDIT	Rule <*Rule name*> is deleted.
MAPS-1102	LOG | AUDIT	Rule <*Rule name*> is modified.
MAPS-1110	LOG | AUDIT	MAPS policy created.
MAPS-1111	LOG | AUDIT	MAPS policy deleted.
MAPS-1112	LOG | AUDIT	MAPS policy cloned.
MAPS-1113	LOG | AUDIT	MAPS policy activated.
MAPS-1114	AUDIT | LOG	Rule <*Rule name*> added to Policy <*Policy name*>.
MAPS-1115	AUDIT | LOG	Rule <*Rule name*> deleted from Policy <*Policy name*>.
MAPS-1120	LOG | AUDIT	Group <*Group name*> created.
MAPS-1121	LOG | AUDIT	Group <*Group name*> deleted.
MAPS-1122	LOG | AUDIT	Group <*Source group name*> cloned to <*Target group name*>.
MAPS-1123	LOG | AUDIT	Group <*Group name*> modified.
MAPS-1124	LOG | AUDIT	Flow <*Flow name*> imported.
MAPS-1125	LOG | AUDIT	Flow <*Flow name*> deimported.
MAPS-1130	LOG | AUDIT	MAPS actions configured on the switch.

TABLE 134 Configuration change events (Continued)

Event ID	Type	Description
MAPS-1200	AUDIT | LOG	Fabric Watch thresholds are converted to MAPS policies.
MAPS-1201	LOG | AUDIT	MAPS has started monitoring the system and Fabric Watch monitoring is disabled.
MAPS-1202	LOG | AUDIT	MAPS is disabled.
MAPS-1204	LOG | AUDIT	Port toggle action is successful on raslog port.
MAPS-1205	LOG | AUDIT	MAPS aborted port toggle action on port raslog.
NSM-1001	LOG	10 Gbe port online.
PMGR-1001	LOG | AUDIT	Attempt to create switch ID succeded.
PMGR-1003	LOG	Attempt to delete switch ID succeded.
PMGR-1005	LOG	Attempt to move port on slot succeded.
PMGR-1007	LOG	Attempt to change switch succeded.
PMGR-1009	LOG	Attempt to change the base switch to software succeded.
PMGR-1011	LOG	Attempt to move port to switch succeeded.
SEC-1319	LOG	FCS Policy changed.
SEC-3051	AUDIT | LOG	License key added or removed.
SNMP-1005	AUDIT | LOG	SNMP configuration hanged.
SNMP-1006	AUDIT | LOG	Indicates that the specified SNMP configuration group was reset to the factory default.
SULB-1004	AUDIT | LOG	Firmware commit has completed.
XTUN-2007	LOG	FCIP Tunnel Circuit created.
XTUN-2020	LOG	FCIP Tunnel deleted.
XTUN-2021	LOG	FCIP Tunnel Circuit deleted.
XTUN-2022	LOG	FCIP Tunnel modified.
XTUN-2024	LOG	FCIP Tunnel Circuit modified.
ZONE-1022	LOG	The effective configuration has changed to <*Effective configuration name*>. <*AD Id*>.
ZONE-1023	LOG	Switch connected to port (<*port number*>) is busy. Retrying zone merge.
ZONE-1024	LOG	Indicates that the cfgSave command has completed successfully.
ZONE-1034	LOG	A new zone database file is created.
ZONE-1042	LOG	The effective configuration has been disabled. <*AD Id*>.
ZONE-1043	LOG	The Default Zone access mode is set to No Access.
ZONE-1044	LOG	The Default Zone access mode is set to All Access.

User Privileges

User privileges

The Management application provides the user administrator with a high level of control over what functions individual users can see and use. The user privileges describes the effect that each user privilege has on the application when placed in one of the three available configurations: no privilege, read-only, and read/write.

User privilege is the Management application's method of providing role-based access control (RBAC) to the software's user administrator.

In the Management application, privileges are assigned to roles and devices are assigned to areas of responsibility (AORs). Privileges and devices are not directly assigned to users; users receive privileges and obtain access to devices from the roles and AORs to which they are assigned. You can assign multiple roles and AORs to a single user.

The following tables define all the privileges in the Management application and the behavior of the application if the privilege is not given, read-only, or read/write.

TABLE 135 Application privileges and behavior

Privilege	Description	No Privilege	Read-Only	Read/Write
Active Session Management	Allows you view active client sessions and disconnect an unwanted user.	Disables the **Active Sessions** command from the **Server** menu.	Enables the **Active Sessions** command from the **Server** menu. Disables all commands and functions on the dialog box except the **Close** and **Help**.	Enables the **Active Sessions** command from the **Server** menu. Enables all commands and functions on the dialog box.
Call Home	Allows you to configure call home centers, devices, and event filters.	Disables the **Call Home** command on the **Monitor > Event Notification** menu.	Enables the **Call Home** command on the **Monitor > Event Notification** menu. Enables the **Add**, **Edit**, **Remove**, **Edit Centers**, and **Show/Hide Centers** buttons as well as the **Enabled** check boxes on the dialog box; however, disables the **OK** and **Apply** buttons on the **Call Home**, **Call Home Event Filter**, and **Configure Call Home Center** dialog box boxes.	Enables the **Call Home** command on the **Monitor > Event Notification** menu. Enables all functions in the dialog box.
Certificate Management	Allows you to access the **Certificate Management** dialog box and manage server truststores.	Disables **Certificates** on the **Options** dialog box.	Enables **Certificates** on the **Options** dialog box. Only viewing of the certificates is supported.	Enables **Certificates** on the **Options** dialog box. Enables all functions in the dialog box.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Configuration Management	Allows you to access the **Configuration Management** dialog box and perform configuration upload and replication. Allows you to access the **COMPASS** dialog box and link, unlink, and synchronize configuration templates to fabrics or groups and monitor for drifts.	Disables **Save**, **Restore**, **Configuration Repository**, and **Schedule Backup** under **Configure > Switch** and the **Configuration** command under **Configure > Switch > Replicate**. Disables **COMPASS** command from the **Configure** menu.	Enables **Configuration Repository** under **Configure > Switch**. Only viewing of saved configuration is supported. Configuration upload and replication are disabled. Enables **COMPASS** command from the **Configure** menu. Allows you to access the **Monitor** tab of the **COMPASS** dialog box; however, you cannot make any changes.	Enables all commands under **Configure > Switch**. Allows you to perform configuration upload, download and restore. Enables **COMPASS** command from the **Configure** menu. Allows you to perform all functions in the **Monitor** tab of the **COMPASS** dialog box.
Dashboard Management	Allows you to access the Dashboard Management.	• **Dashboard** tab cannot be viewed. • User cannot publish Performance widgets and Flow widgets. • Disables **Save as Widget** button from the Performance graph and Flow Vision graph.	Allows you to perform the following operations on the dashboard: • Show or hide the default status and Performance widgets to the dashboard. • Customize Network Scope and Time Scope. • Dashboard Playback operation. • Sharing the dashboard. • Dashboard creation and deletion. Restricts you from performing the following operations on the dashboard: • Add or remove custom widgets. • Publish Performance widgets or Flow widgets to the dashboard.	Allows you to perform all read and write operations on the dashboard.
DCB Management	Allows you to configure DCB devices.	Disables the **DCB** command from the **Configure** menu.	Enables the **DCB** command from the **Configure** menu. Disables all commands and functions on the dialog box except the **Close** and **Help**.	Enables the **DCB** command from the **Configure** menu. Enables all commands and functions on the dialog box.
Element Manager	Allows you to access the device Element Manager.	Disables the Element Manager command.	Enables the Element Manager command. Allows you to open the Element Manager; however, disables all functions.	Enables the Element Manager command. Allows you to perform all Element Manager functions.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Element Manager - Product Administration	An Element Manager privilege that enables most functions.	Disables the functions described in the Element Manager User Manual for which you do not have rights. Displays the message, "You do not have rights to perform this action."	Same as No Privilege.	Enables the functions described in the Element Manager User Manual.
E-mail Event Notification Setup	Allows you to define the e-mail server used to send e-mail.	Disables **Event Notification E-mail** command on the **Monitor** menu and the **E-Mail Event Notification Setup** button in the **Users** dialog box. Disables the **E-mail** option in the Master Log shortcut menu. Currently asks, "Are you sure you want to assign Event Management privileges to this group that does not otherwise have read/write for: E-mail Event Notification Setup?".	Enables the **Event Notification E-mail** command on the **Monitor** menu and the **E-mail Event Notification Setup** button in the **Users** dialog box. Allows you to open the **E-Mail Event Notification Setup** dialog box; however, disables the **OK** button.	Enables **Event Notification E-mail** command on the **Monitor** menu and the **E-mail Event Notification Setup** button in the **Users** dialog box. Enables all functions in the **E-Mail Event Notification Setup** dialog box.
Event Management	Allows you to define rules with event triggers and actions.	Disables the **Event Policies** menu item.	Enables access to the **Event Policies** menu item and allows existing rules to be selected and viewed. Disables all action buttons on the tab.	Enables access to the **Event Policies** menu item and enables all functions on the tab.
Fabric Watch	Fabric Watch — Allows you to launch Fabric Watch. Port Fencing — Allows you to configure the function that logs ports out of fabrics automatically if they are misbehaving. Frame Monitor — Allows you to monitor frames. Performance Thresholds — Allows you to configure performance thresholds.	Disables the **Fabric Watch** command from the **Monitor** menu.	Enables the **Fabric Watch** commands from the **Monitor** menu. DIsables the functions on the **Port Fencing** dialog box. DIsables the functions on the **Frame Monitor** dialog box. DIsables the functions on the **Configure Thresholds** dialog box.	Enables the **Fabric Watch** commands from the **Monitor** menu. Enables you to launch Fabric Watch. Enables all functions on the **Port Fencing** dialog box. Enables all functions on the **Frame Monitor** dialog box. Enables the functions on the **Configure Thresholds** dialog box.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Fault Management	Allows you to control access to the **SNMP Trap Registration and Forwarding** dialog box, the **Event Storage** option of the **Options** dialog box, the **Syslog Registration and Forwarding** dialog box, as well as the **Export** and **Clear** functions in the **Event Log** dialog box and the **Show** and **Hide** functions in the **Customize Columns** dialog box.	Disables the **SNMP Trap** and **Syslog configuration** commands from the **Monitor** menu. Disables the **Event Storage** option on the **Options** dialog box. If you do not have other read/write privileges to **Options** dialog box functions, disables the **Server > Options** command. Enables the **Logs >** *<Log_Type>* from the **Monitor** menu.	Enables the **SNMP Trap** and **Syslog configuration**, commands from the **Monitor** menu. Enables the **Event Storage** option on the **Options** dialog box. Enables the **Server > Options** command. Only enables the **Cancel** function for the dialog box boxes. Enables the **Logs >** *<Log_Type>* from the **Monitor** menu.	Enables the **SNMP Trap** and **Syslog configuration**, commands from the **Monitor** menu. Enables the following functions from the dialog box boxes: • Configure Management server registration • Configure TRAP or Syslog forwarding • Register other servers as a recipient • Apply changes Enables the **Server > Options** command. Enables the **Event Storage** option on the **Options** dialog box. Enables the following functions from the dialog box: • Configure max events • Configure event purging policy • Apply changes Enables the following functions from the **Master Log** right-click menu: • Clear events • Show events • Hide events • Export events Note that the **Export** command on the **Master Log** right-click menu also requires the Export privilege because it launches the **Export** dialog box. Enables the **Clear** and **Export** buttons on the individual log dialog box boxes.
FCoE Management	Allows you to configure FCoE devices.	Disables the **FCoE** command from the **Configure** menu.	Enables the **FCoE** command from the **Configure** menu. Disables all commands and functions on the dialog box except the **Close** and **Help**.	Enables the **FCoE** command from the **Configure** menu. Enables all commands and functions on the dialog box.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Firmware Management	Allows you to download firmware to selected switches and manage the firmware repository.	Disables the **Firmware Management** command from the **Configure** menu and right-click menu.	Enables the **Firmware Management** command from the **Configure** menu and right-click menu. Disables all commands and functions on the dialog box except the **Close** and **Help**.	Enables the **Firmware Management** command from the **Configure** menu and right-click menu. Enables all commands and functions on the dialog box.
Host Adapter Management	Allows you to configure a host.	Disables the **Element Manager** command on the right-click menu and the **Element Manager > HCM** command on the **Configure** menu.	Disables the **Element Manager** command on the right-click menu and the **Element Manager > HCM** command on the **Configure** menu.	Enables the **Element Manager** command on the right-click menu and the **Element Manager > HCM** command on the **Configure** menu.
L2 ACL	Allows you to configure a layer 2 access control list.	Disables the **Security > L2 ACL** command on the **Configure** menu.	Enables the **Security > L2 ACL** command on the **Configure** menu. Disables all functions on the dialog box.	Enables the **Security > L2 ACL** command on the **Configure** menu. Enables all functions on the dialog box.
License Update	Allows you to update your license. Allows you to control access to the **License** dialog box from the **Help** menu.	Disables the **License** command on the **Help** menu.	Enables the **License** command on the **Help** menu; however, disables the **Update** and **OK** buttons.	Enables the **License** command on the **Help** menu and enables you to change the license key.
Performance	Allows you to configure the performance subsystem, the display of performance graphs, and threshold settings.	Disables entire **Performance** submenu of the **Monitor** menu as well as the right-click **Performance Graph(s)** command on ports and switch products. Disables the **Port Optics** command on the right-click menu. Disables the **Performance** button in the **DCB Configuration** dialog box.	Enables entire **Performance** submenu off the **Monitor** menu as well as the right-click **Performance Graph(s)** command on ports and switch products. Allows you to open the **Performance Setup** dialog box; however, disables the **OK** button. No changes can be made. Allows you to open the **Performance Graphs** dialog box and enables all controls; however, disables the check boxes under the **Set Thresholds** label on the individual port dialog box (double-click a graph).	Enables entire **Performance** submenu of the **Monitor** menu and the right-click **Performance Graph(s)** command on ports and switch products. Enables changes to the **Performance Setup** dialog box. Allows you to open the **Performance Graphs** dialog box and enables all controls. Enables all functions on the individual port dialog box (double-click a graph). Enables the **Port Optics** command on the right-click menu.
Configuration Policy Manager	Allows you to configure policy manager.	Disables**Configuration Policy Manager** command on the **Monitor** menu.	Enables **Configuration Policy Manager** command on the **Monitor** menu. Allows you to open the **Configuration Policy Manager** dialog box; however, disables the **Add**, **Delete**, and **Run** buttons. No changes can be made. Enables you to use the **Edit**, **Report**, and **History** buttons to view content.	Enables **Configuration Policy Manager** command on the **Monitor** menu. Allows you to open the **Configuration Policy Manager** dialog box and enables all controls.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Properties Edit	Allows you to edit many director and switch properties.	Enables the **Properties** command on **Edit** menu and right-click menus. Disables edit function (removes green triangles) from editable property fields. Disables the **Names** command on the **Configure** menu.	Enables the **Properties** command on **Edit** menu and right-click menus. Disables edit function (removes green triangles) from editable property fields. Enables the **Names** command on the **Configure** menu; however, disables all edit functions in the dialog box.	Enables **Properties** command on **Edit** menu and right-click menus. Enables editable properties (marked by a green triangle) in the Product List and the Properties Sheets. Enables the **Names** command on the **Configure** menu and enables all functions in the dialog box.
Reports	Allows you to generate and view the following reports: • Fabric Ports • Fabric Summary	Disables the **View** command and the **Generate** command on the **Reports** menu. If this privilege is removed and the Event Management privilege is assigned then this message appears: <title: <Product> Message> <Warning>Removing the Report privilege does not remove users' ability to generate reports in Event Management. You might also want to consider removing the Event Management privilege as well.	Enables the **View** command on the **Reports** menu. Disables the **Generate** command on the **Reports** menu.	Enables the **View** command and the **Generate** command on the **Reports** menu.
Security	Allows you to enable and configure SANtegrity features.	Disables the **Security** command from the **Configure > Switch > Replicate** menu. Disables the **Security Log** command on the **Monitor > Logs** menu. Disables the **Security Misc** command from the **Server > Options** menu.	Disables the **Security** command from the **Configure > Switch > Replicate** menu. Enables the **Security Log** command on the **Monitor > Logs** menu. Enables the **Security Misc** command from the **Server > Options** menu; however, disables the functions.	Enables the **Security** command from the **Configure > Switch > Replicate** menu. Enables the **Security Log** command on the **Monitor > Logs** menu. Enables the **Security Misc** command from the **Server > Options** menu. Enables all functions in the dialog box boxes.
Server Backup	Allows you to control the function that copies (backs up) the application data files to another disk.	Disables the **Backup Now** and **Configure** commands on the Backup icon right-click menu on the application status bar. Disables all controls for Backup on the **Options** dialog box.	Disables the **Configure** command on the Backup icon right-click menu on the application status bar. Disables all controls for Backup on the **Options** dialog box.	Enables the **Backup Now** and **Configure** commands on the Backup icon right-click menu on the application status bar. Enables all functions for Backup on the **Options** dialog box.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Server Software Configuration	Allows you to configure some of the properties of the client and server of the management application.	Disables the **Software Configuration Parameters** folder and subpages in the **Options** dialog box. The configuration cannot be viewed.	Enables the **Software Configuration Parameters** folder and subpages in the **Options** dialog box; however, disables the **OK** and **Apply** buttons when any of the subpages are selected.	Enables the **Software Configuration Parameters** folder and subpages in the **Options** dialog box. Enables all functions when any of those subpages are selected.
Setup Tools	Allows you to define and place commands on product icons and in the **Tools** menu.	Disables the **Setup Tools** command on the **Tools** menu. Any existing **Tools** and/or right-click commands already defined or defined by others are available for use; however, you cannot configure new items. If this privilege is removed and the Event Management privilege is assigned then this message appears: <title: <Product> Message> <Warning>Removing the Log Management privilege does not remove users' ability for Setup Tools in Event Management. You might also want to consider removing the Event Management privilege as well.	Enables the **Setup Tools** command on the **Tools** menu; however, disables the **OK** button.	Enables the **Setup Tools** command on the **Tools** menu. Enables all functions in the **Setup Tools** dialog box.
Technical Support Data Collection	Allows you to capture support data from Fabric OS switches.	Disables the **SupportSave**, **Upload Failure Data Capture**, and **View Repository** commands from the **Monitor > Technical Support** menu and right-click menu.	Enables the **View Repository** command from the **Monitor > Technical Support** menu and right-click menu. Disables the **SupportSave** and **Upload Failure Data Capture** commands from the **Monitor > Technical Support** menu and right-click menu.	Enables the **SupportSave**, **Upload Failure Data Capture**, and **View Repository** commands from the **Monitor > Technical Support** menu and right-click menu. Enables all functions on the dialog box boxes.
User Management	Allows you to create and define users and groups, as well as assign privileges and views to groups.	Disables the **Users** command on the main **Server** menu and the **Users** button on the main tool bar.	Enables the **Users** command on the **Server** menu and the **Users** button on the main tool bar; however, disables the **Add**, **Edit**, and **Remove Users**, **Add and Remove Groups**, and **OK** buttons on the **Users** dialog box. Enables the **Edit Groups** button to display the **Group** dialog box (with **OK** button disabled).	Enables the **Users** command on the **Server** menu and the **Users** button on the main tool bar. Enables all functions on the **Users** dialog box and the secondary **Group** dialog box.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Virtual Network Management	Allows you to perform VMM based host discovery and management.	Disables the **VM Manager** command on the **Discover** menu.	Enables the **VM Manager** command on the **Discover** menu. Disables all functions on the dialog box.	Enables the **VM Manager** command on the **Discover** menu. Enables all functions on the dialog box.
VLAN Manager	Allows you to manage VLAN Management	Disables the **VLAN Manager** command.	Enables the **VLAN Manager** command; however, disables functions on the dialog box.	Enables the **VLAN Manager** command and all functions on the dialog box.
Web Services	Allows you to use Web Services API.			
Zoning Activation (Fabric and offline zone database) **NOTE** You must also have the Zoning Offline and Zoning Online privileges to launch the **Zoning** dialog box. **NOTE** You must also have the LSAN privilege to launch the **Activate LSAN Zones** dialog box from the **Zone Database (DB)** tab of the **Zoning** dialog box.	Allows you to activate a zone configuration selected in the **Zoning** dialog box.	Disables the **Activate**, **Deactivate**, and **Zoning Policies** buttons in the **Zoning** dialog box.	Enables the **Zoning Policies** button; however, you cannot perform any operations within the **Zoning** dialog box. Disables the **Activate** and **Deactivate** buttons in the **Zoning** dialog box.	Enables the **Activate**, **Deactivate**, and **Zoning Policies** buttons in the **Zoning** dialog box.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Zoning Online **NOTE** You must also have the Zoning Activation privilege to enable the Activate button. **NOTE** You must also have the Zoning g Offline privilege to enable the **Save As** function in the in the **Zoning** and **Compare/Merge** dialog box boxes.	Allows you to edit any of the fabric zone databases in the available fabrics within the **Zoning** dialog box from the client side and then save to the switch.	In **Zoning** dialog box, the **Zone DB** list includes online and offline zones; however, if an online zone is selected, the contents are not loaded into the **Zoning** dialog box. To launch offline zones you must have the Zoning Offline privilege. Disables all zone database editing and switch pushing functions.	In **Zoning** dialog box, the **Zone DB** list includes online and offline zones. If you select an online zone, the contents are loaded into the **Zoning** dialog box. To launch offline zones you must have the Zoning Offline privilege. Disables all online zone database editing, activation, and persisting functions. In **Zoning** dialog box, enables the **Cancel** and **Help** buttons and the **Compare** and **Export** functions in the **Zone DB Operation** list. On the **Zone DB** tab, enables the find buttons. On the **Active Zone Config** tab, enables the **Zone Member Display** list and **Report** button. In the **Compare/Merge** dialog box, enables the **Cancel** and **Help** buttons. In the **Potential Members** table, enables all functions in the right-click menu. In the **Zones** table, enables the **Port Label**, **Search**, and **Properties** (not editable) functions in the right-click menu. In the **Zone Configs** table, enables the **Properties** (not editable) function in the right-click menu.	Enables all functions on the **Zoning** dialog box.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Zoning Offline **NOTE** You must also have the Zoning Activation privilege to enable the Activate button. **NOTE** You must also have the Zoning g Online privilege to enable the **Save to Switch**, **Activate**, **Deactivate**, and **Rollback** functions in the **Zoning** dialog box and the **Save** function in the **Compare/Merge** dialog box.	Allows you to edit the zone database in offline mode and save the zone database to the repository or to the switch.	In **Zoning** dialog box, the **Zone DB** list includes offline zones; however, if an offline zone is selected, the contents are not loaded into the **Zoning** dialog box. Only displays the Fabric Zone DB (if you have the Zoning Online privilege) in the **Zone DB** list. Disables the **Save As** function from **Zone DB Operation** list for Fabric Zone DBs. Disables the **Save To** function on the **Active Zone Config** tab.	In **Zoning** dialog box, the **Zone DB** list includes offline zones. If you select an offline zone, the contents are loaded into the **Zoning** dialog box. Disables all offline zone DB editing, activating, and persisting functions. In **Zoning** dialog box, enables the **Cancel** and **Help** buttons and the **Compare** and **Export** functions in the **Zone DB Operation** list. On the **Zone DB** tab, enables the find buttons. On the **Active Zone Config** tab, enables the **Zone Member Display** list and **Report** button. In the **Compare/Merge** dialog box, enables the **Cancel** and **Help** buttons. In the **Potential Members** table, enables all functions in the right-click menu. In the **Zones** table, enables the **Port Label**, **Search**, and **Properties** (not editable) functions in the right-click menu. In the **Zone Configs** table, enables the **Properties** (not editable) function in the right-click menu.	Enables all functions on the **Zoning** dialog box.

TABLE 135 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Zoning - LSAN	Allows you to edit and activate LSAN zones for the LSAN fabrics that are available within the **Zoning** dialog box. Prerequisite: Both the backbone fabrics as well as all directly connected edge fabrics must be added to a resource group and a user with LSAN Zoning privilege must be assigned to this specific resource group.	Disables the **Zoning > LSAN Zoning (Device Sharing)** command on the **Configure** menu. In **Zoning** dialog box, the **Zoning Scope** list does not include *LSAN_<FabricName>* as an entry.	Enables the **Zoning > LSAN Zoning (Device Sharing)** command on the **Configure** menu. In **Zoning** dialog box, the **Zoning Scope** list includes LSAN_<FabricName> as an entry, if discovered. If LSAN_<FabricName> is selected, LSAN zone contents are loaded into the **Zoning** dialog box. Disables LSAN zone functions on all dialog box boxes. Disables all online zone database editing, activation, and persisting functions. In **Zoning** dialog box, enables the **Cancel** and **Help** buttons. In the **Potential Members** table, enables all functions in the right-click menu. In the **LSAN Zones** table, enables the **Search** functions in the right-click menu.	Enables all LSAN zone functions on all dialog box boxes.
Zoning - Set Edit Limits	Allows you to set the number of zoning edit operations that can be performed on a fabric zone database before activating a zone configuration.	Disables the **Zoning > Set Edit Limits** command from the **Configure** menu.	Enables the **Zoning > Set Edit Limits** command from the **Configure** menu. Disables all commands and functions on the dialog box except the **Close** and **Help**.	Enables the **Zoning > Set Edit Limits** command from the **Configure** menu. Enables all commands and functions on the dialog box.

TABLE 136 SAN privileges and application behavior

Privilege	Description	No Privilege	Read-Only	Read/Write
SAN - Discovery Setup	Allows you to configure discovery setup.	Disables **Setup** on the **Discover** menu and toolbar.	Enables **Setup** on the **Discover** menu and toolbar. Allows you to open the **Discover Setup** dialog box; however, disables all functions.	Enables **Setup** on the **Discover** menu and toolbar. Enables all functions in the **Discover Setup** dialog box.
SAN - Fabric Binding	Allows you to configure fabric binding.	Disables the **Fabric Binding** command.	Enables the **Fabric Binding** command; however, disables functions on the dialog box.	Enables the **Fabric Binding** command and all functions on the dialog box.
SAN - Fabric Configuration	Allows you to access the **COMPASS** dialog box and create configuration blocks and templates.	Disables **COMPASS** command from the **Configure** menu.	Enables **COMPASS** command from the **Configure** menu. Allows you to access the **Templates** and **Configuration** tabs of the **COMPASS** dialog box; however, you cannot make any changes.	Enables **COMPASS** command from the **Configure** menu. Allows you to perform all functions in the **Templates** and **Configuration** tabs of the **COMPASS** dialog box.
SAN - Fabric Tracking	Allows you to define the current devices and connections present in a fabric as a baseline and to highlight any changes to that baseline.	Disables the **Track Fabric Changes** and **Accept Changes** commands on the **Monitor** menu and right-click menus of **Fabrics**.	Same as no privilege.	Enables the **Track Fabric Changes** and **Accept Changes** commands on the **Monitor** menu and right-click menus of **Fabrics**.
SAN - FCIP Management	Allows you to configure FCIP tunnels and troubleshooting of IP interfaces (IP performance, IP ping and IP trace route).	Disables the **Configure > FCIP Tunnel** and **Monitor > Troubleshooting > FCIP** commands. Disables the **FCIP Tunnel** command on the Fabric right-click menu.	Enables the **Configure > FCIP Tunnel** and **Monitor > Troubleshooting > FCIP** commands. Only enables the **Cancel** function for the dialog box boxes.	Enables the **Configure > FCIP Tunnel** and **Monitor > Troubleshooting > FCIP** commands. Enables all commands and functions on the associated dialog box boxes. Also enables all commands on the **FCIP Tunnels** tab in the device's **Properties** dialog box.

TABLE 136 SAN privileges and application behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
SAN - FICON Management	Allows you to configure Cascade FICON Fabric and Cascade FICON Fabric Merge. Also allows you to configure block ports and allow/prohibit matrix on active configuration or any offline configurations.	Disables the **Configure Fabric**, **Merge Fabrics** commands on the **Configure > FICON** menu. Disables the **Allow/Prohibit Matrix** command from the **Configure** menu and right-click menu.	Disables the **Configure Fabric**, **Merge Fabrics** commands on the **Configure > FICON** menu. Enables the **Allow/Prohibit Matrix** command from the **Configure** menu and right-click menu. Disables all commands and functions on the **Configure Allow/Prohibit Matrix** dialog box except the **Close** and **Help**.	Enables the **Configure Fabric**, **Merge Fabrics** commands on the **Configure > FICON** menu. Enables the **Allow/Prohibit Matrix** command from the **Configure** menu and right-click menu. Enables all commands and functions on the associated dialog box boxes.
SAN - High Integrity Fabric	Allows you to set Fabric Binding and Insistent Domain IDs.	Disables the **High Integrity Fabric** command from the **Configure** menu.	Enables the **High Integrity Fabric** command from the **Configure** menu. Disables all commands and functions on the dialog box except the **Cancel** and **Help**.	Enables the **High Integrity Fabric** command from the **Configure** menu. Disables all commands and functions on the dialog box.
SAN - Logical Switch Configuration	Allows you to create a new logical switch, assign and remove ports from a logical switch, delete a logical switch, configure a logical fabric, and change the fabric ID of a logical switch. You must be assigned to the 'All Fabrics' resource group to access Logical Switch Configuration feature.	Disables the **Logical Switches** command from the **Configure** menu.	Enables the **Logical Switches** command from the **Configure** menu. Disables all functions from the dialog box except view. Also requires access to All Resources resource group to access the **Logical Switches** dialog box.	Enables the **Logical Switches** command from the **Configure** menu. Enables all commands and functions on the dialog box. Also requires access to All Resources resource group to access the **Logical Switches** dialog box.
SAN - Main Display - MAPS	Allows you to configure MAPS.	Disables MAPS.	Enables the MAPS configuration from the **Configure** menu. Allows to view MAPS violations, **Out of Range Violations** and **Port Health Violations** widgets on the Dashboard, and access MAPS-specific widgets, .	Enables the MAPS configuration from the **Configure** menu. Enables all commands and functions on the dialog box.
SAN - Port Connectivity	Allows you to view all of the port details and connected devices.	Disables the **Port Connectivity** command from the **Monitor** menu and right-click menu.	Enables the **Port Connectivity** command from the **Monitor** menu and right-click menu.	Enables the **Port Connectivity** command from the **Monitor** menu and right-click menu.

TABLE 136 SAN privileges and application behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
SAN - Port Mapping - Host	Allows you to identify all the HBAs that are in the same server.	Disables the **Host Port Mapping** command from the **Discover** menu. Disables the **Server** right-click command on HBAs.	Enables **Host Port Mapping** command from the **Discover** menu and right-click menu; however, disables the **Create**, **Delete**, and **OK** buttons.	Enables **Host Port Mapping** command from the **Discover** menu and right-click menu. Enables all functions in the **Servers** dialog box.
SAN - Port Mapping - Storage	Allows you to construct multi-port storage systems out of individual storage ports.	Disables the **Storage Port Mapping** command from **Discover** menu and right-click menus for Storage products and ports in the tree and map.	Enables the **Storage Port Mapping** command from **Discover** menu right-click menus for Storage products and ports in the tree and map. Allows you to open the **Storage Port Mapping** dialog box; however, disables the **Create**, **Delete**, right and left arrow, and **OK** buttons.	Enables the **Storage Port Mapping** command from **Discover** menu and right-click menus for Storage products and ports in the tree and map. Enables all functions on the **Storage Port Mapping** dialog box.
SAN - Properties - Add/Delete Columns	Allows you to define new properties as well as remove them.	Disables the **Add**, **Edit** and **Delete** buttons on the **Create View** dialog box **Columns** tab. Disables the **Add Column**, **Edit Column**, and **Delete Column** commands on the right-click menu of the **Product List** column headers. Disables the **Add**, **Edit**, and **Delete** commands on the property headers in property sheets.	Same as No Privilege.	Enables the **Add**, **Edit**, and **Delete** properties commands and buttons in the **Create View** and **Edit View** dialog box boxes, the **Product List** column header right-click menu, and the Property Sheet property header right-click menu.
SAN - Routing Configuration	Allows you to configure Routing and domain IDs of phantom switches.	Disables the **Routing Configuration** and **Routing Domain IDs** commands from the **Configure** menu and right-click menu.	Disables the **Routing Configuration** and **Routing Domain IDs** commands from the **Configure** menu and right-click menu.	Enables the **Routing Configuration** and **Routing Domain IDs** commands from the **Configure** menu and right-click menu. Enables all functions in the dialog box boxes.
SAN - SCOM Management	Allows you to manage the SCOM plug-in.	Disables the **Plug-in for SCOM** command from the **Tools** menu.	Disables the **Plug-in for SCOM** command from the **Tools** menu.	Enables the **Plug-in for SCOM** command from the **Tools** menu. Enables all functions in the dialog box boxes.

TABLE 136 SAN privileges and application behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
SAN - SMIA Operations	Allows you to access the CIMOM (Common Information Model Object Manager) server and the SMIA Configuration Tool.	Disables the **Configure SMI Agent** button from the Server Console. Disables the SMIA Configuration Tool Java web start application.	Enables the **Configure SMI Agent** button from the Server Console. Enables the SMIA Configuration Tool Java web start application. However, disables all functions in the dialog box.	Enables the **Configure SMI Agent** button from the Server Console. Enables the SMIA Configuration Tool Java web start application. Enables all functions in the dialog box.
SAN - Storage Encryption Configuration	Allows you to configure storage encryption configuration, including selecting storage devices and LUNs, viewing and editing switch, group, or engine properties, viewing and editing storage device encryption properties, and initiating manual LUN re-keying.	Disables the **Encryption** command from the **Configure** menu.	Enables the **Encryption** command from the **Configure** menu. Disables all functions from the dialog box except view.	Enables the **Encryption** command from the **Configure** menu. Enables the following functions from the dialog box: • Viewing and editing switch, group, or engine properties • Viewing and editing storage device encryption properties • Selecting storage devices and LUNs • Initiating manual LUN re-keying. Disables all other functions from the **Configure Encryption** dialog box.

TABLE 136 SAN privileges and application behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
SAN - Storage Encryption Key Operation	Allows you to configure storage encryption key operation, including selecting storage devices and LUNs, viewing switch, group, or engine properties, viewing storage device encryption properties, initiating manual LUN re-keying, enabling and disabling an engine, zeroizing an engine, restoring a Master Key, and all smart card operations.	Disables the **Encryption** command from the **Configure** menu.	Enables the **Encryption** command from the **Configure** menu. Disables all functions from the dialog box except view.	Enables the **Encryption** command from the **Configure** menu. Enables the following functions from the dialog box: • Viewing switch, group, or engine properties • Viewing storage device encryption properties • Selecting storage devices and LUNs • Initiating manual LUN re-keying. • Enabling and disabling an engine • Zeroizing an engine • Restoring a Master Key • All smart card operations Disables all other functions from the **Configure Encryption** dialog box.

TABLE 136 SAN privileges and application behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
SAN - Storage Encryption Security	Allows you to configure storage encryption security, including creating a new encryption group, adding a switch to an existing group, zeroizing an encryption engine, backing up or restoring a master key, and enabling encryption functions after a power cycle.	Disables all functions from the dialog box except view. The **Encryption** command from the **Configure** menu is enabled and disabled by the Storage Encryption Configuration privilege.	Disables all functions from the dialog box except view. The **Encryption** command from the **Configure** menu is enabled and disabled by the Storage Encryption Configuration privilege.	Enables the **Encryption** command from the **Configure** menu. Enables the following functions from the dialog box: • Ceating a new encryption group • Ading a switch to an existing group • Zeroizing an encryption engine • Backing up or restoring a master key • Enabling encryption functions after a power cycle • Changing key vaults for an encryption group. • Create/edit/delete High Availability (HA) Clusters. • Removing switches from encryption groups. • Enable/disable encryption engines. • Create new master keys (backup and restore of master keys is already listed)
SAN - Troubleshooting	Allows you to run device connectivity check, fabric device sharing check and trace route.	Disables the **Device**, **Fabric Device Sharing**, **Connectivity and Trace Route** commands under **Monitor > Troubleshooting > FC** . Disables the **Configuration Wizard** command under the **Configure** menu.	Disables the **Device Connectivity**, **Fabric Device Sharing**, and **Trace Route** commands under **Monitor > Troubleshooting > FC** .	Enables the **Device Connectivity**, **Fabric Device Sharing**, and **Trace Route** commands under **Monitor > Troubleshooting > FC** . Enables all functions in the dialog box boxes.

TABLE 136 SAN privileges and application behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
SAN - View Management	Allows you to create, edit, and delete views. Selecting from views should always be allowed unless restricted by the assignment of Views in the Group definition in the **Users** dialog box.	Disables the **Create View**, **Copy View**, **Edit View**, **Delete View**, and **Connectivity View** commands in the **View > Manage View** menu and the first tab header on the main desktop. Allows you to select an assigned view but not create or change. Disables the **Create View Automatically** command in the shortcut menu.	Enables the **Create View** and **Edit View** commands in the **View > Manage View** menu and the first tab header on the main desktop; however, disables the **OK** button in the **Create View** and **Edit View** dialog box boxes. Disables the **Copy View**, **Delete View**, and **Connectivity View > Create** and **Refresh** commands. Allows you to select an assigned view but not create or change.	Activates all view commands in the **View > Manage View** menu and the first tab header on the main desktop. Enables all functions in the dialog box boxes.

Roles and Access Levels

The Management application provides preconfigured roles (SAN System Administrator, IP System Administrator, Security Administrator, Zone Administrator, Security Officer, Operator, and Network Administrator); however, the SAN System Administrator can also create roles manually (refer to "Creating a new role" on page 144 for instructions.)

TABLE 137 Application features and role access levels

Feature	Roles with Read/Write Access	Roles with Read-Only Access
Active Session Management	SAN System Administrator, Security Officer	Operator
Call Home	SAN System Administrator, Operator	
Certificate Management	SAN System AdministratorNetwork Administrator, Host Administrator, Security Administrator	Operator
Configuration Management	SAN System AdministratorNetwork Administrator	Operator
DCB Management	SAN System Administrator, Network Administrator	Security Administrator, Security Officer
E-mail Event Notification Setup	SAN System Administrator, Operator	
Element Manager	SAN System Administrator,	
Element Manager - Product Administration	SAN System Administrator,	
Event Management	SAN System Administrator, Network Administrator	Operator
Fabric Watch	SAN System Administrator,	
Fault Management	SAN System AdministratorNetwork Administrator	Operator

TABLE 137 Application features and role access levels (Continued)

Feature	Roles with Read/Write Access	Roles with Read-Only Access
FCoE Management	SAN System Administrator, Network Administrator	Security Administrator, Zone Administrator, Security Officer, Operator
Firmware Management	SAN System AdministratorNetwork Administrator	Operator
Host Adapter Management	SAN System Administrator, Security Officer, Host Administrator	Operator
L2 ACL	SAN System Administrator, Security Administrator	
License Update	SAN System Administrator	Operator
MAPS	SAN System Administrator, , Network Administrator	Operator
Performance	SAN System Administrator, Host Administrator, Network Administrator	Operator
Properties Edit	SAN System Administrator, , Host Administrator	Operator
Reports	SAN System AdministratorNetwork Administrator	Operator
Security	SAN System Administrator, , Security Administrator, Security Officer, Host Administrator	Operator
Server Backup	SAN System Administrator, Product Administrator, Operator	
Server Software Configuration	SAN System Administrator	Operator
Setup Tools	SAN System Administrator	Operator
Technical Support Data Collection	SAN System Administrator	Operator
User Management	SAN System Administrator, Security Officer	Operator
Virtual Network Management	SAN System Administrator	Operator
VLAN Manager	SAN System Administrator	Operator
Web Services	SAN System Administrator	Operator
Zoning - LSAN	SAN System Administrator, Zone Administrator	Operator
Zoning Set Edit Limits	SAN System Administrator	Zone Administrator, Operator
Zoning Activation	SAN System Administrator, Zone Administrator	Operator
Zoning Offline	SAN System Administrator, Zone Administrator	Operator
Zoning Online	SAN System Administrator, Zone Administrator	Operator

TABLE 138 SAN features and role access levels

Feature	Roles with Read/Write Access	Roles with Read-Only Access
SAN- Discovery Setup	SAN System Administrator, Host Administrator	Operator
SAN - Element Manager	SAN System Administrator,	
SAN - Element Manager - Product Operation	SAN System Administrator, Operator	
SAN- Fabric Binding	SAN System Administrator, Security Administrator, Security Officer	Operator
SAN- Fabric Tracking	SAN System Administrator	Operator
SAN- FCIP Management	SAN System Administrator	Operator
SAN- FICON Management	SAN System Administrator	Operator
SAN- High Integrity Fabric	SAN System Administrator, Security Administrator, Security Officer	Operator
SAN- Logical Switch Configuration	SAN System Administrator	
SAN - Main Display - MAPS	IP System Administrator, Network Administrator	Operator
SAN- Port Connectivity	SAN System Administrator	
SAN- Port Mapping - Host	SAN System Administrator, Security Officer, Host Administrator	Operator
SAN- Port Mapping - Storage	SAN System Administrator	Operator
SAN- Properties - Add/Delete Columns	SAN System Administrator, Host Administrator	Operator
SAN- Routing Configuration	SAN System Administrator	Operator
SAN- SCOM Management	SAN System Administrator	
SAN- SMIA Operations	SAN System Administrator	Operator
SAN- Storage Encryption Configuration	SAN System Administrator, Security	Operator
SAN- Storage Encryption Key Operations	SAN System Administrator, Security Administrator, Security Officer	
SAN- Storage Encryption Security	SAN System Administrator, Security Administrator	Operator
SAN- Troubleshooting	SAN System Administrator	
SAN- View Management	SAN System Administrator, Security Administrator, Zone Administrator, Network Administrator, Security Officer, Operator, Host Administrator	

Device Properties

SAN device properties

You can customize the device and fabric **Properties** dialog boxes to display only the data you need by creating user-defined property labels. You can also edit property fields to change information.

Viewing Fabric properties

To view the properties for a fabric, complete the following step.

1. Right-click any fabric and select **Properties**.

   The *Fabric_Name* **Properties** dialog box displays, with information related to the selected fabric.

   To add user-defined property labels, refer to "Adding a property field" on page 1372.

   Fields containing a green triangle ( ) in the lower right corner are editable.

   TABLE 139 Fabric properties

Field/Component	Description
Name	The name specified through the switch Element Manager.
FID Fabric Name	Enter a name for the fabric (up to 128 characters). Supported on seed switches running Fabric OS 7.0 or later.
Seed Switch	The IP address of the seed switch.
AD Enabled	Whether admin domain is enabled on the switch or not.
Status	The operational status.
Switch and AG Count	The number of switches and Access Gateway's in the fabric.
Description	A description of the customer site.
Principal Switch	The IP address of the principal switch.
Active Zone Configuration	Whether active zone configuration is activated on the fabric.
Last Discovery	The date and time of last discovery.
Tracked	Whether the fabric is tracked.
Location	The customer site location.
Contact	The primary contact at the customer site.
Add button	Click to add a user-defined property. For more information, refer to "Adding a property field" on page 1372.
Edit button	Click to edit a user-defined property. For more information, refer to "Editing a property field" on page 1372.
Delete button	Click to delete a user-defined property. For more information, refer to "Deleting a property field" on page 1373.

2. Click **OK** on the *Fabric_Name* **Properties** dialog box to close.

Viewing SAN device properties

To view the properties for a device, complete the following steps.

1. Right-click any product icon and select **Properties**.

   The **Properties** dialog box displays, with information related to the selected device (such as switches, directors, HBAs, trunks, tunnels, and nodes).

   To add user-defined property labels, refer to "Adding a property field" on page 1372.

   Fields containing a green triangle ( ) in the lower-right corner are editable.

   NOTE
   Depending on the device type, some of the properties listed in the following table may not be available for all products.

TABLE 140 Device properties

Field/Component	Description
Addressing Mode	The addressing mode of the switch.
Back to Edge Routing	Whether back to edge routing is supported.
Bandwidth	The bandwidth of the FCIP tunnel.
Capability	The node capability.
Compression	Whether compression is on or off for the FCIP tunnel.
Connected Virtual FCoE Port	The fabric name, switch name, and virtual FCoE port number of the connected virtual FCoE port.
Contact	The primary contact at the customer site.
Contributors	The device contributors.
Device Type	Whether the device is an initiator or target.
Description	A description of the customer site.
Destination IP Address	The IP address of the of the FCIP tunnel destination device.
Discovery Status	The discovery status of the switch. Examples include "Discovered: Seed Switch" and "Discovered: Not Reachable".
Domain ID	The device's domain ID, which is the top-level addressing hierarchy of the domain.
Fabric	The fabric name.
Fabric Name	The name specified through the device Element Manager.
Fabric Watch	Whether Fabric Watch is up or down.
Factory Serial Number	The factory serial number.
Fastwrite	Whether fastwrite is on or off for the FCIP tunnel.
FC Port	The FC port of the FCIP tunnel.
FCoE Capable	Whether the device is Fibre Channel over Ethernet capable.
FCS Role	Whether FCS is supported.
Routing Policy	The routing policy configured on the switch.
Firmware	The firmware version.
GigE Port	The GigE port of the FCIP tunnel.
Host Name	The Host name.

TABLE 140 Device properties (Continued)

Field/Component	Description
HIF Mode	For 7.3.0 or later switches, whether HIF mode is enabled or disabled. Not applicable for 7.2.0 or earlier switches.
IKE Policy #	The IKE policy number. Also includes the following information: • Authentication Algorithm • Encryption Algorithm • Diffie-Hellman • SA Life
IP Address	The device's IP address.
IPSec Policy #	The IPSec policy number. Also includes the following information: • Authentication Algorithm • Encryption Algorithm • SA Life
L2 Capable	Whether the device is Layer 2 capable.
L3 Capable	Whether the device is Layer 3 capable.
L2 Mode	The Layer 2 mode. Options include Access, Converged, or Trunk.
LAG ID	The link aggregation group identifier.
Last Discovery	The date and time of the last discovery.
Location	The customer site location.
MAC address	In a network, the Media Access Control (MAC) address is a unique number that identifies a specific hardware interface. It is a 12-digit hexadecimal number.
Managed By	The management program used to manage the fabric.
Master Port	The master port of the trunk.
Member Ports	The member ports of the trunk.
Model	The model number of the device.
Name	The user-defined name of the switch.
Node Name	The name of the node.
Node WWN	The world wide name of the node.
Physical/Logical	Whether the device is a physical device or a logical device.
Port Count	The number of ports.
Port Type	The port type.
Preshared key configured	Whether the preshared key is configured for the FCIP tunnel.
Reason	The device status.
Remote Switch Name	The remote switch name of the trunk.
Remote Switch IP	The remote switch IP address of the trunk.
Remote Switch WWN	The remote switch world wide name of the trunk.
Remote Slot #	The remote slot number of the trunk.
Remote Master Port	The remote master port of the trunk.
Remote Member Ports	The remote member ports of the trunk.
Sequence number	The sequence number of the switch.
Serial #	The hardware serial number.

TABLE 140 Device properties (Continued)

Field/Component	Description
Slot #	The slot number of the trunk.
Source IP Address	The IP address of the of the FCIP tunnel source device.
Speed (Gb/s)	The speed of the port in gigabits per second.
State	The device's state, for example, online or offline.
Status	The operational status.
Switch Name	The switch name.
Switch IP	The switch IP address.
Switch WWN	The switch world wide name.
Tape Pipelining	Whether tape pipelining is on or off for the FCIP tunnel.
Tunnel ID	The tunnel identifier.
Type	The device type.
Unit Type	The unit type of the node.
Vendor	The product vendor.
# Virtual FCoE port count	The number of virtual FCoE ports on the device. There is a one-to-one mapping of TE ports to virtual FCoE ports. Therefore, the number of virtual session ports is one for directly connected devices.
VLAN #	The VLAN number of the FCIP tunnel.
VLAN Class of Service for Control Connection	The VLAN class of service for the control connection of the FCIP tunnel.
VLAN Class of Service for Data Connection	The VLAN class of service for the data connection of the FCIP tunnel.
VLAN ID	The VLAN identification number.
WWN	The world wide name of the device.
Add button	Click to add a user-defined property. For more information, refer to "Adding a property field" on page 1372.
Edit button	Click to edit a user-defined property. For more information, refer to "Editing a property field" on page 1372.
Delete button	Click to delete a user-defined property. For more information, refer to "Deleting a property field" on page 1373.

2. To view port properties, select one of the following tabs.

   The following port types are available depending on the selected device:

   - FC Ports
   - GigE Ports
   - IP Ports
   - iSCSI Ports
   - Virtual Sessions Ports
   - Virtual FCoE Ports
   - Virtual Machine Ports

3. If you selected the **FC Ports** tab, select the port type:

   - FC

- ICL
- GigE

For a description of the port properties, refer to "Port properties" on page 1361.

4. Click **OK** on the **Properties** dialog box to close the dialog box.

Viewing storage properties

The **Storage Properties** dialog box displays information related to a selected storage device. To view the properties for a storage device, complete the following steps.

1. Select a storage icon.
2. Select **Edit > Properties**.

   The **Properties** dialog box displays.
3. Click the **Storage** tab.

   NOTE
   Some fields may not be available for all products.

TABLE 141 Storage Properties

Field	Description
(**Status**)	Lists two kinds of data: the LUN health and the state of the LUN disks. The colored icon in the lower-left corner indicates the LUN health. In most cases, there is also a number that represents the RAID type. The possible RAID types are 0, 1, 5, or 10, and the number does not display if the RAID type is different from those. The following are examples of generic LUN status: **Normal.** All disks are operating normally and online. **Transitioning.** One or more disks are in a transitioning state. For example, rebuilding or binding. **Faulted/Offline.** One or mores disks is offline or faulted. **Unknown.** Status is not available.
Array	A group of disks designated by the user to be managed by the RAID 5 technique.
Assigned LUNs (Count)	All LUNs assigned (masked) to host ports that currently exist on this storage device.
Assigned LUNs (Size GB)	The total amount of storage space carved into LUNs and assigned (masked) to host ports on the storage device.
Block Size (B)	The size of the individual blocks on the disk, in bytes.
Device Adapter	(IBM ESS products only) Displays one of eight ESS product adapters deployed in pairs, one for each cluster that provides communication between the cluster and storage products.
Disks	The number of disks across which this LUN is striped.
Free LUNs (Count)	All LUNs not assigned (masked) to any host ports (available) that currently exist on this storage device.
Free LUNs (Size GB)	The total amount of storage space carved into LUNs but not assigned (masked) to host ports on the storage device, in gigabytes.
Free Space (Count)	The number of contiguous free space instances not yet carved into LUNs (available to be carved) on the storage device. Typically, there is one free space for each disk group on a storage device.
Free Space (Size GB)	The total amount of storage space not carved into a LUN (available for new LUNs) on the storage device, in gigabytes.
Hosts Assigned	The number of hosts to which this LUN has been assigned.

TABLE 141 Storage Properties (Continued)

Field	Description
Host Spares	The number of disks assigned as host spares in addition to the disks that make up the LUN.
Label	A user-specified label. The default value is the name of the label as specified in the storage product.
Loop	(IBM ESS products only) The physical connection between a pair of product adapters in the ESS product.
LSS ID	Specifies the logical subsystem of an IBM ESS product.
LUN Name	The name of the LUN.
LUN Status	The LUN status (online or offline).
Management Link	The management link status (Up or Down) of the product.
Model #	The model number of the product.
Name (in-band)	The name of the in-band product.
Operational Status	The operational status of the product.
OS Type	The operating system.
Protocol	The LUN protocol.
Size (GB)	The total size of this LUN's storage, in gigabytes.
State	The state of the LUN.
Storage LUN ID	The storage product's LUN ID number for this LUN.
Storage Ports	The total number of storage ports assigned to the server or the port, or bound to the LUN.
Type	The level or type of RAID storage. Possible values are as follows: **0** — Striped disk array without fault tolerance. **1** — Mirroring and duplexing. **2** — Hamming code ECC. **3** — Parallel transfer with parity. **4** — Independent data disks with shared parity disk. **5** — Independent data disks with distributed parity blocks. **6** — Independent data disks with two independent distributed parity schemes. **7** — Optimized asynchrony for high I/O rates as well as high data transfer rates. **10** — Very high reliability combined with high performance. **53** — High I/O rates and data transfer performance. **0+1** — High data transfer performance.
Total (Count)	All LUNs, whether assigned or not, that currently exist on this storage device.
Total (Size GB)	The total amount of storage space on the storage device, in gigabytes.
Unique LUN ID	Identifies the unique LUN identifier.
Volume State	The volume state of the LUN.
Add button	Click to add a user-defined property. For more information, refer to "Adding a property field" on page 1372.
Edit button	Click to edit a user-defined property. For more information, refer to "Editing a property field" on page 1372.
Delete button	Click to delete a user-defined property. For more information, refer to "Deleting a property field" on page 1373.

4. Click **OK** on the **Properties** dialog box to close the dialog box.

Viewing iSCSI properties

The **iSCSI Properties** dialog box displays information related to iSCSI. To view the properties for an iSCSI device, complete the following steps.

1. Right-click a product icon and select **Properties**.

   The **Properties** dialog box displays.

2. Select the **iSCSI** tab.

   NOTE
   Some fields may not be available for all products.

   TABLE 142 iSCSI Properties

Field	Description
Agent	The Caffeine agent version number.
Applications	The applications.
Assigned LUNs	The number of unique LUNs (not LUN paths) masked to this host.
Assigned LUNs Size (GB)	The total size of the unique LUNs (not LUN paths) in gigabytes.
Command Descriptor Block Count	The number of command descriptor blocks on the product.
Comments	Comments regarding the product.
Contact	A contact for the product.
Department	The department.
Description	A description of the product.
Device Type	The product type.
Digest Error Count	The number of digest errors on the product.
Driver	The iSCSI driver.
Driver Version	The iSCSI driver version.
Firmware	The firmware for the product.
Group	The name of the portal group.
Initiator Type	The type of initiator (such as HBA or Software).
Interface	The name of the interface.
IP Address	The product's IP address.
iSCSI Alias	The name of the alias target.
iSCSI Node Name	The node name of the product.
iSCSI Node Type	The node type of the product.
iSCSI Service	The service status; for example, running or not running.
iSNS IP Address	The IP address of the server to which the product is pointed.
ISNS IP Address	A list of the iSNS IP addresses this product has been assigned by the user to query.
iSNS Service	Whether the product is registered with an iSNS server.
Location	The location of the product.
Management Link	The management link status (Up or Down) of the product.
Name (Product)	The name of the product.

TABLE 142 iSCSI Properties (Continued)

Field	Description
OS	The name of the operating system running on the product.
OS Build	The operating system build running on the product.
OS Release	The operating system release running on the product.
Portal Addresses	The list of IP addresses.
Port	The port number.
Protocol Error Count	The number of protocol errors.
Tag	The group tag ID of the portal.
Sessions button	Select to display the **Filer Sessions** dialog box for the product.
Statistics button	Select to display the **Filer iSCSI Statistics** dialog box for the product.
Storage Arrays	The number of arrays containing LUNs masked to the server.
Storage Logins	The number of unique filers to which hosts on this server are logged in.
Target Portals table	Target portals of the product.
Total LUN Size (GB)	The size in gigabytes (GB) of all unique LUNs (not LUN paths) masked to the product.
Vendor	The vendor of the product.

3. Click **OK** on the **Properties** dialog box to close the dialog box.

Viewing port properties

The following port types are available depending on the device:

- FC Ports
- GigE Ports
- IP Ports
- iSCSI Ports

  NOTE
  iSCSI ports that have an FC Address of all zeros are inactive. All others are active.

- Virtual Sessions Ports
- Virtual FCoE Ports

To view a port's properties, right-click a port and select **Properties**, or double-click the port.

The **Port Properties** dialog box displays (Figure 582).

To add user-defined property labels, refer to "Adding a property field" on page 1372.

Fields containing a green triangle ( ) in the lower-right corner are editable.

FIGURE 582Port Properties dialog box

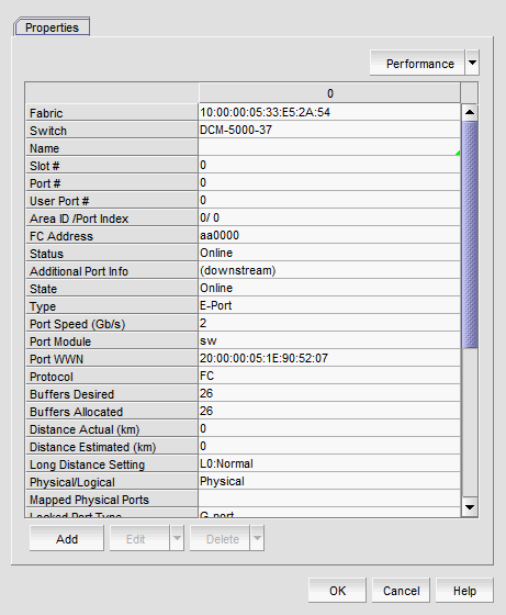

NOTE
Depending on the port type, some of the following properties may not be available for all products.

TABLE 143 Port properties

Field	Description
Additional Port Info	Additional error information relating to the selected port.
Address	The address of the port.
Addressing Mode	The addressing mode of the switch.
Active FC4 Types	The active FC4 types.
Active Tunnels	The number of active tunnels.
Area ID (hex)/Port Index (hex)	The area identifier, in hexadecimal, of the switch-to-product connection.
Associated GE Port	The port number of the associated GE port.
Attached Port #	The port number of the attached product.
Back to Edge Routing Support	Whether back to edge routing is supported.
Bandwidth	The bandwidth of the FCIP tunnel.
Blocked	The configuration of the switch (blocked or unblocked).
Bottleneck Status	Whether the port is bottlenecked or not.
Buffers Desired	The number of buffers desired but not allocated.
Buffers Allocated	The number of buffers allocated.
Capability	The node capability.
Class	The class of the port.
Class of Service	The class of service.
Compression	Whether compression is enabled or disabled.

TABLE 143 Port properties (Continued)

Field	Description
Connected Devices	The number of connected devices. Click the icon in the right side of the field to open the **Virtual FCoE Port <Number> Connected Devices** dialog box.
Connected Switch	The name of the connected switch.
Delete button	Click to delete the ports.
Description	A description of the customer site.
Destination IP Address	The IP address of the of the FCIP tunnel destination device.
Device Type	Whether the device is an initiator or target.
Discovery Status	The discovery status of the switch. Examples include “Discovered: Seed Switch” and “Discovered: Not Reachable”.
Distance Actual (km)	The actual distance (in km) for end port connectivity.
Distance Estimated (km)	The estimated distance (in km) for end port connectivity.
Domain ID	The device's domain ID, which is the top-level addressing hierarchy of the domain.
Encryption	Whether encryption is enabled or disabled.
Fabric	The fabric's IP address.
Fabric Name	The name of the fabric.
Fabric Watch	Whether Fabric Watch is up or down.
Fastwrite	Whether fastwrite is on or off for the FCIP tunnel.
FC Port	The FC port of the FCIP tunnel.
FC Port Count	The number of FC ports on the device.
FCIP Capable	Whether the port is FCIP capable.
FCoE Capable	Whether the device is Fibre Channel over Ethernet capable.
FCS Role	Whether FCS is supported.
Flag (FICON related)	Whether a flag is on or off.
Firmware	The firmware version.
Forward Error Correction (FEC)	Whether FEC is enabled or disabled.
GigE Port	The GigE port of the FCIP tunnel.
GigE Port Count	The number of GigE ports on the device.
Host Name	The Host name.
IKE Policy #	The IKE policy number. Also includes the following information: • Authentication Algorithm • Encryption Algorithm • Diffie-Hellman • SA Life
Inband Management Status	The inband management status (online or offline).
Index	The index of the Virtual FCoE Port.
Interface Count	The interface count.
IP Address	The device's IP address.

TABLE 143 Port properties (Continued)

Field	Description
IPSec Policy #	The IPSec policy number. Also includes the following information: • Authentication Algorithm • Encryption Algorithm • SA Life
iSCSI button	Click to launch the Element Manager.
iSCSI Capable	Whether the port is iSCSI capable or not.
L2 Capable	Whether the device is Layer 2 capable.
L3 Capable	Whether the device is Layer 3 capable.
L2 Mode	The Layer 2 mode. Options include Access, Converged, or Trunk.
LAG ID	The link aggregation group identifier.
Last Discovery	The date and time of the last discovery.
Location	The customer site location.
Locked Port Type	The port type of the locked product.
Long Distance Setting	Whether the connection is considered to be normal or longer distance.
MAC Address	The Media Access Control address assigned to a network adapters or network interface cards (NICs).
Managed By	The management program used to manage the fabric.
Manufacturer Plant	The name of the manufacturer plant.
Master Port	The master port of the trunk.
Member Ports	The member ports of the trunk.
Model	The model number of the device.
Modify button	Click to launch the Element Manager.
Name	The name of the port (up to 128 characters). This field is editable.
Node Name	The name of the node.
Node WWN	The world wide name of the node.
Performance list	Select to launch the dialog box of one of the following performance options: • Real Time Graph • HIstorical Graph • HIstorical Report
Physical/Logical	Whether the port is a physical port or a logical port.
Port #	The number of the port.
Port Address	The address of the port.
Port Count	The number of ports.
Port ID	The identifier of the port.
Port Module	The port's module.
Port NPIV	The number of NPIV ports.
Port Speed (Gb/s)	The port speed, in gigabits per second.
Port State	The port state (online or offline).
Port Status	The port's operational status (online or offline).
Port Type	The port type.

TABLE 143 Port properties (Continued)

Field	Description
Port WWN	The port's world wide name.
Preshared key configured	Whether the preshared key is configured for the FCIP tunnel.
Prohibited	Whether the port is prohibited.
Protocol	The network protocol, for example, Fibre Channel.
Reason	The device status.
Remote Switch Name	The remote switch name of the trunk.
Remote Switch IP	The remote switch IP address of the trunk.
Remote Switch WWN	The remote switch world wide name of the trunk.
Remote Slot #	The remote slot number of the trunk.
Remote Master Port	The remote master port of the trunk.
Remote Member Ports	The remote member ports of the trunk.
Routing Policy	The routing policy configured on the switch.
Sequence number	The sequence number of the switch.
Serial #	The hardware serial number.
Slot #	The location (slot) of the port.
Source IP Address	The IP address of the of the FCIP tunnel source device.
Speed (Gb/s)	The port speed, in gigabits per second.
State	The port state (online or offline).
Status	The port's operational status (online or offline).
Switch Name	The switch name.
Switch IP	The switch IP address.
Switch WWN	The switch world wide name.
Symbolic Name	The symbolic name of the port.
Tag	The tag number of the port.
Tape Pipelining	Whether tape pipelining is on or off for the FCIP tunnel.
Troubleshooting list	Select to launch the dialog box of one of the following troubleshooting options: • IP Ping • IP Traceroute • IP Performance
Tunnel Count	The number of tunnels.
Tunnel ID	The tunnel identifier.
Type	The type of port, for example, U_port.
Unit Type	The unit type of the node.
User Port #	The number of the user port.
Vendor	The product vendor.
# Virtual FCoE port count	The number of virtual FCoE ports on the device. There is a one-to-one mapping of TE ports to virtual FCoE ports. Therefore, the number of virtual session ports is one for directly connected devices.
# Virtual Session Ports	The number of virtual session ports associated with the GE port.

TABLE 143 Port properties (Continued)

Field	Description
VLAN #	The VLAN number of the FCIP tunnel.
VLAN Class of Service for Control Connection	The VLAN class of service for the control connection of the FCIP tunnel.
VLAN Class of Service for Data Connection	The VLAN class of service for the data connection of the FCIP tunnel.
VLAN ID	The VLAN identification number.
WWN	The world wide name of the device.
QSFP Unit Number	The QSFP unit number to which the port belongs to.
Add button	Click to add a user-defined property. For more information, refer to "Adding a property field" on page 1372.
Edit button	Click to edit a user-defined property. For more information, refer to "Editing a property field" on page 1372.
Delete button	Click to delete a user-defined property. For more information, refer to "Deleting a property field" on page 1373.

Viewing VC module properties

To view Virtual Connect (VC) module properties, complete the following steps.

1. Right-click a VC module and select **Properties**.
2. Review the properties for the device.

TABLE 144 Properties tab

Field	Description
Fabric	The name of the fabric.
Name	(Fabric OS modules only) The name of the device.
WwnName	The world wide name of the device.
IP Address	(Fabric OS modules only) The IP address of the device.
Status	The operational status.
Type	The device type - Virtual Connect.
Port Count	The number of ports.
Product Name	The product name.
Serial #	The hardware serial number.
VC Firmware	The downloaded firmware version of the VC Ethernet management module and all VC FC modules managed by the VC Domain.
VC Domain Name	The domain name.
VC Domain Group	The domain group.
IO Bay	
Discovery Status	The discovery status of the VCEM server of this module.
Last Discovery	The last time data collection was performed for this VC module on the VCEM server.

TABLE 145 Port tab

Field	Description
Fabric	The name of the fabric.
Switch	The name of the VC module.
Port #	The port number.
Type	The port type.
Status	The status of the port. For example, LOGGED-IN or NOT-LOGGED-IN.
Port Speed (Gb/s)	The speed of the port in gigabits per second.
Port WWN	The world wide name of the port.
Physical/Logical	Whether the port is Physical or Logical.
NPIV Enabled	Whether the port is NPIV enabled or not.
Connected Switch	The name of the switch connected to the port.

TABLE 146 NPIV WWNs tab

Field	Description
NPIV Port WWN	The world wide name of the NPIV port.
NPIV Node WWN	The world wide name of the NPIV node.
Name	The user-defined name of the NPIV WWN. This is an editable field.
Uplink Port Number	The port number of the uplink.
Uplink Port WWN	The port world wide name of the uplink.
Server Profile	The server profile.
Server Bay	The server bay number.
Virtual Serial Number	The serial number.

3. Click **Close** to close the **Properties** dialog box.

Host properties

You can view device and port properties from the Product List or the map.

You can customize the Host **Properties** dialog boxes by creating user-defined property labels (refer to "Adding a property field" on page 1372).

NOTE
You cannot create user-defined property labels at the adapter level.

You can also edit property fields to change information. Fields containing a green triangle ( ) in the lower right corner are editable.

Viewing adapter port properties

To view adapter port properties, complete the following steps.

1. Right-click an HBA icon and select **Show Ports**.
2. Right-click the port and select **Properties**, or double-click the port.

Fields containing a green triangle ( ) in the lower right corner are editable.

The ***HBA_Port*** **Properties** dialog box displays. Table 147 details the properties of the selected port.

TABLE 147 Adapter port properties

Field	Description
Port Attributes	
Port #	The port number: 0 or 1.
Name	The name that is manually assigned to the port.
Zone Alias	The alternate name of the zone.
Symbolic Name	The symbolic name (nickname) for the HBA port.
HCM Name	The version of the Host Connectivity Manager (HCM) application.
Associated VMs	Virtual machines associated with the HBA port.
Port WWN	The port's world wide name.
Node WWN	The node's (parent device) world wide name.
Factory Port WWN	The world wide name assigned at the factory for the HBA port.
Factory Node WWN	The world wide name assigned at the factory for the HBA.
Media	The type of media; for example, 8G-sw (8 Gbps software).
Product Type	The device port type; for example, N_Port.
Vendor	The port's vendor.
Type	The port type; for example, N_Port.
FC Address	The port's Fibre Channel address.
Attached Port #	The port number of the attached product.
Active FC4 Types	The active FC4 types; for example, SCSI or IP.
Class of Service	The class of the port; for example, Class-2 or Class-3.
Switch	The name of the switch.
Fabric	The name of the Fabric.
VM Port Name	The port name of the virtual machine associated with the host.
Preboot Created	Indicates whether preboot was created on the virtual port.
PCI Function Index	The PCI function number associated with the physical port.
Fabric Assigned Address	The state (enabled or disabled) of the fabric assigned address for the adapter.
WWN Source	The source of the world wide name. Options include: Fabric — The WWN is assigned from the fabric. The fabric assigned address must be enabled. Factory — The WWN is assigned at the factory.
Hyper-V Virtual FC	Indicates if the port is a Hyper-V virtual FC port. For non-virtual ports, the field displays as N/A. **NOTE:** This property is applicable only to Windows server's version 2012 and later. It is a special type of NPIV port that can be presented directly to VMs in a Hyper-V environment.
Configuration	
Configured State	Indicates whether the port is enabled or disabled.
Max Bandwidth	The maximum allowable bandwidth output for the selected port.
Operating State	Indicates whether the port is online or offline.
Configured Speed	The configured port speed.

TABLE 147 Adapter port properties **(Continued)**

Field	Description
Operating Speed	The speed at which the port is operating.
Max Speed Supported	The maximum speed that is supported on the port. For the FC port, the maximum speed is 8 Gbps.
Configured Topology	The configured topology setting: auto, point-to-point, or loop.
Operating Topology	The operating topology setting: auto, point-to-point, or loop.
Boot over SAN	Indicates whether boot over SAN is enabled.
Receive BB Credits	The number of buffer credits received.
Transmit BB Credits	The number of buffer credits transmitted.
Frame Field Size	The frame size, in bytes, of the port.
Hardware Path	The hardware path of the HBA.
Virtual Port Count	The number of virtual ports associated with the HBA.
FEC State	The state of FEC (Forward Error Correction) is an error recovery mechanism.
BB Credit Recovery	The status of Buffer to Buffer Credit Recovery.
Configured BBSCN Count	The count of configured buffe state change number.
Negotiated BBSCN Count	The count of negotiated buffer state change number.
Operating State	Displays details about the state of the following operating parameters: • Beacon State • Link Beacon State • MPIO Mode State • Path Time Out • Logging Level • Target Rate Limit • Default Rate Limit
FC-SP	
Authentication	Indicates whether FC-SP authentication is enabled or disabled.
FCSP Status	Whether FC-SP authentication is being used.
Algorithm	The configured authentication algorithm.
Group	The DH group, which is DH-null (group 0), which is the only option.
Error Status	The health status of the Fibre Channel Security Protocol parameters.
QoS	
Configured QoS State	Indicates whether QoS is enabled or disabled.
Operating QoS State	Indicates whether QoS is on or off.
Total BB Credit	The total number of buffer credits.
Priority Levels	Lists the available priorities (High, Medium, Low).
QoS Percentage	The value represents the bandwidth in percentage for each of the priorities (high, medium, and low) and the three values must equal 100 percent.
IO Execution Throttle	
IO Execution Throttle Max Value	Indicates the maximum value of IO Execution Throttle. Maximum value is 2000.
IO Execution Throttle Operational Value	Indicates the operational value of IO Execution Throttle. Operational value ranges from 0 through 2000.

TABLE 147 Adapter port properties **(Continued)**

Field	Description
IO Execution Throttle Configured Value	Indicates the configured value of IO Execution Throttle. Configuration value ranges from 0 through 2000.
VM	Displays details about the VM of the following operating parameters: • VM Name • State • Status • GUID • Path • Memory Assigned • Hard Drives Count • Processor Count • Uptime • Notes
Add button	Click to add a user-defined property. For more information, refer to "Adding a property field".
Edit button	Click to edit a user-defined property. For more information, refer to "Adding a property field".
Delete button	Click to delete a user-defined property. For more information, refer to "Adding a property field".

3. Click **OK** to close.

Viewing Virtual Machine properties

The host must be running a supported version of hypervisor for the **Virtual Machines** tab to display in the **Host Properties** dialog box as well as the **HBA Virtual Machine** and **HBA Port Properties** dialog boxes. On the **HBA Virtual Machine** and **HBA Port Properties** dialog boxes, the list of virtual machines displayed is filtered to show only virtual machines using the selected HBA or HBA port. If there are no virtual machines present for the HBA or HBA port, the **Virtual Machines** tab is present but without values.

The top table of the **Virtual Machines** tab contains a standard properties table with one column for each virtual machine configured on the host. None of these properties are editable.

The bottom table lists the SAN resources (disks and tapes) assigned to the virtual machine that is currently selected in the top table. The bottom table contents change depending on the virtual machine selection. If there are multiple paths to the SAN resource, the SAN resource is listed in multiple columns, once for each path. None of the properties in the second table are editable.

1. Right-click the VM and select Properties.

Field/Component	Description
Virtual Machine Count	The number of virtual machines.
Virtual Machines table (top)	The standard properties table. **Name — The name of the virtual machine.** **Property UUID — The Universal Unique Identifier of the virtual machine.** **Instance UUID — The Instance UUID/Entity ID of the virtual machine.** **Application Name — The Application name, which is running in the virtual machine. You can map the VM name and enter the Application name only when the vCenter is discovered.** **Network Address — The Network address of the virtual machine.** **IP Address — The virtual machine's IP address.** **VM Hypervisor — The Virtual Machine Manager. (VMM)** OS Type — The type of operating system. **Status** — The status of the virtual machines; for example, Running or Stopped. **Time Started** — The time the virtual machine was started. **vCPU Count** — The number of virtual CPUs. **CPU Resources** — The CPU resources on the virtual machine (for example, MHz reserved). **VM Memory Size** — The memory size of the virtual machine. **Datastore Name** — The datastore name of the virtual machine. **Datastore Location** — The hypervisor name of the device containing the datastore. **Description** — A description of the virtual machine. **vNICs** — The virtual NIC world wide names. **Provisional Storage** — The maximum storage capacity that the VM can use. The maximum is not always guaranteed, however, because some part of the storage may be kept as shared (across multiple VMs). If other VMs require the storage, the shared storage size will reduce. **Not-shared Storage** — The storage that only the selected VM can access and no other VM can use. This is the minimum guaranteed storage available to the VM. **Used Storage** — The current storage capacity being used by the VM. This is a more dynamic property and is expected to change on a more regular basis.

Field/Component	Description
Data Path Count for the Selected Virtual Machine	The number of data paths for the selected virtual machine.
Data Path Count for the Selected Virtual Machine table (bottom)	The SAN resources (disks and tapes) assigned to the virtual machine that is currently selected in the top table. **Device Name** — The end device when there are multiple paths to the same device. **Storage Type** — The type of storage. **Storage Status** — The status of the storage (for example, Online or Offline). **Model** — The model of the storage device. **Serial #** — The serial number of the storage device. **Capacity** — The capacity of the storage device. **Path Name** — The hardware path to the virtual machine. **Fabric** — The fabric that contains this path. **Path Policy** — The policy attribute describes how the server distributes traffic over multiple paths to the target. For example, **Fixed** (uses a specified path unless it is down), **MRU** (uses the path used last, unless it is done), and **Round Robin** (alternate requests on both paths). **Path Status** — The status of the path (for example, Enabled, Active, Preferred or Disabled, Inactive, Not Preferred). **HBA Node WWN** — For an iSCSI storage path, this field displays the iSCSI name for the adapter. If the adapter is not FC or iSCSI, this field is blank. **HBA Port WWN** — For an iSCSI storage path, this field displays the iSCSI name for the adapter port. If the adapter is not FC or iSCSI, this field is blank. **Virtual Node WWN** — The world wide name of the virtual node. **Virtual Port WWN** — The world wide name of the virtual port. **Target Node** — For an iSCSI storage path, this field displays the iSCSI name for the target. If the adapter is not FC or iSCSI, this field is blank. **Target Port** — For an iSCSI storage path, this field displays the iSCSI name for the target port. If the adapter is not FC or iSCSI, this field is blank. **iSCSI Target Address** — The IP address of the iSCSI target. If the adapter is not iSCSI, this field is blank. **iSCSI Target Port** — The TCP port of the storage device. If not specified, the default value of 3260 is used. If the adapter is not iSCSI, this field is blank. **NAS Remote Host** — The host that runs the NFS server. If the storage is not network-attached (NAS), this field is blank. **NAS Remote Path** — The remote path of the NFS and CIFS mount point. If the storage is not network-attached (NAS), this field is blank.

Properties customization

NOTE
Properties customization requires read and write permissions to the Properties - Add / Delete Columns privilege.

You can customize the product and fabric **Properties** dialog boxes by creating user-defined fabric, product, and port properties. You can also edit or delete user-defined properties, as needed.

You can create up to three user-defined property labels from the **Properties** dialog box for each of the following object types: fabric, product, and port properties. Product and fabric property labels created from the **Properties** dialog box display in the Product List and the **Properties** dialog box. Port property labels created from the **Properties** dialog box display in the Product List and the **Properties** dialog box. User-defined properties must be unique across all **Properties** dialog boxes and the Product List.

Property fields containing a green triangle ( ) in the lower right corner are editable. To edit a field with a green triangle ( ), click in the field and make your changes.

Adding a property field

You can add up to three new user-defined properties to the fabric **Properties** dialog box as well as the **Properties** and **Ports** tabs of the device **Properties** dialog box.

To add a user-defined property, complete the following steps.

1. Right-click any product icon and select **Properties**.

   The **Properties** dialog box displays.

2. Select the tab to which you want to add a property, if necessary.
3. Click **Add**.

   The **Add Property** dialog box displays.

4. Enter a label and description for the property.

   The label must be unique and can be up to 30 characters.

   The description can be up to 126 characters.

5. Select **Fabric**, **Port**, or **Property** from the **Type** list, if available.
6. Click **OK**.

   The new property displays in the properties list as well as the Product List. To edit the user-defined property field, click in the field and make your changes.

Editing a property field

NOTE
Properties customization requires read and write permissions to the Properties - Add / Delete Columns privilege.

You can edit any property that you create on the **Properties** dialog box.

Fields containing a green triangle ( ) in the lower right corner are editable. To edit a field with a green triangle ( ), click in the field and make your changes.

To edit a user-defined property, complete the following steps.

1. Right-click any product icon and select **Properties**.

   The **Properties** dialog box displays.

2. Select the tab on which you want to edit a property, if necessary.
3. Click **Edit** > *Property_Label*.

   The **Edit Property** dialog box displays.

4. Change the label and description for the property, as needed.

   The label must be unique and can be up to 30 characters.

   The description can be up to 126 characters.

5. Select **Fabric**, **Port**, or **Property** from the **Type** list, if available.
6. Click **OK**.

Deleting a property field

NOTE
Properties customization requires read and write permissions to the Properties - Add / Delete Columns privilege.

You can delete any user-defined property from the **Properties** dialog box. To delete a user-defined property, complete the following steps.

1. Right-click any product icon and select **Properties**.

   The **Properties** dialog box displays.
2. Select the tab on which you want to delete a user-defined property, if necessary.
3. Click **Delete** > *Property_Label* (where *Property_Label* is the user-defined property you want to delete).
4. Click **Yes** on the confirmation message.

   The property you selected is deleted.

Editing a property field directly

You can edit fields containing a green triangle ( ) in the lower right corner. To edit a field, complete the following steps.

1. Right-click any product icon and select **Properties**.

   The **Properties** dialog box displays.
2. Select the tab on which you want to edit a field.

   Fields containing a green triangle ( ) in the lower right corner are editable.
3. Click in an editable field and change the information.
4. Click **OK**.

Regular Expressions

This appendix presents a summary of Unicode regular expression constructs that you can use in the Management application.

TABLE 148 Characters

Construct	Matches
x	The character x
\\	The backslash character
\0n	The character with octal value 0n (0 <= n <= 7)
\0nn	The character with octal value 0nn (0 <= n <= 7)
\0mnn	The character with octal value 0mnn (0 <= m <= 3, 0 <= n <= 7)
\xhh	The character with hexadecimal value 0xhh
\uhhhh	The character with hexadecimal value 0xhhhh
\t	The tab character ('\u0009')
\n	The newline (line feed) character ('\u000A')
\r	The carriage-return character ('\u000D')
\f	The form-feed character ('\u000C')
\a	The alert (bell) character ('\u0007')
\e	The escape character ('\u001B')
\cx	The control character corresponding to x

TABLE 149 Character classes

Construct	Matches
[abc]	a, b, or c (simple class)
[^abc]	Any character except a, b, or c (negation)
[a-zA-Z]	a through z or A through Z, inclusive (range)

TABLE 149 Character classes

Construct	Matches
[a-d[m-p]]	a through d, or m through p: [a-dm-p] (union)
[a-z&&[def]]	d, e, or f (intersection)
[a-z&&[^bc]]	a through z, except for b and c: [ad-z] (subtraction)
[a-z&&[^m-p]]	a through z, and not m through p: [a-lq-z](subtraction)

TABLE 150 Predefined character classes

Construct	Matches
.	Any character (may or may not match line terminators)
\d	A digit: [0-9]
\D	A non-digit: [^0-9]
\s	A whitespace character: [\t\n\x0B\f\r]
\S	A non-whitespace character: [^\s]
\w	A word character: [a-zA-Z_0-9]
\W	A non-word character: [^\w]

TABLE 151 POSIX character classes (US-ASCII only)

Construct	Matches
\p{Lower}	A lower-case alphabetic character: [a-z]
\p{Upper}	An upper-case alphabetic character:[A-Z]
\p{ASCII}	All ASCII:[\x00-\x7F]
\p{Alpha}	An alphabetic character:[\p{Lower}\p{Upper}]
\p{Digit}	A decimal digit: [0-9]
\p{Alnum}	An alphanumeric character:[\p{Alpha}\p{Digit}]
\p{Punct}	Punctuation: One of !"#$%&'()*+,-./:;<=>?@[\]^_`{|}~
\p{Graph}	A visible character: [\p{Alnum}\p{Punct}]
\p{Print}	A printable character: [\p{Graph}\x]
\p{Blank}	A space or a tab: [\t]
\p{Cntrl}	A control character: [\x00-\x1F\x7F]
\p{XDigit}	A hexadecimal digit: [0-9a-fA-F]
\p{Space}	A whitespace character: [\t\n\x0B\f\r]

TABLE 152 java.lang.Character classes (simple java character type)

Construct	Matches
\p{javaLowerCase}	Equivalent to java.lang.Character.isLowerCase()
\p{javaUpperCase}	Equivalent to java.lang.Character.isUpperCase()

TABLE 152 java.lang.Character classes (simple java character type)

Construct	Matches
\p{javaWhitespace}	Equivalent to java.lang.Character.isWhitespace()
\p{javaMirrored}	Equivalent to java.lang.Character.isMirrored()

TABLE 153 Classes for Unicode blocks and categories

Construct	Matches
\p{InGreek}	A character in the Greek block (simple block)
\p{Lu}	An uppercase letter (simple category)
\p{Sc}	A currency symbol
\P{InGreek}	Any character except one in the Greek block (negation)
[\p{L}&&[^\p{Lu}]]	Any letter except an uppercase letter (subtraction)

TABLE 154 Boundary matches

Construct	Matches
^	The beginning of a line
$	The end of a line
\b	A word boundary
\B	A non-word boundary
\A	The beginning of the input
\G	The end of the previous match
\Z	The end of the input but for the final terminator, if any
\z	The end of the input

TABLE 155 Greedy quantifiers

Construct	Matches
X?	X, once or not at all
X*	X, zero or more times
X+	X, one or more times
X{n}	X, exactly n times
X{n,}	X, at least n times
X{n,m}	X, at least n but not more than m times

TABLE 156 Reluctant quantifiers

Construct	Matches
X??	X, once or not at all
X*?	X, zero or more times
X+?	X, one or more times
X{n}?	X, exactly n times
X{n,}?	X, at least n times
X{n,m}?	X, at least n but not more than m times

TABLE 157 Possessive quantifiers

Construct	Matches
X?+	X, once or not at all
X*+	X, zero or more times
X++	X, one or more times
X{n}+	X, exactly n times
X{n,}+	X, at least n times
X{n,m}+	X, at least n but not more than m times

TABLE 158 Logical operators

Construct	Matches
XY	X followed by Y
X|Y	Either X or Y
(X)	X, as a capturing group

TABLE 159 Back references

Construct	Matches
\n	Whatever the nth capturing group matched
Quotation	
\	Nothing, but quotes the following character
\Q	Nothing, but quotes all characters until \E
\E	Nothing, but ends quoting started by \Q

TABLE 160 Special constructs (non-capturing)

Construct	Matches
(?:X)	X, as a non-capturing group
(?idmsux-idmsux)	Nothing, but turns match flags on–off
(?idmsux-idmsux:X)	X, as a non-capturing group with the given flags on–off
(?=X)	X, through zero-width positive lookahead
(?!X)	X, through zero-width negative lookahead
(?<=X)	X, through zero-width positive lookbehind
(?<!X)	X, through zero-width negative lookbehind
(?>X)	X, as an independent, non-capturing group

Troubleshooting

Application Configuration Wizard troubleshooting

The following section states a possible issue and the recommended solution for Management application Configuration Wizard errors.

Problem	Resolution
Unable to launch the Management application Configuration Wizard on a Windows Vista, Windows 7,or Windows 2008 R2 system	The Windows Vista, Windows 7,or Windows 2008 R2 system enables the User Access Control (UAC) option by default. When the UAC option is enabled, the Configuration Wizard cannot launch. If the Configuration Wizard does not launch, use one of the following options to disable the UAC option: The following are the various ways we can disable UAC in vista: **Disable using msconfig by completing the following steps.** 1 Select **Start > Run**. 2 Type msconfig on the **Run** dialog box and click **OK**. 3 Click the **Tools** tab on the **System Configuration Utility**. 4 Scroll down to and select the **Disable UAC** tool name. 5 Click **Launch**. A command window displays and runs the disable UAC command. When the command is complete, close the window. 6 Close the **System Configuration Utility**. 7 Restart the computer to apply changes. **NOTE:** You can re-enable UAC using the above procedure and selecting the **Enable UAC** tool name in step 4. **Disable using regedit by completing the following steps.** **NOTE:** Before making changes to the registry, make sure you have a valid backup. In cases where you're supposed to delete or modify keys or values from the registry it is possible to first export that key or value(s) to a .REG file before performing the changes. 1 Select **Start > Run**. 2 Type regedit on the **Run** dialog box and click **OK**. 3 Navigate to the following registry key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies\System 4 Right-click the **EnableLUA** value and select Modify. 5 Change the **Value data** field to 0 on the **Edit DWORD Value** dialog box and click **OK**. 6 Close the **Registry Editor**. 7 Restart the computer to apply changes. **NOTE:** You can re-enable UAC using the above procedure and changing the **Value data** field to 1 in step 5.

Browser troubleshooting

The following section states a possible issue and the recommended solution for browser errors.

Problem	Resolution
The **Cancel** button does not work on the **Report via E-mail** dialog box when you use the Mozilla Firefox browser.	Mozilla Firefox Browser does not support window close script. Click the browser Close button to cancel. **NOTE:** The **Cancel** button still displays on all **Report via E-mail** dialog boxes.

Client browser troubleshooting

The following section states a possible issue and the recommended solution for client browser errors.

Problem	Resolution
Downloading Client from a Internet Explorer Browser over HTTPS	If the JNLP file does not launch automatically, use one of the following options: • Complete the following steps. 1 Save the JNLP file to the local host. 2 Launch the JNLP file manually. • In Internet Explorer 7, complete the following steps. 1 Select **Tools > Internet Options**. 2 Click the **Advanced** tab. 3 Clear the **Do not save encrypted pages to disk** check box. If the browser warns you about the security certificate, use the fully qualified hostname to launch the web page.

Discovery troubleshooting

The following section states a possible issue and the recommended solution for discovery errors.

Problem	Resolution
After upgrading to Management application 12.x from 11.x, unable to discover Fabric OS devices. Error "4002" displays when trying to discover Fabric OS devices.	The Management application 12.x uses Java 1.7, which disables the use of certificates with "weak authentication". You must update your Java certificates to resolve the issue.

Fabric tracking troubleshooting

The following section states a possible issue and the recommended solution for fabric tracking errors.

Problem	Resolution
If a switch is replaced by another switch having the same IP address but a different node WWN while fabric tracking is on, the Management application does not update the Product List, Connectivity Map, or switch properties with the new node WWN.	Choose from one of the following options: • Turn fabric tracking off while the switch is replaced. This causes the old switch to be removed and the new switch added. • After the switch is replaced, remove and re-add the fabric in the **Discover Setup** dialog box.

FICON troubleshooting

The following section states a possible issue and the possible cause for FICON errors.

Problem	Causes
FICON not supported on switch error.	FICON Unsupported Configurations: • FICON is not supported on base switches. • FICON is not supported on a logical switch which has an XISL configured. • FICON is not supported if the PID format is 2. • FICON is not supported if 10 bit address is enabled on 8-slot Backbone Chassis for non-default switch. • FICON is not supported if any port address is greater than the maximum port number of the switch. • 48-port blades are not allowed in the Director Chassis for FICON. • FICON is not supported on 48-port blades in the 8-slot Backbone Chassis when Virtual Fabrics is disabled. However, when Virtual Fabrics is enabled in the Backbone Chassis, FICON is supported on the 48-port blade as long as the 48-port blade is part of a logical switch. If the 48-port blade is part of the default switch on the Backbone Chassis, FICON is not supported. • FICON is not supported on Admin Domain-enabled fabrics. • FICON is not supported on 64-port blades.

Firmware download troubleshooting

The following section states a possible issue and the recommended solution for firmware download errors.

Problem	Resolution
If you configured an internal FTP server and the Management application server is running IPv6, firmware download is not supported.	Choose from one of the following options: • If the Management application is running IPv6 only, configure an external FTP server. • If the Management application is running IPv4 and IPv6, configure IPv4 to be the preferred address.
Firmware download using SCP/SFTP does not work because of one of the following issues: • For internal SCP/SFTP server, the application was uninstalled and reinstalled without migration • For external SCP/SFTP server, the SSH handshake keypair is changed - manually - due to an external server reinstall - due to the SCP/SFTP server preference (Options dialog box) being changed from built-in to external (installed on same machine) or vice versa	Clear the SSH (SCP/SFTP) server IP address or hostname from the known_hosts table of the device. • For Fabric OS devices, use the following command: `sw0:FID128:admin>` **sshutil delknownhost** `IP Address/Hostname to be deleted:` ***SSH_server_IP_address*** where ***SSH_server_IP_address*** is the IP address of the SSH server you want to delete. • For Network OS devices running firmware version 3.0 and later, use the following command: `sw0#` **clear ssh-key** ***SSH_server_IP_address*** where ***SSH_server_IP_address*** is the IP address of the SSH server you want to delete. • For Network OS devices running firmware version 2.1.1b, use the following command: `sw0#` **execute-script sshdeleteknownhost** `IP Address/Hostname to be deleted:` ***SSH_server_IP_address*** where ***SSH_server_IP_address*** is the IP address of the SSH server you want to delete.
Firmware download using an external SCP server does not work after reinstalling the application or the external SCP server on the Host machine.	Clear the SCP server IP address or hostname from the known_hosts table of the device using the following command: `sw0# FID10:root> ssh-keygen –R` ***Host_Name*** where ***Host_Name*** is the IP address or host name of the external SCP server.

Launch Client troubleshooting

The following section states a possible issue and the recommended solution if you are unable to launch the remote client.

Problem	Resolution
Remote client does not upgrade from versions prior to 11.0.	The remote client does not automatically upgrade when you select the remote client shortcut of client versions earlier than 11.0. To clear the old client and launch the new remote client version, complete the following steps. 1 Clear the previous version from the Java cache,. a Select **Start > Settings > Control Panel > Java**. The **Java Control Panel** dialog box displays. b Click **View** on the **General** tab. The **Java Cache Viewer** dialog box displays. c Right-click the application and select **Delete**. d Click **Close** on the **Java Cache Viewer** dialog box. e Click **OK** on the **Java Control Panel** dialog box. 2 Launch the remote client. a Open a web browser and enter the IP address of the Management application server in the **Address** bar. If the web server port number does not use the default (443 if is SSL Enabled; otherwise, the default is 80), you must enter the web server port number in addition to the IP address. For example, *IP_Address*:*Port_Number*. The Management application web start screen displays. b Click the Management application web start link. The **Log In** dialog box displays. c Enter your user name and password. The defaults are **Administrator** and **password**, respectively. **NOTE:** Do not enter *Domain**User_Name* in the **User ID** field for LDAP server authentication. d Select or clear the **Save password** check box to choose whether you want the application to remember your password the next time you log in. e Click **Login**. f Click **OK** on the **Login Banner** dialog box. The Management application displays. **NOTE:** When you launch the Management application or navigate to a new view, the **SAN** tab displays with a gray screen over the Product List and Topology Map while data is loading.

Problem	Resolution
Unable to log into the Client (the application does not launch when you use a valid user name and password and exceptions are thrown in the client side).	Use one the following procedures to configure the IP address in the host file. **Windows operating systems** 1 Log in using the 'Administrator' privilege. 2 Select **Start > Run**. 3 Type drivers in the **Open** field and press **Enter**. 4 Go to the 'etc' folder and open the 'hosts' file using a text editor. 5 Add the IP address and host name of the client in the following format: *IP_Address Host_Name*. For example, 127.0.0.1 localhost 6 Save and exit the file. **Unix operating systems** 1 Log in using the 'root' privilege. 2 Open the '/etc/hosts' file using a text editor. 3 Add the IP address and host name of the client in the following format: *IP_Address Host_Name*. For example, 127.0.0.1 localhost 4 Save and exit the file.
Unable to launch the remote client (the SSL setting, web server port number, or server starting point number changed during the server upgrade).	To remove the old link and launch the correct remote client version, complete the following steps. 1 Clear the previous version from the Java cache,. a Select **Start > Settings > Control Panel > Java**. The **Java Control Panel** dialog box displays. b Click **View** on the **General** tab. The **Java Cache Viewer** dialog box displays. c Right-click the application and select **Delete**. d Click **Close** on the **Java Cache Viewer** dialog box. e Click **OK** on the **Java Control Panel** dialog box. 2 Log into the remote client from the browser. a Open a web browser and enter the IP address of the Management application server in the **Address** bar. If the web server port number does not use the default (443 if is SSL Enabled; otherwise, the default is 80), you must enter the web server port number in addition to the IP address. For example, *IP_Address*:*Web_Server_Port_Number*. The Management application web start screen displays. b Click the Management application web start link. The **Log In** dialog box displays. c Enter your user name and password. The defaults are **Administrator** and **password**, respectively. **NOTE:** Do not enter *Domain**User_Name* in the **User ID** field for LDAP server authentication. d Select or clear the **Save password** check box to choose whether you want the application to remember your password the next time you log in. e Click **Login**. f Click **OK** on the **Login Banner** dialog box. The Management application displays. **NOTE:** When you launch the Management application or navigate to a new view, the **SAN** tab displays with a gray screen over the Product List and Topology Map while data is loading.

Names troubleshooting

The following section states a possible issue and the recommended solution for names errors.

Problem	Resolution
Duplicate name error.	If you configured the Management application to only allow unique names and you try to use a name that already exists in the fabric. You can enter a different name for the device or search for the duplicate name using one of the following procedures: • "Searching for a device by name" on page 98 in the **Configure Names** dialog box • "Searching for a device by WWN" on page 99 in the **Configure Names** dialog box • "Searching for a device" on page 314

Patch troubleshooting

The following section states a possible issue and the recommended solution for patch errors.

Problem	Resolution
Unable to launch the SMC on a Windows Vista,Windows 7,or Windows 2008 R2 system	The Windows Vista,Windows 7,or Windows 2008 R2 system enables the User Access Control (UAC) option by default. When the UAC option is enabled, the SMC cannot launch. If the SMC does not launch, use one of the following options to disable the UAC option: The following are the various ways we can disable UAC in vista: **Disable using msconfig by completing the following steps.** 1 Select **Start > Run**. 2 Type msconfig on the **Run** dialog box and click **OK**. 3 Click the **Tools** tab on the **System Configuration Utility**. 4 Scroll down to and select the **Disable UAC** tool name. 5 Click **Launch**. A command window displays and runs the disable UAC command. When the command is complete, close the window. 6 Close the **System Configuration Utility**. 7 Restart the computer to apply changes. **NOTE:** You can re-enable UAC using the above procedure and selecting the **Enable UAC** tool name in step 4. **Disable using regedit by completing the following steps.** **NOTE:** Before making changes to the registry, make sure you have a valid backup. In cases where you're supposed to delete or modify keys or values from the registry it is possible to first export that key or value(s) to a .REG file before performing the changes. 1 Select **Start > Run**. 2 Type regedit on the **Run** dialog box and click **OK**. 3 Navigate to the following registry key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies\System 4 Right-click the **EnableLUA** value and select Modify. 5 Change the **Value data** field to 0 on the **Edit DWORD Value** dialog box and click **OK**. 6 Close the **Registry Editor**. 7 Restart the computer to apply changes. **NOTE:** You can re-enable UAC using the above procedure and changing the **Value data** field to 1 in step 5.

Performance troubleshooting

The following section states a possible issue and the recommended solution for Performance errors.

TABLE 161

Problem	Resolution

An error message with the following text displays:
Real Time statistics collection has failed. Please see master log for details.
	Make sure that the following prerequisites for Performance Monitoring Data collection are met.
	1 To collect performance statistics for any protocol type (FC/FCIP/FCOE/GE), the snmp access control list must have an empty list or the Management server IP must be included in the access control list.
	For example, data collection occurs in the following cases.
	Case 1: Default access control list is empty
	`FCRRouter:admin> snmpconfig --show accesscontrol`
	`SNMP access list configuration:`
	`Entry 0:  No access host configured yet`
	`Entry 1:  No access host configured yet`
	`Entry 2:  No access host configured yet`
	`Entry 3:  No access host configured yet`
	`Entry 4:  No access host configured yet`
	`Entry 5:  No access host configured yet`
	
	Case 2: Management Server IP included in access control list
	`FCRRouter:admin> snmpconfig --show accesscontrol`
	`SNMP access list configuration:`
	`Entry 0:  Access host subnet area 172.26.1.86 (rw)`
	`Entry 1:  No access host configured yet`
	`Entry 2:  No access host configured yet`
	`Entry 3:  No access host configured yet`
	`Entry 4:  No access host configured yet`
	`Entry 5:  No access host configured yet`
	Verification and Troubleshooting.
	To add the server IP address to the access control list, use the following command from the switch CLI:
	`FCRRouter:admin> snmpconfig --set accesscontrol`
	To set the default access control, use the following command from the switch CLI:
	`FCRRouter:admin> snmpconfig --default accesscontrol`

Problem	Resolution

An error message with the following text displays:
Real Time statistics collection has failed. Please see master log for details.
	2 To collect data, the SNMP credentials in the Management application and switch must match. SNMP v1 or v3: The community strings entered in the **Address Properties** dialog box - **SNMP** tab must match the one entered in the switch.
	If you enter 'test' as the SNMP v1 community string in the Management application, then the community string in the switch must be 'test' as well.
	To view the switch SNMP value, use one of the following commands from the switch CLI:
	`HCLSwitch:admin> snmpconfig --show snmpv1`
	`HCLSwitch:admin> snmpconfig --show snmpv3`
	To set the switch SNMP value, use one of the following commands from the switch CLI:
	`HCLSwitch:admin> snmpconfig --set snmpv1`
	`HCLSwitch:admin> snmpconfig --set snmpv3`
	Example
	`HCLSwitch:admin> snmpconfig --set snmpv1`
	`SNMP community and trap recipient configuration:`
	`Community (rw): [test]`
	`Trap Recipient's IP address : [172.26.1.183]`
	`Trap recipient Severity level : (0..5) [4]`
	`Trap recipient Port : (0..65535) [162]`
	`Community (rw): [OrigEquipMfr]`
	`Trap Recipient's IP address : [172.26.24.26]`
	`Trap recipient Severity level : (0..5) [4]`
	`Trap recipient Port : (0..65535) [162]`
	`Community (rw): [custom]`
	`Trap Recipient's IP address : [172.26.1.158]`
	`Trap recipient Severity level : (0..5) [4]`
	`Trap recipient Port : (0..65535) [162]`
	`Community (ro): [custom]`
	`Trap Recipient's IP address : [0.0.0.0]`
	`Community (ro): [common]`
	`Trap Recipient's IP address : [0.0.0.0]`
	`Community (ro): [FibreChannel]`
	`Trap Recipient's IP address : [172.26.1.145]`
	`Trap recipient Severity level : (0..5) [4]`
	`Trap recipient Port : (0..65535) [162]`

Problem	Resolution

An error message with the following text displays:
Real Time statistics collection has failed. Please see master log for details.
	3 To collect GigE port and FCIP statistics, you must enable the FCIP-MIB capability.
	Verification and Troubleshooting
	To verify that FCIP-MIB capability is enabled, use the following command from the switch CLI:
	`FCRRouter:admin> snmpconfig --show mibcapability`
	`FCIP-MIB: YES`
	
	To enabling FCIP-MIB capability, use the following command from the switch CLI:
	`FCRRouter:admin> snmpconfig --set mibcapability`
	`FA-MIB (yes, y, no, n): [yes]`
	`FICON-MIB (yes, y, no, n): [yes]`
	`HA-MIB (yes, y, no, n): [yes]`
	`FCIP-MIB (yes, y, no, n): [yes]`
	`ISCSI-MIB (yes, y, no, n): [yes]`
	
	4 To collect FCIP or GE statistics, you must configure SNMPv3 credentials in the **Address Properties** dialog box - **SNMP** tab.
	Verify that the SNMPv3 credentials are valid. When you discover a switch using 'admin' as the v3 credentials, a new user (for example, User 6) is created with the SNMP user name 'admin'. To verify the SNMP user credentials, use the following command from the switch CLI:
	`sw1:FID128:admin> snmpconfig --show snmpv3`
	
	`SNMPv3 USM configuration:`
	`User 1 (rw): snmpadmin1`
	`Auth Protocol: noAuth`
	`Priv Protocol: noPriv`
	`User 2 (rw): snmpadmin2`
	`Auth Protocol: noAuth`
	`Priv Protocol: noPriv`
	`User 3 (rw): snmpadmin3`
	`Auth Protocol: noAuth`
	`Priv Protocol: noPriv`
	`User 4 (ro): snmpuser1`
	`Auth Protocol: noAuth`
	`Priv Protocol: noPriv`
	`User 5 (ro): snmpuser2`
	`Auth Protocol: noAuth`
	`Priv Protocol: noPriv`
	`User 6 (ro): admin`
	`Auth Protocol: noAuth`
	`Priv Protocol: noPriv`

Problem	Resolution	

An error message with the following text displays:
Real Time statistics collection has failed. Please see master log for details.
	5	To collect data on Virtual Fabric-enabled switches, the Fabric OS user must have access to all Virtual Fabrics. The SNMPv3 user name must be the same as the Fabric OS user name. If the SNMPv3 and Fabric OS user names do not match, data is not collected for the virtual switches with the non-default VF ID. By default, the user 'admin' has access to all Virtual Fabrics.
		To verify the Fabric OS user (verify Role-LF List), use the following command from the switch CLI:
		`sw1:FID128:admin> userconfig --show`
		`Account name: admin`
		`Description: Administrator`
		`Enabled: Yes`
		`Password Last Change Date: Unknown`
		`Password Expiration Date: Not Applicable`
		`Locked: No`
		`Home LF Role: admin`
		`Role-LF List: admin: 1-128`
		`Chassis Role: admin`
		`Home LF: 128`

	6	To collect real time data, I/O must be running in the switch. To view the statistics in the switch, use one of the following command:
		FC Ports command from the switch CLI:
		`portperfshow <interval>`
		Example `Sprint-65:root> portperfshow 5`
		
		FCIP tunnels: command:
		`portshow fciptunnel <Ge port number> <tunnel no> -perf`
		Example `Sprint-65:root> portshow fciptunnel ge0 1 -perf`

An error message with the following text displays:
Real Time statistics collection has failed. Please see master log for details.
	7	To collect performance statistics from a switch, the SNMP security level must be set correctly in the switch. For example, a secLevel of '3' means "No access" which stops the management application from collecting performance statistics from the switch.
		To show the security level respectively, use the following command from the switch CLI:
		`snmpconfig --show secLevel`
		Example
		`snmpconfig --show secLevel`
		`GET security level = 0, SET level = 0`
		`SNMP GET Security Level: No security`
		`SNMP SET Security Level: No security`
		
		To set the security level respectively, use the following command from the switch CLI:
		`snmpconfig --set secLevel`
		Example
		`snmpconfig --set secLevel 0`
		`Select SNMP GET Security Level`
		`(0 = No security, 1 = Authentication only, 2 = Authentication and Privacy, 3 = No Access): (0..3) [0]`

Port Fencing troubleshooting

The following section states a possible issue and the recommended solution for Port Fencing errors.

Problem	Resolution
If you segment a switch from a fabric then rediscover the switch without accepting changes, the **Port Fencing** dialog box displays the switch twice and the port count is doubled.	Right-click on the fabric that the segmented switch (with red minus icon) is part of and select **Accept Changes**.

Professional edition login troubleshooting

The following section states a possible issue and the recommended solution for Professional edition login errors.

TABLE 162 Professional edition login troubleshooting

Problem	Resolution
Login Failed. Only one client allowed. One client session is active or has not yet timed out.	If you closed the client using Windows Task Manager (End Task or Process) or using Unix process ID (kill command), successful relaunch of the application may take up to 2 minutes.

Server troubleshooting

The following section states a possible issue and the recommended solution for server errors.

TABLE 163

Problem	Resolution
Management server exits unexpectedly on Red hat Linux 6.1	A possible cause is low swap space configured on the system. As per the standard recommendation, swap should equal 2 times physical RAM for up to 2 GB of physical RAM, and then an additional 1 times physical RAM for any amount above 2 GB, but never less than 32 MB. Therefore, if M = Amount of RAM in GB and S = Amount of swap in GB, then If M < 2 S = M *2 Else S = M + 2

Server Management Console troubleshooting

The following section states a possible issue and the recommended solution for server management console errors.

TABLE 164

Problem	Resolution
Unable to launch the SMC on a Windows Vista,Windows 7 , or Windows 2008 R2 system	The Windows Vista,Windows 7,or Windows 2008 R2 system enables the User Access Control (UAC) option by default. When the UAC option is enabled, the SMC cannot launch. If the SMC does not launch, use one of the following options to disable the UAC option: The following are the various ways we can disable UAC in vista: **Disable using msconfig by completing the following steps.** 1 Select **Start > Run**. 2 Type msconfig on the **Run** dialog box and click **OK**. 3 Click the **Tools** tab on the **System Configuration Utility**. 4 Scroll down to and select the **Disable UAC** tool name. 5 Click **Launch**. A command window displays and runs the disable UAC command. When the command is complete, close the window. 6 Close the **System Configuration Utility**. 7 Restart the computer to apply changes. **NOTE:** You can re-enable UAC using the above procedure and selecting the **Enable UAC** tool name in step 4. **Disable using regedit by completing the following steps.** **NOTE:** Before making changes to the registry, make sure you have a valid backup. In cases where you're supposed to delete or modify keys or values from the registry it is possible to first export that key or value(s) to a .REG file before performing the changes. 1 Select **Start > Run**. 2 Type regedit on the **Run** dialog box and click **OK**. 3 Navigate to the following registry key: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Policies\System 4 Right-click the **EnableLUA** value and select Modify. 5 Change the **Value data** field to 0 on the **Edit DWORD Value** dialog box and click **OK**. 6 Close the **Registry Editor**. 7 Restart the computer to apply changes. **NOTE:** You can re-enable UAC using the above procedure and changing the **Value data** field to 1 in step 5.

Problem	Resolution
Unable to launch the SMC on a Windows Vista or Windows 7 system continued	**Disable using the Group Policy by completing the following steps.** You can perform this procedure on you local machine using Local Group Policy editor or for many computers at the same time using the Active Directory-based Group Policy Object (GPO) editor. To disable using the Local Group Policy editor, complete the following steps. 1 On your local Vista computer, select **Start > Run**. 2 Type gpedit.msc on the **Run** dialog box and click **OK**. 3 Browse to **Computer Configuration > Windows Settings > Security Settings > Local Policies > Security Options** in the Group Policy editor. 4 In the right pane scroll to the User Access Control policies (at the bottom of the pane). 5 Right-click the **Behavior of the elevation prompt for Administrators in Admin Approval Mode** policy and select Properties. 6 Select the **No Prompt** option and click **OK**. 7 Right-click the **Detect application installations and prompt for elevation** policy and select **Properties**. 8 Select the **Disabled** option and click **OK**. 9 Right-click the **Run all administrators in Admin Approval Mode** policy and select **Properties**. 10 Select the **Disabled** option and click **OK**. 11 Close the Group Policy editor. 12 Restart the computer to apply changes. To disable using the Active Directory-based GPO editor, complete the following steps. 1 On a Vista computer that is a member of a domain, select **Start > Run**. 2 Type gpedit.msc on the **Run** dialog box and click **OK**. 3 Browse to the required GPO that is linked to the OU or domain where the Vista computers are located, then edit it 4 Browse to **Computer Configuration > Windows Settings > Security Settings > Local Policies > Security Options** in the Group Policy editor. 5 In the right pane scroll to the User Access Control policies (at the bottom of the pane). 6 Right-click the **Behavior of the elevation prompt for Administrators in Admin Approval Mode** policy and select **Properties**. 7 Select the **No Prompt** option and click **OK**. 8 Right-click the **Detect application installations and prompt for elevation** policy and select **Properties**. 9 Select the **Disabled** option and click **OK**. 10 Right-click the **Run all administrators in Admin Approval Mode** policy and select **Properties**. 11 Select the **Disabled** option and click **OK**. 12 Close the Group Policy editor. 13. Restart the computer to apply changes.

Supportsave troubleshooting

The following section states a possible issue and the recommended solution for supportsave errors.

Problem	Resolution
Cannot capture support save information.	Capture support show by running the batch file from the *Install_Home*/bin/supportshow.bat from Windows and UNIX systems. 1 *Open Install_Home*\bin\supportsave.bat. 2 Edit file `supportsave dbuser dbpasswd [tareget-dir] [pause-option].` **NOTE:** Unreachable switches increase the time needed to collect supportSave data.

Technical support data collection troubleshooting

The following section states a possible issue and the recommended solution for technical support data collection errors.

TABLE 165

Problem	Resolution
Technical support data collection using SCP/SFTP does not work because of one of the following issues: • For internal SCP/SFTP server, the application was uninstalled and reinstalled without migration • For external SCP/SFTP server, the SSH handshake keypair is changed - manually - due to an external server reinstall - due to the SCP/SFTP server preference (Options dialog box) being changed from built-in to external (installed on same machine) or vice versa	Clear the SSH (SCP/SFTP) server IP address or hostname from the known_hosts table of the device. • For Fabric OS devices, use the following command: `sw0:FID128:admin>` **sshutil delknownhost** `IP Address/Hostname to be deleted:` *SSH_server_IP_address* where *SSH_server_IP_address* is the IP address of the SSH server you want to delete. • For Network OS devices running firmware version 3.0 and later, use the following command: `sw0#` **clear ssh-key** *SSH_server_IP_address* where *SSH_server_IP_address* is the IP address of the SSH server you want to delete. • For Network OS devices running firmware version 2.1.1b, use the following command: `sw0#` **execute-script sshdeleteknownhost** `IP Address/Hostname to be deleted:` *SSH_server_IP_address* where *SSH_server_IP_address* is the IP address of the SSH server you want to delete.
Technical support data collection using an external SCP server does not work after reinstalling the application or the external SCP server on the Host machine.	Clear the SCP server IP address or hostname from the known_hosts table of the device using the following command: `sw0# FID10:root> ssh-keygen -R` *Host_Name* where *Host_Name* is the IP address or host name of the external SCP server.

View All list troubleshooting

The following section states a possible issue and the recommended solution for **View All** list errors.

TABLE 166

Problem	Resolution
View All list does not display.	The **View All** list does not display until you discover a fabric. To discover a fabric, refer to "Discovering fabrics" on page 35.
View All list does not display and there are discovered fabrics. **Example** If you create a new view 'V1' that has one fabric 'F1' and you display the new view in the SAN tab (select **V1** from the **View All** list). Then you delete the fabric F1 from Discovery, the **View All** list no longer displays and the following messages displays: View loaded, no devices present in the current view. Refer to the Troubleshooting Guide in Help (F1) for assistance.	To select another view, select **View > Manage View > Display View >** *View_Name*.

Wireless troubleshooting

After discovery, the Management application inspects the trap listener and syslog recipient configuration on wireless controllers. If there is a problem with the registration, the Management application changes the "registration success" master log event to a warning event with additional message text. The following section states the possible issues and the associated Master Log message that displays:

Problem	Master log warning message
The Management application successfully registers itself as SNMP trap recipient on the wireless controller; however, trap generation is disabled on the device.	Server <address> is successfully registered as SNMP Trap recipient to the switch <device_address>; but trap generation is disabled on switch.
The Management application successfully registers itself as syslog recipient on the wireless controller; however, logging is disabled on the device.	Server <address> is successfully registered as Syslog recipient to the switch <device_address>; but logging is disabled on switch.
The Management application successfully registers itself as syslog recipient on the wireless controller as the secondary or ternary recipient (primary slot is already occupied). The wireless controller sends syslog messages to secondary or ternary recipients only if primary recipient is not reachable. Therefore, even though the Management application is registered as a syslog recipient it may not receive any messages.	Server <address> is successfully registered as Syslog recipient to the switch <device_address>; but <address> is not the primary syslog recipient.
The Management application successfully registers itself as syslog recipient on the wireless controller; however, it is non-primary recipient and logging is disabled on device.	Server <address> is successfully registered as Syslog recipient to the switch <device_address>; but logging is disabled on switch and <address> is not the primary syslog recipient.

Zoning troubleshooting

The following section states some possible issues and recommended solutions for zoning errors.

Problem	Resolution
Cannot perform zoning on a new switch.	You must use telnet (or the **Product Type and Access** tab in the **Add Properties** dialog box) to change the default password on the new switch before you can use the Management application to perform zoning.
When configuring a large zone configuration a switch displays offline during discovery.	If a large zone configuration is configured in a fabric, switches may temporarily display as being offline during discovery. Wait for the next discovery cycle and click the **Refresh** button on the toolbar.
When activating a large zone configuration on a two-switch fabric on UNIX platforms, an error message displays stating "Failed to perform the requested zoning action: Failed to zone due to exception."	Although the error message states that the requested zoning action failed, the zone configuration will be correctly activated. Wait for the next zoning polling to occur. This issue only occurs on UNIX systems.
Zoning activation message displays for a long time, but zone configuration is not activated.	Telnet zoning can take a long time. To improve speed, open the **Discover Setup** dialog box (**Discover > Setup**) and add the IP address for the device to the **Selected Individual Addresses** list.
Out of memory error caused by running a zoning report for a large zone configuration (1 MB) in a medium-sized SAN due to a third party tool.	You must increase the client memory allocation by completing "Configuring memory allocation settings" on page 117.

Database Fields

Database tables and fields

NOTE
The primary keys are marked by an asterisk (*)

TABLE 167 ACH_CALL_CENTER

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
NAME	Name of the Call Center.	varchar	256

TABLE 168 ACH_CALL_CENTER_CONFIG

Field	Definition	Format	Size
KEY_ *	Key to identify the specific configuration of the Call Center.	varchar	256
CALL_CENTER_ID *	ID of the Call Center.	int	
VALUE	Value of specific configuration identified by Key of the Call Center.	varchar	256

TABLE 169 ACH_DEVICE_KEY

Field	Definition	Format	Size
CORE_SWITCH_ID	Core switch record reference.	int	
DEVICE_KEY	Device key provided on adding the device to ESRS on processing the first call home event. It is used to create HMACs for future REST calls for the same device.	char	500
CREATION_TIME	Time when the device key was created.	timestamp	

TABLE 170 ACH_EVENT

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
REASON_CODE	Reason code of the event.	varchar	256
FRU_CODE	FRU code of the event.	varchar	256
DESCRIPTION	Description of the event.	varchar	256
SEVERITY	Severity of the event.	int	

TABLE 170 ACH_EVENT (Continued)

Field	Definition	Format	Size
TYPE	Type of the event.	varchar	256
CONTRIBUTOR_PATTERN	Indicates the Contributor pattern to be used for searching the event contributor in event description. In some cases, FOS uses same message id for different events (e.g MAPS events). To increase the filtering capability of Call Home events, this contributor pattern string is used along with message id. If the event has unique message id, then contributor pattern string will be empty.	varchar	64

TABLE 171 ACH_EVENT_FILTER_MAP

Field	Definition	Format	Size
FILTER_ID *	ID of the event filter.	int	
EVENT_ID *	Event ID which needs to be associated with the filter.	int	

TABLE 172 ACH_FILTER

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of the event filter.	varchar	256
DESCRIPTION	Description of the event filter.	varchar	256

TABLE 173 ACH_INFO

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
SWITCH_WWN	WWN of the switch.	varchar	23
FILTER_ID	If an event filter is assigned to the switch - the filter ID if no filter is assigned - null.	int	
CALL_CENTER_ID	ID of the call center to which the switch is assigned.	int	
SUPPORT_SAVE	1 = Support save is enabled for the switch. 0 = Support save is disabled for the switch.	smallint	
MANAGED_ELEMENT_ID	Managed element Id for the device. Default value is -1.	int	

TABLE 174 AD_GROUP

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
NAME	Name of the active directory group.	varchar	256
EMAIL	Active Directory Group Email Address.	varchar	1024
SOURCE_SERVER_NETWORK_ADDRESS	The LDAP Server Network Address from which the Active directory group is fetched.	varchar	255

TABLE 175 ADAPTER_DRIVER_FILE_DETAILS

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
FILE_NAME	Name of the driver file	varchar	64
MAJOR_VERSION	Major version of the driver file	smallint	
MINOR_VERSION	Minor version of the driver file	smallint	
MAINTENANCE	Maintenance version of the driver file	smallint	
PATCH	Patch details of the driver file	varchar	32
SUPPORTED_OS	Holds multiple flavors of the OS	varchar	1024
OS_ARCHITECTURE	Supported OS architecture	varchar	32
IMPORTED DATE	Imported date of the driver file	timestamp with time zone	
RELEASE DATE	Release date of the driver file	timestamp with time zone	
LOCATION	Location of the adapter driver file in the repository	varchar	1024

TABLE 176 ADAPTER_PORT_CONFIG_DETAILS

Field	Definition	Format	Size
CONFIG_ID	Configuration ID	int	
PROPERTY_ID	Adapter port property ID	int	
VALUE	User configured adapter port property value	varchar	256

TABLE 177 ADAPTER_PORT_CONFIG_PROPERTY

Field	Definition	Format	Size
ID	Adapter port property ID	int	
NAME	Holds the name of the adapter port property	varchar	64
VALUE_LIST	Holds possible values for each adapter port property	varchar	2048
DEFAULT_VALUE	Holds the default value of the port property	varchar	256
DATA_TYPE	Holds the data type of the port property. 0 - Boolean 1 - Integer 2 - String	int	
GROUP_NAME	Holds the group name of the port property. Possible values is Boot_OVER_SAN, FC, FC_SP.	varchar	256

TABLE 178 AOR_DEVICE_GROUP_MAP

Field	Definition	Format	Size
AOR_ID	ID of the AOR.	int	
DEVICE_GROUP_ID	The Product Group which is in the AOR.	int	

TABLE 179 AOR_DEVICE_MAP

Field	Definition	Format	Size
AOR_ID	ID of AOR	int	
DEVICE_ID	The DEVICE ID can be IP Product or ServerIron ID which is in the AOR	int	

TABLE 180 AOR_FABRIC_MAP

Field	Definition	Format	Size
AOR_ID	ID of AOR	int	
FABRIC_ID	FABRIC ID which is in the AOR	int	

TABLE 181 AOR_HOST_MAP

Field	Definition	Format	Size
AOR_ID	ID of AOR	int	
HOST_ID	HOST ID which is in the AOR	int	

TABLE 182 AOR_INM_PORT_GROUP_MAP

Field	Definition	Format	Size
AOR_ID	ID of AOR	int	
PORT_GROUP_ID	IP of port group	int	

TABLE 183 AOR_VIP_SERVER_MAP

Field	Definition	Format	Size
AOR_ID	The column holds ID of an AOR. It is Foreign Key and refers to ID column of AOR table	int	
VIP_SERVER_ID	The column holds ID of VIP Server. It is Foreign Key and refers to ID column of VIP_SERVER table	int	

TABLE 184 AUTO_TRACE_DUMP

Field	Definition	Format	Size
CORE_SWITCH_ID		int	
ENABLED	The enabled or disabled state of automatic trace dump transfer on the switch	smallint	
PROTOCOL	The protocol Unknown(0)/FTP(1)/SCP(2) to be used for transfer	smallint	
IP_ADDRESS	The IP address of the host	varchar	64
USER_NAME	User name	varchar	64
LOCATION	Location of the directory where trace dump files are to be saved	varchar	1024
PASSWORD	User password	varchar	64

TABLE 185 AVAILABLE_FLYOVER_PROPERTY

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of the available property to be included in the flyover display.	varchar	40
TYPE	Indicates the flyover property type. Product property is 0, Connection property is 1, User Defined property is 2, Cee Product property is 3, Cee Connection property is 4, Host property is 5.	smallint	
DEFAULT_SELECTION	Value 1 in the column indicates default selected product/connection property and 0 indicates not included in the default list.	smallint	

TABLE 186 BIRTREPORT_PARAMETER

Field	Definition	Format	Size
ID	The primary key of the table.	int	
RUN_ID	References the ID column in the BIRTREPORT_RUN_TEMPLATE table.	int	
PARAMETER-TYPE	Control type of the parameter. • 1 - Text Box • 2 - List Box • 3 - Radio Button	int	
PARAMETER_NAME	Name of the parameter in the report template design.	varchar	128
PROMPT_TEXT	Text Label for the parameter. This value will be displayed on the GUI.	varchar	256
DATA_TYPE	Data type of the parameter. Possible values are: • 1 - String • 2 - Float • 3 - Decimal • 4 - Date and Time • 5 - Boolean • 6 - Integer • 7 - Date • 8 - Time	int	
PARAMETER_VALUE	Value of the Parameter.	varchar	256

TABLE 187 BIRTREPORT_RUN_TEMPLATE

Field	Definition	Format	Size
ID	The primary key of the table.	int	
SCHEDULE_ID	References the ID column in the BIRTREPORT_SCHEDULE_CONFIG table.	int	

TABLE 187 BIRTREPORT_RUN_TEMPLATE (Continued)

Field	Definition	Format	Size
REPORT_TEMPLATE_TITLE	Report Template title. This name is the same as the title name in the REPORT_TEMPLATE table. There is no foreign key relation here as the user may delete and add a template but the schedule should still hold good if looked up by title. Also title is unique in the REPORT_TEMPLATE table.	varchar	256
NAME	Name of the generated report file.	varchar	256

TABLE 188 BIRTREPORT_SCHEDULE_CONFIG

Field	Definition	Format	Size
ID	The primary key of the table.	int	
DEPLOYMENT_ID	References the ID column in the DEPLOYMENT_CONFIGURATION table.	int	
NAME	Name of the schedule.	varchar	128
REPORT_STORE_LOCATION	Path to the location where the generated report files are stored.	varchar	256
OVERWRITE	0 and 1 are allowed values.1 indicates overwrite is true. I.e., every run of the schedule will overwrite the previous output.0 indicates archive. I.e., every run of the schedule will create a new folder in the store location with timestamp to ensure that output of all the runs will be archived.	int	
FORMAT_TYPE	Possible values are 0, 1, and 2. • 0 indicates output will be in HTML • 1 indicates PDF • 2 indicates CSV	int	
CREATED_BY	Indicates which user created the schedule.';	int	
EMAIL_DELIVERY	Indicates if the generated report needs to be emailed or not. 0 = Not required, 1 = Required. Default value is 0	int	
FOLDER_DELIVERY	Indicates if the generated report needs to be placed in the specified report_store_location or not. 0 = Not required, 1 = Required. Default value is 0.	int	
EMAIL_RECIPIENTS	Indicates the generated report email recipients.	character	255
EMAIL_FROM	Indicates the email from field.	character	255
EMAIL_REPLYTO	Indicates the email reply to field.	character	255
EMAIL_SUBJECT	Indicates the email subject.	character	255
EMAIL_PROLOGUE	Indicates the email body prolog.	character	255
EMAIL_EPILOGUE	Indicates the email body epilog.	character	255
LAST_MODIFIED_TIME	Indicates when the schedule was last modified.	timestamp	
CREATED_TIME	Indicates when the schedule was created.	timestamp	

TABLE 189 BIRTREPORT_SCHEDULE_PARAMETER

Field	Definition	Format	Size
ID	The primary key of the table.	int	
BIRTREPORT_SCHEDULE_TEMPLATE_ID	Id of birtreport_schedule_template table.	int	
BIRTTEMPLATE_PARAMETER_ID	Id of birttemplate_parameter table.	int	
PARAM_VALUE	Value of the parameter.	text	
PARAM_DISPLAY_VALUE	Displays value of the parameter.	varchar	256

TABLE 190 BIRTREPORT_SCHEDULE_TEMPLATE

Field	Definition	Format	Size
ID	The primary key of the table.	int	
SCHEDULE_ID	Id of birtreport_schedule_config table.	int	
REPORT_TEMPLATE_ID	Id of report_template table.	int	
PARAMETERIZED	Possible values are 0 and 1. 0 - Indicates NotParameterized 1 - Indicates Parameterized	int	

TABLE 191 BIRTTEMPLATE_PARAMETER

Field	Definition	Format	Size
ID	The primary key of the table.	int	
REPORT_TEMPLATE_ID	Id of report_template table.	int	
PARAM_NAME	Name of the parameter.	varchar	128
PARAM_REQUIRED	Indicates if the report template parameter is required to be specified or not. 0 = Not required, 1 = Required. Default value is 0.';	int	
PARAM_DESCRIPTION	Indicates the report template parameter description.	varchar	128
PARAM_PICKER_TYPE	Indicates the picker type for the report template parameter. 0 = None, 1 = Fabric, 2 = Switch. Default value is 0.	int	

TABLE 192 BOOT_IMAGE_DRIVER_MAP

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
MAJOR_VERSION	Major Version bit from Boot Image file	smallint	
MINOR_VERSION	Minor Version bit from Boot Image file	smallint	

TABLE 192 BOOT_IMAGE_DRIVER_MAP (Continued)

Field	Definition	Format	Size
MAINTENANCE	Maintenance Version bit from Boot Image file	smallint	
PATCH	Patch Version bit from Boot Image file	varchar	32
MD5_HASH	MD5 hash value for Boot Image file	varchar	64
SUPPORTED_DRIVERS	Compatible HCM Drivers delimited by comma	varchar	256

TABLE 193 BOOT_IMAGE_FILE_DETAILS_

Field	Definition	Format	Size
ID		int	
DRIVER_MAPPING_ID		int	
BOOT_IMAGE_NAME	Name of Boot Image file	varchar	64
MAJOR_VERSION	Major Version bit from Boot Image file	smallint	
MINOR_VERSION	Minor Version bit from Boot Image file	smallint	
MAINTENANCE	Maintenance Version bit from Boot Image file	smallint	
PATCH	Patch Version bit from Boot Image file	varchar	32
IMPORTED_DATE	Imported date of Boot Image file	timestamp	
RELEASE_DATE	Release date of Boot Image file	timestamp	
RELEASE_NOTES_LOCATION	Release notes location in Management application Repository	varchar	1024
LOCATION	Boot Image file location in Management application Repository	varchar	1024

TABLE 194 BOOT_LUN_SEQUENCE

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
NAME	Name of the Boot LUN Sequence	varchar	64
FABRIC_ID	PK of the owning fabric	INT	

TABLE 195 BOOT_LUN_SEQUENCE_DETAIL

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
BOOT_LUN_SEQ_ID	PK of the owning Boot LUN Sequence	char	23
PORT_WWN	WWN of the port in the Boot LUN Sequenc	int	
LUN_NUM	LUN number of the port in the Boot LUN Sequence	int	
SEQUENCE_NUM	Sequence number of the port in the Boot LUN Sequence		

TABLE 196 CAPABILITY_

Field	Definition	Format	Size
NAME *	Name of the capability.	varchar	256
DESCRIPTION	Optional detailed description about the capability.	varchar	512

TABLE 197 CARD

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
CORE_SWITCH_ID *	Core switch DB ID.	int	
SLOT_NUMBER	The number of the physical slot in the chassis where the blade is plugged in. For fixed blades, SlotNumber is zero.	smallint	
TYPE	ID of the blade to identify the type.	smallint	
EQUIPMENT_TYPE	The type of the blade. It is either SW BLADE or CP BLADE.	varchar	32
STATE	State of the blade, such as ENABLED or DISABLED.	varchar	32
POWER_STATE	State of power supply to the blade.	varchar	16
ATTN_STATE		varchar	32
SERIAL_NUMBER	Factory serial number of the blade.	varchar	32
PART_NUMBER	The part number assigned by the organization responsible for producing or manufacturing the blade.	varchar	32
TRUNKING_SUPPORTED	1 = trunking is supported on this blade.	smallint	
FICON_DISABLED	1 = FICON is disabled on this blade.	smallint	
IP_ADDRESS	IP address of first Ethernet management port for a given slot with intelligent blade.	char	64
SUBNET_MASK	Mask of first Ethernet man.agement port for a given slot with intelligent blade.	varchar	64
DEFAULT_GATEWAY	Gateway IP address Ethernet management for a given slot with intelligent blade.	varchar	64
PRIMARY_FW_VERSION	Primary firmware version of applications on this blade. Applicable only for AP_BLADE.	varchar	48
SECONDARY_FW_VERSION	Secondary firmware version applications on this blade. Applicable only for AP_BLADE.	varchar	48
FCIP_CIRCUIT_CAPABLE	The blade is capable of creating FCIP Circuits. • 1 = true. • 0 = false. • Default value is 0.	smallint	
FCIP_LICENSED	FCIP Advanced Extension Licensing is available. • 1 = available. • 0 = not licensed. • -1 = not supported. • Default value is -1.	smallint	
MAX_FCIP_TUNNELS	The maximum number of tunnels that can be created in this slot. • -1 = not supported. • Default value is -1.	int	
MAX_FCIP_CIRCUITS	Describes the maximum number of circuits that can be created in this slot. • -1 = not supported. • Default value is -1.	int	

TABLE 197 CARD (Continued)

Field	Definition	Format	Size
CP_BLADE_INDEX	CP blade index. Default value is -1.	smallint	
CP_HA_STATE	CP's HA state information like Active/Stand by.	varchar	128
ETHERNET_IPV6_ADDRESS	IPV6 address of Ethernet management port for the blade.	varchar	64
ETHERNET_IPV6_GATEWAY	IPV6 Gateway address of Ethernet management port for the blade.	varchar	64
NUMBER_OF_PORTS		int	
HEADER_VERSION	The OEM or vendor-assigned version number.	int	
GIGE_MODE	Determines the port operating mode for GE ports. • 0 - 1G • 1 - 10G • 2 - Dual mode • 3 - Failover mode Default value -1 means it is not applicable.	smallint	
PHYSICAL_SLOT_NUMBER	For chassis based NOS devices, we are storing the logical slot number in the property slot_number, as this is what is displayed in the if_name of the interfaces. This property will be used to store the physical slot number which is unique within the given chassis.	smallint	
FEATURES_ENABLED	Holds bit mask which represents the features that are enabled in this card/blade. Each bit would represents a specific feature. Bit mask supported : IP_EXTENSION_MODE_ENABLED (131072)	int	

TABLE 198 CARD_CAPABILITY

Field	Definition	Format	Size
CARD_ID *	DB ID of the card.	int	
CAPABILITY_ *	Name of the capability detected on the card.	varchar	256
ENABLED	1 = the capability is enabled on the card. Default value is 0.	int	

TABLE 199 CED_APPLICATION

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of the application. Application represents a collection of active zones in a fabric.	varchar	24
FABRIC_ID	ID of the fabric for which the application is created.	int	

TABLE 200 CED_APPLICATION_MEMBER

Field	Definition	Format	Size
APPLICATION_ID*	Auto-generated DB CED_Application table ID.	int	
ZONE_ID*	Auto-generated DB Zone table ID which joins as a member of the application.	int	

TABLE 201 CED_USER_PREFERENCE

Field	Definition	Format	Size
USER_NAME*	User Name carried from _USER table.	varchar	128
FABRIC_ID*	Fabric ID carried from Fabric table.	int	
APPLICATION_ID	CED application ID representing the group of end devices to be displayed in the fabric.	int	

TABLE 202 CEE_PORT

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
GIGE_PORT_ID	FK to GIGE_PORT	int	
VIRTUAL_SWITCH_ID	FK to owning VIRTUAL_SWITCH	int	
IF_INDEX	Interface index	int	
IF_NAME	Interface name	varchar	256
IF_MODE	Gige port mode (L2, L3, none)	varchar	8
L2_MODE	L2 mode (hybrid, trunk, access)	varchar	32
VLAN_ID	List of VLAN this port belongs to	text	
LAG_ID	LAG ID this port belongs to	int	
IP_ADDRESS	Port''s configured IP address	varchar	128
MAC_ADDRESS	Port''s MAC address	varchar	64
PORT_SPEED	Speed in Gb/sec. The default value is 0.	int	
ENABLED	State. The default value is 0.	smallint	
OCCUPIED	The default value is 0.	smallint	
LAST_UPDATE		bigint	
MAC_ACL_POLICY	stores the MAC ACL policy information of the port	varchar	64
NET_MASK	Netmask of the IPAddress of the port	varchar	128
PROTOCOL_DOWN_REASON	Reason for the port''s operational state being down	varchar	512
QOS_TYPE	QoS Type (Cee-Map, TrafficClass Map, FCoE map)	varchar	32
QOS_NAME	Name of the QoS Map set on the port	varchar	64
DOT1X_ENABLED	Indicate if 802.1x authentication is enabled on this port. The default value is 0.	smallint	
PORT_ROLE	Thisfield is used to store the port role value. The value will be populated by the NosSwitchAssetCollector. This field valid values include ISL or Edge. Default value is empty string.	varchar	32

TABLE 202 CEE_PORT (Continued)

Field	Definition	Format	Size
AMPP_PROFILE_MODE	Specifies whether the interface is set to AMPP profile mode.	smallint	
CONNECTED_STORAGE_TYPE	This column holds the type of the storage device (NAS/iSCSI/Others) connected to this interface. This is applicable only when "EDGE_TYPE" is set as "Storage(3)" otherwise this gets the default value "Others(0)". 0- Others 1-NAS 2-iSCSI	int	
EDGE_TYPE	Indicates the type of device connected to this interface. -1 - Not Applicable 0 - connected to device with unknown type 1 - connected to managed Brocade branded AP 2 - connected to standalone Brocade branded AP 3 - connected to a storage device 4 - connected to a Server	int	

TABLE 203 CFG_BACKUP_ARCHIVE

Field	Definition	Format	Size
CFG_BACKUP_ARCHIVE_ID		int	
DEVICE_ID	IP Product DB ID from which the configuration has been retrieved.	int	
USER_ID	Unique DB ID of user who initiated this config upoload.	int	
PRODUCT_TYPE	Indicates the type of product from which the config is retrieved for example NetIron XMR/MLX.	varchar	32
VERSION	Version of the configs downloaded for each product.	num	(8,0)
LOCATION		varchar	255
DATE_TIME	The date and time at which the configuration has been backedup. The date and time will be saved in the following format "Mon May 10 17:59:13 PDT 2010".	varchar	64
FILE_NAME		varchar	64
IS_BASELINE	Indicates if the configuration file is selected by user as baselined configuration or not.	num	(1,0)
DESCRIPTION	Brief comments and description about this configuration.	varchar	1024
IMAGE_VERSION	The firmware version on the product while the config is retrieved.	varchar	64
CLI_TEMPLATE_REPORT_EXECUTION_ID	DB ID of the cli template report execution.	int	
CLI_TEMPLATE_REPORT_EXECUTION	Result of the cli template execution.	int	

TABLE 203 CFG_BACKUP_ARCHIVE (Continued)

Field	Definition	Format	Size
CONFIG_DATA	The actual configuration data of the IP or DCB product.	txt	
CONFIG_TYPE	Configuration Type DCB_RUNNING=1, DCB_STARTUP=2, IP_STARTUP=3, IP_RUNNING=4	smallint	
DRIFT_STATUS	Indicates whether this config backup is deviated or not with respect to the active baseline at at the time of adding this config backup to the repository or not. If there is no baseline for the product, this column will be set to NO_BASELINE(-1). The possible values NO_DEVIATION=0, DEVIATED=1, NO_BASELINE=-1	int	

TABLE 204 CFM_MA

Field	Definition	Format	Size
MA_DB_ID	Id of the management association.	int	
MD_DB_ID	Id of the management domain.	int	
MA_INDEX	Id of the management association as given by the device.	int	
MA_NAME	Name of the management association.	varchar	255
CCM_INTERVAL	Continuity check message interval value.	smallint	
MIP_CREATION_POLICY	Mip policy interval value.	smallint	
MPLS_SERVICE_DB_ID	MPLS service id.	int	
VLAN_SERVICE_DB_ID	VLAN service id.	int	

TABLE 205 CFM_MD

Field	Definition	Format	Size
MA_DB_ID	Id of the management domain.	int	
Device_DB_ID	Id of the device.	int	
MD_INDEX	Id of the management domain as given by the device.	int	
MD_NAME	Name of the management domain.	varchar	255
MD_LEVEL	Level of the management domain.	smallint	

TABLE 206 CFM_MEP

Field	Definition	Format	Size
MEP_DB_ID	Maintenance end point id.	int	
MA_DB_ID	Maintenance association id.	int	
INTERFACE_DB_ID	Id of the interface.	int	

TABLE 206 CFM_MEP (Continued)

Field	Definition	Format	Size
IDENTIFIER	Maintenance end point identifier under the association.	smallint	
DIRECTION	Maintenance end point direction.	smallint	
VLAN_ID	Id of the Vlan.	smallint	

TABLE 207 CHANGELOG

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
APPLIED_AT		varchar	25
DESCRIPTION		varchar	255

TABLE 208 CLI_TEMPLATE

Field	Definition	Format	Size
CLI_TEMPLATE_ID		int	
USER_ID		int	
NAME		varchar	256
TYPE	The template type. Product Monitoring: 2, Global configuration: 1 in CLI Configuration.	num	(2,0)
CLI_CMD		varchar	
DESCRIPTION		varchar	512
DEVICE_USERNAME		varchar	256
DEVICE_PASSWORD		varchar	256
DATE_TIME		varchar	64
DEVICE_ENABLE_USERNAME		varchar	256
DEVICE_ENABLE_PASSWORD		varchar	256
CLI_FILTER		varchar	
HAS_PARAMETERS		num	(1,0)
PROMPT_ADDITIONAL_TARGET	The flag to indicate whether or not to prompt for additional targets during deployment. 1 = Prompt for additional targets. 0 =Do not prompt for additional target.	smallint	
PARAMETERS	Stores Parameter name and values in XML Format.	text	
PARAMETER_MODE	The flag to indicate whether the same parameter has to be applied for all targets or different values to be applied for each target. 0 - Same value for all targets. 1 - Different values for each targets.	smallint	

TABLE 208 CLI_TEMPLATE (Continued)

Field	Definition	Format	Size
IS_EXAMPLE	The flag to indicate whether or not if the template is example template. 0 - User Defined (not example) template. 1 - Example Template.	smallint	
VALIDATE_CLI_RESPONSE	The flag to indicate whether or not if the CLI responses for each of the CLI commands are to be validated for this template. 0 - Dont validate CLI Responses. 1 - Validate CLI Responses.	smallint	
PROMPT_ADDITIONAL_PARAMETERS	The flag to indicate whether or not to prompt Parameter tab during manual deployment for changes. 0 - Dont prompt Parameter tab. 1 - Prompt Parameter tab.	smallint	
SCHEDULE_ENABLED	The flag to indicate whether or not the CLI Template is scheduled. 0 - Scheduled deployment is turned OFF. 1 - Scheduled Deployment is turned ON.	smallint	
MODULE	Stores the module or feature name of the CLI commands. This is used for example CLI templates.	varchar	256

TABLE 209 CLIENT_VIEW

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
USER_NAME	The Management application user name.	varchar	128
NAME	Client view name.	varchar	255
DESCRIPTION	Client View description.	varchar	255

TABLE 210 CLIENT_VIEW_COLUMN

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
NAME	Name of the column. It is used as column header in product list and property name in property sheet(SAN and IP)	varchar	255
ENTITY_CATEGORY	Holds the type of the entity to whom the column name belongs to like Port, Fabric, IPProduct, VCSInterface, etc'	varchar	128
COLUMN_INDEX	Used to differentiate user defined columns and static columns. For static it is 0 and for user defined columns it is 1,2,3.	small int	
DESCRIPTION	Holds description of the column.	varchar	255
ICON_ID	Holds Icon Id for the column. Currently it is unused.	int	
VISIBLE	Indicates whether the columns are visible. 0 - Not Visible, 1 - Visible	smallint	
EDITABLE	Indicates whether the columns are editable. 0 - Not Editable, 1 - Editable.	smallint	

TABLE 211 CLIENT_VIEW_MEMBER

Field	Definition	Format	Size
CLIENT_VIEW_ID *	Foreign key to CLIENT_VIEW table.	int	
FABRIC_ID *	Foreign key to FABRIC table.	int	

TABLE 212 CLIENT_VIEW_MEMBER_HOST

Field	Definition	Format	Size
CLIENT_VIEW_ID	Primary key of CLIENT_VIEW table	int	
HOST_ID	Primary key of DEVICE_ENCLOSURE table	int	

TABLE 213 CLUSTER

Field	Definition	Format	Size
ID *	Arbitrary integer to identify the cluster.	int	
NAME	User-assigned name to identify the cluster. Names should be unique to avoid user confusion, but the database does not enforce uniqueness.	varchar	64
IP_ADDRESS	The primary hostname or IP address for managing the cluster as a single entity. The definition of primary depends on the clustering technology.	varchar	64

TABLE 214 CLUSTER_MEMBER

Field	Definition	Format	Size
CLUSTER_ID	Identifies the cluster containing a member.	int	
DEVICE_ENCLOSURE_ID	Identifies a member of the cluster.	int	32

TABLE 215 CNA_ETH_PORT

Field	Definition	Format	Size
ID	ID of the Eth port	int	
ETH_DEV	Ethernet device	varchar	64
ETH_LOG_LEVEL	Log level for the Ethernet device. Possible values are 0 - Log Invalid 1 - Log Critical 2 - Log Error 3 - Log Warning 4 - Log Info	int	
NAME	Name of the port	varchar	256
MAC_ADDRESS	MAC Address	varchar	64
IOC_ID	IO controller ID. The default value is 0.	varchar	64
HARDWARE_PATH	Hardware path of the port	varchar	256
STATUS	Status of the Eth port. The default value is -1.	varchar	64
CNA_PORT_ID	ID of the parent port	int	

TABLE 215 CNA_ETH_PORT (Continued)

Field	Definition	Format	Size
CREATION_TIME	CNA Eth port record creation time. This tells when the port was first discovered.	timestamp	
CURRENT_MAC_ADDRESS	User definable Mac address which is by default same as built in Mac address	varchar	64
MAX_BANDWIDTH	Maximum bandwidth	varchar	64
PCIF_INDEX	Pci function Iidex	varchar	64
MAX_PCIF	Maximum number of Pci functions.	smallint	
MIN_BANDWIDTH	Minimum guaranteed bandwidth. Value will be in Gbps (0 to 10).	int	
MTU	Maximum transmission unit in bytes	int	

TABLE 216 CNA_PRODUCT_CONNECTIVITY

Field	Definition	Format	Size
CNA_PORT_ID	CNA Port identifier.	int	
INTERFACE_ID	Interface Identifier.	int	

TABLE 217 CNA_ETH_PORT_CONFIG

Field	Definition	Format	Size
ID	Unique autogenerated db id.	int	
CNA_PORT_ID	Foreign key, related cna eth port config with the CNA port.	int	
CNA_ETH_PORT_ID	Nullable foreign key, related cna eth port config with the CNA eth port.	int	
PCIF_INDEX	PCI Function Index eg 2/1/1(adapter number/physical port number/port index).	varchar	64
CURRENT_MAC_ADDRESS	Current MAC address of the port.	varchar	64
MAX_BANDWIDTH	Maximum guaranteed bandwidth. Value will be in Gbps (1 to 10).	varchar	64
MIN_BANDWIDTH	Minimum guaranteed bandwidth. Value will be in Gbps (0 to 10).	int	
PORT_NUMBER	Physical port number of adapter.	int	
PORT_TYPE	Type of this port. For example, ETH.	varchar	64
CREATION_TIME	Creation time of this DB record.	timestamp	
CONFIGURATION_STATUS	Indicates current configuration status of the port. Possible values are: -1 is Invalid 0 is Added 1 is deleted	int	

TABLE 218 CNA_PORT

Field	Definition	Format	Size
ID	Primary key autogenerated ID	int	
PORT_NUMBER	Port number of the CNA port	int	
PORT_WWN	Port WWN of the port	char	23
NODE_WWN	Node WWN of the port	char	23
PHYSICAL_PORT_TYPE	Port type CNA/FC	varchar	32
NAME	Name of the port	varchar	256
MAC_ADDRESS	MAC address of the port.	varchar	64
MEDIA	Media of the port	varchar	64
CEE_STATE	State of the port.	varchar	64
HBA_ID	ID of the port.	int	
CREATION_TIME	CNA port record creation time. This tells when this port was first discovered.	timestamp	
FACTORY_PORT_WWN	Factory configured Port WWN defined for the CNA port in HCM	varchar	23
FACTORY_NODE_WWN	Factory configured Node WWN defined for the CNA port in HC	varchar	23
PREBOOT_CREATED	Flag to identify vports created during preboot and will accept string values True/false/empty	varchar	23
ALARM_WARNING	List of Alarm and Warning flags (Temperature below low threshold, supply voltage exceed high threshold etc.) for degrading SFP. If the SFP is having good health, both these flags would be None.	varchar	2048

TABLE 219 COLLECTOR

Field	Definition	Format	Size
NAME *	Name of the collector registered with the collection framework.	varchar	256
CLASS_NAME	Java class name which serves as the collector.	varchar	256
DESCRIPTION	Collector description, usually not used.	varchar	512

TABLE 220 COLLECTOR_MIB_OBJECT_ENTRY

Field	Definition	Format	Size
COLLECTOR_MIB_OBJECT_ENTRY_ID	Primary key autogenerated ID.	int	
COLLECTOR_ID	ID of the PERF_COLLECTOR.	int	
MIB_OBJECT_ID	MIB_OBJECT table DB ID.	int	

TABLE 221 COLLECTOR_SNMP_EXPRESSION_ENTRY

Field	Definition	Format	Size
COLLECTOR_SNMP_EXPRESSION_ENTRY_ID	Primary key autogenerated ID.	int	

TABLE 221 COLLECTOR_SNMP_EXPRESSION_ENTRY (Continued)

Field	Definition	Format	Size
COLLECTOR_ID	ID of the PERF_COLLECTOR.	int	
EXPRESSION_ID	Id of the SNMP_EXPRESSION.	int	

TABLE 222 COLLECTOR_TARGET_ENTRY

Field	Definition	Format	Size
COLLECTOR_TARGET_ENTRY_ID	Primary key autogenerated ID.	int	
COLLECTOR_ID	ID of the PERF_COLLECTOR.	int	
TARGET_ID	Target ID of the SNMP collector data. For device level collector it will use deviceId, and for port level it will use interfaceId.	int	
PROP_STR	Property string of the PERF_COLLECTOR.	varchar	8192
COLLECTOR_TARGET_ENTRY_TYPE	Target type of the SNMP collector data. for device level collector the target type is 0, for port level it is 1.	int	

TABLE 223 CONFIG_BLOCK

Field	Definition	Format	Size
ID	ID of the block.	int	
NAME	Name of the block.	varchar	255
DESCRIPTION	Description of the block.	varchar	1024
USE_REGEX	Indicates whether the block start is built with regular expression or not. 0 = Does not contain Regular expression 1 = Contains regular expression	smallint	
BLOCK_START	Block start string to match one or more block starts in the device config. Can be built with regular expression to match more than one block.	varchar	1024
BLOCK_END	Block end string. Used the first match to form config block from start to end.	varchar	1024
CATEGORY	Category of the Block. 0 = User defined 1 = Predefined.	smallint	

TABLE 224 CONFIG_CONDITION

Field	Definition	Format	Size
ID	Condition ID.	int	
NAME	Name of the condition.	varchar	255
DESCRIPTION	Description of the condition.	varchar	1024
REMEDIATION	Remediation details for failed conditions.	text	

TABLE 224 CONFIG_CONDITION (Continued)

Field	Definition	Format	Size
USE_REGEX	Indicates whether the condition lines are built with regular expression or not. 0 = Does not contain Regular expression 1 = Contains regular expression	smallint	
MATCH	The device config should Match or Not match with condition. 0 = Not Match 1 = Match.	smallint	
CONDITION_STR	The condition string to match the device config. Unlimited length. Each line in configuration will be matched in the whole config or a block if the order_lines is 0. Else all lines will be matched together.	text	
ORDER_LINES	Indicates whether the condition_str lines order should be matched in the config or block. 0 = Lines order check is not required. 1 = Lines order should be matched.	smallint	
CATEGORY	Category of the Condition. 0 = User defined, 1 = Predefined.	smallint	

TABLE 225 CORE_SWITCH

Field	Definition	Format	Size
ID*	Auto generated ID for this table.	int	
IP_ADDRESS	IP Address of the switch that is represented by this record. Could be either IPV4 or IPV6 address.	varchar	128
WWN	WWN of the core switch.	char	23
NAME	Switch name if available otherwise stores the wwn of the switch.	varchar	64
TYPE	Stores the switch type, the sw_bd_type of the switch.	smallint	
MODEL	Holds the switch model, whether its Brocade, Mcdata or unknown . Value 2 is for Brocade and 3 is for McData	smallint	
FIRMWARE_VERSION	Firmware version of the switch.	varchar	128
VENDOR	Vendor information for the switch.	varchar	256
MAX_VIRTUAL_SWITCHES	Maximum number of virtual switches supported.	smallint	
NUM_VIRTUAL_SWITCHES	Total number of virtual switch present.	smallint	
REACHABLE	Determines whether the switch is reachable from the Management application. 1 is reachable and 0 is unreachable	smallint	
UNREACHABLE_TIME	Time when the switch becomes unreachable.	timestamp	
OPERATIONAL_STATUS	Chassis operational status like FRU, Power Supply etc..	varchar	128

TABLE 225 CORE_SWITCH (Continued)

Field	Definition	Format	Size
CREATION_TIME	Core switch record creation time. This tells us when the intial discovery has happened.	timestamp	
LAST_SCAN_TIME	Last scan time tells the time when the last time the switch was polled.	timestamp	
LAST_UPDATE_TIME	Last update time tells the time when the last update to the database record happened.	timestamp	
SYSLOG_REGISTERED	Determies whether the switch is registered for sending syslog traps. • 1 is registered • 0 is not registered.	smallint	
CALL_HOME_ENABLED	Determines whether the call home feature is enabled..	smallint	
SNMP_REGISTERED	Determines whether the switch is registered for sending SNMP traps . • 1 is registered • 0 is not registered.	smallint	
USER_IP_ADDRESS	Applicable for McDATA switches and VCS clusters. McDATA Switches - This column is used to store the IP address which user provides for those M-model switches for which seed switch is unable to return IP address. VCS clusters - This column is used to store the IP address given by the user during discovery time. For profile based discovery, we will populate the seed switch IP address selected by the system.	varchar	128
NIC_PROFILE_ID	Nic Profile ID refers to the entry in the NIcProfile table that has IP Address of the Management application which is used as Syslog or SNMP reciepients.	int	
MANAGING_SERVER_IP_ADDRESS	IP address(v4/v6) of the Management application server which is currently managing the M-model switch. Used for M-EOS switch only. It does not apply to Fabric OS switches.	varchar	128
VF_ENABLED	Determines whether Virtual Fabric is enabled on the switch. • 1 is enabled • 0 is disabled	smallint	
VF_SUPPORTED	Determines whether virtual fabric is supported on the switch. • 1 is supported • 0 is unsupported	smallint	
MANAGED_ELEMENT_ID	A unique managed element ID for this physical switch. Also a foreign key reference to the MANAGED_ELEMENT table.	int	

TABLE 225 CORE_SWITCH (Continued)

Field	Definition	Format	Size
NAT_PRIVATE_IP_ADDRESS	NAT private IP Address. Feature available from NMS DC Eureka release onwards. During a successful NAT translation the Private IP that gets translated will be stored in this field. The new translated IP Address will be stored in the existing IP_ADDRESS field. All the NAT look up will be done using the NAT Private IP Address.	varchar	128
ALTERNATE_IP_ADDRESS	Alternate IP address of the switch. Feature available from Eureka release onwards. During fabric discovery the column will be populated based on the values in the fabricinfo.html. If Management application server is IPV6 capable, then we store the switchetherIP NVP else we store the switchetherIPV6. So could be either IPV4 or IPV6 address. If there exists any NAT translation, translated IP will be used.	varchar	128
MAC_ADDRESS	Stores the VCS Mac Address. The value will be populated by the FabricCollector. Default value is empty string. The management interface Mac Address will be stored here.	varchar	64

TABLE 226 CORE_SWITCH_CAPABILITY

Field	Definition	Format	Size
CORE_SWITCH_ID *	DB ID.	int	
CAPABILITY_ *	Name of the capability detected on the core switch.	varchar	256
ENABLED	1 = the capability is enabled on the core switch. Default value is 0.	int	

TABLE 227 CORE_SWITCH_CHECKSUM

Field	Definition	Format	Size
CORE_SWITCH_ID *	DB ID.	int	
CHECKSUM_KEY *	Checksum type.	varchar	32
CHECKSUM	Checksum value.	varchar	16

TABLE 228 CORE_SWITCH_COLLECTION

Field	Definition	Format	Size
CORE_SWITCH_ID *	Core switch ID.	int	
COLLECTION_NAME *	Collector name.	varchar	256
LAST_CORE_SW_MODIFICATION	Last core switch modification time.	timestamp	

TABLE 229 CORE_SWITCH_DETAILS

Field	Definition	Format	Size
CORE_SWITCH_ID*	Primary key for the table.	int	
ETHERNET_MASK	Ethernet mask of the core switch IP address.	char	64
FC_MASK	FC IP Address ethernet mask.	char	64

TABLE 229 CORE_SWITCH_DETAILS (Continued)

Field	Definition	Format	Size
FC_IP	Fibre Channel IP address.	char	64
FC_CERTIFICATE	FC IP Address.	smallint	
SW_LICENSE_ID	License ID of the chassis.	char	23
SUPPLIER_SERIAL_NUMBER	Supplier serial number for the switch.	varchar	32
PART_NUMBER	Partnumber of the switch	varchar	32
CHECK_BEACON	Denotes if Switch Beacon is enabled or not on the switch. 1 = beacon is turned on; otherwise, 0.	smallint	
TIMEZONE	Timezone of the switch.	varchar	32
MAX_PORT	Number of maximum ports physically allowed on the switch.	smallint	
CHASSIS_SERVICE_TAG	Chassis service tag for the switch.	varchar	32
BAY_ID	Bay ID of the switch.	varchar	32
TYPE_NUMBER	Type number is more of details for the type, Ex: SLKWRM.	varchar	32
MODEL_NUMBER	Model number is the same as the model number like Brocade 8000, Brocade VDX 6710.	varchar	256
MANUFACTURER	Manufacturer for the switch.	varchar	32
PLANT_OF_MANUFACTURER	Plant of the manufacturer for the switch.	varchar	32
SWITCH_SERIAL_NUMBER	This is the factory serial number.	varchar	32
ACT_CP_PRI_FW_VERSION	Stores Active CP primary firmware version.	varchar	128
ACT_CP_SEC_FW_VERSION	Stores Active CP secondary firmware version.	varchar	128
STBY_CP_PRI_FW_VERSION	Standby CP primary firmware version.	varchar	128
STBY_CP_SEC_FW_VERSION	Standby CP secondary firmware version.	varchar	128
TYPE	Type of the switch, basically the sw bd type stored in the core switch.	smallint	
EGM_CAPABLE	EGM license supported or not. • 1 is supported • 0 is not supported.	smallint	
SUB_TYPE	Sub Type of the switch. DCX+ and DCX-4S+ has values as 1, otherwise 0.	varchar	32
PARTITION	Partitions supported in the switch.	smallint	
MAX_NUM_OF_BLADES	Required by SMIA to populate Brocade_Chassis.MaxNumOfBlades property. It indicates the max no of blades that can be present in a chassis.	smallint	

TABLE 229 CORE_SWITCH_DETAILS (Continued)

Field	Definition	Format	Size
VENDOR_VERSION	Required by integrated SMI agent to populate Brocade_Product.Version property.	varchar	32
VENDOR_PART_NUMBER	Required by integrated SMI agent to populate Brocade_Product.SKUNumber property.	varchar	32
SNMP_INFORMS_ENABLED	Flag to denote whether SNMP informs option in the switch is enabled or disabled. Default value is 0.	smallint	
RNID_SEQUENCE_NUMBER	RNID sequence number for the switch.	varchar	32
CONTACT	Contact details of the switch. Syscontact from the RFC 1213 Mib.	varchar	256
LOCATION	Location details for the switch. Syslocation from RFC 1213.	varchar	256
DESCRIPTION	Description about the switch. Sysdescr from RFC 1213	varchar	256
FIRMWARE_VERSION	Firmware version of the switch.	varchar	128
CHASSIS_PACKAGE_TYPE	A value indicating the type of chassis.	int	
DOMAIN_NAME	Denotes the domain name configured in switch.	varchar	64
IP_ADDRESS_PREFIX	Required to populate the prefix of IPv6 address. Applicable only for IPv6 address.		
DOMAIN_NAME	Denotes the domain name configured in switch.		
FRAME_LOG_SIZE	The number of entries in the framelog.	int	
FRAME_LOG_ENABLED	Indicates if framelog is enabled on the switch. 0 = disabled, 1 = enabled.	smallint	
MAPS_ENABLED	Boolean flag to indicate if the switch is MAPS enabled or not. Enabled: 1, Disabled: 0.	smallint	

TABLE 230 CRYPTO_HOST

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
CRYPTO_TARGET_CONTAINER_ID	Foreign key reference to the CRYPTO_TARGET_CONTAINER that contains this host.	int	
VI_NODE_WWN	Node WWN of Virtual Initiator that represents this host.	char	23
VI_PORT_WWN	Port WWN of Virtual Initiator that represents this host.	char	23
HOST_PORT_WWN	Physical (real) host''s Port WWN	char	23
HOST_NODE_WWN	Physical (real) host''s Node WWN	char	23

TABLE 231 CRYPTO_LUN

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
CRYPTO_TARGET_CONTAINER_ID	Foreign key reference to the CRYPTO_TARGET_CONTAINER that contains the host for which these LUNs are configured.	int	
SERIAL_NUMBER	The LUN serial number, used to identify the physical LUN.	varchar	256
ENCRYPTION_STATE	Boolean. • True (1) if LUN is being encrypted. • False (0) if cleartext. The default value is 0.	smallint	
STATUS	Not currently used but left in for possible future use. Replaced by LUN_STATE. The default value is 0.	smallint	
REKEY_INTERVAL	The number of days that data encryption keys should be used before automatically generated a new key. 0 = infinite, i.e., no re-keying.	int	
VOLUME_LABEL_PREFIX	A user-configured string used to construct the Brocade-specific volume label on encrypted tapes. Ignored for disk LUNs.	varchar	256
LAST_REKEY_DATE	The last time a data encryption key was generated for this LUN. REKEY_INTERVAL days after this date, a new key will be generated.	timestamp	
LAST_REKEY_STATUS	The success or failure of the most recent re-keying operation, if any. This field is not currently used, but is left in the hope that FOS will support it in the future. Only valid for disk LUNs. The default value is 0.	smallint	
LAST_REKEY_PROGRESS	Indicates whether a re-key operation is in progress. • 0 = no re-keying in progress. • > 0 = percentage done of re-keying operation in progress. Only valid for disk LUNs. The default value is 0.	smallint	
CURRENT_VOLUME_LABEL	If a tape session is in progress, this is the volume label for the currently mounted tape. Only valid for tape LUNs.	varchar	2048
PRIOR_ENCRYPTION_STATE	Not used. When configuring a new disk LUN, this field indicates whether the LUN is already encrypted (1) or cleartext (0). This information does not need to be persisted. Only valid for disk LUNs.	smallint	
ENCRYPTION_FORMAT	If ENCRYPTION_STATE is true, ENCRYPTION_FORMAT indicates the type of encryption. • 0 = cleartext • 1 = DF-compatible • 2 = native	smallint	
ENCRYPT_EXISTING_DATA	Not used. When configuring a disk LUN that was previously cleartext and is to be encrypted, this property tells the switch whether or not to start a re-keying operation to encrypt the existing LUN data. This property does not need to be persisted.	smallint	

TABLE 231 CRYPTO_LUN (Continued)

Field	Definition	Format	Size
DECRYPT_EXISTING_DATA	Not used. When configuring disk LUN that was previously encrypted and is to become cleartext, this property tells the switch whether or not to start a re-keying operation to decrypt the existing LUN data. This property does not need to be persisted. This feature is no longer supported in FOS.	smallint	
KEY_ID	Hex-encoded binary key vault ID for the current data encryption key for this LUN. This ID may be used to locate the data encryption key in the key vault.	varchar	64
CRYPTO_HOST_ID	Foreign key reference to the CRYPTO_HOST that uses this LUN.	int	
LUN_NUMBER	The Logical Unit Number of the LUN, as seen by the LUNs host. This may be an integer (0 - 65565) or a WWN-format 8-byte hex number.	varchar	64
BLOCK_SIZE	The LUN''s Logical Block Size, in bytes. Only valid for disk LUNs.	int	
TOTAL_BLOCKS	The total number of logical blocks in the LUN. Multiplying BLOCK_SIZE by TOTAL_BLOCKS gives the LUN size in bytes.	int	
LUN_STATE	LUN operational status, such as OK or disabled for various reasons. For possible values see the enum definition in CryptoClientConstants. The default value is 0.	int	
LUN_FLAGS	Bitmap of LUN options. The only option currently used is bit 0 (least significant) which indicates that the LUN must be manually enabled because it has been disabled due to inconsistent metadata detected. The default value is 0.	bigint	
ENCRYPTION_ALGORITHM	Stores the Encryption Algorithm used to encrypt/decrypt data on the LUN	varchar	512
KEY_ID_STATE	State of the Key ID retrieval from Key Vault. The default value is 0.	smallint	
REKEY_SESSION_NUMBER	Unique Rekey session number. The default value is 0.	int	
PERCENTAGE_COMPLETE	Percentage of rekey completion. The default value is 0.	int	
REKEY_ROLE	Rekey Role definition. The default value is 0.	smallint	
CURRENT_LBA	Current Logical Block address for the rekey session in progress. The default value is 0.	int	
LUN_STATE_STRING	Lun state string.	varchar	2048
NEW_LUN	This field specifies that when a LUN is added to its container, indicate that it''s a new LUN to be used in SRDF Configuration. Feature available only from 6.4 release onwards and for RSA key vaults. CryptoLun collector and CryptoLunBean fills in this value. The default value is -1.	smallint	

TABLE 231 CRYPTO_LUN (Continued)

Field	Definition	Format	Size
NEW_LUN_TYPE	This field indicates the role of the LUN configured in the SRDF mode. The values could be R1, R2 or UNKNOWN. Feature available only from 6.4 release onwards and for RSA key vaults. CryptoLuncollector fills in this value.	varchar	64
DISABLE_WRITE_EARLY_ACK	This variable indicates whether write early acknowledgement is enabled (if value is 0) or disabled (if value is 1). The value of this variable is considered only for tape LUNs. This value is applicable only for the FOS 6.3.2 version and above.	smallint	
DISABLE_READ_AHEAD	This variable indicates whether read ahead is enabled (if value is 0) or disabled (if value is 1). The value of this variable is considered only for tape LUNs. This value is applicable only for the FOS 6.3.2 version and above.	smallint	
TIME_LEFT_FOR_AUTO_REKEY	The time left until next auto rekey, starts from the time last key for LUN was generated. This field is not updated every minute in DB. Its value is same as last_rekey_date + re_key_interval. As per current CAL implementation, will get only last_rekey_date when rekey is in progress. Otherwise it will be 0. CAL is providing "time left for auto rekey" attribute, and this value is stored in DB.	bigint	
THIN_PROVISION_LUN	Indicates whether the LUN is a Thin Provisioning LUN or not. The different Thin Provisioning values are 0(Unknown), 1(No), 2(Yes).	int	

TABLE 232 CRYPTO_SWITCH

Field	Definition	Format	Size
SWITCH_ID*	Primary key. The value is the same as the primary key of a record in the VIRTUAL_SWITCH table	int	
ENCRYPTION_GROUP_ID	Foreign key to the ENCRYPTION_GROUP table. Identifies the Encryption Group that this switch belongs to. Null indicates the switch is not part of an Encryption Group.	int	
GROUP_LEADER_POSITION	No longer used. Previously indicated whether this switch is the group leader. Use GROUP_LEADER_ID in the ENCRYPTION_GROUP table instead.	smallint	
TAPE_ENCRYPTION	No longer used. Previously enabled or disabled tape encryption at the switch level. This feature has been removed from Fabric OS. Default value is 0.	smallint	
TAPE_KEY_POLICY	No longer used. Previously used to configure a separate data encryption key per volume or per group. This feature has been removed from Fabric OS. Default value is 0.	smallint	
PRIMARY_VAULT_LINK_STATUS	The status of the link key for the primary key vault. Link keys are used only for NetApp LKM key vaults. For possible values, see the enum definition in the DTO class. Default value is 0.	smallint	

TABLE 232 CRYPTO_SWITCH (Continued)

Field	Definition	Format	Size
BACKUP_VAULT_LINK_STATUS	The status of the link key for the backup key vault. Link keys are used only for NetApp LKM key vaults. For possible values, see the enum definition in the DTO class. Default value is 0.	smallint	
CP_CERTIFICATE	The public key certificate, in PEM format, of the switch's Control Processor module. This certificate is exchanged with other switches to establish secure communication between switches in an Encryption Group.	varchar	4096
KAC_CERTIFICATE	The public key certificate, in PEM format, of the switch''s Key Archive Client module. This certificate is installed on key vaults to establish secure communication between this switch and the key vault. For firmware versions below 7.1.0 it will be in PEM format (encoded) and for firmware versions 7.1.0 and above it will be in p12 format (encoded).	varchar	8192
PRIMARY_VAULT_CONNECTIVITY_STATUS	The status of the network connection between this switch and the primary key vault. For possible values, see the enum definition in the DTO class. Default value is 0.	smallint	
BACKUP_VAULT_CONNECTIVITY_STATUS	The status of the network connection between this switch and the backup key vault. For possible values, see the enum definition in the DTO class. Default value is 0.	smallint	
UN_ERASE_ENABLED	This variable indicates whether LUN Erase feature is enabled or not on the switch. The value 1 means LUN Erase is enabled on the switch.The Value 0 means LUN Erase is not enabled on the switch.	smallint	

TABLE 233 CRYPTO_TARGET_CONTAINER

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
ENCRYPTION_ENGINE_ID	Foreign key reference to the ENCRYPTION_ENGINE that owns this container.	int	
NAME	A user-supplied name for the container.	varchar	64
VT_NODE_WWN	The Node WWN of the Virtual Target that represents the real physical target device.	char	23
VT_PORT_WWN	The Port WWN of the Virtual Target that represents the real physical target device.	char	23
FAILOVER_STATUS	Indicates whether this container''s target is being encrypted by the encryption engine on which the container is configured (value 0) or by another encryption engine in the HA Cluster (value 1). Default value is 0..	smallint	
FAILOVER_STATUS_2	Failover status from the HA Cluster peer.	smallint	

TABLE 233 CRYPTO_TARGET_CONTAINER (Continued)

Field	Definition	Format	Size
DEVICE_STATUS	The physical target storage device operational status when the virtual initiator last attempted to access the target. For possible values, see the enum definition in the DTO class. Default value is 0.	smallint	
DEVICE_TYPE	Indicates whether the target storage device is a disk (0) or tape (1) device. Default value is 0.	smallint	
TARGET_PORT_WWN	The Port WWN of the physical target storage device associated with this container.	char	23
TARGET_NODE_WWN	The Node WWN of the physical target storage device associated with this container.	char	23
CONTAINER_FIELD_DATA	Container metadata information	varchar	256
CONFIGURATION_STATUS	Configuration status. Default value is 0.	smallint	
FRONT_END_N_PORT_NUMBER	Indicates N_Port number where CISCO Fabric will be connected when BES is in AG Mode. Default value is -1.	smallint	

TABLE 234 CUSTOM_FAVORITES_OBJECT_LIST

Field	Definition	Format	Size
FAVORITE_ID	Represents the ID in FAVORITES table	int	
OBJECT_ID	Represents the member's ID of the custom favorites. It can be port/tunnel/EE monitor ID.	int	

TABLE 235 DASHBOARD

Field	Definition	Format	Size
ID	Dashboard ID.	int	
NAME	Name of the dashboard.	varchar	128
DESCRIPTION	Description of the dashboard.	varchar	256
CREATED_BY	References the ID column of the USER_ table. Foreign Key USER_(ID) who created the dashboard. For out of dashboards the column will be 2 to indicate system user.	int	
CREATION_TIME	Time when dashboard was created.	timestamp	
LAST_OPENED_TIME	Time when dashboard was last opened.	timestamp	
INSTALLATION_TYPE	Indicates the dashboard is SAN Only (0) / IP Only (1) / SAN_IP (2).	int	
SHARED	Indicates whether the dashboard is shared. 0 - Not Shared 1 - Shared.	int	
LAYOUT_TYPE	Specifies the layout type of the dashboard. 0-Grid layout 1-Freeform layout.	int	

TABLE 236 DASHBOARD_CANVAS

Field	Definition	Format	Size
ID	Dashboard Canvas ID.	int	
NAME	Name of the Dashboard canvas.	varchar	128
DESCRIPTION	Description of the dashboard canvas.	varchar	512

TABLE 237 DASHBOARD_CANVAS_PREFERENCE

Field	Definition	Format	Size
ID	Dashboard preferences like user ID, Scope ID etc are stored per dashboard.	int	
USER_ID	FK USER_.ID. ID of the user who own the dashboard.	int	
SCOPE_ID	FK USERDEFINED_NETWORK_SCOPE.ID. This value will be populated when user selects the predefined scope.	int	
SCOPE_TYPE	FK SCOPE_TYPE.ID. This value will be populated when user select user defined network scope.	int	
DASHBOARD_ID	FK DASHBOARD.ID. The ID of the dashboard to which the preference is applied.	int	
DASHBOARD_CANVAS_ID	FK DASHBOARD_CANVAS.ID. The ID of the Canvas in which the dashboard is shown	int	
VISIBLE	Visibility of the dashboard. 0 - Not Visible 1 - Visible.	smallint	
TIME_SCOPE	Time Scope of the Dashboard.	int	

TABLE 238 DASHBOARD_PROVIDER

Field	Definition	Format	Size
CLASS_NAME	The fully defined class name of the Provider class. This is stored per widget Provider class.	varchar	128
REFRESH_INTERVAL	Refresh Interval of the Widget in seconds. Default is 5 seconds.	int	
PROVIDER_GROUP	The Group to which the Provider belong to. Similar providers will have same group name.	varchar	128
PROVIDER_ORDER	The order of execution passed to the Job Executor framework. Provider belong to same group will have different order number. Default: 0	int	

TABLE 239 DASHBOARD_TEMPLATE

Field	Definition	Format	Size
ID	Auto generated primary key.	serial	
NAME	Name of the template.	varchar	128
DESCRIPTION	Description of the template.	varchar	512

TABLE 239 DASHBOARD_TEMPLATE (Continued)

Field	Definition	Format	Size
CREATED_BY	Foreign key reference to the ID column of the USER_ table. Indicates the user who created the configuration.	int	
CREATION_TIME	Time when the template was created.	timestamp	
TEMPLATE_TYPE	Template type. 0- SAN only, 1- IP only, 2 â€“ SAN + IP.	int	
IS_SHARED	Whether the template shared to other users. 0- not shared, 1 -shared.	smallint	
LAYOUT	Template Layout specification.	jsonb	
TAGS	Tag names associated with the template.	jsonb	
IS_SYSTEM	System template or custom defined template	smallint	

TABLE 240 DASHBOARD_TEMPLATE_PREFERENCES

Field	Definition	Format
ID	Auto generated primary key.	serial
USER_ID	Foreign key reference to the ID column of the user_ table.	int
DASHBOARD_TEMPLATE_ID	Foreign key reference to the ID column of the DASHBOARD_TEMPLATE table.	int
PREFERENCES	Preferences details.	jsonb

TABLE 241 DASHBOARD_TEMPLATE_PREFERENCES_FILTER_MAP

Field	Definition	Format
DASHBOARD_TEMPLATE_PREFERENCES_ID	Foreign key reference to the ID column of the DASHBOARD_TEMPLATE_PREFERENCES table.	int
FILTER_ID	Foreign key reference to the ID column of filter table.	int

TABLE 242 DASHBOARD_TEMPLATE_PREFERENCES_SCOPE_MAP

Field	Definition	Format
DASHBOARD_TEMPLATE_PREFERENCES_ID	Foreign Key to DASHBOARD_TEMPLATE_PREFERENCES (id).	int
NETWORK_SCOPE_ID	Name of the Network Scope id.	int
NETWORK_SCOPE_TYPE_ID	Foreign key reference to the ID column of network_scope_type table.	int

TABLE 243 DASHBOARD_TEMPLATE_WIDGET_MAP

Field	Definition	Format
ID	Auto generated primary key.	serial

TABLE 243 DASHBOARD_TEMPLATE_WIDGET_MAP

Field	Definition	Format
DASHBOARD_TEMPLATE_ID	Foreign key reference to the ID column of the DASHBOARD_TEMPLATE table.	int
WIDGET_ID	Foreign key reference to the ID column of the dashboard_widget table.	int

TABLE 244 DASHBOARD_WIDGET

Field	Definition	Format	Size
ID	ID of the dashboard widget. Auto incremented.	int	
TITLE	Name of the dashboard widget.	varchar	255
DESCRIPTION	Description of the dashboard widget.	varchar	512
EDITABLE	Indicates whether the widget attributes are editable. 0 - Not Editable, 1 - Editable.	smalint	
CATEGORY	Dashboard widget category. Used for categorizing the widgets based on the type. Possible values are 1 - General, 2 - Performance, 3 - Starlifter (future).	int	
PROVIDER_CLASS_NAME	Provides the mapping between widget and the summary provider. Fully qualified class name of the summary provider implementation for the widget. The class should implement SummaryProvider interface.	varchar	128
UI_PANEL_CLASS_NAME	Provides the mapping between widget and UI panel. Fully qualified class name of the dashboard widget user interface class. The class should extend from AbstractGadget.	varchar	128
SUMMARY_CLASS_NAME	Provides the mapping between widget and the summary. Fully qualified class name of the summary implementation for the widget. The class should implement Summary interface.	varchar	128
time_scope_supported	References the ID column of the DASHBOARD_PROVIDER table. Provides the mapping between widget and the summary provider. Fully qualified class name of the summary provider implementation for the widget. The class should implement SummaryProvider interface.	int	
network_scope_supported	Indicates whether the widget supports Time Scope. 0 - Not Supported 1 - Supported 2 - Partial'	int	
installation_type	Indicates the widgets is SAN Only (0) / IP Only (1) / SAN_IP (2)'	int	
shared_provider	Can the provider be shared? 0 - Not Shared 1 - Shared.	int	
PORT_TYPE	Types of ports to use for widgets. 0 - All Ports, 1 - ISL Ports, 2 - Host Ports, 3 - Storage Ports';	int	

TABLE 244 DASHBOARD_WIDGET (Continued)

Field	Definition	Format	Size
COLUMN_SPAN	Specifies the number of columns the widget will span.	int	
ROW_SPAN	Specifies the number of rows the widget will span.	int	

TABLE 245 DASHBOARD_WIDGET_PREFERENCE

Field	Definition	Format	Size
ID	Auto incremented widget preference ID.	int	
WIDGET_ID	Foreign Key to DASHBOARD_WIDGET(ID).	int	
USER_ID	Foreign Key to USER_ (ID).	int	
DASHBOARD_ID	Foreign Key to DASHBOARD(ID).	int	
VISIBLE	Indicates whether the widget is visible for the user in the dashboard. 0 - Not Visible, 1 - Visible.	smallint	
STATE	State of the widget. Possible values are 0 - Normal, 1 - Maximized, 2 - Collapsed.	int	
WIDTH	Width of the widget.	int	
HEIGHT	Height of the widget.	int	
ROW_INDEX	Row position of the widget. -1 for an out-of-box widget defined but not shown.	int	
COLUMN_INDEX	Column position of the widget. -1 for an out-of-box widget defined but not shown.	int	
CANVAS_ID	Foreign Key to DASHBOARD_CANVAS.ID	int	
INDEX	Specifies the index where the widget has to be placed in the dashboard using Freeform layout.	int	
COLUMN_SPAN	Specifies the number of columns the widget will span.	int	
ROW_SPAN	Specifies the number of rows the widget will span.	int	

TABLE 246 DEFAULT_FAVORITES

Field	Definition	Format	Size
ID	Name of the favorite.	int	
NAME	The topnumber of ports (5,10,15,20).	varchar	64
TOP_N	Types of ports (FC/FCIP/GE) and -End Monitors.	varchar	40
SELECTION_FILTER	The time interval in which the graph is shown.	varchar	40
FROM_TIME	Always null. The default favorite is not customized.	varchar	40
CUSTOM_LAST_VALUE	Always null. The default favorite is not customized.	int	
CUSTOM_TIME_UNIT	Always null. The default favorite is not customized.	varchar	40
CUSTOM_FROM	Always null. The default favorite is not customized.	timestamp	
CUSTOM_TO	The default five minutes granularity.	timestamp	
GRANULARITY	Always null.	varchar	40
THRESHOLD	The measure Tx MBps or Rx MBps based on DEFAULT_FAVORITES.NAME	int	

TABLE 246 DEFAULT_FAVORITES (Continued)

Field	Definition	Format	Size
MAIN_MEASURE	The Additional measures based on the FAVORITE.MAIN_MEASURE	varchar	40
ADDITIONAL_MEASURE	The Additional measures based on the FAVORITE.MAIN_MEASURE	int	

TABLE 247 DEFAULT_WIDGET_PREFERENCE

Field	Definition	Format	Size
ID	Auto incremented Dashboard Widget Preference ID.	int	
dashboard_id	Foreign Key to DASHBOARD(ID).	int	
widget_id	Foreign Key to DASHBOARD_WIDGET(ID).	int	
installation_type	Indicates the widgets is SAN Only (0) / IP Only (1) / SAN_IP (2).	int	
visible	Indicates whether the widget is visible for the user in the dashboard. 0 - Not Visible, 1 - Visible.	int	
state	State of the widget. Possible values are 0 - Normal, 1 - Maximized, 2 - Collapsed.	int	
width	Width of the widget.	int	
height	Height of the widget.	int	
row_index	Row position of the widget. -1 for an out-of-box widget defined but not shown.	int	
column_index	Column position of the widget. -1 for an out-of-box widget defined but not shown.	int	
INDEX	Specifies the index where the widget has to be placed in the dashboard using Freeform layout.	int	
COLUMN_SPAN	Specifies the number of columns the widget will span	int	
ROW_SPAN	Specifies the number of rows the widget will span.	int	

TABLE 248 DEPLOYMENT_CONFIGURATION

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
NAME	Name of the configuration	varchar	255
CONFIGURATION_TYPE	Indentifies the save configuration type. • 1 - Not applicable • 1 - Running • 2 - Startup • 3 - Running & Startup	smallint	
DEPLOY_OPTION	Identifies the deployment options. • 1-Deploy Now • 2-Save & Deploy • 3-Save deployment only • 4-Scheduled	smallint	

TABLE 248 DEPLOYMENT_CONFIGURATION (Continued)

Field	Definition	Format	Size
DEPLOYMENT_HANDLER_ID	Foreign Key references DEPLOYMENT_HANDLER (ID). Identifies the handler to use for the configuration	int	
SCHEDULE_ENABLED	1 indicates that the schedule is applied to the configuration	smallint	
SNAPSHOT_ENABLED	1 indicates that snapshot is applied to the configuration	smallint	
CLI_TEMPLATE_ID	Identifies the CLI template details. -1 if SNAPSHOT_ENABLED is False	int	
SNAPSHOT_SETTING	Identifies the setting type • 1-Presnapshot • 2-Postsnapshot • 3-Pre & Post snapshot • -1 if SNAPSHOT_ENABLED is False	smallint	
POST_DEPLOYMENT_DELAY	Post deployment delay in seconds	int	
CREATED_BY	User who created the configuration	varchar	255
LAST_MODIFIED_BY	User who last modified the configuration. When the configuration is first created	varchar	255
MANAGEMENT_FLAG	True if deployment should be managed by Deployment Manager Module and this will be displayed in Deployment Manager UI	smallint	
DESCRIPTION	Used to describe the deployment configuration	varchar	255
IS_FCP_POLICY	0 - Indicates policy monitor created through configuration policy manger 1 - Indicates policy monitor created through Fabric configuration policy(FCP)	smallint	

TABLE 249 DEPLOYMENT_HANDLER

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
MODULE	Identifies the unique deployment module	varchar	64
SUB_MODULE	Identifies sub-module	varchar	64
MODULE_DISPLAYNAME	Display text for module name.	varchar	128
HANDLER_CLASS	Fully qualifies name of handler class for the module. This class has to implement <DeploymentHandler> interface	varchar	255
CLIENT_ACTION_HANDLER_CLASS	Fully qualifies module-specific client class which implements <DeploymentDelegateActionsHandler> interface. Framework will delegate edit, duplicate, delete actions to this class	varchar	255

TABLE 249 DEPLOYMENT_HANDLER (Continued)

Field	Definition	Format	Size
PRIVILEGE_ID	Comma separated privilege IDs	varchar	64
MODULE_DISPLAYNAME	Dsplay text for module name.	varchar	128

TABLE 250 DEPLOYMENT_PRODUCT_STATUS

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
DEPLOYMENT_STATUS_ID	Foreign Key references DEPLOYMENT_STATUS (id). Identifies the execution cycle for the deployment.	int	
DEPLOYMENT_TIME	Time when this product deployment occurred.	timestamp	
PRODUCT_ID	This record will be per product. Hence this will have the id of the product.	int	
PRODUCT_TYPE_ID	Foreign Key references TARGET_TYPE(id). This identifies the PRODUCT_ID. (Whether it is switch, device, etc).	int	
STATUS	Indicated the product deployment status 1-Aborted 2-Succesful 3-Partial Failure 4-Failed	smallint	
MESSAGE	Message to be displayed in the report.	txt	
ERROR_CODE	Error code, can be used for i18n	int	

TABLE 251 DEPLOYMENT_REPORT_TEMPLATE

Field	Definition	Format	Size
DEPLOYMENT_HANDLER_ID	Foreign Key references DEPLOYMENT_HANDLER(id).	int	
HEADER	Stores header content of deployment report. This could be plain text or HTML or XML	text	
FOOTER	Stores the footer content of deployment report. This could be plain text or HTML or XML.	text	

TABLE 252 DEPLOYMENT_STATUS

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
DEPLOYMENT_CONFIGURATION_ID	Foreign Key References DEPLOYMENT_CONFIGURATION(id). Identifies the deployment configuration	int	
DEPLOYMENT_TIME	Start Time of the deployment (UTC)	timestamp with time zone	

TABLE 252 DEPLOYMENT_STATUS (Continued)

Field	Definition	Format	Size
STATUS	Overall status of the deployment. 1-In Progress 2-Success 3-Failure 4-Partially failed	smallint	
DEPLOYED_BY	User who deployed the configuration	varchar	255
STATUS_MESSAGE	Overall Success/Failure status description	txt	
TRIGGER_SOURCE	Maintains the source from which this deployment was triggered such as Event Action <Event policy name>, Manual and Scheduled etc.	varchar	128

TABLE 253 DEPLOYMENT_TARGET_MAP

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
DEPLOYMENT_CONFIGURATION_ID	Foreign Key References DEPLOYMENT_CONFIGURATION (id) Identifies the deployment configuration this row is applied	int	
TARGET_ID	Identifies the target. It will NOT have mapping to any product table like device, etc	varchar	255
TARGET_TYPE_ID	Foreign Key references TARGET_TYPE (id) Identifies the target type	int	
TARGET_PARENT_ID	Identifies the parent of the target. If, switch, device, port group, device group it will be same as target id. If it is port/interfaces the parent id will be the switch id	int	

TABLE 254 DEVICE

Field	Definition	Format	Size
DEVICE_ID	Primary key for this table.	int	
IP_ADDRESS	IP address of this device.	varchar	255
ALIAS_NAME	Device alias name.	varchar	512
HOST_NAME	Best matching host name obtained through the device IP address.	varchar	512
SYS_NAME	An administratively-assigned name for this device.	varchar	255
SYS_CONTACT	The textual identification of the contact person for this device, together with information on how to contact this person.	varchar	255
DESCRIPTION	A textual description of the device.	varchar	512
SYS_LOCATION	The physical location of this device.	varchar	255
COMMUNITY_STR_GET	SNMP GET community string to query the device.	varchar	512

TABLE 254 DEVICE (Continued)

Field	Definition	Format	Size
COMMUNITY_STR_SET	SNMP SET community string of this device.	varchar	512
SYS_OID	The vendor's authoritative identification of this device ie., System Object Identifier.	varchar	255
SUPER_USER_PASSWORD	Super user password configured in the device.	varchar	512
TABLE_SUBTYPE	Device table subtype defined by INM BizObject framework.	varchar	32
LOCAL_USER_NAME	Local user name configured in the device for CLI access.	varchar	512
LOCAL_PASSWORD	Password to access the telnet interface.	varchar	512
TELNET_PASSWORD	Password to access the Telnet interface.	varchar	512
RADIUS_USER_NAME	User name for RADIUS access.	varchar	512
RADIUS_PASSWORD	Password for RADIUS access.	varchar	512
TAC_USER_NAME	User name for TACACS access.	varchar	512
TAC_PASSWORD	Password for TACACS access.	varchar	512
TACPLUS_USER_NAME	User name for TACACS+ access.	varchar	512
TACPLUS_PASSWORD	Password for TACACS+ access.	varchar	512
IS_ROUTER	Flag to identify whether the device is router or not.	num	(1,0)
IS_SLB	Flag to identify whether the device supports server load balancing or not.	num	(1,0)
FIRST_SEEN_TIME		varchar	64
LAST_SEEN_TIME	Time when the device is getting discovered by recent collection.	varchar	64
LAST_PROBE_TIME		varchar	64
LAST_PROBE_STATUS		varchar	64
IS_SFLOW_CAPABLE	Flag to identify whether the device is SFlow capable or not.	num	(1,0)
SNMPV3_RO_AUTH_TYPE	SNMP V3 read only authentication type.	varchar	1
SNMPV3_RO_AUTH_USERNAME	SNMP V3 read only authentication user name.	varchar	512
SNMPV3_RO_AUTH_PASSWORD	SNMP V3 read only authentication password.	varchar	512
SNMPV3_RO_PRIV_PROTOCOL	SNMP V3 read only privacy protocol.	varchar	1
SNMPV3_RO_PRIV_PASSWORD	SNMP V3 read only privacy password.	varchar	512
SNMPV3_RW_AUTH_TYPE	SNMP V3 read write authentication type.	varchar	1
SNMPV3_RW_AUTH_USERNAME	SNMP V3 read write authentication user name.	varchar	512
SNMPV3_RW_AUTH_PASSWORD	SNMP V3 read write authentication password.	varchar	512

TABLE 254 DEVICE (Continued)

Field	Definition	Format	Size
SNMPV3_RW_PRIV_PROTOCOL	SNMP V3 read write privacy protocol.	varchar	1
SNMPV3_RW_PRIV_PASSWORD	SNMP V3 read write privacy password.	varchar	512
LOCAL_USERNAME_PORT_CFG	Agent user name configured in device used for port configuration.	varchar	512
LOCAL_PASSWORD_PORT_CFG	Agent password configured in device used for port configuration.	varchar	512
LOCAL_USERNAME_READ_ONLY	Local user name for read only access.	varchar	512
LOCAL_PASSWORD_READ_ONLY	Local password for read only access.	varchar	512
RADIUS_USERNAME_PORT_CFG	RADIUS user name configured in device used for port configuration.	varchar	512
RADIUS_PASSWORD_PORT_CFG	RADIUS password configured in device used for port configuration.	varchar	512
RADIUS_USERNAME_READ_ONLY	RADIUS user name configured in device used for read only access.	varchar	512
RADIUS_PASSWORD_READ_ONLY	RADIUS password configured in device used for read only access.	varchar	512
TAC_USERNAME_PORT_CFG	TACACS username for port configuration.	varchar	512
TAC_PASSWORD_PORT_CFG	TACACS password for port configuration.	varchar	512
TAC_USERNAME_READ_ONLY	TACACS username for read only access.	varchar	512
TAC_PASSWORD_READ_ONLY	TACACS password for read only access.	varchar	512
TACPLUS_USERNAME_PORT_CFG	TACACS+ username for port configuration.	varchar	512
TACPLUS_PASSWORD_PORT_CFG	TACACS+ password for port configuration.	varchar	512
TACPLUS_USERNAME_READ_ONLY	TACACS+ username for read only access.	varchar	512
TACPLUS_PASSWORD_READ_ONLY	TACACS+ password for read only access.	varchar	512
ENABLE_PASSWORD_PORT_CFG	Enable password configured in device used for port configuration.	varchar	512
ENABLE_PASSWORD_READ_ONLY	Enable password for read only access.	varchar	512
ADMIN_STATUS	Device admin status.	smallint	
ADMIN_STATUS_DURATION	Time duration of the admin status without any change.	int	
ADMIN_STATUS_LAST_UPDATED	Time when the admin status updated last.	bigint	
MEMO_LAST_UPDATED	Time when the memo got updated last.	bigint	
MEMO	Memo updated by the user for this device.	varchar	4096

TABLE 254 DEVICE (Continued)

Field	Definition	Format	Size
TACPLUS_ENABLE_USERNAME	TACACS+ enable user name.	varchar	512
TACPLUS_ENABLE_PASSWORD	TACACS+ enable password.	varchar	512
OPER_STATUS	Device operational status.	smallint	
OPER_STATUS_LAST_UPDATED	Time when the device operational status got updated recently.	bigint	
LLDP_CHASSIS_ID_SUBTYPE	Chassis ID subtype returned by lldp MIB.	smallint	
LLDP_CHASSIS_ID	Chassis ID returned by lldp MIB.	bytea	
IS_FDP_ENABLED	Flag to identify whether Foundry Discovery Protocol is enabled or not.	num	(1,0)
IS_CDP_ENABLED	Flag to identify whether Cisco Discovery Protocol is enabled or not.	num	(1,0)
VENDOR	Vendor of this device.	varchar	64
IS_FOUNDRY	Flag to identify whether the device is Foundry product or not.	num	(1,0)
MANAGED_ELEMENT_ID	A unique managed element ID for this IP switch. Also a foreign key reference to the MANAGED_ELEMENT table.	int	
NODE_WWN	The managed element node WWN if one exists, or null/empty otherwise.	varchar	23
SYSLOG_REGISTERED	This flag is to indicate whether the device is registered DCM as its syslog destination server. • 0 indicates not registered. • 1 indicates registered.	num	1
TRAP_REGISTERED	This flag is to indicate whether the device is registered DCM as its SNMP trap destination server. • 0 indicates not registered. • 1 indicates registered.	num	1
PORT_COUNT	Record the number of presented physical ports on the device.	int	
SERIAL_NUMBER	Record the serial number of the device. If there is no serial number, an empty string will be stored.	varchar	32
CATEGORY	This flag is to classify the device category • 0 is for unknown • 1 is for fixed configuration device • 2 is for chassis device • 3 is for stack device (logical)	int	

TABLE 254 DEVICE (Continued)

Field	Definition	Format	Size
MPLS_MANAGE_STATE	This flag is to classify the device mpls managing state • 0 indicates unknown state for catching all • 1 indicates not applicable; if the IP Product is not XMR/MLX, it will be set to this value. • 2 indicates MPLS unmanaged state; in PP or PPE edition, XMR/MLX product will be set to this value. • 3 indicates MPLS managed state; only XMR/MLX product in EE edition will be set to this value.	int	
LICENSE_PORT_COUNT	It records the number of the ports that presented in the device.	int	
SUB_CATEGORY	This column is used to classify device sub category for DCB switches. Column helps to identify whether the DCB switch is an Elara/Frisco or DCX with Europa blade etc. Value 0 indicates that this is a pure IP device and hence that is the default value. Value 1 indicates that this is an Elara DCB device. The values will be populated by the DCB collector during the discovery of the DCB switch.	int	
LICENSED_FEATURES	This column is used to persist the feature based software licenses existing on the device. This represents bitmask as an integer value, where each bit represents a unique feature.	int	
IS_DCB_SWITCH	This column is used to flag whether the device is a DCB Switch or not. Value 0 indicates that this is not a DCB switch device and hence that is the default value and value 1 indicates that this is a DCB device. The values will be populated by the DCB collector during the discovery of the DCB switch.	num	(1,0)
PRODUCT_FAMILY	Record the product family as "BI", "EI", "FGS/FLS/STK". Make it string field to accommodate dynamic group database search.	varchar	32
NETCONF_TRANSPORT	The transport protocol used to connect to this device through Netconf. Possible values are: • 0=Netconf not supported by this device • 1=SSH • 2=HTTPS • 3=HTTP • 4=WING_HTTPS • 5=WING_HTTP	smallint	
POE_CAPABLE	The POE capability of device. Possible values are: • 0 = POE is not supported by this device • 1 = POE is supported with IEEE 802.3af standard by this device • 2 = POE plus is supported with IEEE 802.3at standard by this device	smallint	

TABLE 254 DEVICE (Continued)

Field	Definition	Format	Size
CLUSTER_MODE	This column is used to determine whether VCS Cluster is in Standalone mode or Cluster mode. The values will be populated by the VCS collector during the discovery of the VCS switch. The default value of -1 means that its a non VCS device. Following Enum will be defined as NON_VCS(-1), STANDALONE(0), CLUSTER(1).	smallint	
CLUSTER_TYPE	This column is used to determine whether VCS is in Fabric Cluster or Logical Chassis. The values are populated by the VCS collector during the discovery of the VCS switch. The default value -1 means that its a non-VCS device. Following are the values and their types: • 0 - Unknown • 1 - Standalone • 2 - Fabric Cluster • 3 - Logical Chassis	smallint	
IS_VCS_CAPABLE	This column is used to determine whether the device is a VCS device. The default value 0 means that the device is not VCS capable and value 1 means that the device is VCS capable.	smallint	
TRACKING	This column helps to identify that whether the device is left/joined the cluster membership. The value will be a bit mask value where 2^1 will be treated as left, 2^2 treated as joined. The default value will be -1.	smallint	
VCS_ID	This column is used to store the VCS ID of the device. The value will be populated by the NosSwitchAssetCollector during the discovery of the VCS Cluster. The non zero value will be stored as VCS ID. Default value is -1.	smallint	
VCS_LICENSED	Indicates whether the cluster device has VCS license or not. Possible values are 0 for not applicable, 1 for licensed, 2 for not licensed. 0 is default. Clusters with 2 or less nodes will have value of 0 as all those clusters are automatically licensed. Clusters with 3 or more nodes will have values 1 or 2 depending on whether the license was acquired or not.		
RBRIDGE_ID	This column is used to store the Rbridge ID of the device. The value will be populated by the NosSwitchAssetCollector during the discovery of the VCS member. The non zero value will be stored as Rbridge ID. Default value is -1.	int	
IS_PRINCIPAL_SWITCH	This column is used to determine whether VCS member is a Principal switch or not. Value 1 indicates that this is a principal switch and 0 indicates that this not a Principal switch. The values will be populated by the VCS collector during the discovery of the VCS switch. The default value of 0 means that its a principal switch.	smallint	

TABLE 254 DEVICE (Continued)

Field	Definition	Format	Size
IS_NETCONF_REACHABLE	This column is used to determine whether the device is netconf reachable. The value will be populated by the NosSwitchAssetCollector. The value of 0 means not reachable, 1 means reachable port and -1 means unknown status. Default value is -1	smallint	
FABRIC_WATCH_STATUS	Switch status based on components.	smallint	
FABRIC_WATCH_STATUS_REASON	Component reason for switch status.	varchar	1028
MAC_ADDRESS	The mac address to identify the wireless controller or AP. This will be empty string for all other devices.	varchar	64
MANAGED_AP_COUNT	Its the number of APs that the controller managed.	int	
CONTROLLER_CLUSTER_MODE	Cluster mode of the controller: Active, Standby and None. -1 : NA, 0 : None, 1 : Active, 2 : Standby.	int	
CONTROLLER_CLUSTER_NAME	This is controller cluster name.	varchar	65
CONTROLLER_CLUSTER_PEER_IP	IP addresses of the controller cluster peer.	varchar	128
WIRELESS_TYPE	To filter the APs from the product. 0 : NonAP, 1 : managed Brocade branded AP, 2 : standalone Brocade branded AP.	int	
BRIEF_PRODUCT_FAMILY	Shorter name for the product family.	varchar	32
USER_DEFINED_VALUE_1	User defined value used for product.	varchar	256
USER_DEFINED_VALUE_2	User defined value used for product.	varchar	256
USER_DEFINED_VALUE_3	User defined value used for product.	varchar	256
CLUSTER_MEMBER_STATE	Indicates the state of the member in Fabric Cluster and logical chassis. States can be Online, Offline, Rejoining etc.. For all other devices this column will be empty.	varchar	64
MODE	Denotes the mode of the device. 1 denotes it is in non-AG mode and 2 denotes that the device is in AG mode.-1 denotes not applicable and will be set to all non NOS devices.	int	
USER_IP_ADDRESS	'This column is used to store the IP address given by the user during discovery time. For profile based discovery, we will populate the seed switch IP address selected by the system.	varchar	255
VCS_VF_ENABLED	This represents where the device has VCS virtual fabric support or not. The possible values are -1,0 and 1. Default value -1 applicable for all IOS device and NOS firmware version less than 4.1.The value 0 will be set if the device supports VCS VF and its disabled. The value 1 will be set if its enabled on the device.	int	

TABLE 254 DEVICE (Continued)

Field	Definition	Format	Size
FEATURES_ENABLED	Holds bit mask which represents the features that are enabled in this device. Each bit would represents a specific feature.	int	
VIRTUAL	Represents the given device is software based device(virtual) or classic hardware based device(physical). For eg., Virtual ADX and Vyatta are software based Brocade platforms. 0 - Physical, 1- Virtual	smallint	
SNMP_VERSION	SNMP version in which the device is discovered. The possible values are: v1, v2, v2c, v3.	varchar	4

TABLE 255 DEVICE_CONNECTION

Field	Definition	Format	Size
ID	The primary key.	int	
FABRIC_ID	ID of the fabric which the device port is connected to. It refers the ID column of FABRIC table.	int	
DEVICE_PORT_ID	ID of end device port. It refers ID field of DEVICE_PORT table.	int	
SWITCH_PORT_ID	ID of switch port which end device port is connected. In case device port is connected through AG, this port refer the switch port which AG is connected. It refers to the ID field of SWITCH_PORT table.	int	
AG_PORT_ID	In case of AG, this would refer to F Port of the AG which end device port is connected. If the device port is directly connected to switch port or NPIV connection then it would be -1. It refers to the ID field of SWITCH_PORT table for the access gateway switches.	int	
CREATION_TIME	Time at which the device connection record is created.	timestamp	
LAST_UPDATED_TIME	Time at which connection properties are modified for this record.	timestamp	
MISSING	Indicates if the device connection is missing or not.	int	
MISSING_TIME	Time from which the device connection has been missing.	timestamp	
TRUSTED	Indicates if the device connection is trusted or not.	int	
QSFP_UNIT_NUMBER	Refers to the QSFP unit number for ports that have QSFP. Default value is -1 for all other ports. In case of wedge, the unit number for ports 48-51 will be 0, for 52-55 will be 1 and so on.	int	

TABLE 256 DEVICE_ENCLOSURE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of the Device enclosure.	varchar	256

TABLE 256 DEVICE_ENCLOSURE (Continued)

Field	Definition	Format	Size
TYPE	Type of Device enclosure - Storage Array/Server.	varchar	32
ICON	Type of Icon.	int	
OS	Operating System.	varchar	256
APPLICATIONS	Application which created device enclosure.	varchar	256
DEPARTMENT	Department using this device enclosure.	varchar	256
CONTACT	Contact person details.	varchar	256
LOCATION	Location of physical setup.	varchar	256
DESCRIPTION	Description if any.	varchar	256
COMMENT_	Comments if any.	varchar	256
IP_ADDRESS	IP Address if assigned by user.	varchar	128
VENDOR	Vendor name.	varchar	256
MODEL	Device enclosure Model.	varchar	256
SERIAL_NUMBER	Serial Number given for the entity.	varchar	256
FIRMWARE	Firmware running on the device which is not applicable for device enclosure logical entity.	varchar	256
USER_DEFINED_VALUE 1	User-defined custom value.	varchar	256
USER_DEFINED_VALUE 2	User-defined custom value.	varchar	256
USER_DEFINED_VALUE 3	User-defined custom value.	varchar	256
HCM_AGENT_VERSION	Version of the HCM agent running on the host	varchar	32
OS_VERSION	Operating system version for the enclosure	varchar	256
CREATED_BY	Module which created this enclosure: 0->Manual, 1->HBA, 2->VM, 3->AUTO_ENCLOSURE. Default value is 0.	int	
TRACK_CHANGES	Flag to enable/disable tracking. Default value is 0.	smallint	
LAST_UPDATE_TIME	Last time at which the host information was updated.	timestamp	
LAST_UPDATE_MODUL E	Module which updated the host information.	smallint	
TRUSTED	Flag to mark the enclosure trusted. Default value is 0.	smallint	
CREATION_TIME	Time when enclosure was created. Default is 'now()'.	timestamp	
MISSING	Flag to indicate missing enclosure. Default value is 0.	smallint	
MISSING_TIME	Time when the enclosure is found to be missing.	timestamp	
HOST_NAME	Host Name corresponding to the Device Enclsoure.	varchar	256

TABLE 256 DEVICE_ENCLOSURE (Continued)

Field	Definition	Format	Size
SYSLOG_REGISTERED	SysLog flag that indicates if syslog has been enabled or not.	smallint	
VIRTUALIZATION	If this enclosure is a host, this column indicates whether the host is running a virtualization hypervisor. 0 = unknown 1 = no supported hypervisor present 2 = VMware ESX 3 = Microsoft Hyper-V. Default value is 0.	smallint	
MANAGED_ELEMENT_ID	A unique managed element ID for a managed host.If the device enclosure is manually created (does not represent a managed host) then the field is null. Also a foreign key reference to the MANAGED_ELEMENT table.	int	
MANAGED_BY	1 - Manual - (user created not managed condition) - Default. 2 - Host Adapter 3 - VMM 4.Both Host Adapter and VMM';	smallint	
QUEUE_DEPTH	Queue Depth can be used to control FCP exchange resource allocation. Queue depth can range from 0 to 254 and default value is 32.	int	
AUTO_DELETE_ALLOWED	0 - not allowed, 1 - allowed. This column represents whether the enclosure can be deleted by HOST and VM Discovery. If the value is 1, the enclosure can be deleted. For fabric discovery - value must be set to 1, For Host and VM Discovery - value must be set to 0, For Manual - value is set based on user input.	smallint	

TABLE 257 DEVICE_ENCLOSURE_MEMBER

Field	Definition	Format	Size
ENCLOSURE_ID*	DEVICE_ENCLOSURE table ID.	int	
DEVICE_PORT_WWN*	WWN Of Device Port.	char	23
DEVICE_PORT_ID	Device_Port table ID.	int	

TABLE 258 DEVICE_FDMI_DETAILS

Field	Definition	Format	Size
DEVICE_NODE_ID	Device node id for the FDMI device node. This column refers to the device_node tables primary key	int	
SERIAL_NUMBER	Holds the serial number of the device available via FDMI	varchar	128
FIRMWARE_VERSION	Holds the firmware version of the device available via FDMI ex: 2.1.0.2	varchar	64
DRIVER_VERSION	Holds the driver version of the device available via FDMI, ex: 2.1.0.2	varchar	64

TABLE 258 DEVICE_FDMI_DETAILS (Continued)

Field	Definition	Format	Size
MANUFACTURER	Holds the manufacturer of the device available via FDMI, ex : Brocade	varchar	64
MODEL	Holds the model of the device available via FDMI, ex : Brocade-825	varchar	64
HARDWARE_VERSION	Holds the hardware version of the device available via FDMI, ex: Rev-C	varchar	64
MODEL_DESCRIPTION	Holds the model description of the device available via FDMI, ex : Brocade-825	varchar	64
NODE_NAME	Holds the node name of the device available via FDMI, ex : 20:00:00:05:1e:7c:64:94	varchar	64

TABLE 259 DEVICE_GROUP

Field	Definition	Format	Size
DEVICE_GROUP_ID	Primary key for this table.	int	
NAME	Name of this device group.	varchar	128
USER_ID	User ID corresponds to the user who created the device.	int	
DESCRIPTION	Device group description.	varchar	255
IS_PUBLIC	Flag to identify whether this group is shared across users.	num	(1,0)
IS_INTERNAL	Flag to identify this group is internal.	num	(1,0)
TABLE_SUBTYPE	Table subtype defined by BizObject framework	varchar	32
IS_AP_GROUP	Flag to identify whether this group is access point device group.	num	(1,0)
IS_SENSOR_GROUP	Flag to identify whether this group is sensor device group.	num	(1,0)
VIEW_MASK	Flag to decide whether to show the device group in topology or not.	num	(1,0)
GROUP_TYPE	This flag is to classify the device group type: • 0 is the default and reserved for internal temporary group • 1 is for System Device Group • 2 is for MPLS System Device Group • 3 is for User Defined Device Group	int	

TABLE 260 DEVICE_GROUP_ENTRY

Field	Definition	Format	Size
DEVICE_GROUP_ID	Database ID of the DEVICE_GROUP instance which the device is member of.	int	
DEVICE_GROUP_ENTRY_ID	Unique database auto generated identifier.	int	
DEVICE_ID	Database ID of the member DEVICE.	int	

TABLE 261 DEVICE_GROUP

Field	Definition	Format	Size
DEVICE_GROUP_ID	Primary key for this table.	int	
NAME	Name of this device group.	varchar	128
USER_ID	User ID corresponds to the user who created the device.	int	
DESCRIPTION	Device group description.	varchar	255
IS_PUBLIC	Flag to identify whether this group is shared across users.	num	(1,0)
IS_INTERNAL	Flag to identify this group is internal.	num	(1,0)
TABLE_SUBTYPE	Table subtype defined by BizObject framework	varchar	32
IS_AP_GROUP	Flag to identify whether this group is access point device group.	num	(1,0)
IS_SENSOR_GROUP	Flag to identify whether this group is sensor device group.	num	(1,0)
VIEW_MASK	Flag to decide whether to show the device group in topology or not.	num	(1,0)
GROUP_TYPE	This flag is to classify the device group type: • 0 is the default and reserved for internal temporary group • 1 is for System Device Group • 2 is for MPLS System Device Group • 3 is for User Defined Device Group	int	

TABLE 262 DEVICE_MAC_GROUP_MAPPING

Field	Definition	Format	Size
MAC_GROUP_DB_ID	Foreign Key Reference to the MAC_GROUP table. Part of Primary key.	int	
DEVICE_ID	Foreign Key reference to DEVICE table. Part of Primary key;	int	

TABLE 263 DEVICE_NODE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
FABRIC_ID	Fabric DB ID to which this device node belongs.	int	
WWN	Device node WWN.	char	23
TYPE	Initiator or target or both or unknown. The possible values are Initiator, Target, Initiator+Target, Unknown(Initiator or Target)	varchar	32
DEVICE_TYPE	0 = physical 1 = virtual 2 = NPV 3 = iSCSI 4 = both physical & virtual	smallint	
SYMBOLIC_NAME	Device node symbolic name.	varchar	256

TABLE 263 DEVICE_NODE (Continued)

Field	Definition	Format	Size
FDMI_HOST_NAME	Device node FDMI host name.	varchar	128
VENDOR	Device node vendor.	varchar	64
CAPABILITY_		varchar	16
TRUSTED	1 = the node is trusted for "fabric tracking. Default value is 0.	smallint	
CREATION_TIME	Timestamp when the record is created by the Management application server.	timestamp	
MISSING	1 = the device node is missing from the fabric. Default value is 0.	smallint	
MISSING_TIME	Time when the device node missed.	timestamp	
PROXY_DEVICE	One of the device ports of this device node has translated domain. That device port is set as the Proxy Device and this Device Node is treated as virtual by assigning a value of 1 to this field. Default value is 0.	smallint	
AG	1 = the device node is actually an AG connected to a switch in the fabric. Default value is 0.	smallint	
PREVIOUS_MISSING_STATE	Default value is 0.	smallint	
SIMULATED	Indicates whether the device is simulated by flow vision or not: 0 - Not simulated 1 - Simulated by flow vision	smallint	

TABLE 264 DEVICE_PORT

Field	Definition	Format	Size
NODE_ID	Reference to the ID of the Device Node of which this device port is a part of.	int	
DOMAIN_ID	Stores the Domain ID of the switch to which this device port is connected to.	int	
WWN	Stores the Device Port WWN	char	23
SWITCH_PORT_WWN	Stores the switch port wwn to which this device port is physically connected to. However If the device is connected to an AG, this will contain the switch port WWN till the AG impact is applied by the application. If AG impact fails to be applied this will continue to have the switch port wwn instead of the AG port wwn.	char	23
NUMBER	Stores the port number of this device port.	smallint	
PORT_ID	Stores the FDMI host name.	varchar	6
TYPE	Stores the Vendor of this device.	varchar	32
SYMBOLIC NAME	Stores the Symbolic Name.	varchar	256
FC4_TYPE		varchar	64

TABLE 264 DEVICE_PORT (Continued)

Field	Definition	Format	Size
COS	Stores the Class of Service.	varchar	16
IP_PORT		varchar	63
HARDWARE_ADDRESS	Stores the Hardware Address.	varchar	32
TRUSTED	Denotes if the device port is trusted or not.	smallint	
CREATION_TIME	The creation time of this record.	timestamp	
MISSING	Denotes if this device port is missing or not.	smallint	
MISSING_TIME	Denotes the time from which the device port is missing. Applicable only if the device is missing.	timestamp	
NPV_PHYSICAL	Denotes if this is physical device port or a logical NPIV port.	smallint	
EDGE_SWITCH_PORT_WWN	EDGE_SWITCH_PORT_WWN will be the same as the SWITCH_PORT_WWN except in the case of devices behind the AG. This field will be updated by the name server info collector, added for the feature support of AG WWN N port mapping. This is a null able field. It is used to determine which mapping is used by the AG.	char	23
LOGGED_TO_AG	Indicates if the device is connected with an AG. Not null field and default value is 0, device not connected to AG	smallint	
AG_NODE_WWN	If the device is connected with an AG, the AG switch WWN will be populated. Not null field and default value is empty	char	23
AG_N_PORT_WWN	If the device is connected with an AG, N-Port WWN of AG which is connected to switch will be populated from the N2F and N2WWN map	char	23
MISSING_REASON	The device missing reason.	varchar	1024

TABLE 265 DEVICE_PORT_GIGE_PORT_LINK

Field	Definition	Format	Size
DEVICE_PORT_ID	The primary key of the DevicePort	int	
GIGE_PORT_ID	The primary key of the GigEPort.	int	
DIRECT_ATTACH	Indicates whether the device port is directly attached to the CEE 10G TE port.	smallint	
VIRTUAL_FCOE_PORT_ID	The value of virtual_fcoe_port_id in the Device_Port_Gige_Port_Link table is applicable only for NOS devices. For FOS devices, the virtual_fcoe_port_id value, will be null, as currently in the Management application that mapping data is not collected. Hence the default value is null.	int	
LAG_ID	LAG interface ID which associates port channel with end device. This will be null if device port is associated with physical gige port.	int	

TABLE 266 DEVICE_PORT_MAC_ADDRESS_MAP

Field	Definition	Format	Size
DEVICE_PORT_ID	The primary key of the device port	int	
MAC_ADDRESS	Mac address of the device	varchar	64

TABLE 267 DISK_USAGE

Field	Definition	Format	Size
ID	Primary key of the table. Autogenerated.	int	
TIME_THRESHOLD_CROSSED	Holds the timestamp at which the disk space was analyzed and found to have crossed the threshold (both low to high and vice versa).	timestamp	
MEMORY_UTILIZATION	Holds the disk usage as a percentage. Value varies from 0 to 100.	double precision	
THRESHOLD_TYPE	Denotes whether disk space usage crosses above or below threshold limit. 1 if it is goes above threshold, 0 if it goes below threshold (in previous instance it was above threshold).	smallint	
FREE_SPACE	Holds the free disk space at the particular time in bytes.	bigint	
TOTAL_SPACE	Holds the total disk space at the particular time in bytes.	bigint	

TABLE 268 ENCRYPTION_ENGINE

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SWITCH_ID	Foreign key reference to both the VIRTUAL_SWITCH table and the CRYPTO_SWITCH table (both switch tables use the same primary key values). Identifies the switch that contains this encryption engine.	int	
SLOT_NUMBER	For chassis switches, the slot or blade that contains the encryption engine. Always 0 for pizza-box switches with a single embedded encryption engine.	smallint	
STATUS	Not used. Previously used to indicate the engine''s operational status. Replaced by EE_STATE. The default value is 0.	smallint	
HA_CLUSTER_ID	Foreign key reference to an HA_CLUSTER record. Identifies the HA Cluster that this engine belongs to. Null if this engine does not belong to an HA Cluster.	int	
SYSTEM_CARD_STATUS	Indicates whether a System Card is currently inserted in the Encryption Engine, and whether the card is valid or not. This feature is not yet supported. The default value is 'disabled'.	varchar	256
WWN_POOLS_AVAILABLE	Not used. Previously used to indicate the number of WWN pools remaining for allocation on this encryption engine. This feature is no longer supported.	int	

TABLE 268 ENCRYPTION_ENGINE (Continued)

Field	Definition	Format	Size
STATE	Administrative state for this engine. 0 = disabled, 1 = enabled. The default value is 0.	smallint	
SP_CERTIFICATE	The public key certificate, in PEM format, for the Security Processor within the Encryption Engine. Used to create link keys for Decru LKM key vaults.	varchar	4096
EE_STATE	The operational status of this Encryption Engine. For possible values, see the enum defintion in the DTO class The default value is 0.	int	
HA_CLUSTER_STATUS	Stores the status of the HA Cluster to which the engine is a pair participant The default value is 0.	smallint	
ROUTING_MODE		smallint	
MEDIA_TYPE		char	50
LINK_IP_ADDRESS	Local EE - BP Link IP Address, If there are no links the IP Address could be 0.0.0.0	varchar	64
LINK_NET_MASK	Local EE - BP Link IP new mask	varchar	64
LINK_GW_IP_ADDRESS	Local EE- BP Gateway Address	varchar	64
LINK_MAC_ADDRESS	Local EE Link MAC Address	varchar	64
INK_MTU	Local EE Link MTU. The default value is -1.	int	
LINK_STATE	Local EE State says whether link is down or up	varchar	256
REBALANCE_REQUIRED	This field indicates whether a rebalance operation is required on the Encryption Engine. It can take two values, One(1) indicating that rebalance is required on the Encryption Engine and zero(0) indicating that no rebalance is required on the Encryption Engine. Encryption Engine is said to be unbalanced when the disk and Tape containers are not evenly balanced on the hosting engine. The default value is 0.	smallint	

TABLE 269 ENCRYPTION_GROUP

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
NAME	User-assigned name for this encryption group.	varchar	64
LEADER_SWITCH_ID	'Foreign key reference to both the VIRTUAL_SWITCH table and the CRYPTO_SWITCH table (both switch tables use the same primary key values). Identifies the switch that currently provides central configuration and reporting capabilities for the encryption group. This column may be null if the group leader is not in a discovered fabric.	int	

TABLE 269 ENCRYPTION_GROUP (Continued)

Field	Definition	Format	Size
LEADER_SWITCH_WWN	The Node WWN of the current group leader switch. Each encryption group has one group leader switch.	char	23
DEPLOYMENT_MODE	Indicates Transparent (0) or NonTransparent (1) deployment mode. Only Transparent mode is currently supported. All switches in the Encryption Group share the same deployment mode. Transparent mode uses re-direction zones to preserve existing zoning of physical hosts and targets. Non-transparent mode requires zoning changes to zone physical hosts with Virtual Targets and to zone Virtual Initiators with physical targets. The default value is 0.	smallint	
FAILBACK_MODE	Indicates Automatic (0) or Manual (1) failback. Failback occurs when a previously unavailable Encryption Engine comes back online. In Auto mode, the restored Encryption Engine resumes encrypting all traffic for target containers configured on the Encryption Engine. In manual mode, encryption continues running on the backup encryption engines until manually changed. The default value is 0.	smallint	
SYSTEM_CARD_REQUIRED	Boolean value that indicates whether a System Card (smart card) must be inserted in the Encryption Engine to enable the engine after power-up. This feature is not yet supported. The default value is 0.	smallint	
ACTIVE_MASTER_KEY_STATUS	The operational status of the "master key" or "Key Encryption Key (KEK)" used to encrypt Data Encryption Keys in a key vault. Not used for Decru LKM key vaults. 0 = not used, 1 = required but not present, 2 = present but not backed up, 3 = okay. The default value is 0.	smallint	
ALT_MASTER_KEY_STATUS	The operational status of an alternate "master key" used to access older data encryption keys. Not used for Decru LKM key vaults. 0 = not used, 1 = not present, 3 = okay. The default value is 0.	smallint	
QUORUM_SIZE	The number of authentication cards required to approve certain secure operations. This feature is not yet supported. The default value is 0.	smallint	
RECOVERY_SET_SIZE	No longer used. Previously used to indicate the number of smart cards used to back up a Master Key. The number of cards is now specified when the backup is created, and not persisted in the database. The default value is 0.	smallint	

TABLE 269 ENCRYPTION_GROUP (Continued)

Field	Definition	Format	Size
KEY_VAULT_TYPE	Indicates the type of key vault used by switches in this Encryption Group. 0 = Decru Lifetime Key Manager (LKM), 1 = RSA Key Manager (RKM), 2 = Brocade internal key storage (for demo use only). The default value is 0.	smallint	
PRIMARY_KEY_VAULT_ID	Foreign key reference to the KEY_VAULT record that describes the primary key vault for this Encryption Group. Null if no primary key vault is configured.	int	
BACKUP_KEY_VAULT_ID	Foreign key reference to the KEY_VAULT record that describes the backup key vault for this Encryption Group. Null if no backup key vault is configured.	int	
GROUP_LEADER_STATUS	Stores the status of the Group leader node	int	
SRDF_MODE	This field denotes whether the SRDF support is enabled or not. Feature available only from 6.4 release onwards and for RSA key vaults. EncryptionGroup collector and EncryptionGroupBean fills in this value. The default value is -1.	smallint	

TABLE 270 ENCRYPTION_GROUP_MEMBER

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
ENCRYPTION_GROUP_ID	Foreign key reference to the ENCRYPTION_GROUP record that identifies the encryption group that this member switch belongs to	int	
MEMBER_IP_ADDRESS	The management IP address (IPv4, IPv6, or hostname) of the member switch	varchar	128
MEMBER_WWN	the node WWN of the member switch	char	23
MEMBER_STATUS	The reachability status of the member switch as seen by the group leader switch. For possible values see the enum definition in the DTO class	smallint	

TABLE 271 ENCRYPTION_KMIP_PARAMETERS

Field	Definition	Format	Size
ENCRYPTION_GROUP_ID	Foreign key reference to the ENCRYPTION_GROUP record that describes the group ID of this Encryption Group.	int	
HA_MODE	Indicates the configured High Availability mode for the encryption group. Possible values are noHA, opaque, transparent, and NA.	varchar	32

TABLE 271 ENCRYPTION_KMIP_PARAMETERS (Continued)

Field	Definition	Format	Size
AUTHENTICATION_MODE	Indicates the configured User Authentication mode for the encryption group. Possible values are None, Username, UserPass, and NA.	varchar	32
CERTIFICATE_TYPE	Indicates the configured Certificate Type for the encryption group. Possible values are self, CASign, and NA.	varchar	32

TABLE 272 ENCRYPTION_TAPE_POOL

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SWITCH_ID	No longer used. Tape pools used to belong to specific switches, but are now shared by all switches in an encryption group	int	
ENCRYPTION_ENGINE_ID	No longer used. Tape pools used to belong to specific encryption engines, but are now shared by all encryption engines in an encryption group	int	
ENCRYPTION_GROUP_ID	Foreign key reference to the ENCRYPTION_GROUP record that describes which encryption group this tape pool belongs to	int	
TAPE_POOL_NAME	User-supplied name or number for the tape pool. This is the same name or number specified in the tape backup application. Numbers are stored in hex	varchar	64
TAPE_POOL_OPERATION_MODE	Specifies which type of encryption should be used by tape volumes in this tape pool. 0 = Native, 1 = DF-compatible	smallint	
TAPE_POOL_POLICY	Specifies whether tape volumes in this tape pool should be encrypted. 0 = encrypted, 1 = cleartext	smallint	
KEY_EXPIRATION	Number of days each data encryption key for this tape pool should be used. After the configured number of days, a new data encryption key is automatically generated for any further tape volumes in this pool. 0 = no expiration	int	
TAPE_POOL_LABEL_TYPE	Indicates whether the TAPE_POOL_NAME field is a name or a number. 0 = name, 1 = number	smallint	

TABLE 273 ETHERNET_CLOUD

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SWITCH_ID	The unique id of the switch this cloud is associated to.	int	

TABLE 274 ETHERNET_ISL

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SOURCE_PORT_ID	The unique id of the source port.	int	

TABLE 274 ETHERNET_ISL (Continued)

Field	Definition	Format	Size
DEST_PORT_ID	The unique id of the destination port.	int	
MISSING	Flag to identify whether the ethernet isl link is missing from the switch.	smallint,	
MISSING_TIME	Time when the ethernet isl link is missing from the switch.	timestamp	
TRUSTED	Is this ethernet isl link is trusted.	smallint,	
CREATION_TIME	Time when the ethernet isl link record is created.	timestamp	

TABLE 275 EVENT

Field	Definition	Format	Size
ID*	Unique generated database identifier for an event.	int	
ME_ID	Unique managed element ID used to refer the product that is associated with the event.	int	
SEVERITY	Indicates the severity of the event. Possible values : Emergency- 0, Alert- 1, Critical- 2, Error- 3,Warning- 4,Notice- 5, Info- 6,Debug- 7,Unknown- 8.	int	
AREA	Indicates the Area from which the event has occurred. Possible values : Unknown- 0, SAN- 1, IP- 2, Application Events -3, SAN+IP- 4.	smallint	
ACKNOWLEDGED	Indicates whether the user has acknowledged the event or not. Possible values: Unacknowledged-0 , Acknowledged-1.	smallint	
SOURCE_NAME	This field indicates the name of the source that triggered the event. This could be the name of the source switch or name of the Management application server in the case of application events.	varchar	255
SOURCE_ADDR	'Indicates the IP Address of the source that triggered the event. This could be the IP address of the source switch or IP address of the Management application server in the case of application events.	varchar	50
EVENT_ORIGIN_ID	Database ID of the event origin such as Trap, Syclog etc referring to EVENT_ORIGIN metadata.	int	
EVENT_CATEGORY_ID	Database ID of the event category referring to EVENT_CATEGORY metadata.	int	
EVENT_MODULE_ID	Database ID of the event module referring to EVENT_MODULE metadata.	int	
EVENT_DESCRIPTION_ID	Indicates the identifier of the event description in the EVENT_DESCRIPTION table.	int	
LAST_OCCURRENCE_HOST_TIME	Indicates the the Management application server timestamp when this event occurred last.	timestamp	
EVENT_COUNT	Indicates the number of occurrences of the event. Count indicates the number of times the same event occurred in a given ten minute window.	int	
RESOLVED	This field indicates whether an event is resolved due to another matching event or based on user action. Possible values: Unresolved – 0, Resolved – 1.	smallint	

TABLE 275 EVENT (Continued)

Field	Definition	Format	Size
ACKED_TIME	Indicates the timestamp when the event was acknowledged.	Timestamp	
FIRST_OCCURRENCE_HOST_TIME	Indicates the the Management application server timestamp when the event occurred for the first time.	timestamp	10
EVENT_AUDIT	'Indicates whether this is an audit event or not.	varchar	255
EVENT_KEY	Unique key for the event. This is a string message key represents message ID from events originated from switch or the predefined message Id for application events in the Management application.	varchar	
EVENT_ACTION_ID	Reference to the ID in the EVENT_POLICY table. Represents the event action policy that was responsible for generating this event.	int	
DEVICE_GROUP_ID	Reference to the DEVICE_GROUP_ID in the DEVICE_GROUP table.	int	
PORT_GROUP_ID	Reference to the ID in the PORT_GROUP table.	int	
SPECIAL_EVENT	'Indicates whether the event is marked as special event or not. Not a Special event-0, Special event-1.	smallint	
CALLHOME_EVENT	Indicates whether the event is marked as call home event. 0 - not a call home event 1- call home event	smallint	

TABLE 276 EVENT_CALL_HOME

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
EVENT_ID	Database ID of the EVENT instance.	int	
EVENT_NUMBER	Indicates the Event Number for the event from the Events.html of the associated product .	int	
FRU_CODE	Indicates the Field Replaceable Unit code of the Call Home event.	int	
REASON_CODE	Indicates the reason code of the Call Home event.	int	
FRU_POSITION	Indicates the FRU position of the Call Home event.	int	

TABLE 277 EVENT_CATEGORY

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
DESCRIPTION	Holds the event categories. Possible values : Unknown- 0, Product Event- 1, Link Incident Event- 2 , Product Audit Event- 3, Product Status Event- 4, Security Event- 5 , User Action Event- 6, Management Server Event- 7.	varchar	50

TABLE 278 EVENT_DESCRIPTION

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
DESCRIPTION	Holds the description of the Event.	text	1024

TABLE 279 EVENT_DETAILS

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
EVENT_ID	Database ID of the EVENT instance.	int	
FIRST_OCCURRENCE_SWITCH_TIME	Indicates the first occurrence switch timestamp of the event.	timestamp	
LAST_OCCURRENCE_SWITCH_TIME	Indicates the last occurrence switch timestamp of the event.	timestamp	
CONTRIBUTORS	Indicates the contributing factor for the event resulted due to a status change of the switch.	varchar	512
OPERATIONAL_STATUS	Indicates the operational Status of the product associated with the event.	varchar	255
NODE_WWN	Unique World Wide Number for the product.	varchar	23
PORT_WWN	Unique World Wide Number for the port for which the event was generated.	varchar	23
OID	Indicates the Object ID of the Trap or Syslog.	varchar	50
VIRTUAL_FABRIC_ID	Indicates the Virtual Fabric id of the switch which triggered the event.	smallint	
UNIT	Indicates the Unit number of the Chassis from which the event was triggered.	smallint	
SLOT	Indicates the blade or the slot number in which the port is present.	int	
PORT	indicates the switch port number for which the event was generated.	int	
PRODUCT_ADDRESS	Indicates the IP Address of the Product from which the event is originated.	varchar	
RAS_LOG_ID	Indicates the RASLOG Id of the RASLOG event.	varchar	20
INTERFACE_TYPE	Indicates the type of the interface – Possible Values: Ethernet Port-0, FC Port-1.	smallint	
USER_NAME	Captures the user information from audit Syslog messages.	Varchar	512
PORT_NAME	Shows the PortName for the corresponding port.	Varchar	255
MAC_ADDRESS	'Indicates the MAC address of the Access Point from which this event is received. If the event is received from the wireless controller or any other products, this will be empty.';	varchar	64

TABLE 279 EVENT_DETAILS (Continued)

Field	Definition	Format	Size
ANNOTATED_BY	The user name who annotates the event. In Events context, Acknowledging/Un-Acknowledging event is called annotation. For Auto acknowledged events, SYSTEM User is set in this column.	varchar	128
ANNOTATIONS	Annotations of the event. User can add annotations through Acknowledge/Un-Acknowledge process. Events can be annotated automatically through auto acknowledge option as part of event action. Custom notes can be added manually to all events in master log without performing acknowledgement or un-acknowledgement and those notes are also stored.	varchar	1024

TABLE 280 EVENT_FWD_FILTER

Field	Definition	Format	Size
ID		int	
NAME	Filter Name	varchar	32
DESCRIPTION		varchar	256
TYPE	Filter Type (SNMP/ SYSLOG)	smallint	
APPLICATION_ENABLED	If Application Events enabled	smallint	
SNORT_ENABLED	If Snort Messages enabled	smallint	
PSUDO_ENABLED	If Pseudo Events enabled	smallint	
REGULAR_EXP	Common filtering message for Syslog Forwarding	varchar	512
SEVERITY	Emergency(0), Alert(1), Critical(2), Error(3), Warning(4), Notice(5), Info(6), Debug (7). Traps with selected severity and those with higher severity will be forwarded.	smallint	

TABLE 281 EVENT_INSTANCE

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
EVENT_POLICY_ID	Foreign Key to Event_Policy Table	int	
EVENT_KEY	A String Key string which identifies a specific instance of an Event.	varchar	64
STRING_PATTERN	A Regular expression pattern string which can be used to match an Event instance.	varchar	1024
CATEGORY	A small integer which identifies the Category of an Event instance. 0 - Unknown 1 - Product Event 2- Link Incident Event 3 - Product Audit Event 4- Product Status Event 5 - Security Event 6- User Action Event 7- Management Server Event. The default value is 0.	smallint	

TABLE 281 EVENT_INSTANCE (Continued)

Field	Definition	Format	Size
SEVERITY	The Severity of the Event that is logged per Event Policy 0- Unknown 1- Emergency 2- Alert 3- Critical 4- Error 5- Warning 6- Notice 7- Info 8- Debug. The default value is 0.	smallint	
SEQUENCE_NUMBER	The sequence number of an event instance that"s specific to the policy. The default value is 0.	smallint	
MSG_IDS	Stores the Message ID(s) configured for Custom Event Type	varchar	512

TABLE 282 EVENT_MODULE

Field	Definition	Format	Size
ID	The default value is 0.	int	
DESCRIPTION		varchar	256

TABLE 283 EVENT_NOTIFICATION

Field	Definition	Format	Size
ID*		int	
STATUS	Status of Event Notification. value will be 0 if disabled, 1 otherwise. Default value is 0.	smallint	
SERVER_NAME	E-mail (SMTP) server name.	varchar	256
REPLY_ADDRESS	Reply E-mail address.	varchar	50
SEND_ADDRESS	E-mail address for which a Test E-mail notification is to be sent.	varchar	512
SMTP_PORT	SMTP Port number. Default value is 25.	int	
USER_NAME	User name for authentication.	varchar	256
PASSWORD	Password for authentication.	varchar	256
NOTIFICATION_INTERVAL	Time interval between successive event notifications.	int	
NOTIFICATION_UNIT	Time interval Unit: 0 = Seconds 1 = Minutes 2 = Hours Default value is 0.	smallint	
TEST_OPTION	Time interval Unit: 0 = Send test to configured e-mail address. 1 = Send test to all enabled users. Default value is 0.	smallint	
SSL_ENABLED	Default value is 0.	smallint	

TABLE 284 EVENT_ORIGIN

Field	Definition	Format	Size
ID	0 - Unknown 1 - Trap 2 - Syslog 3 - Snort 4 - Pseudoevent 5 - Application Events 6 - Others	int	
DESCRIPTION		varchar	50

TABLE 285 EVENT_POLICY

Field	Definition	Format	Size
ID		int	
TYPE	Even Policy Type 0 - Pseudo Event 1 - Event Action	smallint	
NAME	The Name of the Event Policy	varchar	256
DESCRIPTION	The Description of the Event Policy	varchar	1024
STATUS	Administrative status of the Event Policy 0 - disabled 1- enabled	smallint	
LAST_MODIFIED_TIME	The Severity of the Event that is logged per Event Policy 0- Unknown 1- Emergency 2- Alert 3- Critical 4- Error 5- Warning 6- Notice 7- Info 8- Debug';	timestamp	
SEVERITY	The Event Policy Sub Type Escalation (0), Resolving (1), Flapping (2),Repeating (3). The default value is 0.	smallint	
MESSAGE		varchar	256
SUB_TYPE	The Event Policy Sub Type Escalation (0), Resolving (1), Flapping (2),Repeating (3)	smallint	
EVENT_ORIGIN	0- SNMP Trap 1- Application Event 2- Pseudo Event 3- Snort 4- Pseudo Event 5- Custom Event	smallint	
PROPERTIES	The property string which is used to define various parameters that are associated with an Event Policy such as flapping time and durations etc	varchar	2048
AUTOACK_ANNOTATIONS		text	

TABLE 286 EVENT_POLICY_SOURCE_ENTRY

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
EVENT_POLICY_ID	Foreign Key to Event_Policy Table	int	

TABLE 286 EVENT_POLICY_SOURCE_ENTRY (Continued)

Field	Definition	Format	Size
MANAGEMENT_ELEMENT_ID	A soft reference key to the Management Element ID. Do not maintain it as a foreign key constraints. The default value is 0.	int	
INTERFACE_ID	A soft reference key to the Interface ID. Do not maintain it as a foreign key constraints. The default value is 0.	int	
DEVICE_GROUP_ID	A reference key to the Device Group Do not maintain it as a foreign key constraints. The default value is 0.	int	
PORT_GROUP_ID	A reference key to the Port Group Do not maintain it as a foreign key constraints. The default value is 0.	int	
SOURCE_SELECTION_TYPE	Option selected to give Source Information • 0- IPAddress/Node wwn/Name provided • 1- Source selected from available list of sources. The default value is 0.	smallint	
IP_ADDRESS	IP address of source	varchar	1024
WWN	Node WWN of source	varchar	1024
SOURCE_NAME	Source Name	varchar	1024
AUTOACK_ANNOTATIONS	The Annotations to be added to the matching event.	varchar	1024

TABLE 287 EVENT_PROCESSOR_MAP

Field	Definition	Format	Size
PROCESSOR_CLASS_NAME	The fully qualified processor class name which will be invoked for the corresponding event id in this table. Column added as part of the Event Processing Framework	varchar	1024
EVENT_ID	The Event Id is the Trap OID on which the corresponding processor needs to act up on . Column added as part of the Event Processing Framework	varchar	1024

TABLE 288 EVENT_RULE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of the Event Rule.	varchar	255
TYPE	Event Rule Type: • 0 = Port Offline • 1 = PM Threshold crossed • 2 = Security Violation • 4 = Event	int	
DESCRIPTION	Description about the Event Rule.	varchar	512
OPERATOR1	AND operator used to append the rule.	varchar	12

TABLE 288 EVENT_RULE (Continued)

Field	Definition	Format	Size
EVENT_TYPE_ID	The Selected Event type ID from the Event type combo box.	int	
OPERATOR2	AND operator used to append the rule.	varchar	12
MESSAGE_ID	Message ID provided by the user.	varchar	20
OPERATOR3	AND operator used to append the rule.	varchar	12
IP_ADDRESS	Source IP Address.	varchar	1024
OPERATOR4	AND operator used to append the rule.	varchar	12
WWN	Source WWN.	varchar	1024
OPERATOR5	AND operator used to append the rule.	varchar	12
COUNT	Count of the specified event.	int	
OPERATOR6	AND operator used to append the rule.	varchar	12
DURATION	Duration of the specified event.	bigint	
STATE	State of the rule: • 0 = Disabled • 1 = Enabled	smallint	
SEVERITY_LEVEL	Event severity level. Default value is 4.	int	
SOURCE_NAME	Name of the source.	varchar	1024
DESCRIPTION_CONTAINS	Description pattern about the rule.	varchar	255
LAST_MODIFIED_TIME	Rules last edited time.	timestamp	
SELECTED_TIME_UNIT	Timestamp unit of the selected rule: • 0 = second • 1 = Minutes • 2 = Hours Default value is 1.	smallint	

TABLE 289 EVENT_RULE_ACTION

Field	Definition	Format	Size
ID*		int	
RULE_ID	The rule ID present in the Event_Rule Table.	int	
NAME	Name of the Event Rule Action: • Launch Script = for launch script • Send E-mail = for send e-mail • Raise Event = for broadcast message	varchar	255
TYPE	Name of the action: • script = for Launch Script • e-mail = for E-mail • message = for Broadcast message	varchar	30
FIELD1	Data for the selected action.	varchar	512
FIELD2	Data for the selected action.	varchar	512
FIELD3	Data for the selected action.	varchar	512

TABLE 289 EVENT_RULE_ACTION (Continued)

Field	Definition	Format	Size
FIELD4	Data for the selected action.	varchar	512
STATE	State of the Action: • 0 = Action Disabled • 1 = Action Enabled Default value is 0.	smallint	

TABLE 290 FABRIC

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
SAN_ID	Foreign key to SAN table; usually 1 since there is only one SAN.	int	
SEED_SWITCH_WWN	WWN of the virtual switch used as seed switch to discover the fabric.	char	23
NAME	User-assigned fabric name.	varchar	256
CONTACT	User-assigned "contact" for the fabric.	varchar	256
LOCATION	User-assigned "location" for the fabric.	varchar	256
DESCRIPTION	User-assigned fabric description.	varchar	256
TYPE	Denotes the type of Fabric. 0 = legacy fabric, 1 = base fabric, 2 = logical fabric, 3 = partial fabric, 4 = ethernet fabric. Default value is 0.	smallint	
SECURE	1 = it is a secured fabric. Default value is 0.	smallint	
AD_ENVIRONMENT	1 = there are user-defined ADs in this fabric. Default value is 0.	smallint	
MANAGED	1 = it is an actively "monitored" fabric; otherwise, it is an "unmonitored" fabric. Default value is 1.	smallint	
MANAGEMENT_STATE	Bit map to indicate various management indications for the fabric. Default value is 0.	smallint	
TRACK_CHANGES	1 = changes (member switches, ISL and devices) in the fabric are tracked. Default value is 0.	smallint	
STATS_COLLECTION	1 = statistics collection is enabled on the fabric. Default value is 0.	smallint	
CREATION_TIME	When the fabric record is inserted, i.e., created. Default value is 'now()'.	timestamp	
LAST_FABRIC_CHANGED	Time when fabric last changed.	timestamp	
LAST_SCAN_TIME	Last Scan time for the fabric i.e. when the switch was scanned for changes.	timestamp	
LAST_UPDATE_TIME	Time when fabric was last updated. Default value is 'now()'.	timestamp	

TABLE 290 FABRIC (Continued)

Field	Definition	Format	Size
ACTIVE_ZONESET_NAME	Name of the zone configuration which is effective / active in that fabric.	varchar	256
USER_DEFINED_VALUE_1	User-defined custom value.	varchar	256
USER_DEFINED_VALUE_2	User-defined custom value.	varchar	256
USER_DEFINED_VALUE_3	User-defined custom value.	varchar	256
PRINCIPAL_SWITCH_WWN	WWN of the principal switch of the fabric	char	23
ZONE_TRANSACTION_TIMEOUT	Number of seconds that a ZONE_TRANSACTION can be idle Default value is 180.	int	
FABRIC_MODEL	Default value is 1.	smallint	
LAST_FAILURE_TIMESTAMP	Denotes the last failure timestamp.	timestamp	
LAST_SUCCESSFUL_TIMESTAMP	Dentoes the last successful timestamp.	timestamp	
ENHANCED_TI_ZONE_SUPPORT	Holds the value if the fabric has enhanced TI Zone support or not. Default: 0 Values: 0|1.	smallint	
FABRIC	The fabric name persisted on switches running FOS 7.0 and later. Not to be confused with NAME, which is store on the Management application only.	varchar	128
STATUS	Overall operational status of the fabric. 0 is unknown, 1 is healthy, 2 is marginal, 3 is down, 5 is Reachable, 6 is unreachable, 7 is Degraded link.	int	
TRACKING_STATUS	This represents bitmask as an integer value which represents missing or untrusted state of fabric members, ISLs, SANConnections, device Nodes and device ports. 1 is missing switch/ISL in fabric, 2 is untrusted switch or ISL in fabric, 4 is missing initiator or port in fabric, 8 is untrusted initiator or port in fabric, 16 is missing target or port in fabric, 32 is untrusted target or port in fabric.	int	
BOTTLENECK_STATUS	Holds bottleneck status of fabric. Default is 0, Values are 0 or 1.	int	
VCS_LICENSED	Indicates whether the fabric has VCS license or not. Possible values are 0 for not applicable, 1 for licensed, 2 for not licensed. 0 is default. Fabrics representing clusters with 2 or less nodes will have value of 0 as all those are automatically licensed. Fabrics representing clusters with 3 or more nodes will have values 1 or 2 depending on whether the license was acquired or not.	int	
HAS_NOS_AG	Denotes whether fabric has NOS AG connected to it or not. 0 denotes that the fabric has no NOS AG connected to it and 1 denotes a NOS AG is connected to the fabric. -1 denotes that it is not applicable and will be set to NOS clusters.	int	

TABLE 291 FABRIC_CHECKSUM

Field	Definition	Format	Size
FABRIC_ID *	Fabric ID, foreign key to the FABRIC table.	int	
CHECKSUM_KEY *	Type of checksum, e.g. device data or zone data.	varchar	32
CHECKSUM	Actual checksum value.	varchar	16

TABLE 292 FABRIC_COLLECTION

Field	Definition	Format	Size
FABRIC_ID *	Fabric ID, foreign key to the FABRIC table.	int	
COLLECTOR_NAME *	Name of the collector, e.g., NameServerInfoCollector, TopologyCollector, ZoneInfoCollector, ActiveZoneInfoCollector.	varchar	256
SEED_SWITCH_IP	IP address of the switch which serves as the seed switch. This is the switch from which above mentioned fabric level collectors get their information.	varchar	128
LAST_SEED_SW_ MODIFICATION	Timestamp of the seed switch, when the particular HTML page was changed last. Note that this is not when the last time collection was done.	timestamp	

TABLE 293 FABRIC_CONFIGURATION

Field	Definition	Format	Size
ID	Auto generated primary key.	int	
NAME	Name of the configuration block.	varchar	128
DESCRIPTION	Description of the configuration block.	varchar	512
CREATED_BY	Foreign key reference to the ID column of the USER_ table. Indicates the user who created the configuration block.	int	
CREATION_TIME	Time when configuration block was created.	timestamp	
LAST_UPDATE_TIME	Time when configuration block was updated.	timestamp	
STATUS	Indicates if config block values are populated. 0 - Empty, 1- Partial, 2- Fully populated.	smallint	

TABLE 294 FABRIC_CONFIGURATION_DRIFTS

Field	Definition	Format	Size
ID	Auto generated primary key.	int	
FABRIC_CONFIG_SETTIN G_TYPE_ID	Foreign key to FABRIC_CONFIGURATION_SETTING_TYPE(ID). Indicates the type of configuration.	int	
FABRIC_CONFIG_MONIT OR_ID	Refers to FABRIC_CONFIGURATION_MONITOR (ID). The configuration drifts should be retained over a period of time even if the monitor is deleted.	int	
VIRTUAL_SWITCH_ID	Foreign key reference to the ID column of virtual group.	int	
DRIFT_OBJECT	Object containing the configuration drift details.	bytea	

TABLE 294 FABRIC_CONFIGURATION_DRIFTS (Continued)

Field	Definition	Format	Size
LAST_UPDATE_TIME	Time when the drift was found.	timestamp	
CHECK_STATUS	Indicates if the drift check against the switch is successful or not. 0- successful, 1- failed.	smallint	

TABLE 295 FABRIC_CONFIGURATION_ENTRY

Field	Definition	Format	Size
ID	Auto generated primary key.	int	
FABRIC_CONFIG_ID	Foreign Key to FABRIC_CONFIGURATION (ID). Refers to the configuration block that this configuration entry belong to.	int	
FABRIC_CONFIG_SETTING_TYPE_ID	Foreign key to FABRIC_CONFIGURATION_SETTING_TYPE(ID). Indicates the type of configuration.	int	
CONFIG_KEY	Name of the configuration key.	varchar	128
CONFIG_VALUE_OBJECT	Payload object containing the configuration parameter values to be applied on switches.	byte	

TABLE 296 FABRIC_CONFIGURATION_POLICY

Field	Definition	Format	Size
ID	Auto generated primary key.	int	
NAME	Name of the configuration policy.	varchar	128
DESCRIPTION	Description of the configuration policy.	varchar	512
FABRIC_CONFIG_TEMPLATE_ID	Foreign Key to FABRIC_CONFIGURATION_TEMPLATE (ID).	int	
POLICY_MONITOR_CONFIG_ID	Foreign Key to DEPLOYMENT_CONFIGURATION (ID). ID of the policy monitor policy. This can be NULL if the user does not add any policy.	int	
CREATED_BY	Foreign key reference to the ID column of the USER_ table. Indicates the user who created the configuration policy.	int	
CREATION_TIME	Time when configuration policy was created.	timestamp	
LAST_UPDATE_TIME	Time when configuration policy was updated.	timestamp	

TABLE 297 FABRIC_CONFIGURATION_MONITOR

Field	Definition	Format	Size
ID	Auto generated primary key.	int	
LINKED_TIME	Time when the template was linked to a fabric or a switch group.	timestamp	
FABRIC_CONFIG_TEMPLATE_ID	Foreign key reference to FABRIC_CONFIGURATION_TEMPLATE (ID).	int	
USER_DEFINED_NETWORK_SCOPE_ID	Foreign Key to USERDEFINED_NETWORK_SCOPE (ID). This can be NULL if the user does not link the fabric configuration template to a user defined network scope.	int	
FABRIC_ID	Foreign key reference to FABRIC (ID).This can be NULL if the user does not link the fabric configuration template to a fabric.	int	

TABLE 298 FABRIC_CONFIGURATION_SETTING_TYPE

Field	Definition	Format	Size
ID	Auto generated primary key.	int	
SETTING_TYPE	Indicates the switch configuration settings such as FTP Server, Syslog destination, etc.	varchar	512

TABLE 299 FABRIC_CONFIGURATION_TEMPLATE

Field	Definition	Format	Size
ID	Auto generated primary key.	int	
NAME	Name of the configuration template.	varchar	128
DESCRIPTION	Description of the configuration template.	varchar	512
CREATED_BY	Foreign key reference to the ID column of the USER_ table. Indicates the user who created the configuration template.	int	
CREATION_TIME	Time when configuration template was created.	timestamp	
LAST_UPDATE_TIME	Time when configuration template was updated.	timestamp	

TABLE 300 FABRIC_CONFIGURATION_TEMPLATE_ENTRY

Field	Definition	Format	Size
ID	Auto generated primary key.	int	
FABRIC_CONFIG_TEMPLATE_ID	Foreign Key to FABRIC_CONFIGURATION_ TEMPLATE (ID).	int	
FABRIC_CONFIG_ID	Foreign key to FABRIC_CONFIGURATION (ID). Refers to the fabric configuration block.	int	

TABLE 301 FABRIC_CONFIG_DRIFTS_LASTCHECK_DETAILS

Field	Definition	Format	Size
FABRIC_CONFIG_MONITOR_ID	Foreign key reference to FABRIC_CONFIGURATION_MONITOR (ID.	int	
VIRTUAL_SWITCH_ID	Foreign key reference to VIRTUAL_SWITCH (ID).	int	
LASTCHECK_TIMESTAMP	Last check time stamp.	timestamp	
LASTCHECK_STATUS	Last check status. Currently, 0 = Successful, 1 = Failed. Default value is 0.	int	
LASTCHECK_STATUS_DETAILS	Last check status details.	varchar	512

TABLE 302 FABRIC_DISCOVERY_POLICY

Field	Definition	Format	Size
FABRIC_ID	The database ID of the fabric that the policy belongs to.	int	
DISCOVER_ALL_MEMBERS	This column indicates if all the members of the fabric can be discovered. 1 means discover all members and 0 means do not discover all members. In the case of 0, the filtering rules comes from FabricDiscoveryPolicyRule table.	smallint	
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	

TABLE 303 FABRIC_DISCOVERY_POLICY_RULE

Field	Definition	Format	Size
ID		serial	
FABRIC_ID	The database ID of the fabric that the policy belongs to.	int	
FILTER	Filter to be applied for this fabric. This could be IP Address or WWN or SwitchType. The Type of the filter comes from the FilterType column. This can be either in included list or excluded list depending on the EXCLUDED column value.	varchar	128
FILTER_TYPE	This column indicates type of the filter. It could take values like 0 for IP Address, 1 for WWN and 2 for SwitchType. Default is IP address type.	smallint	
EXCLUDED	This column indicates if the Filter in the record should be included or excluded. 1 means exclude and 0 means include. Default is to include.	smallint	
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	

TABLE 304 FABRIC_MEMBER

Field	Definition	Format	Size
FABRIC_ID*	Fabric ID, foreign key to FABRIC table.	int	
VIRTUAL_SWITCH_ID*	ID of the virtual switch which is a member of this fabric, foreign key to VIRTUAL_SWITCH table.	int	
TRUSTED	1 = the switch is a trusted member of the fabric. Either found in the initial discovery or user subsequently entrusted the switch by user action. Default Value is 0.	smallint	
CREATION_TIME	When the switch became a member. Default Value is 'now()'.	timestamp	
MISSING	1 = it is missing from the fabric. Default Value is 0.	smallint	
MISSING_TIME	When it is missed from the fabric; null if the member is entrusted.	timestamp	
LAST_UPDATE	Last Updated time for the record.	bigint	

TABLE 305 FABRIC_THRESHOLD_SETTING

Field	Definition	Format	Size
FABRIC_ID*	References the ID in FABRIC table	int	
POLICY_ID*	References the ID in THRESHOLD_POLICY table	int	

TABLE 306 FABRIC_VCS_CLUSTER_MAP

Field	Definition	Format	Size
FABRIC_ID	Foreign key to ID in fabric table.	int	
VCS_CLUSTER_ME_ID	Foreign key to ID in ManagedElement table. This is the VCS cluster entry managed_element_id reference.	int	

TABLE 307 FABRIC_ZONING_EDIT_RESTRICTION

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
FABRIC_ID	PK of the owning fabric	int	
CHANGE_COUNT	Count of the maximum changes allowed in active zone config in the fabric.The default value is 0.	int	

TABLE 308 FAVORITES

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of the favorite.	varchar	64
USER_	The application user credentials.	varchar	128
TOP_N	The top number of ports(5,10,15,20).	varchar	40
SELECTION_FILTER	Types of ports (FC/FCIP/GE) and -End Monitors.	varchar	40

TABLE 308 FAVORITES (Continued)

Field	Definition	Format	Size
FROM_TIME	The time interval in which the graph is shown. Time interval can be predefined or custom. If FROM_TIME is Custom, the user can choose the number of minutes/hours/days or specify the time interval.	varchar	40
CUSTOM_LAST_VALUE	The number of minutes/hours/days. It becomes null in two cases. 1. When the value of FROM_TIME is not Custom. 2. When FROM_TIME is Custom, and user chooses the time interval (CUSTOM_FROM and CUSTOM_TO)	int	
CUSTOM_TIME_UNIT	The unit type (Minutes, Hours, Days) of the CUSTOM_LAST_VALUE.	varchar	40
CUSTOM_FROM	The starting time.	timestamp	
CUSTOM_TO	The ending time.	timestamp	
GRANULARITY	The granularity.	varchar	40
THRESHOLD	The reference line.	int	
MAIN_MEASURE	The measure of FC/FCIP/GE.	varchar	40
ADDITIONAL_MEASURE	The additional measures.	int	
CUSTOM_SELECTION_OBJECT_TYPE	Represents the selected filter type. • 0 - Default favorite • 1 - FC Ports • 2 - Device Ports • 3 - ISL Ports • 4 - 10GE Ports • 5 - FCIP Tunnels • 6 - EE Monitors Selected member identifiers are stored in CUSTOM_FAVORITES_OBJECT_LIST table if this favorite is not default.	int	
PLOT_EVENTS	Indicates whether the PM historical chart should overlay the events on the graph. 0 - No, 1 - Yes.	smallint	

TABLE 309 FCIP_CIRCUIT_PORT_MAP

Field	Definition	Format	Size
CIRCUIT_ID		int	
SWITCH_PORT_ID	SWITCH_PORT_ID of one end of the circuit	int	
DP1_SWITCH_PORT_ID	Switch port Id of the DP1 ip address. This field will be -1 for non Skybolt switches or if DP1 ip address not configured in Skybolt circuit.	int	

TABLE 310 FCIP_PORT_TUNNEL_MAP

Field	Definition	Format	Size
SWITCHPORT_ID*	Switch Port ID.	int	
TUNNEL_ID*	FCIP Tunnel ID.	int	

TABLE 311 FCIP_TUNNEL

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
TUNNEL_ID	Tunnel ID for that GigE Port.	smallint	
VLAN_TAG	VLAN Tag on the tunnel (if present). Default value is -1.	int	
SOURCE_IP	Source IP on which the tunnel is created.	char	64
DEST_IP	Destination IP on the other end of tunnel.	char	64
LOCAL_WWN	Local port WWN for the tunnel.	char	23
REMOTE_WWN_RESTRICT	Remote Port WWN for the tunnel.	char	23
COMMUNICATION_RATE	Bandwidth specified for the tunnel.	double precision	
MIN_RETRANSMIT_TIME	FCIP Tunnel Parameter.	int	
SELECTIVE_ACK_ENABLED	FCIP Tunnel Parameter.	smallint	
KEEP_ALIVE_TIMEOUT	FCIP Tunnel Parameter.	int	
MAX_RETRANSMISSION	FCIP Tunnel Parameter.	int	
WAN_TOV_ENABLED	Is WAN TOV enabled. Default value is 0.	smallint	
TUNNEL_STATUS	Tunnel Status (Active/Inactive).	int	
DESCRIPTION	Description for the created tunnel.	varchar	64
FICON_TRB_ID_ENABLED	Whether Ficon_Tape_Read_Block is enabled on that tunnel. Default value is 0.	smallint	
FICON_TT_EMUL_ENABLED	Whether Ficon_Tin_Tir_Emulation is enabled on that tunnel. Default value is 0.	smallint	
FICON_DLA_EMUL_ENABLED	Whether Device_Level_Ack_Emulation is enabled on that tunnel. Default value is 0.	smallint	
FICON_TAPE_WRITE_MAX_PIPE	The Value for FICON_TAPE_WRITE_MAX_PIPE on the tunnel. Default value is -1.	int	
FICON_TAPE_READ_MAX_PIPE	The Value for FICON_TAPE_READ_MAX_PIPE on the tunnel. Default value is -1.	int	
FICON_TAPE_WRITE_MAX_OPS	The Value for FICON_TAPE_WRITE_MAX_OPS on the tunnel. Default value is -1.	int	
FICON_TAPE_READ_MAX_OPS	The Value for FICON_TAPE_READ_MAX_OPS on the tunnel. Default value is -1.	int	

TABLE 311 FCIP_TUNNEL (Continued)

Field	Definition	Format	Size
FICON_TAPE_WRITE_TIMER	The Value for FICON_TAPE_WRITE_TIMER on the tunnel. Default value is -1.	int	
FICON_TAPE_MAX_WRITE_CHAIN	The Value for FICON_TAPE_MAX_WRITE_CHAIN on the tunnel. Default value is -1.	int	
FICON_OXID_BASE	The Value for FICON_OXID_BASE on the tunnel. Default value is -1.	int	
FICON_XRC_EMULATION_ENABLED	Whether Xrc_Emulation is enabled on that tunnel. Default value is 0.	smallint	
FICON_TW_EMUL_ENABLED	Whether Ficon_Tape_Write_Emulation is enabled on that tunnel. Default value is 0.	smallint	
FICON_TR_EMUL_ENABLED	Whether Ficon_Tape_Read_Emulation is enabled on that tunnel. Default value is 0.	smallint	
FICON_DEBUG_FLAGS	FICON_DEBUG_FLAGS for that particular tunnel. Default value is -1.	double precision	
REMOTE_WWN	Configured WWN of the Remote Node.	char	64
CDC	CDC Flag. Default value is 0.	smallint	
ADMIN_STATUS	Admin Status of the Tunnel. Default value is 0.	smallint	
CONTROL_L2_COS	Class of service as defined by IEEE 802.1p for tunnel. Default value is -1.	int	
DSCP_CONTROL	DiffServe marking for control frame. Default value is -1.	int	
TRUNKING_ALGORITHM	Trunking Algorithm. Default value is -1.	int	
EXTENDED_TUNNEL	Indicates if the tunnel is an Extended Tunnel (i.e. new Tunnel type on the switch). Default value is 0.	smallint	
VIRTUAL_SWITCH_ID	Refers to the virtual switch to which the tunnel record belongs to.	int	
CIRCUIT_COUNT	The number of circuits configured on the tunnel. Default value is 1.	smallint	
MISMATCHED_CONFIG_DETAILS	Details of the reasons as to why the tunnel is down.	varchar	2048
LAST_UPDATE	Last update time tells the time when the last update to the database record happened.	bigint	

TABLE 311 FCIP_TUNNEL (Continued)

Field	Definition	Format	Size
SLOT_NUMBER	SLOT_NUMBER on which the VE Port of the tunnel exists. Default value is 0.	int	
FICON_ENABLED	Is Ficon enabled. Default: 0, Values: 0|1. Default value is 0.	smallint	
TPERF_ENABLED	Is Tperf enabled. Default: 0, Values: 0|1. Default value is 0.	smallint	
AUTH_KEY	This is the preshared-key to be used during IKE authentication.	varchar	128
CONNECTED_COUNT	Active connections count. Default value is 1.	smallint	
TUNNEL_STATUS_STRING	Tunnel Status string value from switch for the tunnel.	varchar	256
COMPRESSION_MODE	Compression mode value (0,1,2,3). Default value is 0.	smallint	
TURBO_WRITE_ENABLED	Whether turbo write (fast write) is enabled or not (0,1). Default value is 0.	smallint	
TAPE_ACCELERATION_ENABLED	Whether turbo write (fast write) is enabled or not (0,1). Default value is 0.	smallint	
IPSEC_ENABLED	Default value is 0.	smallint	
PRESHARED_KEY	The preshared key on tunnel.	char	32
QOS_HIGH	QoS high value.	smallint	
QOS_MEDIUM	QoS medium value.	smallint	
QOS_LOW	QoS low value.	smallint	
BACKWARD_COMPATIBLE	Whether the 10G tunnel is backward compatible with previous FOS versions.	smallint	
FICON_TERADATA_READ_ENABLED	Whether Ficon_Teradata_Read_Pipelining is enabled on that tunnel.	smallint	
FICON_TERADATA_WRITE_ENABLED	Whether Ficon_Teradata_Write_Pipelining is enabled on that tunnel.	smallint	
HA_STATUS		int	
IP_EXTN_MODE	IP Extension mode of the tunnel. 0 - Disabled, 1 - Enabled.	int	
IP_COMPRESSION_MODE	Compression mode for the IP traffic. 0-Off, 1 - Advanced, 2 - LZ, 3- Deflate, 4 - Aggressive Deflate, 5- Fast Deflate.	int	
IP_QOS	IP QOS Value array. Ex {50,30,20} High 50% Medium 30% Low 20%, null for IP Extn disabled tunnels.	int	
QOS_DISTRIBUTION_MODE	0- non-IP Ext tunnel 1 -default, 2 - protocol, 3 – priority.	int	

TABLE 311 FCIP_TUNNEL (Continued)

Field	Definition	Format	Size
QOS_DISTRIBUTION_VALUE	Array of QOS distribution values. Will have two values in case of protocol mode and 3 values in priority mode. Ex in default: {20,15,15,30,10,10} Here FC High 20% Medium 15% Low 15% IP High 30% Medium 10% Low 10% in protocol: {60,40} Here FC 60 % IP 40% in priority: {50,30,20} Here High 50% Medium 30% Low 20% for non-IP Ext tunnels: null ';	int	
LOAD_LEVEL	Holds the Load level (Spillover/Failover) settings for the tunnel for Harpoon and Skybolt Switch with FOS (8.0.1 or above) which are set in BNA while creating/Editing a Tunnel. The Load Level can either be failover or spillover.	varchar	64

TABLE 312 FCIP_TUNNEL_CIRCUIT

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
TUNNEL_ID	Tunnel ID to which the circuit belongs to	int	
CIRCUIT_NUMBER	Circuit Number of the Circuit from the switch	smallint	
COMPRESSION_ENABLED	Whether Compression is enabled on that circuit	smallint	
TURBO_WRITE_ENABLED	Whether TurboWrite is enabled on that circuit'	smallint	
TAPE_ACCELERATION_ENABLED	Whether TapeAccelaration is enabled on that circuit	smallint	
IKE_POLICY_NUM	The IKE Policy on the circuit.The default value is -1.	int	
IPSEC_POLICY_NUM	The IPSEC Policy on the circuit'. The default value is -1	int	
PRESHARED_KEY	The preshared Key on the circuit	char(	32
SOURCE_IP	SOURCE_IP of the circuit	varchar	64
DEST_IP	DESTINATION_IP of the circuit	varchar	64
VLAN_TAG	VLAN Tag of the circuit. The default value is -1	int	
SELECTIVE_ACK	Select acknowledgement flag.The default value is 0.	smallint	
QOS_MAPPING	QOS Mapping. The default value is 0.	smallint	

TABLE 312 FCIP_TUNNEL_CIRCUIT (Continued)

Field	Definition	Format	Size
PATH_MTU_DISCOVERY	MTU Discovery Path. The default value is 0.	smallint	
MIN_COMM_RATE	Minimum communication Speed. The default value is 0.	int	
MAX_COMM_RATE	Maximum communication Speed. The default value is 0.	int	
MIN_RETRANSMIT_TIME	Minimum Retransmission Time. The default value is -1	int	
MAX_RETRANSMIT_TIME	Maximum retransmission time. The default value is -1	int	
KEEP_ALIVE_TIMEOUT	Keep Alive timeout. The default value is -1	int	
ADMIN_STATUS	Is admin status enabled. The default value is 0.	smallint	
METRIC	Circuit metric to set priority. The default value is -1	int	
DATA_L2_COS	Class of service as defined by IEEE 802.1p for circuit. The default value is -1.	int	
DSCP_DATA	DiffServe marking for Data Frame. The default value is -1	int	
MAX_RETRANSMISSIONS	Max number of Retransmission attempts on the circuit. The default value is 0.	int	
SLOT_NUMBER	Slot number of the circuit. The default value is 0.	smallint	
VE_PORT_NUMBER	VE port number of the tunnel to which the circuit belongs.	int	
SECURITY_FLAG	Security Flag associated with the circuit. The default value is 0.	int	
DSCP_CONTROL	Diffserve marking for control frame. The default value is 0.	int	
CIRCUIT_STATUS	Status of the circuit. The default value is 0.	smallint	
ENABLED	Is circuit enabled. Default: 0, Values: 0|1. The default value is 0.	smallint	

TABLE 312 FCIP_TUNNEL_CIRCUIT (Continued)

Field	Definition	Format	Size
MISMATCHED_CONFIGURATIONS	If a tunnel is down due to mismatched configurations on local and remote end, this property specifies the list of such mismatched configurations.	varchar	1024
CIRCUIT_STATUS_STRING	Circuit Status string value from switch for the tunnel	varchar	256
L2COS_F_CLASS	The default value is 0.	smallint	
L2_COS_HIGH	The default value is 0.	smallint	
L2_COS_MEDIUM	The default value is 0.	smallint	
L2_COS_LOW	The default value is 0.	smallint	
DSCP_F_CLASS	The default value is 0.	smallint	
DSCP_HIGH	The default value is 0.	smallint	
DSCP_MEDIUM	The default value is 0.	smallint	
DSCP_LOW	The default value is 0.	smallint	
FAILOVER_CIRCUIT	Whether the circuit is configured as failover or not.	smallint	
FAILOVER_GROUP_ID	Represents the failover group id for the circuit 0 - Default Failover Group. 1 - 9 Failover Group numbers for the circuits. -1 - Not supported. For the switches running less than FOS 7.2.	int	
DP1_SOURCE_IP	DP1 source IP address of the circuit. This field will be empty if the circuit belongs to a non Skybolt switch or the DP1 source IP address is not configure in Skybolt.	varchar	64
DP1_DEST_IP	DP1 destination IP address of the circuit. This field will be empty if the circuit belongs to a non Skybolt switch or the DP1 destination IP address is not configure in Skybolt.	varchar	64
DP1_VLAN_TAG	Vlan tag of the DP 1 ip address. This will be -1 for the non Skybolt switches.	int	

TABLE 313 FCIP_TUNNEL_PERFORMANCE

Field	Definition	Format	Size
TUNNEL_ID	Primary key of the Switch Port	int	
SWITCH_ID	The number of octets or bytes that have been transmitted by this port. One second periodic polling of the port. This value is saved and compared with the next polled value to compute net throughput. Note, for Fibre Channel, ordered sets are not included in the count	int	
TX	'The number of octets or bytes that have been transmitted by this port. One second periodic polling of the port. This value is saved and compared with the next polled value to compute net throughput. Note, for Fibre Channel, ordered sets are not included in the count	double precision	
RX	The number of octets or bytes that have been received by this port. One second periodic polling of the port. This value is saved and compared with the next polled value to compute net throughput. Note, for Fibre Channel, ordered sets are not included in the count.	double precision	
TX_UTILIZATION	The computed value of TX based on speed of port	double precision	
RX_UTILIZATION	The computed value of RX based on speed of port	double precision	
DROPPED_PACKETS	Number of TCP packets dropped	double precision	
COMPRESSION	Compression ratio	bigint	
LATENCY	Round trip time (latency) in milliseconds	int	
LINK_RETRANSMITS	Number of segments retransmitted	double precision	
RTT_BY_TIME_OUT	Counter of retransmit packets due to timeout	double precision	
RTT_BY_DUP_ACK	Counter of retransmit packets due to duplicate acknowledgement'	double precision	
DUPLICATE_ACK	Counter of duplicate acknowledgement packets	double precision	
ROUND_TRIP_TIME	Round trip time in milliseconds	double precision	
TCP_OUT_OF_ORDER	Counter of TCP out-of-order.	double precision	
SLOW_START	Counter of slow starts	double precision	
LAST_UPDATE_TIME	'Time when this stats record was updated	timestamp	

TABLE 314 FCOE_DEVICE

Field	Definition	Format	Size
DEVICE_NODE_ID	The primary key of the DeviceNode.	int	
DIRECT_ATTACH	Indicates whether the fcoe device is directly attached to the switch''s TE port or to a cloud.	smallint	
ATTACH_ID	The primary key of the port (if direct attached) or cloud (if not direct attached).	int	
MAC_ADDRESS	Mac address of device.	varchar	64

TABLE 315 FCR_ROUTE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
BB_FABRIC_ID	Backbone fabric DB ID.	int	
FCR_FABRIC_ID	FID assigned to edge fabric.	int	
SWITCH_WWN	WWN of the router switch.	varchar	128
NR_PORT_ID	Route parameter.	int	
FCRP_COST	Route parameter.	int	
EX_PORT_WWN	Ex_port WWN.	varchar	128

TABLE 316 FEATURE

Field	Definition	Format	Size
FEATURE_ID*	ID used to uniquely identify the feature.	int	6
NAME	Name of the feature.	varchar	256
DESCRIPTION	Description for the feature.	varchar	256

TABLE 317 FEATURE_EDITION_MAP

Field	Definition	Format	Size
FEATURE_ID*	ID used to uniquely identify the feature.	int	
EDITION_MASK	Used to associate a feature to the edition (Reserved for future).	int	

TABLE 318 FEATURES_USAGE

Field	Definition	Format	Size
CLIENT_IP	Identifies client IP.	varchar	128
USER_NAME	Identifies the feature used user name.	varchar	128
FEATURE_NAME	Identifies the unique feature(module) name.	varchar	128
SUB_FEATURE_NAME	Identifies the sub module name	varchar	128
OPERATION_NAME	Identifies the major operation or action happened in that feature.	varchar	128
LAST_UPDATED_TIME	Identifies the last updated time stamp.	timestamp	
USAGE_COUNT	Count shows how many times the feature is accessed.	int	
FIRST_UPDATED_TIME	Identifies the first updated time stamp.	timestamp	

TABLE 319 FICON_DEVICE_PORT

Field	Definition	Format	Size
DEVICE_PORT_ID*	Value for the device port to which these FICON properties are applied.	int	
TYPE_NUMBER		varchar	16
MODEL_NUMBER	Ficon device model number, such as S18.	varchar	64
MANUFACTURER	Manufacturer of the device, typically IBM.	varchar	64

TABLE 319 FICON_DEVICE_PORT (Continued)

Field	Definition	Format	Size
MANUFACTURER_PLANT	Plant number where the device is manufactured.	varchar	64
SEQUENCE_NUMBER	Device sequence number.	varchar	32
TAG	FICON device property, e.g., 809a or 809b.	varchar	16
FLAG	FICON device property, e.g., 0x10 (hex).	varchar	8
PARAMS	FICON device property string, e.g., Valid channel port.	varchar	16

TABLE 320 FILTER

Field	Definition	Format	Size
ID	Auto incremented integer ID for this table.	serial	
USER_ID	ID of the user, who created this filter and User_id as foreign key to USER_ table.	int	
NAME	Name of the filter.	varchar	128
NETWORK_TYPE	For SAN-1, IP - 2.	int	
CRITERIA	Inputs from user to filter by using this criteria.	jsonb	
TYPE	Filter type for Flow â€“1, It can be update for Fabric, Switch, Port in future.	int	
SUB_TYPE	Flow filter sub types are None - -1, ITL â€“ 11, ZONEALIAS â€“12, ACTIVEZONE â€“ 13, VTAP -14.	int	
LAST_MODIFIED_TIME	Filter last modified time.	timestamp	

TABLE 321 FIRMWARE_FILE_DETAIL

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
FIRMWARE_NAME	Name of the firmware file.	varchar	64
MAJOR_VERSION	Major version bit from the firmware version.	smallint	
MINOR_VERSION	Minor version bit from the firmware version.	smallint	
MAINTENANCE	Maintenance bit from the firmware version.	smallint	
PATCH	Patch bit from the firmware version.	varchar	64
PHASE	Phase bit from the firmware version.	varchar	64
RELEASE_DATE	Release date of the firmware file.	timestamp	
IMPORTED_DATE	Imported date of the file to the Management application.	timestamp	
FIRMWARE_FILE_SIZE	Firmware file size.	int	
FIRMWARE_LOCATION	Firmware file location in the Management application repository.	varchar	1024
RELEASE_NOTES_LOCATION	Release notes file location in theManagement application repository.	varchar	1024

TABLE 321 FIRMWARE_FILE_DETAIL (Continued)

Field	Definition	Format	Size
FIRMWARE_REPOSITORY_TYPE	Repository type to identify the FTP server: 0 = internal FTP. 1 = external FTP.	smallint	
FIRMWARE_TYPE	Firmware Type specifies whether a FOS or AMP firmware type. 0 - FOS, 1 - AMP.	smallint	

TABLE 322 FIRMWARE_SWITCH_DETAIL

Field	Definition	Format	Size
FIRMWARE_ID*	ID for the firmware file.	int	
SWITCH_TYPE*	Switch type that supports this firmware file.	smallint	
REBOOT_REQUIRED	Reboot required flag for the switch type.	smallint	
NUMFILES	Number of files in the firmware.	int	

TABLE 323 FOUNDRY_DEVICE

Field	Definition	Format	Size
DEVICE_ID	Database ID of the DEVICE instance.	int	
IMAGE_VERSION	Firmware image version currently running in the device.	varchar	128
PRODUCT_TYPE	Product type of the device computed based on sysoid and version of main board. To get the main board version for devices, refer octet 28 of snChasMainBrdId MIB in foundry.mib.	varchar	32
FEATURE_MASK	A bit string representing the software features available in the switch/router. Each bit represent a feature and if the feature available then bit value would be 1 and 0 otherwise. Refer snAgSoftwareFeature MIB in foundry.MIB for various features supported and its corresponding details.	bytea	
IS_PORT_VLAN_ENABLED	'Port VLANs enabled for the product or not.	num	(1,0)
ARCHITECTURE_TYPE	Chassis architecture type. Refer snChasArchitectureType MIB in foundry.mib for possible values.	num	(2,0)
BUILD_LABEL	The image label of the built software.	varchar	64
SSL_SLOT	Slot number of the SSL module.	num	(4,0)

TABLE 324 FOUNDRY_MODULE

Field	Definition	Format	Size
MODULE_ID	Unique generated database identifier.	int	
SERIAL_NUM	Serial number of this module.	varchar	32
DRAM_SIZE	Dynamic RAM size in Kilo bytes. Currently it is not populated and used.	num	(4,0)
BOOT_FLASH_SIZE	Boot flash size in Kilo bytes. Currently it is not populated and used.	num	(4,0)

TABLE 324 FOUNDRY_MODULE (Continued)

Field	Definition	Format	Size
MODULE_TYPE	Type of this module. Refer snAgentBrdMainBrdId in foundry.mib for more details and possible values.	num	(4,0)
CODE_FLASH_SIZE	Code flash size in Kilo bytes. Currently it is not populated and used.	num	(4,0)
EXPANSION_MODULE_TYPE	Expansion board type. Refer snAgentBrdExpBrdId in foundry.mib for more details and possible values.	num	(4,0)
EXPANSION_MODULE_DESCRIPTION	The expansion board description string. Expansion board are those boards attaching on the main board.	varchar	128

TABLE 325 FOUNDRY_PHYSICAL_DEVICE

Field	Definition	Format	Size
PHYSICAL_DEVICE_ID	Unique generated identifier.	int	
SERIAL_NUMBER	The serial number of the chassis.	varchar	32
PRODUCT_TYPE	Product type based on sysoid or architecture type and management module main board id.	varchar	32

TABLE 326 FOUNDRY_PHYSICAL_PORT

Field	Definition	Format	Size
PHYSICAL_PORT_ID	Database ID of PHYSICAL_PORT instance.	int	
CONNECTOR_TYPE	The type of connector that the port offers. Refer snSwPortInfoConnectorType of foundry.mib for more details and possible values.	smallint	
MEDIA_TYPE	The media type for the port. Refer snSwPortInfoMediaType of foundry.mib for more details and possible values.	smallint	
GIG_TYPE		smallint	

TABLE 327 FPORT_TRUNK_GROUP

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
VIRTUAL_SWITCH_ID	Virtual switch ID where this F_Port Trunk Group is defined.	int	
MASTER_USER_PORT	User port number for the master port of this trunk.	smallint	
WWN	WWN of the trunk group.	char	23
TRUNK_AREA	User-assigned area number used to group together F_ports of the trunk.	smallint	

TABLE 328 FPORT_TRUNK_MEMBER

Field	Definition	Format	Size
GROUP_ID*	Foreign key to the PORT_TRUNK_GROUP table.	int	
PORT_NUMBER*	Member user port number.	smallint	
WWN	Member port WWN.	char	23

TABLE 329 FRU

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
CORE_SWITCH_ID		int	
TAG	provides the TAG number of FRU element, requested by SMIA and values filled in by Switch Asset Collector. Field probably contains information such as asset tag or serial number data. This value varies depending on the type of physical package	varchar	64
PART_NUMBER	provides the part number of the FRU element, requested by SMIA and values filled in by Switch Asset Collector. Field probably contains the part number assigned by the organization responsible for producing or manufacturing the physical element	varchar	64
SERIAL_NUMBER	provides the serial number of the FRU element, requested by SMIA and values filled in by Switch Asset Collector	varchar	64
VENDOR_PART_ NUMBER	provides the Vendor-assigned part number of this package, requested by SMIA and values filled in by Switch Asset Collector	varchar	64
VENDOR_SERIAL_ NUMBER	provides the Vendor-assigned serial number of this package, requested by SMIA and values filled in by Switch Asset Collector'	varchar	64
CAN_BE_FRUED	provides whether this element can be removed from the switch, requested by SMIA and values filled in by Switch Asset Collector. Maps to IsRemovable field in the html. The default value is -1.	int	
SLOT_NUMBER	provides the slot number of this FRU element , requested by SMIA and values filled in by Switch Asset Collector.The default value is -1.	int	
MANUFACTURER_DAT E	provides the manufactured date of this FRU element, requested by SMIA and values filled in by Switch Asset Collector	timestamp	
UPDATE_DATE	provides the updated date of this FRU element, requested by SMIA and values filled in by Switch Asset Collector	timestamp	
VERSION		varchar	32
MANUFACTURER	provides the manufacturer of this FRU element ,requested by SMIA and values filled in by Switch Asset Collector	varchar	64
VENDOR_EQUIPMENT _TYPE	provides the vendor equipment type of the FRU element, requested by SMIA and values filled in by Switch Asset Collector	varchar	32

TABLE 329 FRU (Continued)

Field	Definition	Format	Size
OPERATIONAL_STATUS	provides the operational status of the FRU element, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above. The default value is -1.	int	
TOTAL_OUTPUT_POWER	provides the total power output of the power supply FRU element, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above. this field is applicable only for the power supply FRU element. The default value is -1.	bigint	
SPEED	provides the speed of the FAN FRU element, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above. this field is applicable only for the FAN FRU element. The default value is -1.	int	
CREATION_TIME	provides the record creation time, standard columns for Management application and values filled in by Switch Asset Collector	timestamp	
LAST_UPDATE_TIME	provides the record creation time, standard columns for Management application and values filled in by Switch Asset Collector	timestamp	
PREVIOUS_OP_STATUS	provides the previous operational status of FRU element, requested by SMIA and values filled in by Switch Asset Collector. Helps identify the operational status transitions. The default value is -1.	int	
VENDOR	This holds the vendor name information for FRU	varchar	256

TABLE 330 FTP_SERVER

Field	Definition	Format	Size
ID*		int	
TYPE	Type indicates the what type of file. Internal FTP - 0, External FTP - 1, External SCP - 2, Internal SCP/SFTP - 3, External SFTP - 4 and Technical support FTP - 100. Technical Support FTP server configuration is created by user to transfer the technical support files from the Management application repository to specified FTP server. Other server configurations can be seen in Options dialog.	smallint	
IP	FTP server IP address.	varchar	64
USER_NAME	FTP server user name.	varchar	64
PASSWORD	FTP server user password.	varchar	512
ROOT_DIRECTORY	FTP server root directory location.	varchar	1024
PORT	Port on which FTP server is configured.	int	

TABLE 331 GENERATED_REPORT

Field	Definition	Format	Size
ID*		int	
NAME	Report name.	varchar	256

TABLE 331 GENERATED_REPORT (Continued)

Field	Definition	Format	Size
TYPE_ID	Report type.	int	
EFCM_USER	Management application user who has generated this report.	varchar	128
REPORT_OBJECT	Report object BLOB.	bytea	
TIMESTAMP_	Timestamp when the report is generated.	timestamp	
FABRIC_NAME	Fabric Name.	varchar	256

TABLE 332 GIGE_PORT

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
SWITCH_PORT_ID	ID for the GigE Port in SWITCH_PORT.	int	
PORT_NUMBER	GigE Port Number(0 for ge0 and 1 for ge1).	int	
SLOT_NUMBER	Slot number on which the GigE Port is present.	int	
ENABLED	Enabled or disabled. Default value is 0.	smallint	
SPEED	Port speed details. Default value is 0.	bigint	
MAX_SPEED	Port maximum speed supported.	bigint	
MAC_ADDRESS	MAC Address of that port.	varchar	64
PORT_NAME	GigE Port Name.	varchar	64
OPERATIONAL_STATUS	LED status.	int	
LED_STATE	LED status.	smallint	
SPEED_LED_STATE	GigE Port type details.	smallint	
PORT_TYPE	Port type for the GigE Port.	varchar	64
PERSISTENTLY_DISABLED	Whether the GigE Port is persistently disabled.	smallint	
INTERFACE_TYPE		smallint	
CHECKSUM		varchar	16
FCIP_CAPABLE	1 = FCIP capable; otherwise, 0. Default value is 2.	smallint	
ISCSI_CAPABLE	1 = ISCSI capable; otherwise, 0. Default value is 2.	smallint	
REMOTE_MAC_ADDRESS	MAC address of attached port of the 10G GigE Port.	varchar	64
INBAND_MANAGEMENT_STATUS	1 = Inband Management status is enabled; otherwise, 0. Default value is 0.	smallint	
OCCUPIED	Default value is 0.	smallint	
LAST_UPDATE		bigint	

TABLE 333 GIGE_PORT_ETHERNET_CLOUD_LINK

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
CLOUD_ID		int	
SWITCH_PORT_ID	The unique id of the switch TE port that this member connects to.	int	
TRUSTED		smallint	
CREATION_TIME		timestamp	
MISSING		smallint	
MISSING_TIME		timestamp	

TABLE 334 GIGE_PORT_STATS

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
SWITCH_ID	References the ID in CORE_SWITCH table.	int	
PORT_ID	References the ID in SWITCH_PORT table.	int	
CREATION_TIME	The polling time.	timestamp	
TX	Transmit (TX) value in bytes.	double precision	
RX	Receive (RX) value in bytes.	double precision	
TX_UTILIZATION	Transmit utilization (TX%) value in percentage.	double precision	
RX_UTILIZATION	Receive utilization (RX%) value in percentage.	double precision	
DROPPED_PACKETS	Number of dropped packets.	double precision	
COMPRESSION	The compression value.	double precision	
LATENCY	The latency value.	double precision	
BANDWIDTH	The bandwidth value.	double precision	

TABLE 335 GLOBAL_VLAN

Field	Definition	Format	Size
GLOBAL_VLAN_DB_ID	Unique database generated identifier.	int	
NAME	Name for Global VLAN.	varchar	255
CONTEXT_DEVICE_ID	Database ID of the DEVICE instance which is associated with global VLAN.	int	

TABLE 336 GRE_TUNNEL_INTERFACE

Field	Definition	Format	Size
INTERFACE_ID	This column is used to store the id of the interface. The value will be populated by the discovery engine where the corresponding GRE Tunnel Interface details will be persisted in the INTERFACE table.	int	

TABLE 337 HA_CLUSTER

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
NAME	User-supplied name for the HA Cluster.	varchar	64
ENCRYPTION_GROUP_ID	Foreign key reference to the ENCRYPTION_GROUP that contains this HA Cluster.	int	
MEMBER_LIST	A comma-separated list of Encryption Engines in the HA Cluster. Each engine is identified by a switch node WWN, followed by "/", followed by the slot number. The slot number is 0 if the switch does not have removable blades.	varchar	256

TABLE 338 HBA

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
HOST_ID	ID of the Device Enclosure (Host) to which this HBA belongs to.	int	
NAME	User defined name of the HBA	varchar	128
POWER_MODE	Power mode of the HBA	varchar	256
MODEL	Model code of the HBA	varchar	256
MODEL_DESCRIPTION	Model description for the HBA	varchar	256
OPERATING_STATUS	Current operating status of the HBA: 1: Enabled/0: Disabled. The default value is 0.	smallint	
CHIP_REVISION	Revision level of the chip used in the HBA	varchar	64
HARDWARE_PATH	Hardware path for the HBA.	varchar	256
SERIAL_NUMBER	Serial number of the HBA	varchar	64
TEMPERATURE	Temperatur of HBA. Both in Celsius/Fahrenheit	varchar	16
USERNAME	User name to be used for logging into the HBA.	varchar	256
PASSWORD	Password used for logging into the HBA	varchar	256
MANAGEMENT_STATE	Management state bit mask, Managed/Auth failed etc. The default value is -1.	int	
MANAGEMENT_INTERFACE	Management interface bit mask, JSON/WMI/SMI etc . The default value is -1.	int	
DRIVER_VERSION	The version level of the host adapter driver	varchar	256
DRIVER_NAME	The name of the HBA driver	varchar	256
FIRMWARE_VERSION	The version level of the firmware	varchar	256
BIOS_VERSION	The version level of the BIOS	varchar	256

TABLE 338 HBA (Continued)

Field	Definition	Format	Size
PCI_REG_VENDOR_ID	The identifier of the PCI Register''s vendor	varchar	32
PCI_REG_DEVICE_ID	The device ID of the PCI Register	varchar	32
PCI_REG_SUBSYSTEM_ID	The ID of the PCI subsystem	varchar	32
PCI_REG_SUBSYS_VENDOR_ID	The ID of the PCI subsystem vendor.	varchar	32
PCI_REG_LANE_COUNT	The number of PCI lanes, in Gbps, each way between the PCI slot and the adapter. The default value is 8.	int	
PCI_REG_NEG_LANE_COUNT	The set number of PCI lanes that were initially negotiated. The default value is 8.	int	
PCI_REG_GENERATION	PCI generation	varchar	256
TRUSTED	Denotes whether HBA is trusted by user or not. When the host first time discovered, all the HBAs will be trusted by default. If any HBA added later, then it will be in untrusted stated. 0 denotes untrusted and 1 is for trusted.	smallint	
CREATION_TIME	HBA record creation time. This tells us when this HBA was first discovered.	timestamp	
MISSING	Denotes whether HBA is missing or not. 0 denotes present and 1 states that HBA is missing from host.	smallint	
MISSING_TIME	States the missing time of the HBA. This will be null if the HBA is available.	timestamp	
CIM_NAMESPACE	Reflects the CIM namespace used to discover the HBA	varchar	128
CARD_TYPE	FC for HBA, CNA for CNA. The default value is 'FC'.	varchar	32
WWN	WWN of the adapter	varchar	23
HCM_AGENT_VERSION	Version of HCM agent used to managed the HBA	varchar	128
MAC_ADDRESS	Adapter mac address	varchar	64
MAX_SPEED_SUPPORTED	The maximum port speed that is supported on the port, in Gb/s. The default value is 0.	int	
VPD_PRODUCT_DESCRIPTION	Description of the product	varchar	256
VPD_PART_NUMBER	Part Number of the device	varchar	32
VPD_EC_LEVEL	EC Level of the device	varchar	32
VPD_FRU_NUMBER	FRU number of the device	varchar	256
VPD_SERIAL_NUMBER	serial number of the device	varchar	32

TABLE 338 HBA (Continued)

Field	Definition	Format	Size
VPD_PW	PW details of the device	varchar	32
VPD_EDC	EDC details of the device	varchar	32
VPD_MDC	MDC details of the device	varchar	32
VPD_FABRIC_GEOGRAPHY	FABRIC_GEOGRAPHY of the device	varchar	256
VPD_LOCATION	LOCATION of the device	varchar	256
VPD_MANUFACTURER_ID	MANUFACTURER_ID of the device	varchar	256
VPD_PCI_GEOGRAPHY	PCI_GEOGRAPHY of the device	varchar	256
VPD_VENDOR_DATA	VENDOR_DATA of the device	varchar	256
VPD_EXT_CAPABILITY	EXT_CAPABILITY of the device	varchar	256
VPD_OEM	OEM details of the device	varchar	256
VPD_OEM_INFO	OEM related information of the device	varchar	256
MAX_PCIF	Maximum number of Pci functions.	smallint	
CARD_MODE	The mode that the card is operating on.	smallint	
DRIVER_CARD_MODE	It is the same as card type but uses new values applicable for 3.0 and later driver versions. Deprecates the card type field. Possible values are: • HBA/CNA/AnyIO/Mezzanine • HBA/Mezzanine CNA/Mezzanine AnyIO	varchar	32
VENDOR	Adapter vendor name.	varchar	128
HOST_DISCOVERY_REQUEST_ID	ID indicates the host discovery request through which the adapter is discovered.	int	
NAMESPACE	This will used to identify the adapter discovered using what name space. this is auto populated based on HBA discovery.	varchar	128

TABLE 339 HBA_NODE_MAP

Field	Definition	Format	Size
DEVICE_NODE_ID	Primary key from the Device Node table	int	
HBA_ID	Primary key from the HBA table	int	

TABLE 340 HBA_PORT

Field	Definition	Format	Size
DEVICE_PORT_ID	Primary key on the owner Device port table	int	
CONFIGURED_STATE	Indicates whether the port is enabled or disabled. The default value is 0.	smallint	
CONFIGURED_SPEED	The configured speed of the port. E.g. Auto-negotiate	varchar	64

TABLE 340 HBA_PORT (Continued)

Field	Definition	Format	Size
CONFIGURED_TOPOLOGY	The topology setting. The default value is 1.	int	
MAX_SPEED_SUPPORTED	The maximum port speed that is supported on the port, in Gb/s. The default value is 0.	int	
OPERATING_STATE	Indicates whether the link is online or offline. The default value is 0.	smallint	
OPERATING_TOPOLOGY	The topology setting at which the port is operating. The default value is 1.	int	
SUPPORTED_FC4_TYPES	List of supported FC4 types for this port.	varchar	32
SUPPORTED_COS	Supported Class of Service (COS) for this port.	varchar	32
TRUSTED	Denotes whether port is trusted or not. 0 denotes untrusted and 1 is for trusted.	smallint	
CREATION_TIME	HBA port record creation time. This tells us when this HBA port was first discovered.	timestamp	
MISSING	Denotes whether port is missing or not. 0 denotes present and 1 states that port is missing from fabric.	smallint	
MISSING_TIME	States the missing time of the this port.	timestamp	
OPERATING_SPEED	Operating speed of the hba port. The default value is 0.	varchar	64
CNA_PORT_ID	Nullable foreign key, related FC pot with the CNA port	int	
PORT_NWWN	Node WWN for the HBA port	varchar	23
PHYSICAL_PORT_WWN	Physical Ports WWN in case of V port	varchar	128
SWITCH_IP	IP of the switch, HBA port is connected to	varchar	23
PRINCIPAL_SWITCH_WWN	WWN of the principal switch of the fabric, HBA is connected to	varchar	128
HBA_ID	HBA ID of the HBA this port belongs to	int	
PORT_NUMBER	Port number of this HBA port.	smallint	
NAME	Name defined for the HBA port in HCM	varchar	
FACTORY_PORT_WWN	Factory configured Port WWN defined for the HBA port in HCM	varchar	
FACTORY_NODE_WWN	Factory configured Node WWN defined for the HBA port in HCM	varchar	
PREBOOT_CREATED	Flag to identify vports created during preboot	varchar	
MAX_BANDWIDTH	Maximum bandwidth	varchar	64
PCIF_INDEX	Pci function index	varchar	64

TABLE 340 HBA_PORT (Continued)

Field	Definition	Format	Size
MAX_PCIF	Maximum number of Pci functions.	smallint	
SYNTHETIC_FC	Synthetic FC is applicable for Windows only: • 0 - Unknow • 1 - Yes • 2 - No.	int	

TABLE 341 HBA_PORT_DETAIL

Field	Definition	Format	Size
DEVICE_PORT_ID	Device port id acts as the primary key	int	
PERSISTENT_BINDING	Persistent binding value of the port. With persistent binding (on the host), one can bind a LUN to a specific device file, thus making sure devices reappear on the same device files after reboots. 0 – disable 1 – enabled	smallint	
FABRIC_NAME	Principal switch WWN of the Fabric to which the port is associated with.	varchar	64
BOOT_OVER_SAN	Flag to indicate whether boot over SAN is enabled or not.. The default value is 0.	smallint	
BOOT_OPTION	Boot option for the port. Possible values are 0 - AUTO_DISCOVERED_FROM_FABRIC , 1 - FIRST_VISIBLE_LUN, 2 - USER_CONFIGURED_LUN	smallint	
BOOT_SPEED	Boot speed for the port in Gbps. Possible values are 0 - AUTO_NEGOTIATE and 2, 4, 8, 16 Gbps. The default value is 0.	int	
BOOT_TOPOLOGY	Boot topology for the port. Possible values are 0 - Point to Point , 1 - Loop. The default value is 1.	int	
BOOTUP_DELAY	On starting system how long system needs to wait for user action. Configured value ranges 0,1,2,5 and 10 minutes. Default value is 0.	int	
BB_CREDIT	The maximum number of receive buffer. The default value is 8.	int	
FRAME_DATA_FIELD_SIZE	The default value is 512.	int	
HARDWARE_PATH	Indicates whether MPIO is enabled or disabled		
V_PORT_COUNT	Number of logical ports. The default value is 0.	int	
QUEUE_DEPTH	The number of I/O operations that can be run in parallel on a device. The default value is 0.	int	
INTERRUPT_CONTROL_COALESCE	Indicates whether interrupt control is on or off. The default value is 0.	smallint	

TABLE 341 HBA_PORT_DETAIL (Continued)

Field	Definition	Format	Size
INTERRUPT_CONTROL_LATENCY	Sets the interrupt control latency value.. The default value is 0.	int	
INTERRUPT_CONTROL_DELAY	Sets the interrupt control delay value.. The default value is 0.	int	
BEACON_STATE	Indicates whether beaconing is on or off.. The default value is 0.	smallint	
LINK_BEACON_STATE	Indicates whether link beaconing is on or off.. The default value is 0.	smallint	
MPIO_MODE_STATE	Indicates whether multipathing mode is on or off.. The default value is 0.	smallint	
PATH_TIME_OUT	The value between 0 to 60 that specifies the time out session. Note you can only enable or edit the path time out when MPIO is disabled. The default value is 0.	int	
LOGGING_LEVEL	The port logging level. Values include Log Critical, Log Error, Log Warning, and Log Info. The default value is 0.	smallint	
TARGET_RATE_LIMIT	Target rate limit of the port. Possible values are 0 -disabled, 1 - enabled. The default value is 0.	smallint	
DEFAULT_RATE_LIMIT	Default target rate limit of the port speed (1 Gbps). The default value is 0.	int	
VF_MODE	True if the port is in VF (Virtual Fabric) mode.	smallint	
RECIEVE_BUFFER_CREDIT	Receiving buffer-to-buffer credits (BB_credits) for the port.	varchar	64
TRANSMIT_BUFFER_CREDIT	Transmitting buffer-to-buffer credits (BB_credits) for the port.	varchar	64
FCSP_AUTH_STATE	Indicates whether FC-SP authentication is on or off. The default value is 0.	smallint	
FCSP_STATUS	The status of FC-SP authentication. The default value is 'Disabled'.	varchar	32
FCSP_ALGORITHM	The configured authentication algorithm. The default value is 'MD5'.	varchar	64
FCSP_GROUP	The DH Group (DH Null, group 0 is the only option). The default value is 0.	smallint	
FCSP_ERROR_STATUS	The health status of the Fibre Channel Security Protocol parameters	varchar	256
QOS_CONFIGURED_STATE	Indicates whether QoS is enabled or disabled. The default value is 0.	smallint	

TABLE 341 HBA_PORT_DETAIL (Continued)

Field	Definition	Format	Size
QOS_OPERATING_STATE	QOS Operating state. The default value is 'Disabled'.	varchar	256
QOS_TOTAL_BB_CREDIT	The number of receive buffers. The default value is 2.	varchar	16
QOS_PRIORITY_LEVEL	QoS priority levels. Values include High, Medium, and Low	varchar	32
QOS_HIGH_BW_ALLOCATION	Percentage of bandwidth allocation for the High priority level.	varchar	32
QOS_MEDIUM_BW_ALLOCATION	Percentage of bandwidth allocation for the Medium priority level	varchar	32
QOS_LOW_BW_ALLOCATION	Percentage of bandwidth allocation for the Low priority level.	varchar	32
MEDIA	media of port	varchar	64
IOC_ID	IO controller ID	int	
PREBOOT_DISABLED	Boolean value indicating if port was disabled during preboot.. The default value is 0.	smallint	
ALARM_WARNING	A bit mask indicating degrading SFP if the bit mask has any 1s in it. It bit mask is all 0s then SFP is in good state.	varchar	32
IO_EXEC_THROTTLE_MAX	Maximum value is 2000. This feature is available for driver 3.1 and later.	int	
IO_EXEC_THROTTLE_OPERATIONAL	Operation value ranges from 0 - 2000.	int	
IO_EXEC_THROTTLE_CONFIGURED	Configured value ranges from 0 - 2000.	int	
FEC_STATE	State of FEC. The FEC (Forward Error Correction) is an error recovery mechanism that allows the receiver of the corrupted frame to correct the error without referring back to the port which transmitted the frame. Supported on prowler card in FC mode. Applicable values are Online, Offline and Not Supported. Note : Not Supported on (PORT_MEDIA_MEZZANINE_CARD).	varchar	128
BB_CREDIT_RECOVERY_STATUS	Status of Buffer to Buffer Credit Recovery. Supported on FC ports. Applicable values are Online, Offline, Not Applicable, and Disable.	varchar	32
CONFIGURED_BB_SCN_COUNT	Configured value of Buffer to Buffer Credit Recovery state change notification count. Range between 1 to 15.	int	
NEGOTIATED_BB_SCN_COUNT	Buffer to Buffer Credit Recovery state change notification count value set by bcu. Range between 1 to 15.	int	

TABLE 342 HBA_PORT_DEVICE_PORT_MAP

Field	Definition	Format	Size
DEVICE_PORT_ID	ID from the device_port table.	int	
HBA_PORT_ID	DEVICE_PORT_ID from the hba_port table.	int	

TABLE 343 HBA_PORT_FCOE_DETAILS

Field	Definition	Format	Size
DEVICE_PORT_ID		int	
BANDWIDTH	The bandwidth percentage of the FCoE port eg. 10 gb for CNA.	int	
FIP_STATE	FIP (Fibre channel Initialization Protocol) state of the port 0 – disable , 1- enabled.	varchar	64
DISCOVERY_PRIORITY	Discovery priority of the port. Currently not used.	varchar	256
FCF_FCMAP	FC Map value of port. Currently not used.	varchar	256
FCF_FPMA_MAC	FPMA (fabric-provided MAC address) MAC address of port. Currently not used.	varchar	64
FCF_MAC	FCF (FCoE Forwarder) MAC value of port.	varchar	64
FCF_MODE	FCF (FCoE Forwarder) Mode of the port. Currently not used.	varchar	256
FCF_NAMEID	FCF (FCoE Forwarder) Name of the port currently Not used.	varchar	256
FCPIM_MPIO_MODE	Indicates whether multipathing I/O (MPIO) mode is turned on or off. 1- on, 0 - off	smallint	
PORT_LOG_ENABLED	True if port log is enabled.	smallint	
MAX_FRAME_SIZE	The frame size, in bytes, of the FCoE port.	int	
MTU	Maximum transmission unit in bytes of the FCoE port. Default - 2112, 0 - auto	int	
PATH_TOV	The value between 0 and 60 that specifies the time-out session. **NOTE:** You can only enable or edit the path time out when MPIO is disabled	int	
SCSI_QUEUE_DEPTH	The LUN queue depth feature determines how many concurrent IOs the adapter will accept and process per LUN (not at the adapter port level, as with the IO throttle value). Not setting the queue depth to the optimal level can result in poor performance, where outstanding IO queuing can cause bottlenecks. For optimum performance, consider both the configuration settings of the HBA and the physical limits on the storage array. If you set the queue depth too low on the HBA it could lead to under-utilization of storage resources. **NOTE:** The Queue Depth feature is supported for all adapter classes configured in FC or FCoE mode (Windows operating systems only)'	int	
STATE	The state of the FCoE port (online or offline).	varchar	64
SUPPORTED_CLASS	The classes supported on the FCoE port. For example, Class2 and Class3.	varchar	256

TABLE 343 HBA_PORT_FCOE_DETAILS (Continued)

Field	Definition	Format	Size
TRL_SPEED	TRL (Target Rate limit) speed. This will be less than max speed supported by this port.	int	
TRL_STATE	TRL (Target Rate limit) state of the port. Possible values are 0 - disable , 1 - Enable	smallint	
PG_ID	The priority group ID. Possible values are 0-7 (user-definable) and 15.0-15.7 (strict priority).	varchar	32
PRIORITIES	'Lists the available priorities (High, Medium, Low).	varchar	128
FCOE_MAC	FCOE MAC address of the port.	varchar	64
IOC_ID	The IO controller Identifier.	int	

TABLE 344 HBA_REMOTE_PORT

Field	Definition	Format	Size
ID	Autogenerate primary column.	int	
SYMBOLIC_NAME	The symbolic name associated with the remote port.	varchar	256
PORT_WWN	The world wide name of the remote device''s port.	char	23
NODE_WWN	The world wide name of the remote device	char	23
NAME	The name associated with the device	varchar	256
FC_ADDRESS	FC Address for the port in hex	varchar	6
FRAME_DATA_SIZE	The frame size, in bytes, of the device. The default value is 512.	int	
SPEED	Operating speed of the remote port.	int	
STATE	Indicates whether the device is online or offline. The default value is 'Offline'.	varchar	64
SUPPORTED_COS	The types of classes that are supported on the remote port; for example, Class-3	varchar	32
DEVICE_TYPE	The type of the device; for example, Disk or Tape.	varchar	64
BIND_TYPE	The persistent bind type. The default value is 0.	smallint	
TARGET_ID	The identifier of the target device. The default value is 0.	int	
ROLE	The role of the device (target or initiator)	varchar	64
VENDOR	The vendor of the device	varchar	256
PRODUCT_ID	The device''s identifier.	varchar	256
PRODUCT_VERSION	Field which stores information regarging target rate limiting on the remote port	varchar	256
QOS_PRIORITY	QOS Priority on the target. The default value is 'Unknown'.	varchar	64
QOS_FLOW_ID	QOS Flow ID on the target. The default value is 0.	varchar	64
CURRENT_SPEED	Current speed of the remote port, as enforced by TRL. The default value is 0.	varchar	64
TRL_ENFORCED	True if TRL(Target Rate limit) is enforced.	varchar	16

TABLE 344 HBA_REMOTE_PORT (Continued)

Field	Definition	Format	Size
BUS_NO	Channel number in the PCI Bus. The default value is 0.	varchar	32
FCP_IM_STATE	Indicates whether the Fibre Channel Protocol Input Method (FCP-IM) is online or offline.	varchar	128
IO_LATENCY_MIN	Minimum IO Latency value (< 79) in milliseconds and is calculated when HBA port starts SCSI (FCP) exchanges	varchar	32
IO_LATENCY_MAX	IO Latency value in milliseconds and is calculated when HBA port starts SCSI (FCP) exchanges	varchar	32
IO_LATENCY_AVERAGE	Average IO Latency value in milliseconds and is calculated when HBA port starts SCSI (FCP) exchanges	varchar	32
DATA_RETRANSMISSION_SUPPORT	Field to indicate whether the remote port supports data retransmission.0 would mean unsupported and nonzero value implies supported. The default value is 0.	smallint	
REC_SUPPORT	Field to indicate whether the remote port supports the REC ELS command Channel number in the PCI Bus.Zero would mean unsupported and nonzero value implies supported. The default value is 0.	smallint	
TASK_RENTRY_IDENT_SUPPORT	The number of PRLI responses from the target to the initiator and begins when HBA Port starts FCP exchanges.Zero would mean unsupported and nonzero value implies supported. The default value is 0.	int	
CONFIRMED_COMPLETIONS_SUPPORT	The number of confirmed completions on the remote port and begins when HBA Port starts FCP exchanges.Zero would mean unsupported and nonzero value implies supported. The default value is 0.	int	

TABLE 345 HBA_REMOTE_PORT_LUN

Field	Definition	Format	size
ID	Auto generated primary key	int	
HBA_REMOTE_PORT_ID	Primary key of owner row in Remote Port	int	
FCP_LUN	The logical unit number of Fibre Channel Protocol (FCP) device. The default value is 0.	varchar	16
CAPACITY	The capacity of the logical unit. The default value is 0.	int	
BLOCK_SIZE	The block size of the logical unit, in bytes (for example, 512 Bytes). The default value is 0.	int	
VENDOR	The vendor of the device to which the logical unit is assigned	varchar	256
PRODUCT_ID	The product identifier of the device to which the logical unit is assigned	varchar	256
PRODUCT_VERSION	The revision level of the device to which the logical unit is assigned.	varchar	256

TABLE 345 HBA_REMOTE_PORT_LUN (Continued)

Field	Definition	Format	size
PRODUCT_SERIAL_NO	The serial number of the device to which the logical unit is assigned	varchar	256
TARGET_WWN	The world wide name of the target devic	char	23
PHYSICAL_LUN	'If there is a lun connected to a remote port, then it represents a value 1 indicating it is a physical lun otherwise it is a dummy lun with value 0. The default value is 1.	smallint	
LUN_ID	IS lun id	varchar	32

TABLE 346 HBA_TARGET

Field	Definition	Format	size
DEVICE_PORT_ID	Primary key from the Device port table	int	
HBA_REMOTE_PORT_LUN_ ID	Primary key from the HBA Remote port lun table	int	
BOOT_LUN	Flag to indicate of the LUN is bootable. The default value is -1.	smallint	
TRUSTED	Denotes whether target is trusted or not. 0 denotes untrusted and 1 is for trusted.	smallint	
CREATION_TIME	Creation time of the entry	timestamp	
MISSING	Flag to indicate if the remote LUN is missing. The default value is 0.	smallint	
MISSING_TIME	Time at which the LUN is marked missing.	timestamp	
TARGET_ID	The identifier of the target device as reported by each HBA port. The default value is 0.	int	

TABLE 347 HEALTH_STATUS

Field	Definition	Format	Size
ID		serial	
DEPLOYMENT_STATUS_ ID	Identifies the execution cycle for the deployment.	int	
RULE_ID	Policy Monitor rule ID.	smallint	
RULE_DESCRIPTION	Description of what the check is about.	varchar	255

TABLE 348 HEALTH_TARGET_STATUS

Field	Definition	Format	Size
ID		serial	
HEALTH_STATUS_ID	Identifies the execution cycle for the deployment.	int	
TARGET_ID	In case of fabric, this is fabric DB ID; for switch, this is switch DB ID; for host, this is host db ID.	int	

TABLE 348 HEALTH_TARGET_STATUS (Continued)

Field	Definition	Format	Size
TARGET_TYPE		smallint	
STATUS	0 - Failed 1 - Successful	smallint	
MESSAGE	Check result message.	varchar	16384
MESSAGE_TYPE		text	
LEGACY_NAME	Target legacy name.	varchar	256

TABLE 349 HOST_DISCOVERY_OPTION

Field	Definition	Format	Size
ID	Auto generated primary key	int	
DISCOVER_JSON	Flag to indicate JSON agent based discovery. The default value is 1.	smallint	
JSON_USERNAME	Username for the JSON agent	varchar	128
JSON_PASSWD	Password for the JSON agent	varchar	512
DISCOVER_CIM	Flag to indicate CIM based discovery. on/off. The default value is 0.	smallint	
CIM_IMPL	CIM implemenation used. 1: SMI, 2: WMI. The default value is 0.	smallint	
CIM_USERNAME	Username for the CIM based agent	varchar	128
CIM_PASSWORD	Password for the CIM based agent'	varchar	512
CIM_NAMESPACE	CIM Namespace. The default value is 'root/brocade	varchar	128
CIM_PORT	Port number used for the CIM agent. The default value is 5988.	int	
DISCOVER_VM	Flag to indicate VM discovery for a host. On/Off'. The default value is 0.	smallint	
VM_USERNAME	Username to be used for VM discovery	varchar	128
VM_PASSWORD	Password to be used for VM discovery	varchar	512
JSON_PORT	Port Number used for the Json agent. The default value is 34568.	int	
VM_PORT	Port Number used for the VM agent. The default value is 443.	int	
*Application_Name*_USER_NAME	Management application User Name of the user who generated the last operation on the request	varchar	255
*Application_Name*_SERVER_ADDRESS	Management application Server address which generated the last operation on this request	varchar	50

TABLE 350 HOST_DISCOVERY_OPTIONS

Field	Definition	Format	Size
ID	Auto generated primary key.		
DISCOVERY_TYPE	Flag to indicate what type of discovery 1 - JSON, 2 - CIM, 3 - WMI	int	
USERNAME	Username to be used for discovery.	varchar	128
PASSWORD	Password to be used for discovery.	varchar	512
PORT	Port Number used for the discovery. for JSON default port is 34568, for CIM using http 5988 using https 5989, for WMI using WinRM 1.1: http 80 https 443 WinRM 2.0: http 5985 https 5986.	int	
SSL_ENABLED	0- disabled 1- enabled	smallint	
NAMESPACE	Multiple name spaces (comma separated). Auto populated based on discovery e.g. root/brocade, root/emulex	varchar	1024

TABLE 351 HOST_DISCOVERY_REQ_GROUP

Field	Definition	Format	Size
ID	Auto generated primary key	int	
NAME	Unique name for the host request. The default value is ' New Host Group'.	varchar(	256
DISCOVERY_OPTIONS_ID	Primary key from the host discovery options table. Points to the associated discovery options	int	
MANAGEMENT_STATE	Reflects the status of the request E.g. 0-> Completed, 1->Delete Pending. The default value is 0.	int	

TABLE 352 HOST_DISCOVERY_REQUEST

Field	Definition	Format	Size
ID	Autogenerated primary key	int	
HOST_NAME	Hostname: IP address or host name	varchar	256
DEVICE_ENCLOSURE_ID		int	
REQUEST_GROUP_ID	Primary key from the request group table. Null allowed	int	
HOST_DISCOVERY_OPTION_ID	This id is a foreign key to the id in the host_discovery_option table. The default value is -1.	int	
VM_MANAGEMENT_STATE	The status of VM Discovery indicating success or failure. The default value is 0.	int	

TABLE 352 HOST_DISCOVERY_REQUEST (Continued)

Field	Definition	Format	Size
JSON_MANAGEMENT_STATE	The status of HBA discovery using JSON agent, indicating success or failure. The default value is 0.	int	
CIM_MANAGEMENT_STATE	The status of HBA Discovery using CIM, indicating success or failure. The default value is 0.	int	
MANAGEMENT_STATE	Reflects the status of the request E.g. 0-> Completed, 1->Add Pending 2->Delete Pending 3->Edit Pending 4->Delete Failed. The default value is 1.	int	
MANAGEMENT_STATE_DETAILS	This field holds the detailed information on the management state if available. For example, in case of management state 3, this field will have details on all the manually created conflicting enclosures.	varchar	1024

TABLE 353 HOST_DISCOVERY_REQUESTS

Field	Definition	Format	Size
ID	Auto generated primary key.		
HOST_NAME	Hostname: IP address or host name.	varchar	256
DEVICE_ENCLOSURE_ID	Identifier of the Host which this request is associated.	int	
HOST_DISCOVERY_REQ_GROUP_ID	Identifier of the Host discovery request group which this request is associated.	int	
HOST_DISCOVERY_OPTIONS_ID	Reference to the Host discovery options associated with this request.	int	
DISCOVERY STATE	The status of Discovery indicating success or failure. 1 - Success 2 - Authentication failed 3 - Connection Failed 4 - Unknown error 5 - conflicting Manual or auto enclosure found 6 - No Adapter found 7 - SSL certificate issue	int	

TABLE 353 HOST_DISCOVERY_REQUESTS (Continued)

Field	Definition	Format	Size
REQUEST STATE	Reflects the status of the request. 0 - Completed 1 - Add Pending 2 - Edit Pending 4 - Delete Pending 8 - Inactive/deleted request 16 - Redundant request 32 - Collection submission failed 63 - Request internal error 128 - Request DB Update done, waiting for collection 256 - Rediscover pending	int	
REQUEST_STATE_DETAILS	'This field holds the detailed information on the request/discovery state if available. For example, in case of discovery state 5, this field will have details on conflicting manually/auto created enclosures.	varchar	1024

TABLE 354 HYPER_V_VIRTUAL_MACHINE

Field	Definition	Format	Size
ID	Primary Key	int	
VM_NAME	Name of the Virtual Machine.	varchar	256
COMPUTER_NAME	Hyper V host name.	varchar	256
CONFIGURATION_LOCATION	Path where VM configuration data is located.	varchar	256
GUID	Globally Unique ID of the VM.	varchar	256
HARD_DRIVES_COUNT	Count of virtual hard drives in the VM.	int	
MEMORY_ASSIGNED	Amount of memory assigned to the VM.	varchar	256
PATH	Path to the primary disk of the Virtual Machine.	varchar	256
PROCESSOR_COUNT	Number of virtual CPUs of the VM.	int	
STATUS	Status of the VM.	varchar	64
STATE	Operational State of the VM.	varchar	64
NOTES	Notes describing the VM.	varchar	2048
UPTIME	The time since the VM was last powered up.	varchar	512

TABLE 355 HYPER_V_VM_HBA_PORT_MAP

Field	Definition	Format	Size
ID	Primary Key	int	
HYPER_V_VM_ID	ID of the HYPER_VIRTUAL_MACHINE instance.	int	
HBA_PORT_ID	ID of the HBA_PORT instance which is a Hyper V Virtual FC port.	int	

TABLE 356 IFL

Field	Definition	Format	Size
ID*	Primary key for this table. Serial number which is uniquely generated by DB.	int	
EDGE_FABRIC_ID	Edge fabric ID of this IFL link.	int	
EDGE_PORT_WWN	Edge switch port wwn of this IFL link.	varchar	128
BB_FABRIC_ID	Backbone fabric ID of this IFL link.	int	
BB_PORT_WWN	Backbone fabric switch port wwn of this IFL link.	varchar	128
BB_RA_TOV	Backbone fabric resource allocation time out value specified in milliseconds.	int	
BB_ED_TOV	Backbone fabric Error detect time out value specified in milliseconds.	int	
BB_PID_FORMAT	Backbone fabric port identifier format.	smallint	

TABLE 357 IFL_CONNECTION

Field	Definition	Format	Size
ID	The primary key of the table.	int	
EDGE_FABRIC_ID	Edge fabric ID.	int	
EDGE_SWITCH_ID	Edge switch ID, only present when attached fabric is managed.	int	
EDGE_PORT_ID	Edge switch port ID, only present when attached fabric is managed.	int	
EDGE_PORT_WWN	Edge switch port WWN.	varchar	23
BB_FABRIC_ID	Backbone fabric ID.	int	
BB_SWITCH_ID	Backbone FCR virtual switch.	int	
BB_PORT_ID	Backbone switch port ID.	int	
BB_PORT_WWN	Backbone switch port WWN.	varchar	23

TABLE 358 INTERFACE

Field	Definition	Format	Size
INTERFACE_ID		int	
SWITCH_SERVICE_ID		int	
DEVICE_ID		int	
NAME		varchar	255
IDENTIFIER		varchar	128
TABLE_SUBTYPE		varchar	255
TAG_MODE		smallint	
VLAN_TAG_TYPE		int	
UNTAGGED_VLAN_ID	The existing Data type short has been modified to integer. Hence it supports 16 bit additionally.	int	
IF_NAME		varchar	64

TABLE 358 INTERFACE (Continued)

Field	Definition	Format	Size
LLDP_PORT_ID_SUBTYPE		smallint	
LLDP_PORT_ID		bytea	
IS_FDP_ENABLED		num	(1,0)
IS_CDP_ENABLED		num	(1,0)
PORT_STATUS		smallint	
PORT_STATE		smallint	
IF_INDEX	This column is used to store the ifIndex of the interface. The value will be populated by the DCB collector during the discovery of the DCB switch. Since this value is not populated by IP discovery engine, making the field as nullable.	int	
AMPP_PROFILE_MODE	Specifies whether the interface is set to AMPP profile mode.	smallint	
DOT1D_PORT_NUM	To store dot1d port number in DB to reduce SNMP calls to switch from IfIndexUtility	int	
EDGE_TYPE	Indicates the type of device connected to this interface. -1 - Not Applicable 0 - Connected to device with unknown type 1 - Connected to managed Brocade branded AP 2 - Connected to standalone Brocade branded AP. 3 - Connected to a storage device 4 - Connected to a Server	int	
USER_DEFINED_VALUE_1	User defined value used for IP Port.	varchar	256
USER_DEFINED_VALUE_2	User defined value used for IP Port.	varchar	256
USER_DEFINED_VALUE_3	User defined value used for IP Port.	varchar	256
FEATURES_SUPPORTED	Contains the features supported as a bit mask at port level. Possible values are: 1 - Flex port (can be converted to fiber channel or Ethernet port)	int	
CONNECTED_STORAGE_TYPE	This column holds the type of the storage device (NAS/iSCSI/Others) connected to this interface. This is applicable only when "EDGE_TYPE" is set as "Storage(3)" otherwise this gets the default value "Others(0)". 0- Others 1-NAS 2-iSCSI	int	

TABLE 359 INTERFACE_DEPLOYMENT_CONFIG

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
DEPLOYMENT_ID	Deployment configuration ID. Foreign Key for DEPLOYMENT table.	int	
CLEAR_CONFIGURATION	1/0 corresponding to "Clear Assignment" / "Assign Configuration" for interface level configuration.	smallint	
WRITE_TO_DEVICE	1/0 corresponding to Write to device/not write to device for outbound traffic.	smallint	
BINDING_DIRECTION	Represents the binding direction. 0/1 corresponds to IN / OUT direction.	smallint	

TABLE 360 IP_DEVICE_LICENSE

Field	Definition	Format	Size
ID	Primary Key field for the DEVICE_LICENSE	int	
DEVICE_ID	This is the foreign key reference to the Device table	int	
HASH	A unique hash for identifying a license entry in the device. This helps to traverse through the entries with same package name and LID.	varchar	24
PACKAGE_NAME	Name of the license package. Package defines the features enabled by the license. Example:SW-NI-CES-2024-L3U	varchar	64
LICENSE_ID	License ID of the chassis or the line module for which, this entry displays license information.Example: fJucJFgFHG	varchar	24
LICENSE_TYPE	The type of the license, which can be either normal or trial. Values are: permanent(1), trial(2).The default value is 1.	smallint	
EXPIRY_DATE	Expiry Date of the trial license. For normal license, the value is 0.	varchar	19
PRECEDENCE	Defines the priority of a particular trial license among those having the same package and License ID. This is primarily used for determining which license to use, when there are many trial and normal licenses with same package name and LID. The value range is (0..65535)	int	
LICENSE_STATE	This indicates the state of the license. Possible values:invalid(1),unused(2),active(3),expired(4)	smallint	

TABLE 361 IP_INTERFACE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
ETHERNET_PORT_ID	GigE Port ID.	int	
IP_ADDRESS	IP address on the Ip_interface.	varchar	64

TABLE 361 IP_INTERFACE (Continued)

Field	Definition	Format	Size
NET_MASK	Subnet mask for the interface.	varchar	64
MTU_SIZE	MTU Size for that interface.	int	
CHECKSUM	Check Sum.	varchar	64
GIGE_PORT_TYPE	Whether the IP interface is created on a 10G cross port or not. Non-zero value denotes a cross port.	smallint	
VLAN_ID	Vlan id for the Skybolt IP interfaces. For non Skybolt ip interfaces it will be -1.	int	
DP_NUMBER	Indicates the DP number of the GigE port.	int	

TABLE 362 IP_PORT_GROUP

Field	Definition	Format	Size
PORT_GROUP_ID	Unique database generated identifier.	int	
NAME	Name for Port group.	varchar	64
USER_ID	Database ID of the USER_ instance refer a user who created the group.	int	
DESCRIPTION	Description for Port group.	varchar	255
IS_PUBLIC	Represents if the port group is public or not. private-0, public-1.	num	(1,0)
IS_AP_GROUP	Represents if the group created using AP port(s) or not. Non-AP Port group-0, AP Port group-1.	num	(1,0)

TABLE 363 IP_ROUTE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
ETHERNET_PORT_ID	GigE Port ID.	int	
PORT_NUMBER	Port Number related to the GigE Port.	int	
SLOT_NUMBER	Slot Number related to the GigE Port.	int	
NET_MASK	Subnet Mask for the Route.	varchar	64
GATEWAY	Gateway for the Route.	varchar	64
IP_ADDRESS	IP Address created after '"&"' operation of gateway.	varchar	64
METRIC	Metric.	int	
FLAG	Flag.	int	
CHECKSUM	Check Sum.	varchar	64
GIGE_PORT_TYPE	Whether the IP interface is created on a 10G cross port or not. Non-zero value denotes a cross port.	GIGE_PORT_ TYPE	
DP_NUMBER		int	

TABLE 364 IP_SUBNET_VLAN

Field	Definition	Format	Size
VLAN_DB_ID	Database ID of the VLAN instance which is associated with the IP subnet.	int	
IP_ADDRESS	IP address for subnet.	varchar	40
SUBNET_MASK	Subnet Mask of the IP subnet.	varchar	40

TABLE 365 IP_USER_DEFINED_DEVICE_DETAILS

Field	Definition	Format	Size
MAC_ADDRESS	MAC address of the attached end device.	varchar	64
IP_ADDRESS	IP Address of the attached end device.	varchar	256
TYPE	Indicates what type of end device. 0 - Storage, 1 - Server	int	
IP_STORAGE_TYPE	Indicates what type of storage is connected to interface (NAS/iSCSI/Unknown) configured by user. 0-Unknown 1-NAS 2-iSCSI	int	

TABLE 366 IPX_NETWORK_VLAN

Field	Definition	Format	Size
VLAN_DB_ID	Database ID of the VLAN instance which is associated with the IPX network.	int	
NETWORK_NUMBER	Number for IPX network.	varchar	32
FRAME_TYPE	Frame type for IPX. Possible values are 0-Not Applicable, 1-802.2, 2-802.3, 3-Ethernet II and 4-SNAP.	num	(4,0)

TABLE 367 ISL

Field	Definition	Format	Size
ID*	Primary key for the table.	int	
FABRIC_ID	Fabric ID of the associated fabric for this ISL.	int	
SOURCE_DOMAIN_ID	Source domain ID of the ISL.	int	
SOURCE_PORT_NUMBER	Source port number of the ISL.	smallint	
DEST_DOMAIN_ID	Destination or remote domain ID of the ISL.	int	
DEST_PORT_NUMBER	Destination or remote port number of the ISL.	smallint	
COST	The cost of the ISl link.	int	
TYPE	The type of link.	smallint	
TRUSTED	Denotes whether ISL link is trusted or not. • 0 denotes untrusted • 1 denotes trusted.	smallint	
CREATION_TIME	Creation time of the ISL record in the Management application database.	timestamp	

TABLE 367 ISL (Continued)

Field	Definition	Format	Size
MISSING	Denotes whether ISL link is missing or not. • 0 denotes present • 1 states that ISL is missing	smallint	
MISSING_TIME	States the missing time of the this ISL.	timestamp	
missing_reason	The ISL disabled reason. For an ISL either one or both ends might have been disabled. This field will capture the port disable message from both side of ISL. The data is formatted as follows: "<port_wwn>: <disabled_reason> ; <port_wwn>: <disabled_reason>".	varchar	1024
TRUNKED	Determines whether the isl is part of a trunk or not. The value of 0 means not trunked, 1 means this isl is part of a trunk and -1 means not applicable status. Default value is -1.	smallint	

TABLE 368 ISL_CONNECTION

Field	Definition	Format	Size
ID	The primary key of the table.	int	
FABRIC_ID	This is the fabric ID	int	
SOURCE_SWITCH_PORT_ID	The Switch port ID of the Source Switch (local end of the ISL). Maintained as a nullable foreign key to account for ports being moved from one VF to other.	int	
TARGET_SWITCH_PORT_ID	The Switch port ID of the Target Switch (remote end of the ISL). Maintained as a nullable foreign key to account for ports being moved from one VF to other.	int	
COST	Cost of the ISL link.	int	
TYPE	Type of the IS.	int	
TRUSTED	Denotes whether ISL link is trusted or not. 0 denotes untrusted and 1 is for trusted.	int	
MISSING	Denotes whether ISL link is missing or not. 0 denotes present and 1 states that ISL is missing.	int	
MISSING_TIME	Missing timestamp.	timestamp	
MISSING_REASON	The ISL disabled reason. For an ISL either one or both ends might have been disabled. This field will capture the port disable message from both side of ISL. The data is formatted as follows: "<port_wwn>: <disabled_reason> ; <port_wwn>: <disabled_reason>"	varchar	1024
CREATION_TIME	Creation timestamp.	timestamp	
TRUNKED	This column is used to determine whether the isl is part of a trunk or not. The value of 0 means not trunked, 1 means this isl is part of a trunk and -1 means not applicable status. Default value is -1.	int	

TABLE 368 ISL_CONNECTION (Continued)

Field	Definition	Format	Size
MASTER_CONNECTION_ID	This will hold the id of the master ISL connection for a ISL between trunk members. The ISL Connection between masters will have its own ID in this column. Non trunk ISLs will have the default value of -1.	int	
SOURCE_MASTER_PORT	This column will hold the trunk master port for the source port, if the connection is trunked. For the master connection it will have its source por"s port number. For non-trunk connections it will have the default value -1.	int	
TARGET_MASTER_PORT	This column will hold the trunk master port for the target port, if the connection is trunked. For the master connection it will have its target port"s port number. For non-trunk connections it will have the default value -1.	int	
SOURCE_PORT_SPEED	Currently configured speed on the source switch port.	int	
TARGET_PORT_SPEED	Currently configured speed on the target switch port.	int	

TABLE 369 ISL_TRUNK_GROUP

Field	Definition	Format	Size
ID*	Primary key for the table.	int	
VIRTUAL_SWITCH_ID	Foreign key reference to Virtual Switch record associated with the trunk group.	int	
MASTER_USER_PORT	Stores the master user port for the ISL trunk..	smallint	
TRUSTED	Denotes whether ISL trunk group is trusted or not. 0 denotes untrusted and 1 is for trusted.	smallint	
MISSING	Denotes whether ISL trunk group is missing or not. 0 denotes present and 1 states that ISL trunk is missing	smallint	
MISSING_TIME	States the missing time of the this ISL trunk group. If the trunk is not missing then it will be null	timestamp	
MEMBER_TRACKING_STATUS	Member added/removed status of this trunk. This is represented as bitmap value. Each bit is set based on membership state change. Currently only 2 bits from LSB are used. Bit 1 - Member added Bit 2 - Member removed For example if the trunk group has membership change (some members are added and some existing members are removed) then the value would be 3.	int	

TABLE 370 ISL_TRUNK_MEMBER

Field	Definition	Format	Size
GROUP_ID*	Foreign key reference to the trunk group table for this member.	int	
PORT_NUMBER*	Member port number for this trunk..	smallint	
TRUSTED	Denotes whether ISL trunk member is trusted or not. 0 denotes untrusted and 1 is for trusted.	smallint	
MISSING	Denotes whether ISL trunk member is missing or not. 0 denotes present and 1 states that ISL trunk member is missing.	smallint	
MISSING_TIME	We could change this as "States the missing time of the this ISL trunk member. If the member is not missing then it will be null.	timestamp	

TABLE 371 ISNS_CONFIG

Field	Definition	Format	Size
ID	Primary key for the table.	serial	
CLUSTER_ID	NOS cluster database identifier.	int	
ISNS_IP_ADDRESS	ISNS server IP address configured for device.	varchar	128
ESI_TIMEOUT	Configured ESI timeout in seconds. Default is 300 seconds. Min value is 120 seconds and max value is 3600 seconds.	int	
VRF_VALUE	Configured VRF value. For NOS 7.1. release it is preconfigured as 1 and cannot be changed.	int	
CONFIG_CONTENT	Formatted content which contains Discovery Domain sets and Discovery domains configured in cluster. It is formatted by Network Advisor application for easier management.	text	
CONFIG_CHECKSUM	Computed SHA-1 CHECKSUM value for configuration content. This value is used to determine whether the content is changed or not before updating the content in the database.	varchar	64
NO_OF_DD	Number of discovery domains configured in the cluster.	int	
NO_OF_DDS	Number of discovery domain sets configured in the cluster.	int	
NO_OF_ACTIVE_DDS	Number of discovery domain sets active in the cluster.	int	
LAST_MODIFIED_TIME	Last time the configuration is updated in database. It could be by collector which collects the configuration from cluster or updated when user configures through BNA client.	timestamp	

TABLE 372 ISNS_DEVICE

Field	Definition	Format	Size
ID	Primary key for the table.	serial	
CLUSTER_ID	NOS cluster database identifier.	int	
DEVICE_NAME	Fully Qualified Device Name in the form of iqn.yyyy-mm.naming-authority:unique name.	varchar	256

TABLE 372 ISNS_DEVICE

Field	Definition	Format	Size
DEVICE_IP	IP Address of the Device.	varchar	128
DEVICE_TYPE	The type of the device : Initiator or Target.	varchar	64
ENTITY_ID	Key attribute uniquely identifies each Network Entity registered in the ISNS server.	varchar	275

TABLE 373 KEY_VAULT

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
IP_ADDRESS	The IP Address (IPv4, IPv6, or hostname) of the key vault	varchar(	512
PORT_NUMBER	The TCP port number for the key vault	int	
PUBLIC_CERTIFICATE	The key vault"s public key certificate. Switches use this to establish a secure connection to the key vault	varchar(	4096
CERTIFICATE_LABEL	A text name to identify the certificate	varchar(	256
POSITION_	Specifies whether this key vault is the primary key vault or the backup key vault. 0 = primary, 1 = backup.	smallint	
VENDOR_NAME	Indicates the name of the key vault vendor. For non KMIP key vaults, this column will contain value as Not Applicable.	varchar	256

TABLE 374 L2_ACCESS_CONTROL_ENTRY

Field	Definition	Format	Size
ID		serial	
ACCESS_CONTROL_LIST_ID	L2 access control list ID, to which the ACL entry is associated.	int	
SEQUENCE_NUMBER		int	
ACTION	Specifies the action: 0 = Permit 1 = Deny	smallint	
SOURCE_MAC	Source MAC address.	varchar	24
SOURCE__MASK	Source MAC address mask.	varchar	24
DEST_MAC	Destination MAC address.	varchar	24
DEST_MASK	Destination MAC address mask.	varchar	24
VLAN_ID	Specifies the VLAN ID for the L2 ACL entry.	int	
ETHERNET_TYPE		varchar	24
LOG_ENABLE	Specifies whether logging is enabled or not: 1 = Enabled 0 = Not Enabled	smallint	

TABLE 374 L2_ACCESS_CONTROL_ENTRY (Continued)

Field	Definition	Format	Size
SOURCE_TYPE	Indicates the source MAC type (any, host or mac) for DCB Switch L2 ACL entry.	varchar	24
DEST_TYPE	Indicates the destination MAC type (any, host or mac) for DCB Switch L2 ACL entry.	varchar	24

TABLE 375 L2_ACCESS_CONTROL_LIST

Field	Definition	Format	Size
ID		serial	
ACCESS_CONTROL_LIST_TYPE	Specifies the ACL Type: • 0 = Standard • 1= Extended	smallint	
ACCESS_CONTROL_LIST_NUM	L2 ACL number.	int	
ACCESS_CONTROL_LIST_NAME	L2 ACL name.	varchar	255
STARTING_SEQUENCE_NUMBER		int	
INCREMENTAL_VALUE		int	
CLEAR_STATS	• 1 = Clear stats is enabled. • 0 = Clear stats is disabled.	smallint	
OPERATION_TYPE		varchar	10
DEPLOYMENT_ID	Deployment configuration ID. Foreign Key for DEPLOYMENT table.	int	
INT_BINDING_DIRECTION	Represents the interface binding direction. 0/1/2 corresponds to IN /OUT/BOTH direction.	smallint	

TABLE 376 L2_ACL_DEVICE_DEPLOY_MAP

Field	Definition	Format	Size
DEPLOYMENT_ID	Deployment configuration ID. Foreign Key for DEPLOYMENT table.	int	
L2_ACCESS_CONTROL_LIST_ID	L2 Access control List ID for reference to the L2_ACCESS_CONTROL_LIST. Foreign Key for L2_ACCESS_CONTROL_LIST table.	int	

TABLE 377 L2_ACL_INTERFACE_DEPLOY_MAP

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
DEPLOYMENT_ID	Deployment configuration ID. Foreign Key for DEPLOYMENT table.	int	
INBOUND_L2_ACL_ID	L2 Access control List ID of the L2 ACL selected for inbound. Foreign Key for L2_ACCESS_CONTROL_LIST table.	int	
INBOUND_WRITE_TO_DEVICE	1/0 corresponding to Write to device/not write to device for inbound traffic.	smallint	

TABLE 377 L2_ACL_INTERFACE_DEPLOY_MAP (Continued)

Field	Definition	Format	Size
OUTBOUND_L2_ACL_ID	L2 Access control List ID of the L2 ACL selected for outbound. Foreign Key for L2_ACCESS_CONTROL_LIST table.	int	
OUTBOUND_WRITE_TO_DEVICE	1/0 corresponding to Write to device/not write to device for outbound traffic.	smallint	

TABLE 378 L2_NEIGHBOR

Field	Definition	Format	Size
L2_NEIGHBOR_ID		int	
INTERFACE_ID		int	
RMT_IP_ADDRESS		varchar	40
RMT_IF_NAME		varchar	256
LAST_SEEN_TIME		int	
LLDP_REM_CHASSIS_ID_SUBTYPE		smallint	
LLDP_REM_CHASSIS_ID		bytea	
LLDP_REM_PORT_ID_SUBTYPE		smallint	
LLDP_REM_PORT_ID		bytea	
LLDP_REM_CHASSIS_ID_VALUE	To store the MAC or Network address value in ascii format	varchar	40
LLDP_REM_PORT_ID_VALUE	To store the MAC or Network address value in ascii format	varchar	40

TABLE 379 L3_ACL_DEVICE_DEPLOYMENT_MAP

Field	Definition	Format	Size
DEPLOYMENT_ID	Deployment configuration ID.	int	
L3_ACCESS_CONTROL_LIST_ID		int	

TABLE 380 L3_ACL_INT_DEPLOYMENT_MAP

Field	Definition	Format	Size
ID		serial	
DEPLOYMENT_ID	Deployment configuration ID.	int	
INBOUND_L3_ACL_ID	L3 Access control List ID of the L3 ACL selected for inbound.	int	
INBOUND_WRITE_TO_DEVICE	1/0 corresponding to Write to device/not write to device for inbound traffic.	smallint	
OUTBOUND_L3_ACL_ID	L3 Access control List Id of the L3 ACL selected for outbound. Foreign Key for ACCESS_CONTROL_LIST table.	int	
OUTBOUND_WRITE_TO_DEVICE	1/0 corresponding to Write to device/not write to device for outbound traffic.	smallint	

TABLE 381 LAG

Field	Definition	Format	Size
ID	DB ID of LAG(Port-Channel).	int	
VIRTUAL_SWITCH_ID	FK to owning VIRTUAL_SWITCH	int	
LAG_ID	LAG ID	int	
IF_INDEX	Interface index	int	
IF_NAME	Interface name	varchar	256
ENABLED	LAG is enabled=1, disabled=0	smallint	
LAG_MODE	Static or dynamic (1=dynamic, 2=static)	smallint	
ACTIVE	LACP active or passive (1=active, 2=passive) valid if mode=dynamic	smallint	
TYPE	Trunking type (1=standard, 2=brocade, 3=hybrid)	smallint	
IF_MODE	L2 or L3 mode	varchar	8
L2_MODE	Type of L2 mode (default=access	varchar	32
MAC_ACL_POLICY	stores the MAC ACL policy information of the LAG	varchar	64
VLAN_LIST	Comma separated vlan ID list.	text	
MAC_ADDRESS	MAC address of LAG(Port-Channel).	varchar	64
IP_ADDRESS	Primary IPAddress of the LAG	varchar	128
NET_MASK	Netmask of the Primary IPAddress of the LAG	varchar	128
MINIMUM_LINKS	Least number of operationally UP links to declare the port-channel UP. range 1..16.	int	
MTU	Maximum transmission unit in bytes. range 1522..9208.	int	
LOAD_BALANCE	Load balancing details.	varchar	64
VLAG	Specifies whether the lag is a vlag or not.	smallint	

TABLE 382 LAG_MEMBER

Field	Definition	Format	Size
ID	DB ID of LAG member(port).	int	
LAG_ID	FK to owning LAG	int	
NAME	Member name	varcha	64
TYPE	currently not used. The default value is 0.	smallint	
MEMBER_MODE	Dynamic Mode Active/passive. The default value is 0.	smallint	

TABLE 383 LAST_CONFIG_UPDATE_TIME

Field	Definition	Format	Size
ID	Primary key.		
MANAGED_ELEMENT_ID	The managed element id of the device. This is the foreign key to MANAGED_ELEMENT table.	int	

TABLE 383 LAST_CONFIG_UPDATE_TIME (Continued)

Field	Definition	Format	Size
CONFIG_XPATH	The xpath string.	varchar	1024
LAST_UPDATE_TIME	Timestamp returned by the device for this particular xpath.	bigint	

TABLE 384 LAUNCH_IN_CONTEXT_MODULE

Field	Definition	Format	Size
NAME	Unique dialog name used as a module name when launching in context.	varchar	64
DESCRIPTION	Description about the dialog features.	varchar	256
XML_FILE_NAME	The dialog XML XUL file name used to launch the dialog.	varchar	64
PRIVILEGE_ID	This is the comma separated list of privilege IDs required to launch this dialog. This is either the list of values from PRIVILEGE.ID column or -1 if no privilege is required to launch this dialog.	varchar	64
READ_WRITE_ACCESS	Specifies the read or write access privilege required to launch this dialog. 0 = no access is required to launch this dialog. 1 = At least the read-only access is required for the above privilege to launch this dialog. 2 = The read-write access is required for the above privilege to launch the dialog.	int	
EMPTY_DIALOG_ALLOWED	This field indicates whether the dialog can be launched even when there are no fabrics discovered. • 0 = Yes • 1= No	int	
INTERNAL_MODE_DIALOG	The DCFM main client is not visible when the dialog is launched in internal mode. This mode is used when launching from SMIA config tool. • 0- No • 1- Only internal mode • 2- Internal and external	int	
LICENCE_PACKAGE_TYPE	Column to indicate whether the dialog is related to SAN or IP license package type. • 0 = SAN package • 1 = IP Package	int	
OPTIONAL_PARAMS	Comma separated names of all the optional parameters such as WWN.	varchar	256
OPTIONAL_PARAMS_DESC	Comma separated descriptions for the above optional parameters.	varchar	1024

TABLE 385 LICENSE

Field	Definition	Format	Size
ID	Unique Number assigned for the license information.	int	
LICENSE_KEY	License key string which has encoded value of number of products, ports licensed and package which this license isapplicable, etc.	varchar	1024
SERIAL_NO	Unique serial number string that helps to identify the customer or organization which this license is issued for.	varchar	255
CREATION_TIME	Time at which this license key is added	timestamp	
TYPE	Type of license: • 0 - Trial, • 1 - Permanent. The default value is 0.	smallint	
SUB_TYPE	Sub Type of license: • 0 - Base, • 1 - Addon. The default value is 0.	smallint	
VALID	Is this license still considered: • 0 - No, • 1 - Yes. The default value is 1.	smallint	

TABLE 386 LICENSE_DOWNGRADE_DETAILS

Field	Definition	Format	Size
ID	Primary key ID.		
PREVIOUS_LICENSE_INFO	Previous License information during downgrade. The details will have license type, license count like fabric, device, port etc.	varchar	512
NEW_LICENSE_INFO	New License information during downgrade. The details will have license type, license count like fabric, device, port etc.	varchar	512
DOWNGRADE_TIME	Time when License is downgraded.	timestamp	
DOWNGRADED_BY	User who performed license downgrade.	int	
IS_ACTIVE	Takes the value 0 or 1. • 1 - currently active downgrade • 0 - inactive or older downgrade.	smallint	

TABLE 387 LICENSE_FEATURE_MAP

Field	Definition	Format	Size
LICENSE_ID*	Foreign Key (SWITCH_LICENSE.ID) and is part of the primary key.	int	
FEATURE_ID*	Foreign Key (LICENSED_FEATURE.ID) and is part of the primary.	int	
LICENSED_SLOTS	Holds array of licensed slots where each element indicates the licensed slot number.	int	

TABLE 388 LICENSE_RULE

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
NAME	Name of the license rule	varchar	
DESCRIPTION	Description of the rule	varchar	
SCOPE	Scope of the rule - is it applicable to Fabric, switch or ports	varchar	
CATEGORY	Category of the rule - is it used by unknown - 0, asset collection - 1, or 2 - the license manager service	smallint	
ENABLE	Whether the rule needs to be considered or not. 1 - consider, 0 - do not consider for calculation. The default value is 1.	smallint	

TABLE 389 LICENSED_FEATURE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	License feature name, a short text description.	varchar	64
DESCRIPTION	Optional detailed description about the license feature.	varchar	256

TABLE 390 LINK

Field	Definition	Format	Size
LINK_ID	Unique database generated identifier.	int	
TYPE	Type of the link. Currently it is always U.	varchar	1
NAME	Name of the link which is combination of device display name and ifName of the interface which this link associated.	varchar	255

TABLE 391 LOCK

Field	Definition	Format	Size
NAME	The name of this transaction synchronization lock. The name should be upper case and should describe the activity being synchronized, such as MANAGED_ELEMENT_CREATION.	varchar	40
LAST_USED_BY	Identifies the transaction that last updated this lock record, such as IP_DISCOVERY. This field is primarily here just to have something to modify. The new value does not need to be different than the previous value.	varchar	40
LAST_USED_TIME	Optional time when the lock was last modified. Might be useful for debugging someday.	timestamp	

TABLE 392 LSAN_DEVICE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
BB_FABRIC_ID	Backbone fabric DB ID.	int	
FCR_FABRIC_ID	FID assigned to edge fabric.	int	

TABLE 392 LSAN_DEVICE (Continued)

Field	Definition	Format	Size
DEVICE_PORT_WWN	Device port WWN of physical device.	char	23
PHYSICAL_PID	PID of physical device.	char	6

TABLE 393 LSAN_TAG_CONFIG

Field	Definition	Format	Size
ID*	Unique id for FCR LSAN Tags configuration	int	
VIRTUAL_SWITCH_ID	Database identifier of virtual switch which represent FC Router.	int	
TAG_ENABLED	Indicates whether the LSAN tag is enabled or not. Possible values are 0 -false, 1 - true.	smallint	
ENFORCE_TAGS	List of enforcement tags configured in FC router. Enforce tag reduces the resources used in an FC router by limiting the number of LSAN zones that will be enforced in that FC router. There can be maximum of 8 enforce tags per FC router.	varchar	128
SPEED_TAGS	Speed tag configured in FC router. Speed tag allows you to speed up the discovery process by importing the devices into the remote edge fabrics when the devices come online.	varchar	16

TABLE 394 LSAN_PROXY_DEVICE

Field	Definition	Format	Size
FCR_FABRIC_ID*	FID assigned to edge fabric	int	
PROXY_PID*	Proxy device PID	char	6
STATE	State of the device	varchar	128
LSAN_DEVICE_ID*	LSAN_DEVICE record reference	int	

TABLE 395 LSAN_ZONE_DB_CONFIG

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
BB_FABRIC_ID	Backbone fabric db ID.	int	
EDGE_FABRIC_ID	FID assigned to edge fabric.	int	
ZONE_CONTENT	LSAN zone string.	text	
BACKBONE	0= is not a backbone lsan zone configuration. 1= is a backbone lsan zone configuration.	smallint	
FCR_SWITCH_ID	Virtual Switch ID of the FCR switch from which these LSAN Zones are retrieved.	int	

TABLE 396 MCT_CLIENT

Field	Definition	Format	Size
MCT_CLIENT_ID	MCT Client db ID.	int	
RBRIDGE_ID	MCT Client rbridge ID.	int	
CLIENT_NAME	MCT Client name.	varchar	(100)
PORT_ID	MCT Client port foreign key.	int	
OPER_STATE	MCT Client operational state.	smallint	
DEPLOY_STATE	MCT Client deployment state: • Deployed(0) • Undeployed(1)	smallint	
VCN_MEMBER_ID	Virtual Cluster Node member Cluster id foreign key.	int	

TABLE 397 MAC_FILTER

Field	Definition	Format	Size
ID		serial	
MAC_FILTER_NUMBER	MAC Filter number.	int	
FILTER_ACTION	Defined Permit - 0 or Deny -1	smallint	
DESCRIPTION	Description associated with each MAC Filter entry.	varchar	256
SOURCE_MAC_ADDRESS	Source MAC Address.	varchar	24
SOURCE_ADDRESS_MASK	Source MAC address mask.	varchar	24
DEST_MAC_ADDRESS	Destination MAC Address.	varchar	24
DEST_ADDRESS_MASK	Destination MAC address mask.	varchar	24
ETHERNET_TYPE	This column specifies the Ethernet Type. This field can take 0(not used), 1(etype), 2(llc), 3(snap).	smallint	
OPERARTOR	This column specifies the operator. This field can take 0(=), 1(!=), 2(<), 3(>).	smallint	
FRAME_NUMBER	This column specifies the Frame Number. Range is from 0600-FFFF in hex presentation.	int	
OPERATION_TYPE		varchar	10
DEPLOYMENT_ID	Deployment configuration ID. Foreign Key for DEPLOYMENT table.	int	
INT_BINDING_DIRECTION	Represents the interface binding direction. 0/1/2 corresponds to IN / OUT/BOTH direction.	smallint	

TABLE 398 MAC_FILTER_DEV_DEPLOYMENT_MAP

Field	Definition	Format	Size
DEPLOYMENT_ID	Deployment configuration ID.	int	
MAC_FILTER_ID	MAC FILTER Id for reference to the MAC_FILTER.	int	

TABLE 399 MAC_FILTER_INT_DEPLOYMENT_MAP

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
DEPLOYMENT_ID	Deployment configuration ID. Foreign Key for DEPLOYMENT table.	int	
INBOUND_MAC_FILTER_ID	MAC FILTER Id of the MAC Filter selected for inbound. Foreign Key for MAC_FILTER table.	int	
INBOUND_WRITE_TO_DEVICE	1/0 corresponding to Write to device/not write to device for inbound traffic.	smallint	

TABLE 400 MAC_GROUP

Field	Definition	Format	Size
ID	Sequence number of the records	int	
NAME	Name of the mac group, for internal representation of the group.	varchar	128
MAC_GROUP_ID	The ID will represent any one of the following. 1 A cluster wide unique identifier defined in Network OS switch to be used in mac-based GVLAN classification at access port. Allowed range is 0 through 500 both inclusive. 2 INTERNAL GROUP ID - A dummy group ID (-1), used to represent one or more mac addresses that can be associated with a GVLAN. 3 CUSTOM MAC GROUP ID: Reserved for future usage if the Management application provides an option through UI to create MAC GROUPs.	Int	
TYPE	This indicates if the mac group is internal to BNA or mapped to device. 0- internal mac group, 1- external mac group defined in network OS switch, 2- User defined mac groups.	int	

TABLE 401 MAC_GROUP_MEMBER

Field	Definition	Format	Size
ID	Primary Key. Sequence number of the records.	int	
MAC_GROUP_DB_ID	Foreign Key Reference to the MAC_GROUP table.	int	
MAC_ADDRESS	Mac Address that belongs to the Mac group.	varchar	64
MASK	Mask applied on the mac address.	varchar	64

TABLE 402 MANAGED_ELEMENT

Field	Definition	Format	Size
ID	An ID that is unique across managed elements of all types: SAN physical switches, SAN logical switches, IP switches, and hosts. Also the primary key for the MANAGED_ELEMENT table.	int	
PLACEHOLDER	Not used. iBatis/Abator requires at least one non-serial column to generate correct objects. The default value is 0.	int	

TABLE 403 MAPS_EVENT

Field	Definition	Format	Size
ID	The primary key of the table.	int	
HOST_TIME	The time at which the server processed the event.	timestamp	
CATEGORY	The violations category. i.e. Port Health, Fabric Health, etc.	int	
VIOLATION_TYPE	The type of the violation. i.e. CRC, ITW.	int	
MANAGED_ELEMENT_ID	The managed element corresponding to this event.	int	
ORIGIN_FABRIC_ID	The fabric from which the event originated. Retaining this id as historical data.	int	
SWITCH_PORT_ID	Nullable foreign key. The FC port for which the event occurred. This will only be populated for port events.	int	
FCIP_CIRCUIT_ID	Nullable foreign key. The FCIP tunnel circuit for which the event occurred. This will only be populated for FCIP tunnel events.	int	
FRU_NAME	For switch policy status events, the object name is provided in the event and indicates the name of the FRU affected. i.e. PS 1, Fan 2. As this FRU object name is only provided for one category of events, making the column nullable.	varchar	32
VM_ID	Nullable foreign key. The VM for which the event occurred. This will only be populated for vCenter events.	int	
FLOW_DEFINITION_ID	Nullable foreign key. The NP flow definition for which the event occurred. This will only be populated for flow events	int	
INTERFACE_ID	Denotes interface_id for NOS interface related interface violations. This will only be populated for NOS Ethernet related port Events, for other FC ports of NOS switch_port_id will be populated.	int	
SUB_FLOW_KEY	The subflow source, destination address along with ingress port are received as part of the event varbind. This will only be populated for sub flow events.	varchar	256
FCIP_TUNNEL_ID	Nullable foreign key. The FCIP tunnel on which the violation has occurred. Populated only for FCIP tunnel violations related events, Null will be populated for other violations.	int	
PORT_TYPE	Indicates port type of the port when the violation occurred on it. This will be populated only for port violations. For other violations, it will be empty.	varchar	64
COLLECTION_NAME	Indicates name of the collection. This will be populated only for violations which are triggered due to rules specified in collection(s). It is applicable only for AMP. For other violations, it will be null.	varchar	128

TABLE 404 MAPS_EVENT_DETAILS

Field	Definition	Format	Size
ID	The primary key of the table.	int	
MAPS_EVENT_ID	The corresponding maps_event.	int	

TABLE 404 MAPS_EVENT_DETAILS (Continued)

Field	Definition	Format	Size
SWITCH_TIME	The switch timestamp from the event.	timestamp	
RULE_NAME	The name of the threshold rule.	varchar	32
RULE_CONDITION	The threshold condition in string format. i.e. CRC > 30	varchar	128
TIME_BASE	The time base for the threshold. 0 - None, 1 - Minute, 2 - Hour, 3 - Day	int	
ACTIONS	A bit map for the actions configured for the rule. 0 - None, 1 - RASLOG, 2 - SNMP, 4 - Email, 8 - Fence Port, 16 - SW_ST_DOWN, 32 - SW_ST_MARGINAL.	int	
CURRENT_VALUE	The current value of the measure that triggered the violation.	varchar	32
SWITCH_ENABLED_ACTIONS	MAPS actions enabled on the switch at the time the violation occurred.	int	
SEVERITY	Specifies the severity of the Maps Rule. 1 - Critical, 2 â€ Error, 3 â€“ Warning, 4 â€“ Info. Default Value = 3.	int	

TABLE 405 MAPS_EVENT_CAUSE_ACTION

Field	Definition	Format	Size
VIOLATION_TYPE	The type of the violation. i.e. CRC, ITW, as defined in MapsConstants.	int	
ACTION	Description of the recommended action for the MAPS violation.	varchar	4096

TABLE 406 MAPS_POLICY

Field	Definition	Format	Size
ID	The primary key of the table.	int	
VIRTUAL_SWITCH_ID	The id of the virtual switch.	int	
NAME	The name of the MAPS policy.	varchar	32
IS_ACTIVE	Indicates if the policy is the active policy on the switch. 0 - No, 1 - Yes.	int	
IS_DEFAULT	Indicates if the policy is a default policy on the switch. 0 - No, 1 - Yes.	int	

TABLE 407 MARCHING_ANTS

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
THRESHOLD1_VALUE	The marching ants low boundary threshold value (T1).	int	
THRESHOLD2_VALUE	The marching ants high boundary threshold value (T2).	int	

TABLE 408 MEASURE

Field	Definition	Format	Size
ID	Primary key column.	int	
MEASURE_TYPE	Measure type 0 - SNMP MIB, 1-- Expression 2-- EE Monitor counter 3-- HBA counter,4 - Custom, 5 - NetworkPatroller, 6 - NOS Ethernet Port optics etc.	smallint	
INDEX_TYPE	Identifies the index type for a given SNMP MIB or Expression measure. Various index type supported are 0 - UNKNOWN, 1 - SCALAR, 2 - IF_INDEX, 3 - ETHER_STATS_INDEX, 4 - CONN_UNIT_INDEX, 5 - FCIP_LINK_TABLE_INDEX, 6 - CUSTOM, 7 - SW_TEMP_SENSOR_INDEX, 8 - SW_FAN_SENSOR_INDEX, 9 - SW_POWER_SENSOR_INDEX. For non-SNMP measures like EE Monitors, Ping statistics etc. index type is not applicable. In that case index type would be updated as 0 - Unknown.	smallint	
NAME	Name of the measure.	varchar	64
DETAIL	For SNMP MIB, stores the OID, for expression, stores the expression formula.	varchar	1024
UNIT	Unit string thats used for displaying the chart.	varchar	64
DESCRIPTION	Description for the measure. Default: 1	varchar	512
IS_SYSTEM	Indicates whether this is a system built in measure, for system expressions , user cannot delete it.	smallint	

TABLE 409 MESSAGE_RECIPIENT

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
DESCRIPTION	Description about recipient.	varchar	256
IP_ADDRESS	IP Address of the recipient.	varchar	128
PORT	Port number of the recipient.	int	
RECIPIENT_TYPE_ID	Recipient Type (Syslog or SNMP).	int	
ENABLED	If forwarding to destination is enabled.	smallint	
SOURCE_ADDRESS_ADDED	If source address is added as another varbind in trap. -1 for Syslog i.e RECIPIENT_TYPE_ID: 2. Default value is -1.	smallint	
REPEATER_ENABLED	If filtering is disabled. -1 for Syslog i.e RECIPIENT_TYPE_ID: 2. Default value is -1.	smallint	
VERSION	Snmp version(v1/v2/v3)	varchar	8

TABLE 410 MIGRATION_HISTORY

Field	Definition	Format	Size
ID	The primary key of the table.	int	
SOURCE_RELEASE	Source release name, version number and patch number. Example *Management_Application* 11.3.0a.	varchar	128
SOURCE_RELEASE_BUILD_NUMBER	Source release build number.	int	
TARGET_RELEASE	Target release name and version. Example *Management_Application* 12.1.0.	varchar	128
TARGET_RELEASE_BUILD_NUMBER	Target release build number.	int	
MIGRATION_TIME	Date and Time at which this migration completed.	timestamp	

TABLE 411 MODULE

Field	Definition	Format	Size
MODULE_TYPE_ID	Primary key for this table.	int	
MODULE_TYPE	Type of the module.		
NAME	Name of the module configured in this device.		
DESCRIPTION	Description of the module.	varchar	128
NUM_PORTS	Number of ports present in this module.	num	(4,0)
TABLE_SUBTYPE	Identifies the table name which more properties/attributes about this module stored. Possible value is FOUNDRY_MODULE.	varchar	32
IS_PRESENT	Identifies the module is present or not. Not Present-0, Present-1.	num	(1,0)
IS_MANAGEMENT_MODULE	Identifies the module is management module or not. Other module-0, Management module-1.	num	(1,0)
NUM_CPUS	The number of CPUs present in the module.	smallint	
HW_REVISION	The vendor-specific hardware revision string. Refer entPhysicalHardwareRev of RFC4133-ENTITY-MIB.mib for more details.	varchar	64
FW_REVISION	The vendor-specific firmware revision string. Refer entPhysicalFirmwareRev of RFC4133-ENTITY-MIB.mib for more details.	varchar	64
SW_REVISION	The vendor-specific software revision string. Refer entPhysicalSoftwareRev of RFC4133-ENTITY-MIB.mib for more details.	varchar	64

TABLE 411 MODULE (Continued)

Field	Definition	Format	Size
MODULE_STATUS	Specifies the status of the module. Possible values are 0, 2, 3, 4, 8, 9, 10, 11. 0 - moduleEmpty, 2 - moduleGoingDown, 3 - moduleRejected, 4 - moduleBad, 8 - moduleConfigured, 9 - moduleComingUp, 10 - moduleRunning, 11 - moduleBlocked.	int	
REDUNDANT_STATUS	Specifies the redundant status of the module. Possible values are 1, 2, 3, 4, 5. 1 - other, 2 - active, 3 - standby, 4 - crashed, 5 - comingUp. Non management modules always return value as other. Management module returns the rest of the states.	int	
OPERATIONAL_STATUS	Specifies the operational status of the module. Possible values for NI series products are as below: • CARD_STATE_NOT_PRESENT, • CARD_STATE_INIT, • CARD_STATE_BOOT, • CARD_STATE_LP_SYNC, • CARD_STATE_INTERACTIVE, • CARD_STATE_SW_LOADED, • CARD_STATE_STRIPE_SYNC, • CARD_STATE_UP, • CARD_STATE_DOWN, • CARD_STATE_POWERED_OFF, • CARD_STATE_RECOVERY, • CARD_STATE_REBOOT, • CARD_STATE_SYNC_FID. Empty string indicates, module has not been inserted to the chassis or not applicable for this product.	varchar	64

TABLE 412 MODULE_SLOT_PRESENT

Field	Definition	Format	Size
MODULE_SLOT_PRESENT_ID	Unique database generated identifier.	int	
MODULE_ID	Database ID of the MODULE instance.	int	
SLOT_ID	Database ID of the SLOT instance.	int	

TABLE 413 MPLS_ADMIN_GROUP

Field	Definition	Format	Size
MPLS_ADMIN_GROUP_DB_ID	Unique database generated identifier.	int	
NAME	The group name that this administrative group is associated with.	varchar	255
ID	Identifies the administrative group.	int	
DEVICE_ID	Database ID of the DEVICE instance from which this admin group is retrieved.	int	

TABLE 414 MPLS_ADMIN_GROUP_INTERFACE_RELATION

Field	Definition	Format	Size
MPLS_ADMIN_GROUP_INTERFACE_RELATION_DB_ID	Unique database generated identifier.	int	
MPLS_ADMIN_GROUP_DB_ID	Database ID of the MPLS_ADMIN_GROUP instance.	int	
INTERFACE_ID	Database ID of the INTERFACE instance.	int	

TABLE 415 MPLS_LSP

Field	Definition	Format	Size
MPLS_LSP_DB_ID	Unique database generated identifier.	int	
TABLE_SUBTYPE	Refers to the subtype of the LSP where additional attributes/properties of different type of LSP stored. The possible values are MPLS_RSVP_LSP, MPLS_LSP.	varchar	32
NAME	Name of the Label Switched Path.	varchar	255
DESTINATION_IP_ADDRESS	Destination IP Address of the egress LSR associated with this tunnel instance.	varchar	255
OPER_STATUS	Actual operational status of this tunnel, which is typically but not limited to, a function of the state of individual segments of this tunnel. Up-1, Down-2, Testing-3, Unknown-4, Dormant-5, Not present-6, Lower Layer Down-7.	smallint	
DEVICE_ID	Database ID of the DEVICE Instance from which this LSP retrieved.	int	

TABLE 416 MPLS_PATH

Field	Definition	Format	Size
MPLS_PATH_DB_ID	Unique database generated identifier.	int	
NAME	Name of the MPLS Path as configured in the device. Refer mplsTunnelHopPathOptionName of RFC3812 for more details.	varchar	255
DEVICE_ID	Database ID of the DEVICE Instance from which this path information retrieved.	int	

TABLE 417 MPLS_PATH_HOP

Field	Definition	Format	Size
MPLS_PATH_HOP_DB_ID	Unique database generated identifier.	int	
HOP_INDEX	Index of the MPLS hop.	int	

TABLE 417 MPLS_PATH_HOP (Continued)

Field	Definition	Format	Size
HOP_IP_ADDRESS	The Tunnel Hop Address for this tunnel hop.	varchar	255
HOP_TYPE	Denotes whether this tunnel hop is routed in a strict or loose fashion. Possible Values are Strict-1 and Loose-2.	smallint	
MPLS_PATH_DB_ID	Database ID of the MPLS_PATH Instance which this hop is part of.	int	

TABLE 418 MPLS_RSVP_LSP

Field	Definition	Format	Size
MPLS_LSP_DB_ID	Unique database generated identifier.	int	
IS_ENABLED	Represents whether the LSP is enabled. Enabled-1, Disabled-0.	num	(1,0)
IS_BYPASS	Represents if the LSP is a Bypass LSP or not. Not a Bypass-0, Bypass LSP-1. Currently ByPass LSPs are not supported. So the value will be always 0.	num	(1,0)
FROM_IP_ADDRESS	Represents the Source IP Address of the LSP.	varchar	255
METRIC	Represents the metric of the LSP used by the routing protocols to determine the relative preference among several LSPs towards a given destination. Accepts a range of 1 - 65535.	int	
PATH_SELECT_MODE	Specifies the path selection mode to use. Refer mplsLspPathSelectMode MIB of foundry.mib for more details and possible values.	smallint	
PATH_SELECT_PATH	The user-selected pathname when the Path Selection mode is either Manual or Unconditional. If the device returns null or empty string, this value would be primary.	varchar	255
REVERT_TIMER	The number of seconds to wait after the primary or selected path comes up before traffic reverts to that path. A value of 0 indicates that it will switch immediately after the current working path goes down. The range of values supported are 0-65535.	int	
TIE_BREAKING_MODE	Specifies the tie breaking mode for selecting the Constrained Shortest Path First(CSPF) equal-cost paths. Possible values are Random-1, LeastFill-2 and MostFill-3.	smallint	
IS_USE_LSP_FOR_OSPF_SHORTCUTS	Indicates that this LSP allows shortcut between nodes in an AS. OSPF includes the LSP in its SPF calculation. Possible values are Not Allowed-0 and Allowed-1.	num	(1,0)
IS_USE_LSP_FOR_ISIS_SHORTCUTS	Flag to indicate if the LSP is to be used by ISIS destinations.	num	(1,0)
LSP_FOR_ISIS_SHORTCUTS_LEVEL	ISIS level to which the LSP is advertised into	int	
LSP_FOR_ISIS_SHORTCUTS_RELATIVE_METRIC	Add or subtract relative metric.	int	
IS_LSP_FOR_ISIS_SHORTCUTS_ANNOUNCE	Flag that indicates if the LSP is to be announced into ISIS domain.	num	(1,0)
LSP_FOR_ISIS_SHORTCUTS_ANNOUNCE_METRIC	If announced into ISIS domain metric used by the LSP.	int	

TABLE 419 MPLS_RSVP_LSP_ACTUALLY_ROUTED_HOP

Field	Definition	Format	Size
MPLS_RSVP_LSP_ACTUALLY_ROUTED_HOP_DB_ID	Unique database generated identifier.	int	
HOP_INDEX	Index of actually routed hop.	varchar	255
HOP_IP_ADDRESS	The Tunnel Hop Address for this tunnel hop.	int	
MPLS_LSP_DB_ID	Database ID of the MPLS_RSVP_LSP instance which this hop is part of.	int	

TABLE 420 MPLS_RSVP_LSP_ADMIN_GROUP

Field	Definition	Format	Size
MPLS_RSVP_LSP_ADMIN_GROUP_DB_ID	Unique database generated identifier.	int	
AFFINITY_TYPE	Represents the affinity type of the MPLS Admin Group. Possible values are Unknown-0, Include Any-1, Include All-2 and Exclude Any-3.	smallint	
MPLS_ADMIN_GROUP_DB_ID	Database ID of the MPLS_ADMIN_GROUP Instance.	int	
MPLS_RSVP_LSP_ADMIN_GROUP_CONTAINER_DB_ID	Database ID of the MPLS_RSVP_LSP_ADMIN_GROUP_CONTAINER instance.	int	

TABLE 421 MPLS_RSVP_LSP_ADMIN_GROUP_CONTAINER

Field	Definition	Format	Size
MPLS_RSVP_LSP_ADMIN_GROUP_CONTAINER_DB_ID	Unique database generated identifier.	int	
MPLS_RSVP_LSP_PARAMETERS_DB_ID	Database ID of the MPLS_RSVP_LSP_PARAMETERS instance.	int	
MPLS_RSVP_LSP_FRR_PARAMETERS_DB_ID	Database ID of the MPLS_RSVP_LSP_FRR_PARAMETERS instance.	int	

TABLE 422 MPLS_RSVP_LSP_FRR_PARAMETERS

Field	Definition	Format	Size
MPLS_RSVP_LSP_FRR_PARAMETERS_DB_ID	Unique database generated identifier.	int	
BANDWIDTH	Specifies the bandwidth constraint for the MPLS Fast Reroute Path. The value 0 means that the detour route uses a best-effort value for bandwidth.	int	
HOP_LIMIT	Represents the limit for the number of hops the LSP can traverse. Accepted range is 0 - 255.	num	(3,0)
IS_FACILITY_BACKUP	Specifies whether the request for Facility backup is enabled or not. If the FRR mode is facility then this value will be 1. 0 otherwise.	num	(1,0)
SETUP_PRIORITY	The setup priority for MPLS Fast Reroute.Allowed range between 0-7.	num	(1,0)

TABLE 422 MPLS_RSVP_LSP_FRR_PARAMETERS (Continued)

Field	Definition	Format	Size
HOLD_PRIORITY	'The hold priority for MPLS Fast Reroute. Allowed range between 0-7.	num	(1,0)
MPLS_LSP_DB_ID	Database ID of the MPLS_RSVP_LSP instance which these reroute parameters associated with.	int	

TABLE 423 MPLS_RSVP_LSP_PARAMETERS

Field	Definition	Format	Size
MPLS_RSVP_LSP_PARAMETERS_DB_ID	Unique database generated identifier.	int	
IS_ADAPTIVE	Indicates if the LSP supports adaptive mechanism or not. Non Adaptive-0, Adaptive-1.	num	(1,0)
BFD_TRANSMIT	This object specifies the minimum interval, in milliseconds, that the local system would like to use when transmitting The Bidirectional Forwarding Detection(BFD) Control packets. Accepts a range of 50-30000.	int	
BFD_RECEIVE	This object specifies the minimum interval, in milliseconds, between received Bidirectional Forwarding Detection (BFD) Control packets the local system is capable of supporting. Accepts a range of 50-30000.	int	
BFD_MULTIPLIER	The Bidirectional Forwarding Detect time multiplier. Accepts a range of 3-50.	int	
COS	The Class of Service for this LSP. Allowed range is 0-7 and 255. 255 means COS is not explicitly configured.	smallint	
HOP_LIMIT	Represents the limit for the number of hops the LSP can traverse. Accepted range is 0-255.	num	(3,0)
IS_CSPF	Specifies whether the Constrained Shortest Path First (CSPF) calculation is enabled. Possible values are Disabled-0, Enabled-1.	num	(1,0)
MTU	Specifies the Maximum IP Packet Size of the packets without being fragmented. Valid range is 0-65535.	num	(4,0)
SETUP_PRIORITY	The setup priority of the tunnel. Valid range is 0-7.	num	(1,0)
HOLD_PRIORITY	The holding priority of the tunnel. Valid range between 0-7.	num	(1,0)
IS_RECORD_ROUTES	Specifies whether the route is actually recorded route or not. Not Recorded-0 and Recorded-1.	num	(1,0)
REOPTIMIZE_TIMER	The number of seconds from the beginning of one reoptimization attempt to the beginning of the next attempt. Valid range is 300-65535 seconds. 0 is also accepted.	int	
MPLS_LSP_DB_ID	Database ID of the MPLS_RSVP_LSP instance which these parameters are associated with.	int	
MPLS_RSVP_LSP_PATH_DB_ID	Database ID of the MPLS_RSVP_LSP_PATH instance which these parameters are associated with.	int	

TABLE 424 MPLS_RSVP_LSP_PATH

Field	Definition	Format	Size
MPLS_RSVP_LSP_PATH_DB_ID	Unique database generated identifier.	int	
PATH_TYPE	The type of path that is active, i.e., a primary path, a standby path, or a generic secondary path. Possible values are Other-1, Primary-2, Standby-3 and Secondary-4.	smallint	
IS_STANDBY	Specifies whether the path is standby or not. Currently it is unused and value is always 0 (Not standby)	num	(1,0)
MPLS_LSP_DB_ID	Database ID of the MPLS_RSVP_LSP instance.	int	
MPLS_PATH_DB_ID	Database ID of the MPLS_PATH instance.	int	

TABLE 425 MPLS_RSVP_LSP_TUNNEL_RESOURCE

Field	Definition	Format	Size
MPLS_RSVP_LSP_TUNNEL_RESOURCE_DB_ID	Unique database generated identifier.	int	
MAX_RATE	Specifies the maximum data rate (kilo bits/secs) of the packet travelling over the LSP.	int	
MEAN_RATE	Specifies the mean data rate (kilo bits/secs) of the packet travelling over the LSP.	int	
MAX_BURST	The maximum burst size in bytes that the LSP can send at the maximum rate.	int	
MPLS_RSVP_LSP_PARAMETERS_DB_ID	Database ID of the MPLS_RSVP_LSP_PARAMETERS instance.	int	

TABLE 426 MPLS_SERVICE

Field	Definition	Format	Size
MPLS_SERVICE_DB_ID	Unique database generated identifier.	int	
NAME	Specifies the name of the MPLS Service.	varchar	255
VCID	Virtual Circuit Identifier of the MPLS Service.	bigint	
MPLS_SERVICE_TYPE	The type of the MPLS Service. Local VLL-1, Remote VLL-2, VLL-3, VPLS-4, Admin Group-8, Path-9, RSVP LSP-10.	smallint	
VLL_MODE	Specifies the Virtual Local Loop (VLL) Mode. Possible values are Unknown-0, Raw-1 and Taggged-2.	smallint	
STATUS	Status of the MPLS Service. All Peers Up-1, All Peers Down-2, Some Peers Down-3, Undefined-0.	smallint	
CONFLICTS	The type of Conflict. Possible values are None-0, Name Mismatch-1, VLL Mode Mismatch-2, Peer Incomplete-4, No Endpoints-8, Peer Missing-16, Duplicate VCID-32, Unknown-65535.	int	
LAST_UPDATED_TIME	Time when this service record last updated in the database.	num	(20,0)

TABLE 427 MPLS_SERVICE_DEVICE_RELATION

Field	Definition	Format	Size
MPLS_SERVICE_DEVICE_RELATION_DB_ID	Unique database generated identifier.	int	
MPLS_SERVICE_DB_ID	Database ID of the MPLS_SERVICE instance.	int	
DEVICE_ID	Database ID of the DEVICE instance.	int	
TABLE_SUBTYPE	Specifies the type of MPLS Service Relation with Device. Possible values are VLL_DEVICE_RELATION and VPLS_DEVICE_RELATION.	varchar	32
NAME	Name of the MPLS Service.	varchar	255
COS	This value indicates the Class Of Service for this endpoint (VLL/VPLS). Allowed range is 0-7 and 255. 255 means COS is not explicitly configured.	smallint	
MTU	Represents the maximum packet size configured on the VLL/VPLS instance.	int	

TABLE 428 MPLS_SERVICE_ENDPOINT_RELATION

Field	Definition	Format	Size
MPLS_SERVICE_ENDPOINT_RELATION_DB_ID	Unique database generated identifier.	int	
MPLS_SERVICE_DEVICE_RELATION_DB_ID	Database ID of the MPLS_SERVICE_DEVICE_RELATION instance.	int	255
INTERFACE_ID	Database ID of the INTERFACE instance associated with this end point (VLL/VPLS).	int	
TABLE_SUBTYPE	The Type of the MPLS Service endpoint relation. Possible values are VLL_ENDPOINT_RELATION and VPLS_ENDPOINT_RELATION.	varchar	32
TAG_MODE	This value indicates the vlan mode for this endpoint. Possible values are Tagged-1, Untagged-2.	smallint	
VLAN_ID	Specifies the Outer VLAN ID value of this endpoint (VLL/VPLS).	smallint	
OPER_STATUS	Operational status of the endpoint. Possible values are Up-1, Down-2.	num	(2,0)
TAG_TYPE	The type of tagging supported. Possible values are Untagged-1, Dual-2 and Inner VLAN/ISID-3. ISID applicable only when dual tagging enabled for VPLS.	smallint	
INNER_VLAN_ID	This value indicates the inner tag for this endpoint. If tagging type is dual, then it returns the inner vlan id of the end point (VLL/VPLS). If tagging type is ISID and Untagged this value will be 0.	smallint	

TABLE 429 MPLS_SERVICE_PEER_RELATION

Field	Definition	Format	Size
MPLS_SERVICE_PEER_RELATION_DB_ID	Unique database generated identifier.	int	
MPLS_SERVICE_DEVICE_RELATION_DB_ID	Database ID of the MPLS_SERVICE_DEVICE_RELATION instance.	int	

TABLE 429 MPLS_SERVICE_PEER_RELATION (Continued)

Field	Definition	Format	Size
PEER_DEVICE_ID	Database ID of the peer device.	int	
PW_INDEX	The pseudo-wire service index Mask.	int	
PEER_IP	The IP of the Peer Device of the PW/PE maintenance protocol entity.	varchar	255
OPER_STATUS	Operational Status of the peer with the MPLS Service. Refer PwOperStatus MIB of foundry.mib for more details and possible values.	smallint	

TABLE 430 MRP_RING

Field	Definition	Format	Size
MRP_RING_ID	Auto generated database ID for MRP ring.	int	
RING_ID	User configured unique ring number for MRP ring, Valid values are 1 – 255.	num	(8,0)
RING_NAME	Represents name of MRP ring.	varchar	255
STATUS	Computed Status of MRP ring. Status is computed based on MRP_RING devices. Possible status values are Normal-1, Warning-2 and Critical-3	smallint	
LAST_UPDATED	Time when this ring record was last updated in the database.	bigint	

TABLE 431 MRP_RING_DEVICE

Field	Definition	Format	Size
MRP_RING_DEVICE_DB_ID	Auto generated database ID for device in MRP Ring.	int	
MRP_RING_ID	Database ID of MRP ring.	int	
DEVICE_ID	Database ID of member device.	int	
PORT_VLAN_DB_ID	The database ID of the port VLAN. The master VLAN in the topology group used by this ring. If a topology group is used by MRP, the master VLAN controls the MRP settings for all VLANs in the topology group.	int	
MRP_RING_NAME	User configured name for the ring.	varchar	255
TOPO_GRP_ID	Topology group ID.	int	
STATE	Whether MRP is enabled or disabled on the device. Disabled-1, Enabled -2.	smallint	
ROLE	Represents role of device in MRP topology. Master-2, Member-3.	smallint	
HELLO_TIME	The interval, in milliseconds, at which the Forwarding port on the ring's master node sends Ring Health Packets (RHPs).	int	
PRE_FWD_TIME	The number of milliseconds a MRP interface that has entered the Pre forwarding state will wait before changing to the Forwarding state.	int	
PRI_PORT_INTERFACE_ID	Interface database ID for the Primary port of the device.	int	

TABLE 431 MRP_RING_DEVICE (Continued)

Field	Definition	Format	Size
PRI_PORT_STATE	State of device's primary port. Other-1, Pre-Forwarding- 2, Forwarding-3, Blocking-4, Disabled-5.	smallint	
PRI_PORT_TYPE	Type of device's primary port. Other-1, Regular port-2, Tunnel port-3.	smallint	
PRI_PORT_ACTIVE_INTERFACE_ID	Interface database ID of an primary active port, which is sending RHPs.	int	
SEC_PORT_INTERFACE_ID	Interface database ID for the Secondary port of the device.	int	
SEC_PORT_STATE	State of device's secondary port. Other-1, Pre-Forwarding- 2, Forwarding-3, Blocking-4, Disabled-5.	smallint	
SEC_PORT_TYPE	Type of device's secondary port. Other-1, Regular port-2, Tunnel port-3.	smallint	
SEC_PORT_ACTIVE_INTERFACE_ID	Interface database ID of an secondary active port, which is receiving RHPs.	int	
RHP_TX	The number of RHPs sent on the active interface.	bigint	
RHP_RC	The number of RHP packets received on the interface.	bigint	
STATE_CHANGED	The number of MRP interface state changes that have occurred.	int	
TC_BPDU_RC	The number of Topology Change RHPs received on the interface. A Topology Change RHP indicates that the ring topology has changed.	int	
STATUS	Computed status of device in MRP Ring. Possible status values are Normal-1, Warning-2 and Critical-3.	smallint	
LAST_UPDATED	Time when this record was last updated in the database.	bigint	

TABLE 432 N2F_PORT_MAP

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
VIRTUAL_SWITCH_ID	Virtual switch ID of AG for N to F_port mapping, foreign key to VIRTUAL_SWITCH table.	int	
N_PORT	Port number of port type N_Port which is being mapped, One N_Port can be mapped to multiple F_ports.	smallint	
F_PORT	Port number of port type F_Port which is being mapped.	smallint	

TABLE 433 NETWORK_VLAN

Field	Definition	Format	Size
VLAN_DB_ID	Database ID of the VLAN instance which is associated with the Network.	int	

TABLE 434 NETWORK_SCOPE_TYPE

Field	Definition	Format	Size
ID	The primary key of the table.	int	
NAME	Name of the Scope.	varchar	128
DESCRIPTION	Description of the Scope.	varchar	512
HANDLER_CLASS_NAME	Fully defined Handler Class for the predefined SCOPE.	varchar	128

TABLE 435 NIC_PROFILE

Field	Definition	Format	Size
ID*		int	
NAME	The name of the network interface in the format network interface name / host address.	varchar	255
IP_ADDRESS	The host address of the interface.	varchar	128

TABLE 436 NPORT_WWN_MAP

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
VIRTUAL_SWITCH_ID	AG switch reference on which the Nport wwn mapping resides.	int	
N_PORT	N Port through which AG is connected to the edge switch	smallint	
DEVICE_PORT_WWN	Device Port which is mapped to the N port. This device could be offline device as well.	char	23

TABLE 437 NP_FLOW_DEFINITION

Field	Definition	Format	Size
ID	The primary key of the table.	int	
NAME	The name of the table.	varchar	20
VIRTUAL_SWITCH_ID	The id for the virtual switch.	int	
SRCDEV	Comma separated list of source device ports.	varchar	1024
DSTDEV	Comma separated list of destination device ports.	varchar	1024
SRCPORT	Comma separated list of source switch ports.	varchar	1024
DSTPORT	Comma separated list of destination switch ports.	varchar	1024
BIDIR	This specifies if traffic in both direction has to be monitored, where, 0 - false, 1 - true.	smallint	
SFID	Source fabric ID. In case of FCR flow creation it will have FCR fabric Id in this field. In XISL it will have virtual fabric Id. For learning flow it will have *.	varchar	12
DFID	Destination fabric ID. In case of FCR flow creation it will have FCR fabric Id in this field. In XISL it will have virtual fabric Id. For learning flow it will have *.	varchar	12
SRCDOMAIN	Source domain ID	int	

TABLE 437 NP_FLOW_DEFINITION (Continued)

Field	Definition	Format	Size
DSTDOMAIN	Destination domain ID	int	
LUNID	Comma separate list of LUN IDs	varchar	1024
OXID	FC Originator Exchange ID for the frame.	varchar	1024
QOS	Quality of Service, can be comma separated values of: 1 - low, 2 - medium, 3 - high.	varchar	1024
"OPTION"	A bitmask for options with following bit mapping: Noactive (0th bit) = 2^0 = 1 Noconfig (1st bit) = 2^1 = 2 NoZoneCheck (2nd bit) = 2^2 = 4	int	
SCSICMD	SCSI command frame types.	varchar	32
TYPE	Frame type value	varchar	32
RCTL	Routing control byte.	varchar	32
PROTOCOL_TYPE	Protocol types.	varchar	32
FRAME_OFFSET	Generic frame offset in format of byte offset, mask, value.	varchar	1024
"SIZE"	Size of the frame payload. Range: 64 bytes to max 2112 bytes, 0 for random size.	int	
PATTERN	String to specify the pattern of the payload.	varchar	32
LAST_UPDATED_TIME	Last updated time	timestamp	
MONITOR_FEATURE	Flow Monitor feature state. 0 - not selected, 1 - deactivated, 2 - selected and activated.	int	
GENERATOR_FEATURE	Flow generator feature state. 0 - not selected, 1 - deactivated, 2 - selected and activated.	int	
MIRROR_FEATURE	Flow mirror feature state. 0 - not selected, 1 - deactivated, 2 - selected and activated.	int	
IS_PREDEFINED	Flag which identifies if the flow definition is one of the pre-defined flow definitions on the switch.	smallint	
MIRROR_PORT	Mirror Port to which the External Analyzer will be connected and used for creation of Local Flow Mirroring	varchar	32
SRC_SWITCH_PORT_ID		int	
DST_SWITCH_PORT_ID		int	
SRC_DEVICE_PORT_ID		int	
DST_DEVICE_PORT_ID		int	
SRC_HBA_ID		int	
DST_HBA_ID		int	
SRC_VM_ID		int	
DST_VM_ID		int	
MIRROR_SWITCH_PORT_ID		int	
SRC_DEVICE_NODE_ID		int	
DST_DEVICE_NODE_ID		int	

TABLE 437 NP_FLOW_DEFINITION (Continued)

Field	Definition	Format	Size
REMOTE_MIRROR_PORT	Remote mirror port details for RFM flow, it holds the AMP devices AF port. It will be empty for all other non RFM flows.	varchar	32
MAPS_PORT_GROUP	it holds the MAPS Port group name for the RFM flow. It will be empty for all other non RFM flows.	varchar	64
VM_ENTITY_ID	Entity Id of Virtual Machine from which this flow is initiated.	varchar	256

TABLE 438 NP_SUB_FLOW

Field	Definition	Format	Size
ID	The primary key of the table.	int	
FLOW_DEFINITION_ID	The id of the flow definition	int	
FEATURE	Feature this sub flow is associated with. Feature can be one of the following: Monitor - 0, Generator - 1, Mirror - 2	int	
SRCDEV	Source device port.	varchar	32
DSTDEV	Destination device port.	varchar	32
SRCPORT	Switch Source port.	varchar	32
DSTPORT	Switch Destination port.	smallint	
BIDIR	This specifies if traffic in both direction has to be monitored, where, 0 - false, 1 - true		
SFID	Source fabric ID. In case of FCR flow creation it will have FCR fabric Id in this field. In XISL it will have virtual fabric Id. For learning flow it will have *.	varchar	12
DFID	Destination fabric ID. In case of FCR flow creation it will have FCR fabric Id in this field. In XISL it will have virtual fabric Id. For learning flow it will have *.	varchar	12
SRCDOMAIN	Source domain ID	int	
DSTDOMAIN	Destination domain ID	int	
LUNID	LUN ID.	varchar	32
LAST_UPDATED_TIME	Last updated time	timestamp	
IS_MISSING	Is the sub flow no more available on the switch? 0 - false, 1 - true.	smallint	
OXID	FC Originator Exchange ID for the frame..	int	
RXID	FC Responder Exchange ID for the frame.	int	
CS_CTL	Frame header CS_CTL..	int	
"SIZE"	Size of the frame payload. Range: 64 bytes to max 2112 bytes, 0 for random size.	int	
PATTERN	String to specify the pattern of the payload.	varchar	32
MIRROR_PORT	Mirror Port to which the External Analyzer will be connected and used for creation of Local Flow Mirroring.	varchar	32

TABLE 438 NP_SUB_FLOW (Continued)

Field	Definition	Format	Size
KEY	Will contain the subflow details -Source Device , Destination Device and Ingress/Egress Port details of the sub flow, which will be used to get the violation count for the sub flow.	varchar	256
SUB_FLOW_KEY	The subflow source, destination address along with ingress port are received as part of the event varbind. This will only be populated for sub flow events.	varchar	64
SRC_VIRTUAL_SWITCH_ID		int	
SRC_SWITCH_PORT_ID	It holds the source device connected switch port reference for the AMP sub-flows. It refers the source switch port id from switch_port table.	int	
DST_VIRTUAL_SWITCH_ID		int	
DST_SWITCH_PORT_ID	It holds the destination device connected switch port reference for the AMP sub-flows. It refers the source switch port id from switch_port table.	int	
SRC_DEVICE_PORT_ID	It refers the source device port id from device_ port table.	int	
DST_DEVICE_PORT_ID	It refers the destination device port id from device_ port table.	int	
SUB_FLOW_ORIGIN		varchar	32
SRC_HBA_ID		int	
DST_HBA_ID		int	
SRC_VM_ID		int	
DST_VM_ID		int	
MIRROR_SWITCH_PORT_ID		int	
SRC_DEVICE_NODE_ID		int	
DST_DEVICE_NODE_ID		int	
VM_ENTITY_ID	Entity Id of Virtual Machine from which this flow is initiated.	varchar	256

TABLE 439 OUI_GUESSED_DEVICE_MAP

Field	Definition	Format	Size
OUI*	Vendor OUI.	char	6
TYPE	Guessed device type for this vendor.	varchar	32

TABLE 440 OUI_VENDOR

Field	Definition	Format	Size
OUI*	Vendor OUI, 6-digit hexadecimal number which can have leading digits as zero.	char	6
VENDOR	Vendor name.	varchar	64

TABLE 440 OUI_VENDOR (Continued)

Field	Definition	Format	Size
VENDOR_CATEGORY	Default is 'none'.	varchar	32
USER_MODIFIED		int	

TABLE 441 PASSWORD_HISTORY

Field	Definition	Format	Size
USER_NAME		varchar	128
PASSWORD_UPDATED_DATETIME	The date and time the user updated password recently.	timestamp	
PREVIOUS_PASSWORD	User''s Previous password	varchar	512

TABLE 442 PBR_INTERFACE_CONFIG

Field	Definition	Format	Size
ID	Primary key.	int	
POLICY_ID	PBR policy ID.	int	
INTERFACE_NAME	Name of the ingress interface.	varchar	32
IP_TYPE	Defines the ip type, v4 or v6, of the pbr policy the interface is bound to. Value 1 indicates IPv4 policy. Value 2 indicates IPv6 policy.	int	
ALLOW_VLAN	Indicates if the packets arriving at the ingress ports allow all VLANs or not.	smallint	

TABLE 443 PERF_COLLECTOR

Field	Definition	Format	Size
COLLECTOR_ID	Primary key autogenerated ID.	int	
NAME		varchar	128
STATUS	Status of the collector. if the value is set to "E" means the collector is enabled state and "D" means disabled.	char	1
TYPE	Target type of the SNMP collector data. The target type for, • device level collector is 0 • port level collector is 1.	num	(2,0)
POLLING_INTERVAL	Time interval in seconds; indicates the frequency with which the collector will poll the device to get the data.	int	
CREATED_TIME_SECONDS	Collector created time.	int	
PROPS_STR	Serialized string for the threshold or rearm property.	varchar	256

TABLE 444 PBR_NEXT_HOP

Field	Definition	Format	Size
ID	Primary key.	int	
RULE_ID	PBR rule id.	int	
NEXT_HOP_SEQUENCE	The sequence of the next hop entry that corresponds to a rule within a route map. The sequence of 1 indicates it is the first next hop to be tried for that rule. This is a running integer.	int	
HOP_TYPE	The Next hop type. 1 indicates INTERFACE, 2 indicates IP_ADDRESS, 3 indicates FLOOD VLAN.	smallint	
HOP_VALUE	Depending on the hop type the value can be an IP Address, Vlan id or Interface name.	varchar	64
PRESERVE_VLAN	0 indicates do not preserve vlan, 1 indicates preserve vlan.	smallint	

TABLE 445 PBR_POLICY

Field	Definition	Format	Size
ID	Primary key.	int	
POLICY_NAME	The name of the pbr policy.	varchar	81
IP_POLICY_TYPE	Defines the ip type of the policy, v4, v6 or both v4 and v6. Type v4 will have a value 1, type v6 will have a value 2, both will have a value 3.	smallint	
OPERATION_TYPE	Indicates the action to take on the policy. 1 means ADD, 2 means EDIT, 3 means DELETE.	smallint	
DEPLOYMENT_ID	ID of the deployment_configuration table entry.	int	

TABLE 446 PBR_RULE

Field	Definition	Format	Size
ID	Primary key.	int	
POLICY_ID	PBR policy ID.	int	
RULE_NAME	Name of the pbr rule.	varchar	127
ROUTE_MAP_SEQUENCE	The sequence of the route-map entry that corresponds to a rule within a route-map. The sequence of 1 indicates it is the first rule within the route map. This number will be incremented for every rule entry within a route-map.	int	
OPERATION_TYPE	Indicates the action to take on the rule. 1 is ADD, 2 is EDIT, 3 is DELETE.	smallint	

TABLE 447 PBR_RULE_ACL_LIST

Field	Definition	Format	Size
ID	The primary key of the table.	int	
RULE_ID	PBR_RULE id. This is a foreign key.	int	
ACL_MATCH_SEQUENCE	The sequence of the matching acl entry that corresponds to a rule within a route map. The sequence of 1 indicates it is the first matching acl for that rule. This is a running integer.	int	

TABLE 447 PBR_RULE_ACL_LIST (Continued)

Field	Definition	Format	Size
ACL_NAME	Name of the ACL for the rule.	varchar	81
ACL_TYPE	Indicates the ACL type. Value of 4 denotes IPV4, Value of 6 denotes IPV6.	smallint	

TABLE 448 PHANTOM_PORT

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
WWN	The Wwn of the phantom port.	char	23
VIRTUAL_SWITCH_ID	The id of the phantom switch.	int	
PORT_NUMBER	The port number of the phantom port. The default value is -1.	smallint	
PORT_ID	The portId of the phantom port. The default value is 000000.	varchar	8
SPEED	The speed of the phantom port. The default value is 0.	int	
MAX_SPEED	The max speed of the phantom port. The default value is 0.	int	
TYPE	The portType of the phantom port.The default value is 'Unknown'.	varchar	16
REMOTE_NODE_WWN	The remote node wwn(for E-ports only). Attached port device info must be retrieved from DevicePort table.	char	23
REMOTE_PORT_WWN	The remote port wwn(for E-ports only). Attached port device info must be retrieved from DevicePort table.	char	23
PHANTOM_TYPE	The phantom type of the port, either front or xlate	int	
BB_FABRIC_ID	Denotes the Backbone Fabric ID.	int	

TABLE 449 PHYSICAL_DEVICE

Field	Definition	Format	Size
PHYSICAL_DEVICE_ID	Unique generated database identifier.	int	
DEVICE_ID	Database identifier of the DEVICE instance.	int	
DESCRIPTION	System description of the device.	varchar	255
NUM_SLOTS	Number of slots present in the device.	num	(4,0)
TABLE_SUBTYPE	Table name where additional properties/attributes about this physical device stored. Possible value is FOUNDRY_PHYSICAL_DEVICE.	varchar	32
UNIT_NUMBER	Unit number in the stack if it is stackable device . For non-stacking device it will be always 0.	num	(2,0)
UNIT_NEIGHBOR1	Stacking neighbor's unit(left) number for the stackable devices. If there is no neighbor unit/non stackable devices, then set to 0.	num	(2,0)

TABLE 449 PHYSICAL_DEVICE (Continued)

Field	Definition	Format	Size
UNIT_NEIGHBOR2	Stacking neighbor's unit(left) number for the stackable devices . If there is no neighbor unit/non stackable devices, then set to 0.	num	(2,0)
UNIT_PRESENT	Used to identify the stack unit is present in the chassis or not. Present-1 and Not Present-2.	num	(1,0)
UNIT_TYPE	Indicates unit type in the stack. This column stores the model type of a stackable unit such as "ICX6610-48P". For non-stacking device it will be empty.	varchar	64
IMAGE_VERSION	Image version of the unit in the stack. For non-stacking device it will be always empty.	varchar	128
UNIT_ROLE	Indicates unit role in the stack. Possible values: 1 - other, 2 - active, 3 - standby, 4 - member, 5 - standalone. For non-stacking device it will be always -1.	int	
UNIT_PRIORITY	Indicates unit priority. Possible values 0 to 255. For non-stacking device it will be always -1.	int	
UNIT_STATE	Used to identify unit state in the stack. Possible values: 1 - local, 2 - remote, 3 - reserved, 4 - empty. For non-stacking device it will be always -1.	int	
MAC_ADDRESS	Is to get the mac address of individual unit involved in SPX configuration.	varchar	64

TABLE 450 PHYSICAL_INTERFACE

Field	Definition	Format	Size
INTERFACE_ID	Primary key for this table.	int	
PHYSICAL_PORT_ID	Foreign key which refers PHYSICAL_PORT table.	int	
SPEED_IN_MB	Interface speed in Mega Bytes.	int	
PHYSICAL_ADDRESS	MAC address of this interface.	varchar	64
LINK_ID	Foreign key which refers LINK table.	int	
DUPLEX_MODE	Interface duplex mode. Full/Half/Auto.	smallint	
IS_STACKING_INTERFACE	This flag is to indicate whether the interface is stacking interface or peri port. 0 indicates non-stacking, 1-indicates stacking interface, 2-indicates peri port.	num	(1,0)

TABLE 450 PHYSICAL_INTERFACE (Continued)

Field	Definition	Format	Size
IS_PORT_PRESENT	This flag is to indicate whether the port is presented in the device. 0 = Unknown 1 = Present 2 = Not present	int	
PHYSICAL_DEVICE_ID	For DCB switch, this is the core switch id. For IP products, this is the physical_device_id in physical_device table.	int	
UNIT_NUMBER	This is the unit number of which the interface is located for IP stacking products. If it is not applicable, the value is -1.	int	
SLOT_NUMBER	This is the slot number of which the interface is located for the devices and switches. If it is not applicable, the value is -1.	int	
PORT_NUMBER	This is the port number of the interface.	int	
PORT_TYPE	This column is used to store the port type of the interface. The value will be populated by the NosSwitchAssetCollector during the discovery of the VCS switch.The value of 0 means its edge port, and 1 means its trill port. Default value is 0.	smallint	
UNIT_TYPE	Indicates unit type in the stack. This column stores the model type of a stackable unit such as "ICX6610-48P". For non-stacking device it will be empty	varchar	64
IMAGE_VERSION	Image version of the unit in the stack. For non-stacking device it will be always empty.	varchar	
UNIT_ROLE	Indicates unit role in the stack. Possible values: 1 - other, 2 - active, 3 - standby, 4 - member, 5 - standalone. For non-stacking device it will be always -1'	int	
UNIT_PRIORITY	Indicates unit priority. Possible values 0 to 255. For non-stacking device it will be always -1	int	
UNIT_STATE	Used to identify unit state in the stack. Possible values: 1 - local, 2 - remote, 3 - reserved, 4 - empty. For non-stacking device it will be always -1	int	

TABLE 451 PHYSICAL_PORT

Field	Definition	Format	Size
PHYSICAL_PORT_ID	Database unique generated identifier.	int	
PORT_NUM	Port number from interface identifier.	smallint	
MODULE_ID	Database ID of the module which this port is present.	int	
IS_PORT_PRESENT	This flag is to indicate whether the port is presented in the device. Unknown-0, Present-1 and Not present -2.	smallint	
TABLE_SUBTYPE	PHYSICAL_PORT table sub type.	varchar	32

TABLE 452 PM_COLLECTOR_MEASURE_SETTING

Field	Definition	Format	Size
COLLECTOR_ID	ID of the data_collector.	int	
MEASURE_ID	ID of the measure.	int	

TABLE 453 PM_COLLECTOR_TARGET_SETTING

Field	Definition	Format	Size
COLLECTOR_ID	ID of the data_collector.	int	
TARGET_TYPE	Target type associated to the collector. Possible values are 12 - IP_DEVICE_GROUP and 14 - VIRTUAL_GROUP. To identify the exact target type, combination of TARGET_TYPE and TARGET_ID values are used.	smallint	
TARGET_ID	Target Id associated to the collector. Possible values are 1 - ALL_IOS_PRODUCTS, 2 - ALL_NOS_PRODUCTS, 3 - ALL_IP_TRUNK, 4 - ALL_TRILL_TRUNK, 5 - ALL_PHYSICAL_PORT, 6 - ALL_SAN_FC_PORT, 7 - ALL_SAN_TE_PORT, 8 - ALL_SAN_FCIP_TUNNEL, 9 - ALL_SAN_PRODUCT, 10 - ALL_SAN_EE_MONITOR.	int	
ME_ID	ME_ID of the target.	int	
INDEX_MAP	Stores the index_map value in case of an expression.	varchar	8192

TABLE 454 PM_COLLECTOR_TIME_SERIES_MAPPING

Field	Definition	Format	Size
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
TARGET_NAME	Time series data master table name. It could be either TIME_SERIES_DATA_1 or TIME_SERIES_DATA_2.	varchar	63

TABLE 455 PM_DASHBOARD_WIDGET

Field	Definition	Format	Size
DASHBOARD_WIDGET_ID	Primary key column.	int	
TIME_SCOPE	Time in unit of seconds, for which the data has to be fetched from DB going back from now applicable for top N, distribution, and top Flow, time series.	int	
REFRESHING_INTERVAL	The widget refreshing interval in seconds, in 11.3 we will fix it at 600 (10 mins) and not expose it to user.	int	
MONITOR_TYPE	The widget refreshing interval in seconds, in 11.3 we will fix it at 600 (10 mins) and not expose it to user.	int	
MEASURE_TYPE	TYPE of the user selection measure.	int	

TABLE 455 PM_DASHBOARD_WIDGET (Continued)

Field	Definition	Format	Size
CREATE_USER_ID	ID of the user who created the widget definition.	int	
CREATE_TIME	Widget definition created server time.	timestamp	
MODIFY_USER_ID	ID of the user who last modifed the widget definition.	int	
MODIFY_TIME	Widget definition last modified server time.	timestamp	
TOP_OR_BOTTOM_N	The Top N setting for the Top N, Bottom N and Top XXX monitor TYPE, for other monitor TYPE, this field set to default value. Default is 0.	int	
LEVEL1_ENABLED	Enable / disable the threshold check for first percentage band. This value is applicable only for Top N, Top Flow widgets. Default is 0.	smallint	
LEVEL1_VALUE	Limit value for the first percentage band. Default is 0.	double precision	
LEVEL1_COLOR	RGB color for the first percentage band.	int	
LEVEL2_ENABLED	Enable / disable the second threshold check. This value is applicable only for Top N, Top Flow widgets. Default is 0.	smallint	
LEVEL2_VALUE	Limit value for the second percentage band. Default is 0.	double precision	
LEVEL2_COLOR	RGB color for the second percentage band.	int	
LEVEL3_ENABLED	Enable / disable the third threshold check. This value is applicable only for Top N, Top Flow widgets. Default is 0.	smallint	
LEVEL3_VALUE	Limit value for the third percentage band. Default is 0.	double precision	
LEVEL3_COLOR	Limit value for the third percentage band.	int	
LEVEL4_ENABLED	Enable / disable the fourth threshold check. This value is applicable only for Top N, Top Flow widgets. Default is 0.	smallint	
LEVEL4_VALUE	Limit value for the fourth percentage band. In case of Top N, Top Flow widgets only three percentage bands are available. This value is not applicable. Default is 0.	double precision	
LEVEL4_COLOR	RGB color for the fourth percentage band.In case of Top N, Top Flow widgets we will use this column to store the color value for other percentage band.	int	
LEVEL5_ENABLED	Enable / disable the fifth threshold check. This value is applicable only for Top N, Top Flow widgets. Default is 0.	smallint	
LEVEL5_VALUE	Limit value for the fifth percentage band. In case of Top N, Top Flow widgets only three percentage band s are available. This value is not applicable. Default is 0.	double precision	
LEVEL5_COLOR	RGB color for the fifth percentage band. In case of Top N, Top Flow widgets only three percentage bands are available. This value is not applicable.	int	

TABLE 455 PM_DASHBOARD_WIDGET (Continued)

Field	Definition	Format	Size
CATEGORY	Category of the dashboard monitor. 0 - System defined, 1 - User defined, 2 - Promoted from Historical Graph, 3 - Promoted from Real time Graph. Default is 0..	smallint	
GRAPH_ENABLED	Enable or disable the time series graph for Top n or Bottom n widgets. 0 = disabled, 1 = enabled.	smallint	
FILTER_CRITERIA	Stores the filter criteria to be applied on the selected measure for products or ports.	varchar	256
FILTER_VALUE	Stores the measure value to be used in the filter criteria.	double precision	
USE_DASHBOARD_SCOPE	Use Dashboard scope or widget scope. 0 - Widget scope, 1 - Dashboard scope.	smallint	
PORT_TYPE	Types of ports to use for port measure widgets: 0 - All Ports, 1 - ISL Ports, 2 - Host Ports, 3 - Storage Ports.	smallint	

TABLE 456 PM_DATA_COLLECTOR

Field	Definition	Format	Size
ID	Primary key column.	int	
NAME	The name of the collector definition.	varchar	128
STATUS	Status of the collector. 0 - disabled and 1 - enabled. Default - 0.	smallint	
TYPE	Target type of the snmp collector data. for device level collector the target type is 0, for port level it is 1.	smallint	
POLLING_INTERVAL	Time interval in seconds; indicates the frequency with which the collector will poll the device to get the data.	int	
CREATED_TIME	Collector created time.	timestamp	
CREATE_USER_ID	The user id who has created this collector.	int	
ENABLE_THRESHOLD	Widget definition created server time.	smallint	
THRESHOLD	Stores the threshold value.	double precision	
REARM	Stores the rearm value.	double precision	
THRESHOLD_OP	Stores the threshold operator value.	varchar	10
REARM_OP	Stores the rearm operator value.	varchar	10
IS_REARM_ABS	Whether or not the rearm. Default - 0.	smallint	
THRESHOLD_SEVERITY	The severity for the threshold event.	smallint	
REARM_SEVERITY	The severity for the rearm event.	smallint	
IS_SYSTEM	Indicates whether this is a system built in collector, user cannot delete it. Default - 1.	smallint	

TABLE 457 PM_MEASURE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
DESCRIPTION	The description of the .	varchar	64
NAME	Name of the measure.	varchar	32

TABLE 458 PM_STATS_AGING_POLICY

Field	Definition	Format	Size
ID	Auto generated unique Identifier. Primary key for the table.	int	
RAW_SAMPLE_AGE	The maximum time in seconds for retaining records in the PM stats raw sample tables (TIME_SERIES_DATA_1 or TIME_SERIES_DATA_2) in database.	int	64
THIRTY_MIN_SAMPLE_AGE	The maximum time in seconds for retaining records in the PM stats 30min sample tables (TIME_SERIES_DATA_1_30MIN and TIME_SERIES_DATA_2_30MIN) in database.	int	
TWO_HOUR_SAMPLE_AGE	The maximum time in seconds for retaining records in the PM stats 2hour sample tables (TIME_SERIES_DATA_1_2HOUR and TIME_SERIES_DATA_2_2HOUR) in database.	int	
ONE_DAY_SAMPLE_AGE	The maximum time in seconds for retaining records in the PM stats 1day sample tables (TIME_SERIES_DATA_1_1DAY, TIME_SERIES_DATA_2_1DAY) in database.	int	
POLICY_TYPE	Type of the aging policy. 100 is Default aging; 101 is Raw samples to 1 day.	int	
ACTIVE	State of the aging policy. 1 is Active, 0 is Inactive.	int	32

TABLE 459 PM_WIDGET_MEASURE_TYPE

Field	Definition	Format	Size
Type	Primary key column.	int	
NAME	Storing the NAME of the measure.	varchar	64

TABLE 460 PM_WIDGET_MEASURE_TYPE_ENTRY

Field	Definition	Format	Size
WIDGET_ID	The id of the widget definition.	int	
MEASURE_TYPE	stores measure type id of the widget, a widget could map to multiple measure types.	int	

TABLE 461 PM_WIDGET_MONITOR_TYPE

Field	Definition	Format	Size
Type	Primary key column.	int	
NAME	Storing the NAME of the monitor type.	varchar	64

TABLE 462 PM_WIDGET_TARGET_ENTRY

Field	Definition	Format	Size
WIDGET_ID	The ID of the widget definition.	int	
TARGET_TYPE	0 - Device 1 - Port	smallint	
TARGET_ID	Stores device ID if taret_TYPE is Device, or interface DB ID if target TYPE is port.	int	

TABLE 463 PM_WIDGET_TIME_SERIES_ENTRY

Field	Definition	Format	Size
WIDGET_ID	The ID of the widget definition.	int	
COLLECTOR_ID	ID of the PERF_COLLECTOR.	int	
TARGET_TYPE	0 - Device 1 - Port	smallint	
TARGET_ID	Stores device ID if taret_TYPE is Device, or interface DB ID if target TYPE is port.	int	
MEASURE_ID	Measure table DB ID.	int	
MEASURE_INDEX	Index value for a MIB variable. For scalar value it will be empty.	varchar	256

TABLE 464 PM_WIDGET_TOP_N_COLLECTOR_ENTRY

Field	Definition	Format	Size
WIDGET_ID	The ID of the widget definition.	int	
COLLECTOR_ID	ID of the PERF_COLLECTOR.	int	
MEASURE_ID	Measure table DB ID.	int	
DIRECTION	The direction of the port measure. 0 - default (not used) 1 - receiving 2 -transmitting	smallint	

TABLE 465 PM_WIDGET_USER_ENTRY

Field	Definition	Format	Size
WIDGET_ID	The ID of the widget definition.	int	
USER_ID	ID of the user who is using the widget definition.	int	

TABLE 466 POE_THRESHOLD

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
TYPE	This field indicates if the threshold is defined for product and port level measure. • 0 = product level • 1 = port level	smallint	
DEVICE_ID	This is the foreign key reference key to the Device Table.	int	
INTERFACE_ID	This is the foreign key reference key to the Interface Table.	int	
ENABLED	Flag to indicate of defined threshold is enabled or not. • 0 = disabled • 1 = enabled	smallint	
VALUE	Value of the measure at which threshold is defined.	double precision	
INTERVAL	Time interval at which threshold is triggered.	int	
MEASURE	Product and port level poe measure definition.	smallint	
CONDITION	Condition like >< = to the defined threshold value at which threshold is triggered • 0 > (Greater Than) • 1 >= (Greater Than or Equal) • 2 < (Less Than) • 3 < = (Less Than or Equal) • 4 = (Equal to) • 5 != (Not Equal To)	smallint	
SEVERITY	Severity level of defined threshold on port and product Poe measures.	int	

TABLE 467 POE_THRESHOLD_EVENT

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
TIME_STAMP	This field indicates the time at which a particular threshold was triggered.	bigint	
THRESHOLD_ID	This is the foreign key reference key to the POE_THRESHOLD Table.	int	
EVENT_VALUE	Value of the measure at which threshold was triggered.	double precision	

TABLE 468 POLICY_RULE

Field	Definition	Format	Size
ID	Primary key for the table.	int	
NAME	Rule name.	varchar	255
DESCRIPTION	Rule description.	varchar	1024

TABLE 468 POLICY_RULE (Continued)

Field	Definition	Format	Size
TYPE	Policy Monitor rule type. For e.g. 1-zoning status check, 2-orphaned zone check, 3-fan in ratio check, 4- event registration check... 10-Configuration Rule, etc.	smallint	
CATEGORY	Policy Monitor rule category. For example, 0 - pre-defined, 1 - user-defined for configuration rule check.	smallint	

TABLE 469 POLICY_RULE_MAP

Field	Definition	Format	Size
ID	Primary key for the table.	int	
DEPLOYMENT_CONFIGURATION_ID	Foreign key reference to DEPLOYMENT_CONFIGURATION.ID.	int	
POLICY_RULE_ID	Foreign key reference to POLICY_RULE.ID.	int	
ATTRIBUTES	Attributes for pre-canned rules, this is name value pair string, use & as delimiter. For example minConn=2&maxConn=5.	varchar	255

TABLE 470 PORT_BOTTLENECK_CONFIG

Field	Definition	Format	Size
SWITCH_PORT_ID	The database ID of the switch port that the configuration belongs to.	int	
BOTTLENECK_DETECT_ENABLED	Flag indicates if bottleneck detection is enabled or not. The default value is 0.	smallint	
ALERTS_ENABLED	Flag indicates if bottleneck detection alerts is enabled or not.The default value is -1.	smallint	
CONGESTION_ THRESHOLD	Value of bottleneck detection congestion threshold in percent. The default value is -1.	double precision	
LATENCY_THRESHOLD	Value of bottleneck detection latency threshold in percent. The default value is -1.	double precision	
LATENCY_SEVERITY	The factor by which throughput must drop in a second in order for that second to be considered affected by latency bottlenecking. Range (1 to 1000).	int	
LATENCY_TIME	The minimum fraction of a second that must be affected by latency in order for that second to be considered affected by latency bottlenecking. Range (0 to 1).	double precision	
WINDOW_	Value of bottleneck detection latency window in millisecond. The default value is 0.	int	
QUIET_TIME	Value of bottleneck detection quiet time in millisecond. The default value is 0.	int	
CREATION_TIME	Creation time of the record.	timestamp	
LAST_UPDATE_TIME	Last update time of the record.	timestamp	

TABLE 471 PORT_BOTTLENECK_STATUS

Field	Definition	Format	Size
SWITCH_PORT_ID	The database ID of the switch port that the status belongs to.	int	
STATUS	Flag indicates bottleneck status of the switch port.	smallint	

TABLE 472 PORT_COMMISSION_CIMOM_SERVER

Field	Definition	Format	Size
ID	Primary key for the table.	int	
DESCRIPTION	User defined description of the CIMOM Server.	varchar	1024
NETWORK_ADDRESS	IPv4 or IPv6 address or Host name of the CIMOM server.	varchar	64
CIM_NAMESPACE	Name of the namespace where this CIM_FCPort CIM Class is located.	varchar	256
PORT	Port number which CIMOM server is listening.	int	
SSL_ENABLED	Protocol used for connecting CIMOM server. Default protocol will be HTTPS. If HTTPS is not reachable fall back to HTTP. Supported values 0 - HTTP, 1 - HTTPS	int	
USER_ID	User Id to be used for authenticating CIMOM Server.	varchar	128
PASSWORD	Password to be used for authenticating. Stored in encrypted format.	varchar	512
STATUS	Status before and after contacting the CIMOM Server. Possible values are 0 - OK, 1- Not Contacted Yet , 2 - Credentials Updated, 3 - Credentials Failed, 4 - Not Reachable.	int	
LAST_CONTACTED_TIME	Last time CIMOM server contacted. This will be updated when we test the reachablity of the CIMOM Server and when we perform port decommission/recommission.	timestamp	
ERROR_MESSAGE	Detailed error message. Applicable when communication to CIMOM Server failed.	varchar	2048

TABLE 473 PORT_FENCING_POLICY

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of the policy. The length of the field should be 62 because M-Model switch supports only maximum 62 characters.	varchar	62
TYPE	• 0 = ISL Protocol • 1 = Link • 2 = Security	smallint	
THRESHOLD_LIMIT	Threshold Limits for M-Model Switch.	int	
THRESHOLD_DURATION	Duration In minutes for M-Model Switch.	int	

TABLE 473 PORT_FENCING_POLICY (Continued)

Field	Definition	Format	Size
DEFAULT_POLICY	• 1 = the default port fencing policies. • 0 = the non-default policies. The default port fencing policies are: For ISL - Default Protocol Error Policy For Link Violation type - Default Link Level Policy For Security - Default Security Policy	smallint	
B_THRESHOLD_LIMIT	Threshold Limits for B-Model Switch (Not Supported).	int	
B_THRESHOLD_DURATION	Duration in minutes for B-Model Switch (Not Supported).	int	

TABLE 474 PORT_FENCING_POLICY_MAP

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
POLICY_ID	Foreign key to ID column of PORT_FENCING_POLICY table.	int	
LEVEL	• 0 = All Fabric • 1 = Fabric • 2 = Core Switch Group • 3 = Switch • 4 = Port Type • 5 = Port List	smallint	
SUB_LEVEL	• 1 = E_Port • 2 = F_Port • 3 = FL_Port, Fabric WWN, Switch WWN	char	23
NODE	WWN of Node which policy assigned.	char	23
INHERITANCE	Directly assigned or inherited from root level. • 0 = Directly assigned • 1 = Indirectly assigned	smallint	

TABLE 475 PORT_PROFILE

Field	Definition	Format	Size
ID	Auto generated id for the created profile	int	
SWITCH_ME_ID	Incase of a VCS discovery in M/C mode this is the cluster meid. Incase of VCS discovery in F/C mode, this is a member ID. Incase of a VDX this is a member ID	int	
NAME	Name of the port profile	varchar	255
SWITCH_PORT_MODE	Mode for vlan configuration it can be 0=access 1= trunk 2= converged	smallint	
STATE	Profile state 0=created 1=activate	smallint	
ACL_PROFILE	If port Profile has an acl profile or not. 0=NO 1=YES	smallint	
QOS_PROFILE	If port Profile has an qos profile or not. 0=NO 1=YES	smallint	
VLAN_PROFILE	If port Profile has an vlan profile or not.0=NO 1=YES	smallint	
VLAN_DETAILS	This column lists the way vlan is configured 0=NONE 1=ALL 2=SELECTED 3=EXC EPT	smallint	

TABLE 475 PORT_PROFILE (Continued)

Field	Definition	Format	Size
DEFAULT_PROFILE	This flag determines if this profile is a default profile.0=NO 1=YES	smallint	
ACL_NAME	Name of the mapped ACL	varchar	255
ACTIVATED	Profile activated 0= false 1=true	smallint	
FCOE_PROFILE	If port Profile has an fcoe profile or not. • 0=NO • 1=YES	smallint	
FCOE_MAP_NAME	Name of the FCoE Map.	varchar	255

TABLE 476 PORT_PROFILE_DOMAIN

Field	Definition	Format	Size
ID	Sequence number of the records.	int	
ME_ID	Foreign Key Reference to ID field of MANAGED_ELEMENT table.	int	
NAME	Name of the port profile domain.	varchar	255
DEFAULT_DOMAIN	This flag determines if this domain is a default domain. 0 - NO 1 - YES	smallint	

TABLE 477 PORT_PROFILE_DOMAIN_MAP

Field	Definition	Format	Size
PROFILE_DOMAIN_ID	Foreign Key Reference to ID field of PORT_PROFILE_DOMAIN table.	int	
PROFILE_ID	Foreign Key Reference to ID field of PORT_PROFILE table.	int	

TABLE 478 PORT_PROFILE_INTERFACE_MAP

Field	Definition	Format	Size
ID	Auto generated ID for the created profile	int	
PROFILE_ID	DB id of the port profile	int	
SWITCH_ME_ID	Managed element id of the cluster and its cluster members and a stand alone calisto	int	
INTERFACE_ID	ID of the interface table	int	
SWITCH_PORT_ID	Db id of the Switch port with matching interface	int	

TABLE 479 PORT_PROFILE_MAC_MAP

Field	Definition	Format	Size
ID	Auto generated ID for the created profile	int	
PROFILE_ID	DB id of the port profile	int	
MAC	Mac address mapped to the port profile	varchar	32
NAME	User assigned name to the mac	varchar	256

TABLE 480 PORT_PROFILE_QOS_MAP

Field	Definition	Format	Size
ID	Auto-generated ID for the created profile	int	
PROFILE_ID	DB ID of the port profile	int	
DCB_MODE	If the mode is dcb or non dcb. 1 : DCB.0: NON-DCB	smallint	
ETHERNET_MODE	This integer indicates if pause of priority flow control is set.0: PAUSE 1: PFC	smallint	
PAUSE_TX	Is PAUSE TX is enabled 0=NO 1=YES	smallint	
PAUSE_RX	Is PAUSE RX is enabled 0=NO 1=YES	smallint	
COS_COS	Name of the cos 2 cos map set in the NON DCB mode	varchar	256
TRAFFIC_CLASS	Name of the traffic class map set in the NON DCB mode	varchar	256
CEE_MAP	Name of the cee map set in the DCB mode	varchar	256
COS	Default COS value for QoS Profile can range from 0-7 if set	int	
TRUST_COS	Is trust cos enabled 0=NO 1=YES	smallint	

TABLE 481 PORT_PROFILE_QOS_PFC_MAP

Field	Definition	Format	Size
ID	Auto generated id for the created profile	int	
PROFILE_ID	DB id of the port profile	int	
COS0_TX	TX setting for this cos field 0: NO 1: YES	smallint	
COS0_RX	RX setting for this cos field 0: NO 1: YES	smallint	
COS1_TX	TX setting for this cos field 0: NO 1: YES	smallint	
COS1_RX	RX setting for this cos field 0: NO 1: YES	smallint	
COS2_TX	TX setting for this cos field 0: NO 1: YES	smallint	
COS3_TX	TX setting for this cos field 0: NO 1: YES	smallint	
COS3_RX	RX setting for this cos field 0: NO 1: YES	smallint	
COS4_TX	TX setting for this cos field 0: NO 1: YES	smallint	
COS4_RX	RX setting for this cos field 0: NO 1: YES	smallint	
COS5_TX	TX setting for this cos field 0: NO 1: YES	smallint	
COS5_RX	RX setting for this cos field 0: NO 1: YES	smallint	
COS2_RX	RX setting for this cos field 0: NO 1: YES	smallint	
COS6_TX	TX setting for this cos field 0: NO 1: YES	smallint	
COS6_RX	RX setting for this cos field 0: NO 1: YES	smallint	
COS7_TX	TX setting for this cos field 0: NO 1: YES	smallint	
COS7_RX	RX setting for this cos field 0: NO 1: YES	smallint	

TABLE 482 PORT_PROFILE_VLAN_MAP

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
PROFILE_ID	DB id of the port profile.	int	
VLANID	Configured vlan id on the switch.	int	
VLAN_TYPE	Type of this vlan, 0 : access 1: trunk.	int	
MAC_GROUP_DB_ID	Nullable Foreign Key Reference to ID field of MAC_GROUP table. In case of VLAN_TYPE 3, MAC_GROUP table entry created with empty GROUP_ID with TYPE 3 and MAC_GROUP_MEMBER have the mac address details. In case of VLAN_TYPE 4, MAC_GROUP table entry created with valid GROUP_ID and TYPE(3).	int	
CTAG_ID	This will be populated only if VLAN_TYPE is 6 and 7. This as an Incoming customer tag (c-tag) associated with a GVLAN and its applicable only for trunk mode. If the TLS (Transparent LAN Service) is enabled in the device, a pre-defined range of values are used for C-Tag.	text	

TABLE 483 PORT_VLAN

Field	Definition	Format	Size
VLAN_DB_ID	Database ID of the VLAN instance which is associated with the port.	int	
STP	STP value for VLAN. Possible values are 0-Disabled, 1-STP, 2-RSTP, 3-MSTP, 4-PVST and 5-RPVST.	num	(1,0)
TVF	0 means Disabled, 1 means Enabled. When TVF is enabled, packets are allowed to be forwarded without any form of CPU intervention including MAC learning and MAC destination lookups. When it is disabled (VLAN CPU Protection), it allows traffic intended for pure Layer2.	smallint	
VLAN_ID	The existing Data type short has been modified to integer. Hence it supports 16 bit additionally.	smallint	
QOS	Quality of service for port VLAN.	smallint	
GLOBAL_VLAN_DB_ID	Database ID of the GLOBAL_VLAN instance which is associated with the port.	int	
STP_INSTANCE_ID	Database ID of the STP_INSTANCE instance which is associated with the port.	int	
ADMIN_STATUS	• 0 = Disabled • 1 = Enabled	smallint	
FCOE_ENABLED	Signifies whether this VLAN is the default FCoE VLAN on the DCB switch.	smallint	
PVLAN_TYPE	pvlan_type value for vlan.0- Normal VLAN. The following are PVLAN Types applicable for NOS4.0 and above.1- Primary PVLAN, 2- RSPAN, 3 -TLS..	int	

TABLE 483 PORT_VLAN (Continued)

Field	Definition	Format	Size
PRIMARY_VLAN_ID	Private VLAN domain is built with one primary VLAN and one or more secondary VLANs. This column represents primary VLAN ID associated with this secondary Isolated/Community VLAN (if PVLAN_TYPE column value is 2 or 3) in private VLAN domain. For primary VLAN (if PVLAN_TYPE column value is 1) in private VLAN domain and normal VLAN (if PVLAN_TYPE column value is 0) , then default value (i.e 0) will be populated.	int	
TLS_ID	It represents Transparent LAN Service id and its supported range is 1-1000 in GVLAN. The TLS and VLAN are one-one mapping.	int	

TABLE 484 PRIVILEGE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Privilege Name.	varchar	128
AREA	Privilege Area. 0= Application 1= SAN 2= IP	smallint	

TABLE 485 PRODUCT_APP

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
MENU_TEXT	Name of the product menu.	varchar	256
PROP1_KEY	First condition name to be satisfied by a selected product to launch a particular tool.	varchar	256
PROP1_VALUE	First condition value to be satisfied by a selected product to launch a particular tool.	varchar	256
PROP2_KEY	Second condition name to be satisfied by a selected product to launch a particular tool.	varchar	256
PROP2_VALUE	Second condition value to be satisfied by a selected product to launch a particular tool.	varchar	256
TOOL_ID	The tool to be used for launching the application.	int	
PARAMETERS	Link to that application.	varchar	256
IP_SELECTED	Selected IP Address option.	smallint	
WWN_SELECTED	Selected WWN option.	smallint	

TABLE 486 PROTOCOL_VLAN

Field	Definition	Format	Size
VLAN_DB_ID	Database ID of the VLAN instance which is associated with the protocol.	int	
PROTOCOL	Protocol for VLAN. Possible values are 1-IP, 2-IPX, 3-AppleTalk, 4-DECnet, 5-NetBIOS, 6-Other and 7-IPv6.	num	(4,0)

TABLE 487 PURGED_SWITCH

Field	Definition	Format	Size
WWN		char	23
NAME		varchar	64
VIRTUAL_FABRIC_ID		smallint	
USER_NAME		varchar	64
PASSWORD		varchar	128
IP_ADDRESS		varchar	128
PORT_NUMBER		smallint	
RETRY_COUNT		smallint	
TIMEOUT		smallint	
VERSION		varchar	6
READ_COMMUNITY_STRING		varchar	128
WRITE_COMMUNITY_STRING		varchar	128
SNMP_USER_NAME		varchar	128
CONTEXT_NAME		varchar	128
AUTH_PROTOCOL		varchar	16
AUTH_PASSWORD		varchar	64
PRIV_PROTOCOL		varchar	16
PRIV_PASSWORD		varchar	64
TYPE		smallint	
AUTO_SNMP	If automatic snmp configuration is enabled for this switch. 0 is manual, 1 is automatic.	smallint	
SWITCH_TYPE	The switch model type.	smallint	

TABLE 488 QRTZ_BLOB_TRIGGERS

Field	Definition	Format	Size
TRIGGER_NAME*	Name of the trigger.	varchar	200
TRIGGER_GROUP*	Name of the trigger group.	varchar	200

TABLE 488 QRTZ_BLOB_TRIGGERS (Continued)

Field	Definition	Format	Size
BLOB_DATA	The Scheduler info.	bytea	
SCHED_NAME	DCMScheduler.	varchar	120

TABLE 489 QRTZ_CALENDARS

Field	Definition	Format	Size
CALENDAR_NAME*	Name of the Calendar.	varchar	200
CALENDAR	Calendar object.	bytea	
SCHED_NAME	DCMScheduler.	varchar	120

TABLE 490 QRTZ_CRON_TRIGGERS

Field	Definition	Format	Size
TRIGGER_NAME*	Name of the trigger.	varchar	200
TRIGGER_GROUP*	Name of the trigger group.	varchar	200
CRON_EXPRESSION	The CRON trigger Expression (ex:"0 0 12 * * ?" - meaning:Fire at 12pm (noon) every day).	varchar	120
TIME_ZONE_ID	Given "cron" expression resolved with respect to the TimeZone.	varchar	80
SCHED_NAME	DCMScheduler.	varchar	120

TABLE 491 QRTZ_FIRED_TRIGGERS

Field	Definition	Format	size
ENTRY_ID*	Fired instance ID.	varchar	95
TRIGGER_NAME	Name of the trigger.	varchar	200
TRIGGER_GROUP	Name of the trigger group.	varchar	200
IS_VOLATILE	Whether the job should not be persisted in the JobStore for re-use after the program restarts.	boolean	
INSTANCE_NAME	Trigger instance name.	varchar	200
FIRED_TIME	The trigger fired time.	num	(13,0)
STATE	The fired trigger job state.	varchar	16
JOB_NAME	Name of the job.	varchar	200
JOB_GROUP	Name of the job group.	varchar	200
IS_STATEFUL	Whether the job implements the interface StatefulJob.	boolean	
REQUESTS_RECOVERY	True or false.	boolean	
SCHED_NAME	DCMScheduler.	bigint	

TABLE 492 QRTZ_JOB_DETAILS

Field	Definition	Format	Size
JOB_NAME*	Name of the job.	varchar	200
JOB_GROUP*	Name of the job group.	varchar	200
DESCRIPTION	Description of the job (optional).	varchar	200

TABLE 492 QRTZ_JOB_DETAILS (Continued)

Field	Definition	Format	Size
JOB_CLASS_NAME	The instance of the job that will be executed.	varchar	200
IS_DURABLE	Whether the job should remain stored after it is orphaned.	boolean	
IS_VOLATILE	Whether the job should not be persisted in the JobStore for re-use after program restarts.	boolean	
IS_STATEFUL	Whether the job implements the interface StatefulJob.	boolean	
REQUESTS_RECOVERY	Instructs the scheduler whether or not the job should be re-executed if a "recovery" or "fail-over" situation is encountered.	boolean	
JOB_DATA	To persist the job-related and application-related informations.	bytea	
SCHED_NAME	DCMScheduler.	varchar	120

TABLE 493 QRTZ_JOB_LISTENERS

Field	Definition	Format	Size
JOB_NAME*	Name of the job.	varchar	80
JOB_GROUP*	Name of the job group.	varchar	80
JOB_LISTENER*	Job listener action class instance.	varchar	80

TABLE 494 QRTZ_LOCKS

Field	Definition	Format	Size
LOCK_NAME*	Resource identification name assigned by user.	varchar	40
SCHED_NAME	DCMScheduler.	varchar	120

TABLE 495 QRTZ_PAUSED_TRIGGER_GRPS

Field	Definition	Format	Size
TRIGGER_GROUP*	Name of the trigger group.	varchar	200
SCHED_NAME	DCMScheduler.	varchar	120

TABLE 496 QRTZ_SCHEDULER_STATE

Field	Definition	Format	Size
INSTANCE_NAME*	Instance of the scheduler.	varchar	200
LAST_CHECKIN_TIME	Last fired time in milliseconds.	bigint	
CHECKIN_INTERVAL	Repeat interval.	bigint	
RECOVERER	Misfire instruction.	varchar	80
SCHED_NAME	DCMScheduler.	varchar	120

TABLE 497 QRTZ_SIMPLE_TRIGGERS

Field	Definition	Format	size
TRIGGER_NAME*	Name of the trigger	varchar	200
TRIGGER_GROUP*	name of the trigger group	varchar	200

TABLE 497 QRTZ_SIMPLE_TRIGGERS (Continued)

Field	Definition	Format	size
REPEAT_COUNT	number of times to repeat	bigint	
REPEAT_INTERVAL	interval for first and second job	bigint	
TIMES_TRIGGERED	Number of times the corresponding trigger fired	bigint	
SCHED_NAME	DCMScheduler.	varchar	120

TABLE 498 QRTZ_SIMPROP_TRIGGERS

Field	Definition	Format	size
SCHED_NAME		varchar	120
TRIGGER_NAME		varchar	200
TRIGGER_GROUP		varchar	200
STR_PROP_1		varchar	512
STR_PROP_2		varchar	512
STR_PROP_3		varchar	512
INT_PROP_1		int	
INT_PROP_2		int	
LONG_PROP_1		bigint	
LONG_PROP_2		bigint	
DEC_PROP_1		numeric	13,4
DEC_PROP_2		numeric	13,4
BOOL_PROP_1		boolean	
BOOL_PROP_2		boolean	

TABLE 499 QRTZ_JTRIGGER_LISTENERS

Field	Definition	Format	Size
TRIGGER_NAME*	Name of the trigger.	varchar	80
TRIGGER_GROUP*	Name of the trigger group.	varchar	80
TRIGGER_LISTENER*	The listener action.	varchar	80

TABLE 500 QRTZ_TRIGGERS

Field	Definition	Format	Size
TRIGGER_NAME*	Name of the trigger.	varchar	200
TRIGGER_GROUP*	Name of the trigger group.	varchar	200
JOB_NAME	Name of the job.	varchar	200
JOB_GROUP	Name of the job group.	varchar	200
IS_VOLATILE	Whether the trigger should be persisted in the JobStore for re-use after program restarts.	boolean	

TABLE 500 QRTZ_TRIGGERS (Continued)

Field	Definition	Format	Size
DESCRIPTION	A description for the trigger instance - may be useful for remembering/displaying the purpose of the trigger, though the description has no meaning to Quartz.	varchar	250
NEXT_FIRE_TIME	The next fire time in milliseconds.	bigint	
PREV_FIRE_TIME	The previous fired time in milliseconds.	bigint	
TRIGGER_STATE	The state of the trigger (viz. Error, wait,etc.)	varchar	16
TRIGGER_TYPE	The type of the trigger (Simple,cron).	varchar	8
START_TIME	The job start time.	bigint	
END_TIME	The job end time (-1 means infinite).	bigint	
CALENDAR_NAME		varchar	200
MISFIRE_INSTR	Instructs the scheduler to execute the misfired job.	smallint	
JOB_DATA	Persists the job-related info.	bytea	
SCHED_NAME	DCMScheduler.	varchar	120

TABLE 501 QUERY_BASED_DEVICE_GROUP

Field	Definition	Format	Size
DEVICE_GROUP_ID		int	
QUERY_TEXT		text	
GROUP_CRITERIA	Holds the dynamic device group criteria XML value.	text	

TABLE 502 QUORUM_CARD_GROUP_MAPPING

Field	Definition	Format	Size
ID*		int	
ENCRYPTION_GROUP_ID	Foreign key reference to the ENCRYPTION_GROUP for which an authorization card is registered.	int	
SMART_CARD_ID	Foreign key reference to the SMART_CARD that is registered as an authorization card for the encryption group.	int	

TABLE 503 RAS_LOG

Field	Definition	Format	Size
MSG_ID*	Message ID of the event.	varchar	15
MODULE_ID	Module ID of the event.	varchar	10
SEVERITY	Severity of the event.	varchar	10
CAUSE	Probable root cause for the event.	varchar	4096
ACTION	Recommended action for the event.	varchar	4096
OLD_MSG_ID	Old message ID.	varchar	45
TYPE	Indicates the product type associated to the RASLOG. 1 - Fabric OS 2 - Network OS	int	

TABLE 504 RECIPIENT_TYPE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
TYPE	Type of the recipient (Syslog or SNMP).	varchar	20

TABLE 505 RECOVERY_CARD_GROUP_MAPPING

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
ENCRYPTION_GROUP_ID	Foreign key reference to the ENCRYPTION_GROUP for which a recovery card is registered.	int	
SMART_CARD_ID	Foreign key reference to the SMART_CARD that is registered as a recovery card for the encryption group.	int	
POSITION_	The position of the card within the recovery card set. 1 = first card, 2 = second card, etc.	int	

TABLE 506 REPORT_TYPE

Field	Definition	Format	Size
ID*	Meta Data for available reports.	int	
NAME	Report name.	varchar	128
DESCRIPTION	Report type description.	varchar	256

TABLE 507 REPORT_TEMPLATE

Field	Definition	Format	Size
ID	The primary key of the table.	int	
NAME	Name of the report and the report names must be descriptive. For example, Wired Device Report.	varchar	256
TITLE	The title of the report that briefly describes the report contents. This title will also be used for the report header and menu item. Title should be unique. For example, Wired Products List.	varchar	256
CREATED_TIME	Timestamp of when the report was created.	timestamp	
CREATED_BY	Foreign key to the user table, to identify which user created the report.	int	
REPORT_TYPE	0 = Precanned template which will not be deleted or edited, 1 = Editable report which can be deleted as well, 2 = Not Editable report but can be deleted.	int	
REPORT_DEFINITION	XML representation of the report.	text	
PARAMETERIZED		int	
CATEGORY	Indicates the report template category. 0 = SAN domain, 1 = IP domain. Default value is 0.	int	
SHARED	Indicates if the report template is shared or not. 0 = Not shared, 1 = Shared. Default value is 0.	int	

TABLE 508 REPORT_TEMPLATE_FAVORITE

Field	Definition	Format	Size
FAVORITED_BY_USER_ID	Indicates the user ID of the user who favorited this report template.	int	
FAVORITED_REPORT_TEMPLATE_ID	Indicates the favorited report template ID.	int	

TABLE 509 REPORT_DRILLDOWN_TEMPLATE

Field	Definition	Format	Size
ID*	The primary key of the table.	int	
REPORT_TEMPLATE_ID	References the ID column in the REPORT_TEMPLATE table.	int	
NAME	Name of the report. Names should be descriptive so users will know exactly what kind of report they will be running or scheduling. E.g. Wired Device Report.	varchar	256
REPORT_DRILLDOWN_DEFINITION	XML representation of the report.	text	

TABLE 510 RESOURCE_FABRIC_MAP

Field	Definition	Format	Size
RESOURCE_GROUP_ID*	Resource group ID.	int	
FABRIC_ID*	Fabric ID, which is in the resource group.	int	

TABLE 511 RESOURCE_GROUP

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Resource group name.	varchar	128
DESCRIPTION	Resource group description.	varchar	512

TABLE 512 RESOURCE_HOST_MAP

Field	Definition	Format	Size
RESOURCE_GROUP_ID	Resource Group ID	int	
HOST_ID	HOST_ID,which is in the resource group	int	

TABLE 513 ROLE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Role name.	varchar	128

TABLE 513 ROLE (Continued)

Field	Definition	Format	Size
DESCRIPTION	Role description.	varchar	512
HIDDEN	Field to identify whether the role is Hidden from users or not. Values: • 0= Not Hidden • 1= Hidden Currently, only ''All Users'' Role is hidden and other roles are visible to user. Default value is 0.	smallint	

TABLE 514 ROLE_PRIVILEGE_MAP

Field	Definition	Format	Size
ROLE_ID*	User role ID.	int	
PRIVILEGE_ID*	Privilege ID.	int	
PERMISSION	Privilege permission: 1 = RO 2 = RW 0 = No privilege Default value is 0.	smallint	

TABLE 515 RULE_BLOCK_MAP

Field	Definition	Format	Size
POLICY_RULE_ID	Foreign key reference to POLICY_RULE.ID.	int	
CONFIG_BLOCK_ID	Foreign key reference to CONFIG_BLOCK.ID.	int	

TABLE 516 RULE_CONDITION_MAP

Field	Definition	Format	Size
POLICY_RULE_ID	Foreign key reference to POLICY_RULE.ID.	int	
CONFIG_CONDITION_ID	Foreign key reference to CONFIG_CONDITION.ID.	int	

TABLE 517 RULE_LOGICAL_EXPRESSION_MAP

Field	Definition	Format	Size
POLICY_RULE_ID	Policy rule ID.	int	
LOGICAL_EXPRESSION_XML	Configuration Rule Logical Expression XML.	text	

TABLE 518 SAN

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of this SAN.	varchar	256
CONTACT	Contact person for this SAN.	varchar	256
LOCATION	Location of this SAN.	varchar	256
DESCRIPTION	Description.	varchar	256

TABLE 518 SAN (Continued)

Field	Definition	Format	Size
STATS_COLLECTION	1 = statistics collection is enabled; otherwise, 0. Default value is 0.	smallint	
CREATION_TIME	time at which this record was created. Default value is 'now()'.	timestamp	
LAST_UPDATE_TIME	time when this was last updated. Default value is 'now()'.	timestamp	

TABLE 519 SAN_CONNECTION

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
SOURCE_SWITCH_ID	Foreign key to VIRTUAL_SWITCH table. This is the virtual switch ID of AG	int	
SOURCE_PORT_ID	Foreign key to SWITCH_PORT table. This is the switch port id of N-port	int	
SOURCE_PORT_WWN	WWN of the AG N port	varchar	32
SOURCE_PORT_TYPE	Type of source port	varchar	16
SOURCE_USER_PORT_NUMBER	User port number of AG N port	smallint	
DESTINATION_SWITCH_ID	Foreign key to VIRTUAL_SWITCH table. This is the virtual switch ID of edge switch	int	
DESTINATION_PORT_ID	Foreign key to SWITCH_PORT table. This is the switch port id of F-port	int	
DESTINATION_PORT_WWN	WWN of the F port	varchar	23
DESTINATION_PORT_TYPE	Type of destination port	varchar	16
DESTINATION_USER_PORT_NUMBER	User port number of F-port	smallint	
FABRIC_ID	Foreign key to FABRIC table	int	
TRUSTED	Indicates if the connection is trusted	smallint	
MISSING	Indicates if the connection is missing	smallint	
MISSING_TIME	Timestamp when the connection went missing	timestamp	
LAST_UPDATE_TIME	Last update time for this record	timestamp	
CREATION_TIME	Creation timestamp	timestamp	

TABLE 520 SCOM_HOST

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
HOST	The FQDN or the ip address of the host	varchar	256
DOMAIN	The domain of the SCOM server host	varchar	256
USER_NAME	The domain user to login into the SCOM Server	varchar	64

TABLE 520 SCOM_HOST (Continued)

Field	Definition	Format	Size
PASSWORD	The password to login into the SCOM Server	varchar	64
VERSION	The version of SCOM. Default is 6.1.7221.0 which is SCOM 2007 R2. The default value is '6.1.7221.0' .	varchar	32
TOKEN_ID	Unique ID for each SCOM host	varchar	32
STATUS	Status of Plug-in registration to the SCOM server where 0-registered, 1-unregistered, 2-authentication failed, 3-not reachable	int	

TABLE 521 SECURITY_POLICY

Field	Definition	Format	Size
VIRTUAL_SWITCH_ID *	DB ID of virtual_switch.	int	
POLICY_NUMBER*	IPSec Policy Number. The number can range from 1 to 32.	smallint	
POLICY_TYPE*	Type of the Policy. The possible values are IKE or IPSec	smallint	
ENCRYPTION_ALGORITHM	Encryption Algorithm for the policy.The following are the possible Encryption: NONE,DES,3DES,AES-128,AES-256,AES-CM-128 or AES-CM-256.	varchar	32
AUTHENTICATION_ALGORITHM	Authentication Algorithm for the policy: NONE SHA-1 MD5 AES-XCBC	varchar	32
PERFECT_FORWARD_POLICY_ENABLED	Perfect Forward Secrecy for the policy. The possible values are 0 or 1.	smallint	
DIFFIE_HELLMAN_GROUP	Diffie-Hellman Group used in PFS negotiation.	smallint	
SECURITY_ASSOC_LIFE	Association lifetime in seconds.	double precision	
SECURITY_ASSOC_LIFE_IN_MB	Security association lifetime in megabytes.	double precision	

TABLE 522 SELECTED_FLYOVER_PROPERTY

Field	Definition	Format	Size
PROPERTY_ID*	Refers to Flyover_Property ID from AVAILABLE_FLYOVER_PROPERTY table.	int	
USER_NAME*	The name of the user who selected the property to be shown on flyover.	varchar	128
POSITION_	The user preferred position of the selected flyover property.	int	

TABLE 523 SENSOR

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
CORE_SWITCH_ID		int	
SENSOR_ID	Identifies the sensor device , requested by SMIA and values filled in by Switch Asset Collector. Maps to Device Id in the html page. The default value is -1.	int	
CURRENT_READING	Identifies the current temperature reading sensor, requested by SMIA and values filled in by Switch Asset Collector, Maps to value field in the html page. The default value is -1.	bigint	
TYPE	The default value is -1.	int	
SUB_TYPE	The default value is -1.	int	
DESCRIPTION	Provides the description of the temperature sensor, requested by SMIA and values filled in by Switch Asset Collector	varchar	128
STATUS	provides the status of the sensor, requested by SMIA and values filled in by Switch Asset Collector,Values could be 0 or 1. 0 means faulty and 1 is ok.The default value is -1.	int	
OPERATIONAL_STATUS	provides the operational status of the sensor, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above. The default value is -1.	int	
PART_NUMBER	provides the part number of the sensor, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above	varchar	64
SERIAL_NUMBER	provides the serial number of the sensor, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above	varchar	64
VERSION	provides the version of the sensor, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above	varchar	32
CREATION_TIME	provides the record creation time, standard columns for Management application and values filled in by Switch Asset Collector	timestamp	
LAST_UPDATE_TIME	provides the record last updated time, standard columns for Management application and values filled in by Switch Asset Collector	timestamp	
FRU_TYPE	provides the type of the sensor, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above. The values represents FAN,PS, SLOT etc. The default value is -1.	int	

TABLE 523 SENSOR (Continued)

Field	Definition	Format	Size
UNIT_NUMBER	provides the unit number of the sensor, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above . This the gives the index of the unit. For SLOT FRU, this will be slot number. For FAN fru, this will be fan number. The default value is -1.	int	
STATE	provides the state of the sensor, requested by SMIA and values filled in by Switch Asset Collector will be available only from FOS 6.4 switches and above. This gives the value whether the FRU is On or Off . The default value is -1.	int	

TABLE 524 SFLOW_CHECKPOINT

Field	Definition	Format	Size
TABLE_TO_DROP	Staging child tables previously checkpointed (indicating that their aggregation was completed).	varchar	40

TABLE 525 SFLOW_HOUR_SUMMARY

Field	Definition	Format	Size
SLNUM	This column is used to store a counter value to identify the total row count.	bigserial	
TIME_IN_SECONDS	Data collection time in seconds.	int	
DEVICE_ID	ID of the product which sends the sflow traffic.	int	
IN_UNIT	Unit number of the incoming traffic interface. Default value is 0.	smallint	
IN_SLOT	Slot number of the incoming traffic interface.	smallint	
IN_PORT	Port number of the incoming traffic interface.	smallint	
OUT_UNIT	Unit number of the outgoing traffic interface. Default value is 0.	smallint	
OUT_SLOT	Slot number of the outgoing traffic interface.	smallint	
OUT_PORT	Port number of the outgoing traffic interface.	smallint	
IN_VLAN	Vlan ID of the incoming traffic interface.	smallint	
OUT_VLAN	Vlan ID of the outgoing traffic interface.	smallint	
IN_PRIORITY	Priority ID of the incoming traffic interface.	smallint	
OUT_PRIORITY	Priority ID of the outgoing traffic interface.	smallint	
SRC_MAC	MAC address of the source in the received sFlow packet.	byte	
DEST_MAC	MAC address of the destination in the received sFlow packet.	byte	
L3_SRC_ADDR	L3 address of the source in the received sFlow packet.	byte	
L3_DEST_ADDR	L3 address of the destination in the received sFlow packet.	byte	

TABLE 525 SFLOW_HOUR_SUMMARY (Continued)

Field	Definition	Format	Size
L3_PROTOCOL	L3 protocol value in the received sFlow packet. For example, ARP.	int	
IP_TOS	Type of service id in the received sFlow packet.	smallint	
L4_PROTOCOL	L4 protocol value in the received sFlow packet. For example, IGP	smallint	
L4_SRC_PORT	L4 source port number in the received sFlow packet.	int	
L4_DEST_PORT	L4 destination port number in the received sFlow packet.	int	
SRC_SUBNET_BITS	Subnet value of the incoming traffic interface.	smallint	
DEST_SUBNET_BITS	Subnet value of the outgoing traffic interface.	smallint	
LOCAL_AS	Second AS number of the received sFlow packet in case of BGP traffic.	bigint	
SRC_AS	Source AS number of the received sFlow packet in case of BGP traffic.	bigint	
SRC_PEER_AS	Peer AS number of the received sFlow packet in case of BGP traffic.	bigint	
SFLOW_IP_ROUTE_INFO_ID	route_info_id of the received sFlow packet in case of BGP traffic.	int	
IP_FLOW_LABEL	IP flow label value in the received sFlow packet.	int	
SRC_USER	Name of the Source user in the received sFlow packet.	int	
DEST_USER	Name of the destination user in the received sFlow packet.	int	
FRAMES	Number of frames transmitted through the sFlow sample collected.	bigint	
BYTES	Number of bytes transmitted through the sFlow sample collected.	bigint	
TCP_FLAGS	TCP flag value of the received sFlow packet.	smallint	
IN_PORT_TYPE	This column is used to store the port type of the incoming traffic interface. For VCS switch the value of • 0 means its edge port. • 1 means its trill port. For other devices Default value is 0.	smallint	
OUT_PORT_TYPE	This column is used to store the port type of the outgoing traffic interface. For VCS switch the value of • 0 means its edge port. • 1 means its trill port. For other devices Default value is 0.	smallint	

TABLE 526 SFLOW_IP_ROUTE_INFO

Field	Definition	Format	Size
SFLOW_IP_ROUTE_INFO_ID	This column is the primary key for IP routing information.	int	
LOCAL_PREF	Local preference value of routing information in the received sFlow packet.	int	

TABLE 526 SFLOW_IP_ROUTE_INFO (Continued)

Field	Definition	Format	Size
LAST_USED_TIME	Last used time of the routing information.	int	
DST_AS_PATH	Routing path information in the received sFlow packet.	varchar	2048
COMMUNITIES	Communities value in the received sFlow packet.	varchar	1024

TABLE 527 SFLOW_MINUTE_BGP

Field	Definition	Format	Size
SLNUM	This column is used to store a counter value to identify the total row count.	bigserial	
TIME_IN_SECONDS	Data collection time in seconds.	int	
DEVICE_ID	ID of the product which sends the sflow traffic.	int	
SRC_AS	Source AS number of the received sFlow packet in case of BGP traffic.	bigint	
SFLOW_IP_ROUTE_INFO_ID	route_info_id of the received sFlow packet in case of BGP traffic.	int	
IN_VLAN	Vlan ID of the incoming traffic interface.	smallint	
OUT_VLAN	Vlan ID of the outgoing traffic interface.	smallint	
FRAMES	Number of frames transmitted through the sflow sample collected.	bigint	
BYTES	Number of bytes transmitted through the sflow sample collected.	bigint	
IN_PORT_TYPE	Port type of the incoming traffic interface. For VCS member the value of, • 0 means its edge port. • 1 means its fabric port. For other devices Default value is 0.	smallint	
OUT_PORT_TYPE	Port type of the outgoing traffic interface. For VCS member the value of, • 0 means its edge port • 1 means its fabric port. For other devices Default value is 0.	smallint	

TABLE 528 SFLOW_MINUTE_BGP_SLNUM

Field	Definition	Format	Size
MAX_SLNUM	Maximum row count.	bigint	

TABLE 529 SFLOW_MINUTE_L3_SLNUM

Field	Definition	Format	Size
MAX_SLNUM	Maximum row count.	bigint	

TABLE 530 SFLOW_MINUTE_MAC

Field	Definition	Format	Size
SLNUM	This column is used to store a counter value to identify the total row count.	bigserial	
TIME_IN_SECONDS	Data collection time in seconds.	int	
DEVICE_ID	ID of the product which sends the sflow traffic.	int	
FRAMES	Number of frames transmitted through the sflow sample collected.	bigint	
BYTES	Number of bytes transmitted through the sflow sample collected.	bigint	
IN_VLAN	Vlan ID of the incoming traffic interface.	smallint	
OUT_VLAN	Vlan ID of the outgoing traffic interface.	smallint	
SRC_MAC	MAC address of the Source in the received sFlow packet.	bytea	
DEST_MAC	MAC address of the destination in the received sFlow packet.	bytea	
IN_PORT_TYPE	Port type of the incoming traffic interface. For VCS member the value of, • 0 means its edge port • 1 means its fabric port. For other devices Default value is 0.	smallint	
OUT_PORT_TYPE	Port type of the outgoing traffic interface. For VCS member the value of, • 0 means its edge port • 1 means its fabric port. For other devices Default value is 0	smallint	
L3_SRC_ADDR	This column is used to store the L3 address of the source in the received sFlow packet.	bytea	
L3_DEST_ADDR	This column is used to store the L3 address of the destination in the received sFlow packet.	bytea	

TABLE 531 SFLOW_MINUTE_MAC_SLNUM

Field	Definition	Format	Size
MAX_SLNUM	Maximum row count.	bigint	

TABLE 532 SFLOW_MINUTE_SUMMARY

Field	Definition	Format	Size
SLNUM	This column is used to store a counter value to identify the total row count.	bigserial	
TIME_IN_SECONDS	Data collection time in seconds.	int	
DEVICE_ID	ID of the product which sends the sflow traffic.	int	
FRAMES	Number of frames transmitted through the sflow sample collected.	bigint	
BYTES	Number of bytes transmitted through the sflow sample collected.	bigint	

TABLE 532 SFLOW_MINUTE_SUMMARY (Continued)

Field	Definition	Format	Size
IN_PORT_TYPE	Port type of the incoming traffic interface. For VCS member the value of, • 0 means its edge port • 1 means its fabric port. For other devices Default value is 0.	smallint	
OUT_PORT_TYPE	Port type of the outgoing traffic interface. For VCS member the value of, • 0 means its edge port • 1 means its fabric port. For other devices Default value is 0	smallint	

TABLE 533 SFLOW_MINUTE_VLAN

Field	Definition	Format	Size
SLNUM	This column is used to store a counter value to identify the total row count.	bigserial	
TIME_IN_SECONDS	Data collection time in seconds.	int	
DEVICE_ID	ID of the product which sends the sflow traffic.	int	
FRAMES	Number of frames transmitted through the sFlow sample collected.	bigint	
BYTES	Number of bytes transmitted through the sFlow sample collected.	bigint	
IN_VLAN	Vlan ID of the incoming traffic interface.	smallint	
OUT_VLAN	Vlan ID of the outgoing traffic interface.	smallint	
IN_PORT_TYPE	Port type of the incoming traffic interface. For VCS member the value of, • 0 means its edge port. • 1 means its fabric port. For other devices Default value is 0.	smallint	
OUT_PORT_TYPE	Port type of the outgoing traffic interface. For VCS member the value of, • 0 means its edge port. • 1 means its fabric port. For other devices Default value is 0.	smallint	

TABLE 534 SFLOW_MINUTE_VLAN_SLNUM

Field	Definition	Format	Size
MAX_SLNUM	Maximum row count.	bigint	

TABLE 535 SFLOW_REPORT_L3_SOURCE

Field	Definition	Format	Size
SFLOW_REPORT_L3_SOURCE_ID	Primary key autogenerated ID.	int	
REPORT_DEFINITION_ID	Report definition ID.	int	

TABLE 535 SFLOW_REPORT_L3_SOURCE (Continued)

Field	Definition	Format	Size
ADDRESS_GROUP_ID	ACL network group IDs mapped with a report definition.	int	
IP_SUBNET_DEFINITION_ID	Subnet IDs mapped with a Report definition.	int	

TABLE 536 SFLOW_REPORT_L4_SOURCE

Field	Definition	Format	Size
SFLOW_REPORT_L4_SOURCE_ID	Primary key autogenerated ID.	int	
REPORT_DEFINITION_ID	Report definition ID.	int	
SERVICE_PORT_DEFINITION_ID	Service port Id mapped with a report definition.	int	
SERVICE_GROUP_ID	Service group Id mapped with a report definition.	int	

TABLE 537 SFLOW_STAGING

Field	Definition	Format	Size
SLNUM	This column is used to store a counter value to identify the total row count.	bigserial	
TIME_IN_SECONDS	Data collection time in seconds.	int	
DEVICE_ID	ID of the product which sends the sflow traffic.	int	
IN_UNIT	Unit number of the incoming traffic interface. Default value is 0.	smallint	
IN_SLOT	Slot number of the incoming traffic interface.	smallint	
IN_PORT	Port number of the incoming traffic interface.	smallint	
OUT_UNIT	Unit number of the outgoing traffic interface. Default value is 0.	smallint	
OUT_SLOT	Slot number of the outgoing traffic interface.	smallint	
OUT_PORT	Port number of the outgoing traffic interface.	smallint	
IN_VLAN	Vlan ID of the incoming traffic interface.	smallint	
OUT_VLAN	Vlan ID of the outgoing traffic interface.	smallint	
IN_PRIORITY	Priority ID of the incoming traffic interface.	smallint	
OUT_PRIORITY	Priority ID of the outgoing traffic interface.	smallint	
SRC_MAC	MAC address of the source in the received sFlow packet.	bytea	
DEST_MAC	MAC address of the destination in the received sFlow packet.	bytea	
L3_SRC_ADDR	L3 address of the source in the received sFlow packet.	bytea	
L3_DEST_ADDR	L3 address of the destination in the received sFlow packet.	bytea	
L3_PROTOCOL	L3 protocol value in the received sFlow packet. For example, ARP.	int	
IP_TOS	Type of service ID in the received sFlow packet.	smallint	

TABLE 537 SFLOW_STAGING (Continued)

Field	Definition	Format	Size
L4_PROTOCOL	L4 protocol value in the received sFlow packet. For example, IGP.	smallint	
L4_SRC_PORT	L4 source port number in the received sFlow packet.	int	
L4_DEST_PORT	L4 destination port number in the received sFlow packet.	int	
SRC_SUBNET_BITS	Subnet value of the incoming traffic interface.	smallint	
DEST_SUBNET_BITS	Subnet value of the outgoing traffic interface.	smallint	
LOCAL_AS	Second AS number of the received sFlow packet in case of BGP traffic.	bigint	
SRC_AS	Source AS number of the received sFlow packet in case of BGP traffic.	bigint	
SRC_PEER_AS	Peer AS number of the received sFlow packet in case of BGP traffic.	bigint	
SFLOW_IP_ROUTE_INFO_ID	route_info_id of the received sFlow packet in case of BGP traffic.	int	
IP_FLOW_LABEL	IP flow label value in the received sFlow packet.	int	
SRC_USER	Name of the source user in the received sFlow packet.	int	
DEST_USER	Name of the destination user in the received sFlow packet.	int	
FRAMES	Number of frames transmitted through the sflow sample collected.	bigint	
BYTES	Number of bytes transmitted through the sflow sample collected.	bigint	
TCP_FLAGS	Tcp flag value of the received sFlow packet.	smallint	
IN_PORT_TYPE	This column is used to store the port type of the incoming traffic interface. For VCS switch the value of 0 means its edge port, and 1 means its trill port. For other devices Default value is 0.	smallint	
OUT_PORT_TYPE	This column is used to store the port type of the outgoing traffic interface. For VCS switch the value of 0 means its edge port, and 1 means its trill port. For other devices Default value is 0.	smallint	

TABLE 538 SFLOW_STAGING_SLNUM

Field	Definition	Format	Size
MIN_SLNUM	Maximum row count.	bigint	

TABLE 539 SMART_CARD

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
CARD_TYPE	Indicates how this smart card is configured: 0 = authorization card. The default value is 0.	smallint	

TABLE 539 SMART_CARD (Continued)

Field	Definition	Format	Size
CARD_INFO	Additional smart card details. For recovery set cards, the details include the recovery set size and the card"s position within the set; e.g., 2 of 5	varchar	64
CARDCN_ID	A unique name for the card, derived from the card"s serial number and usage	varchar	64
FIRST_NAME	Optional first name of the person responsible for this card.	varchar	64
LAST_NAME	Optional last name of the person responsible for this card	varchar	64
NOTES	User-supplied notes about the card.	varchar	256
PUBLIC_CERTIFICATE	The public key certificate of the card, in PEM format. Used to validate the card and set up a secure communications channel to the card.	varchar	4096
CERTIFICATE_LABEL	User-supplied name for the card''s public key certificate	varchar	256
GROUP_NAME	The name of the Encryption Group used to initialize the card. For recovery set cards, this identifies which group"s master key is backed up on the card.	varchar	64
CREATION_TIME	The date and time that the card was initialized. For recovery set cards, this is the date and time the master key was written to the card. The default value is 'now()'.	timestamp	

TABLE 540 SMIA_SAN_NAME

Field	Definition	Format	Size
NAME	'This will be the principal switch WWN of the fabric.';	varchar	24
ELEMENT_NAME	User friendly name to identify the SAN	varchar	32
IS_PRIMARY_FABRIC	This value will indicate whether principal switch WWN has primary ownership or not. In case of simple FC fabric, the value will be always 1. In case of Meta SAN, Fabric with highest principal switch WWN will have primary ownership (value 1) and other fabric entries within the same SAN will have value as 0.	int	

TABLE 541 SNAPSHOT_PRODUCT_STATUS

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
DEPLOYMENT_STATUS_ID	Foreign Key references DEPLOYMENT_STATUS_ID (id). Identifies the execution cycle for the deployment.	int	
MANAGED_ELEMENT_ID	Associates for which target the status applies to.	int	
SNAPSHOT_TYPE	Indicates the type of snapshot: • 1-Pre snapshot • 2-Post snapshot	int	

TABLE 541 SNAPSHOT_PRODUCT_STATUS (Continued)

Field	Definition	Format	Size
SNAPSHOT_TIME	Time when this pre/post snapshot occurred.		
MESSAGE	Detailed message for snapshot report.	text	

TABLE 542 SNMP_CREDENTIALS

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
VIRTUAL SWITCH_ID	Virtual switch ID for which this instance of the SNMP credentials apply.	int	
RECIPIENT_ID	Recipient in the MESSAGE_RECIPIENT table.	int	
POR)_NUMBER	Port number of the SNMP agent on the switch for get and set requests.	smallint	
RETRY_COUNT	Number of times to retry if get/set request to the SNMP agent times out. Default value is 3.	smallint	
TIMEOUT	Timeout value in seconds for a get/set request to the SNMP agent. Default value is 5.	smallint	
VERSION	SNMP agent version running on the switch, as in SNMPv1 or SNMPv3.	varchar	6
READ_COMMUNITY_ STRING	The SNMP Read-Only Community String is like a password. It is sent along with each SNMP Get-Request and allows (or denies) access to a device. The default value is "public". This is applicable if the agent is configured to operate in SNMPv1.	varchar	128
WRITE_COMMUNITY_ STRING	The SNMP Write-Only Community String is like a password. It is sent along with each SNMP Set-Request and allows (or denies) access to a device. The default value is "private". This is applicable if the agent is configured to operate in SNMPv1.	varchar	128
USER_NAME	A human readable string representing the name of the user. This is applicable if the agent is configured to operate in SNMPv3.	varchar	64
CONTEXT_NAME	Text ID associated with the user, used by the SNMP agent to provide different views. This is applicable if the agent is configured to operate in SNMPv3.	varchar	128
AUTH_PROTOCOL	An indication of whether messages sent or received on behalf of this user can be authenticated and if so, which authentication protocol to use. The supported values for this field are: usmNoAuthProtocol, usmHMACMD5AuthProtocol, and usmHMACSHAAuthProtocol. This is applicable if the agent is configured to operate in SNMPv3.	varchar	16
AUTH_PASSWORD	The localized secret key used by the authentication protocol for authenticating messages. This is applicable if the agent is configured to operate in SNMPv3.	varchar	64

TABLE 542 SNMP_CREDENTIALS (Continued)

Field	Definition	Format	Size
PRIV_PROTOCOL	An indication of whether messages sent or received on behalf of this user can be encrypted and if so, which privacy protocol to use. The current values for this field are: usmNoPrivProtocol and usmDESPrivProtocol. This is applicable if the agent is configured to operate in SNMPv3.	varchar	16
PRIV_PASSWORD	The localized secret key used by the privacy protocol for encrypting and decrypting messages. This is applicable if the agent is configured to operate in SNMPv3.	varchar	64

TABLE 543 SNMP_DATA

Field	Definition	Format	Size
ID	Primary key column.	serial	
MIB_OBJECT_ID	MIB Object ID.	int	
TARGET_TYPE	Target type of the SNMP collector data. The target type for, • device level collector is 0 • port level collector it is 1.	num	(2,0)
TARGET_ID	Target id of the SNMP collector data. for device level collector it will use deviceId, and for port level it will use interfaceId.	int	
VALUE	Value of the OID retrieved from the corresponding target.	double	
TIME_IN_SECONDS	Time, in seconds, at which the record was inserted.	int	
COLLECTOR_ID	Correspoding collector table ID.	int	
MIB_INDEX	Index value for a MIB varaible. For scalar value it will be empty.	varchar	256

TABLE 544 SNMP_DATA_1DAY

Field	Definition	Format	Size
ID	Primary key autogenerated ID.	int	
MIB_OBJECT_ID	The DB ID of MIB_OBJECT.	int	
TARGET_TYP	Target or source type can be, • device - 0 or • interface or ports - 1	num	(2,0)
TARGET_ID	DB ID of the target which can be device or interface.	int	
VALUE	Aggregated value inserted during the aging process.	double precision	
TIME_IN_SECONDS	Time, in seconds, at which the record was inserted.	int	
COLLECTOR_ID	DB ID of the collector object used for collection.	int	
MIB_INDEX	MIB index used for collection if applicable.	char	256

TABLE 545 SNMP_DATA_2HOUR

Field	Definition	Format	Size
ID	The DB ID of MIB_OBJECT.	int	
MIB_OBJECT_ID	The DB ID of MIB_OBJECT.	int	
TARGET_TYPE	Target or source type can be, • device - 0 or • interface or ports - 1	num	(2,0)
TARGET_ID	DB ID of the target which can be device or interface.	int	
VALUE	Aggregated value inserted during the aging process.	double precision	
TIME_IN_SECONDS	Time, in seconds, at which the record was inserted.	int	
COLLECTOR_ID	DB ID of the collector object used for collection.	int	
MIB_INDEX	MIB index used for collection if applicable.	char	256

TABLE 546 SNMP_DATA_30MIN

Field	Definition	Format	Size
ID	Primary key autogenerated ID	int	
MIB_OBJECT_ID	MIB OID used for collection	int	
TARGET_TYPE	Target or source type can be, • device - 0 or • interface or ports - 1	num	(2,0)
TARGET_ID	DB Id of the target which can be device or interface	int	
VALUE	Value collected by the engine	double precision	
TIME_IN_SECONDS	Time at which collection occured in seconds	int	
COLLECTOR_ID	DB Id of the collector object used for collection	int	
MIB_INDEX	MIB index used for collection if applicable	char	256

TABLE 547 SNMP_EXPR_DATA

Field	Definition	Format	Size
ID	Primary key column.	serial	
EXPRESSION_ID	MIB object ID.	int	
TARGET_TYPE	Target type of the SNMP collector data. Th target type for, • device level collector is 0, • for port level collector is 1.	smallint	
TARGET_ID	Target ID of the SNMP collector data. for device level collector it will use deviceId, for port level it will use interfaceId.	int	
VALUE	Value of the OID retrieved from the corresponding target.	double	
TIME_IN_SECONDS	Time when value of the OID was retrieved from the corresponding target.	int	
COLLECTOR_ID	Correspoding collector table ID.	int	

TABLE 548 SNMP_EXPR_DATA_1DAY

Field	Definition	Format	Size
ID	Primary key autogenerated ID.	int	
EXPRESSION_ID	DB ID of the expression object used for collection.	int	
TARGET_TYPE	Target or source type can be, • device - 0 or • interface or ports - 1	smallint	
TARGET_ID	DB Id of the target which can be device or interface.	int	
VALUE	Aggregated value inserted during the aging process.	double precision	
TIME_IN_SECONDS	Time, in seconds, at which the record was inserted in seconds.	int	
COLLECTOR_ID	DB ID of the collector object used for collection.	int	

TABLE 549 SNMP_EXPR_DATA_30MIN

Field	Definition	Format	Size
ID	Primary key autogenerated ID	int	
EXPRESSION_ID	DB ID of the expression object used for collection	int	
TARGET_TYPE	Target/Source type can be device:0 or interface/ports:1'	smallint	
TARGET_ID	DB Id of the target which can be device or interface	int	
VALUE	Value collected by the engine'	double precision	
TIME_IN_SECONDS	Time at which collection occured in seconds	int	
COLLECTOR_ID	DB Id of the collector object used for collection	int	

TABLE 550 SNMP_EXPRESSION

Field	Definition	Format	Size
EXPRESSION_ID	Primary key column.	serial	
NAME	Name of the expression.	varchar	64
DESCRIPTION	Description of the expression.	varchar	512
EQUATION	Equation of the expression.	varchar	1024
UNIT	Unit that is used for displaying the chart.	varchar	64
IS_TRANSIENT	Explicitly identified whether expressions is used for Real time collector or not. A transient expression will not be allowed for user editing.	numeric	(1,0)

TABLE 551 SNMP_PROFILE

Field	Definition	Format	Size
NAME*	A text string representing a set of SNMP agent profile. When created, one or more virtual switches could refer to this profile for its SNMP credentials unless a unique set of SNMP credentials has been defined in SNMP_CREDENTIAL.	varchar	256
PORT_NUMBER	Port number of the SNMP agent on the switch for get and set requests	smallint	
RETRY_COUNT	Number of times to retry if get/set request to the SNMP agent times out. Default value is 3.	smallint	
TIMEOUT	Timeout value in seconds before for a get/set request to the SNMP agent. Default value is 5.	smallint	
VERSION	SNMP agent version running on the switch as in SNMPv1 and SNMPv3	varchar	6
READ_COMMUNITY_STRING	The SNMP Read-Only Community String is like a password. It is sent along with each SNMP Get-Request and allows (or denies) access to device. The default value is "public". This is applicable if the agent is configured to operate in SNMPv1.	varchar	128
WRITE_COMMUNITY_STRING	The SNMP Write-Only Community String is like a password. It is sent along with each SNMP Set-Request and allows (or denies) access to a device. The default value is "private". This is applicable if the agent is configured to operate in SNMPv1	varchar	128
USER_NAME	A human-readable string representing the name of the user. This is applicable if the agent is configured to operate in SNMPv3.	varchar	64
CONTEXT_NAME	A text ID associated with the user, used by SNMP agent to provide different views. This is applicable if the agent is configured to operate in SNMPv3.	varchar	128
AUTH_PROTOCOL	An indication of whether or not messages sent or received on behalf of this user can be authenticated and if so, which authentication protocol to use. The supported values for this field are: usmNoAuthProtocol, usmHMACMD5AuthProtocol, and usmHMACSHAAuthProtocol. This is applicable if the agent is configured to operate in SNMPv3.	varchar	16
AUTH_PASSWORD	The localized secret key used by the authentication protocol for authenticating messages. This is applicable if the agent is configured to operate in SNMPv3.	varchar	64
PRIV_PROTOCOL	An indication of whether or not messages sent or received on behalf of this user can be encrypted and if so, which privacy protocol to use. The current values for this field are: usmNoPrivProtocol and usmDESPrivProtocol. This is applicable if the agent is configured to operate in SNMPv3.	varchar	16

TABLE 551 SNMP_PROFILE (Continued)

Field	Definition	Format	Size
PRIV_PASSWORD	The localized secret key used by the privacy protocol for encrypting and decrypting messages. This is applicable if the agent is configured to operate in SNMPv3.	varchar	64
SNMP_INFORMS_ENABLED	To denote whether SNMP informs option is enabled or disabled Default value is 0.	smallint	

TABLE 552 SNMP_TRAP_CREDENTIAL

Field	Definition	Format	Size
ID	PK for the table to uniquely identify the record	int	
VERSION	to identify the version of Credentials: v1v2c and v3 are the values	varchar	6
COMMUNITY_STRING	to decode the v1/v2c traps	varchar	64
USER_NAME	user access name for v3 trap	varchar	64
AUTH_PROTOCOL	authentication protocol used for v3 traps	varchar	16
AUTH_PASSWORD	authentication password for v3 traps	varchar	64
PRIV_PROTOCOL	privacy protocol used for v3 traps	varchar	16
PRIV_PASSWORD		varchar	64
POSITION_	order of credentials to authenticate v1/v2c or v3 traps	int	

TABLE 553 SNMP_V3_FORWARDING_CREDENTIAL

Field	Definition	Format	Size
ID*		int	
USER_NAME	USM user name.	varchar	64
CONTEXT_NAME	USM context name.	VARCHAR	128
AUTH_PROTOCOL	Authorization protocol.	VARCHA	16
AUTH_PASSWORD	Authorization password.	VARCHAR	64
PRIV_PROTOCOL	Privilege protocol.	VARCHAR	16
PRIV_PASSWORD	Privilege password.	VARCHAR	64

TABLE 554 SOURCE_OBJECT_TYPE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
TYPE_NAME	Type of the object to which the event applies, such as Fabric, Switch or Port.	char	64
DESCRIPTION	Description of the object	varchar	255

TABLE 555 SSL_CERTIFICATE

Field	Definition	Format	Size
SSL_CERTIFICATE_ID		int	
NAME		varchar	255
LOCATION		varchar	255
FILE_NAME		varchar	255
KEY_ID		int	
CERT_TYP		num	(2,0)
START_TIME		num	(20,0)
EXPIRATION_TIME		num	(20,0)
FORMAT		num	(2,0)
DESCRIPTION		varchar	1024
NOTIFICATION_TIME	The time stamp (long format) of the last expiration notification sent	num	(20,0)
NOTIFICATION_SENT	The status of last notification sent. Possible values: Unknown -0, Good 1, Expiring 2, Expired 3	num	(2,0)
NOTIFICATION_REPEAT	Indicates whether repeat expiration notification is enabled for this certificate or not. Possible values: Repeat Disabled - 0, Repeat Enabled - 1	num	(2,0)
SYNC_DEVICE	Indicates whether this certificate is in sync with device or not. Possible values: Need Deploy 0, Imported 1, Deployed 2, Unknown 3.	num	(2,0)
CERTIFICATE	The content of the ssl certificate.	txt	
USER_ID	This field will be populated when the Management application user creates certificate or import certificates from file. User can view this certificate not bound to any vip in SSL certificate dialog	int	

TABLE 556 SSL_CERTIFICATE_VIP_SERVER_MAP

Field	Definition	Format	Size
SSL_CERTIFICATE_ID	Foreign key to SSL_CERTIFICATE_ID in ssl_certificate table	int	
VIP_SERVER_ID	The column holds ID of VIP Server. It is Foreign Key and refers to ID column of VIP_SERVER table	int	

TABLE 557 SSL_KEY

Field	Definition	Format	Size
ssl_key_id		int	
name		varchar	255
location		varchar	255
file_name		varchar	255

TABLE 557 SSL_KEY (Continued)

Field	Definition	Format	Size
key_type		varchar	2
encryption_type		varchar	2
password		varchar	255
description		varchar	1024
strength	The strength of the private key in bits.	int	
private_key	Content of the private key.	txt	
USER_ID	This field will be populated when the Management application user creates certificate or import certificates from file. User can view this certificate not bound to any vip in SSL certificate dialog.	int	

TABLE 558 SSL_KEY_PASSWORD

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
KEY_PASSWORD_ALIAS	Key Password Alias is the alias name used for the encrypted key password. This alias name is used to identify the password in client UI.	varchar	16
KEY_PASSWORD	SSL keys are protected by passwords, and these passwords are used during key import operation from device. The key password is stored encrypted in the tables.	varchar	256

TABLE 559 STATS_AGING

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
FIVE_MIN_VALUE	Configured maximum samples value for the five minute table.	int	
THIRTY_MIN_VALUE	Configured maximum samples value for the thirty minute table.	int	
TWO_HR_VALUE	Configured maximum samples value for the two hour table.	int	
ONE_DAY_VALUE	Configured maximum samples value for the one day table.	int	
MAX_SAMPLES_VALUE	The maximum number of samples value, i.e., 3456.	int	
INTERPOLATE	Whether interpolation is enabled or disabled.	smallint	
POLICY_TYPE	The type of the aging ploicy. • 100 - Default aging (1 day 5 mins samples, 3 days 30 mins samples,7 days 2 hrs sample and 2 years 1 day samples) • 101 - 5 mins to 1 day aging(8 days 5 mins samples and 90 days of 1 day samples)	smallint	
ACTIVE	The active state of the policy.	smallint	

TABLE 560 STP_PORT

Field	Definition	Format	Size
STP_PORT_ID	Prmiary key Identifier.	int	
STP_INSTANCE_ID	Foreign Key Reference to STP_INSTANCE table.	int	
STP	If MSTP is enabled, the value will be 1 else 0.	numeric	(1,0)
INTERFACE_ID	Foreign Key Reference to INTERFACE table	int	
PATH_COST	Port Path Cost.	bigint	
PRIORITY	Port Priority.	bigint	
LINK_TYPE	Link Type. 1- Shared 2 - P2P.	numeric	(1,0)
PORT_FAST	Port Fast. 0 - Disabled 1 - Enabled	numeric	(1,0)
BPDU_FILTER	BPDU Filter. 0 - Disabled 1 - Enabled	numeric	(1,0)
BPDU_GUARD	BPDU guard. 0 - Disabled 1 - Enabled	numeric	(1,0)
EDGE_PORT	Edge port. 0 - Disabled 1 - Enabled	numeric	(1,0)
AUTO_EDGE	Auto edge. 0 - Disabled 1 - Enabled	numeric	(1,0)
ROOT_GUARD	Root guard. 0 - Disabled 1 - Enabled	numeric	(1,0)
HELLO_TIME	Number of seconds between generation of config BPDUs on CIST.	smallint	
VLAN_ID	The PVST and RPVST values needs to mapped with the VLAN.	int	

TABLE 561 STP_INSTANCE

Field	Definition	Format	Size
STP_INSTANCE_ID		int	
INSTANCE_TYPE		num	(2,0)
INSTANCE_ID		num	(4,0)
DEVICE_ID		int	
STP_MODE		num	(2,0)

TABLE 561 STP_INSTANCE (Continued)

Field	Definition	Format	Size
FORWARD_DELAY		num	(2,0)
MAX_AGE		num	(2,0)
HELLO_TIME		num	(2,0)
PRIORITY		num	(6,0)
STP_VERSION		num	(2,0)
RE_ENABLE_PORT_INTERVAL	FOS/NOS Field. Re enable port interval.	int	
RE_ENABLE_PORT_STATE	FOS/NOS Field. Re enable port state.	smallint	
PATH_COST		bigint	
STP	Possible values: • 0 - Disabled • 1 - Enabled	smallint	
CISCO_INTER_OP	Cisco Interoperability Enabled/Disabled.	num	(1,0)
TX_HOLD_COUNT	Transmit HoldCount of the Bridge	smallint	
MAX_HOPS	MST max hop count (1-40)	smallint	
REGION	MST Region.	varchar	255
REVISION	Revision Number for Configuration information.	smallint	

TABLE 562 SWITCH_BOTTLENECK_CONFIG

Field	Definition	Format	Size
VIRTUAL_SWITCH_ID	The database ID of the switch that the configuration belongs to	int	
BOTTLENECK_DETECT_ENABLED	Flag indicates if bottleneck detection is enabled or not	smallint	
ALERTS_ENABLED	Flag indicates if bottleneck detection alerts is enabled or not	smallint	
CONGESTION_THRESHOLD	Value of bottleneck detection congestion threshold in percent	double precision	
LATENCY_THRESHOLD	Value of bottleneck detection latency threshold in percent	double precision	
WINDOW_	Value of bottleneck detection latency window in millisecond	int	
QUIET_TIME	Value of bottleneck detection quiet time in millisecond	int	
CREATION_TIME	Creation time of the record	timestamp	
LAST_UPDATE_TIME	Last update time of the record	timestamp	

TABLE 562 SWITCH_BOTTLENECK_CONFIG (Continued)

Field	Definition	Format	Size
LATENCY_SEVERITY	The factor by which throughput must drop in a second in order for that second to be considered affected by latency bottlenecking. Range (1 to 1000).	int	
LATENCY_TIME	The minimum fraction of a second that must be affected by latency in order for that second to be considered affected by latency bottlenecking. Range (0 to 1).	double precision	

TABLE 563 SWITCH_CONFIG

Field	Definition	Format	Size
ID*		int	
NAME	Name of the switch configurations uploaded from the switch either on demand or through scheduler	varchar	64
SWITCH_ID	ID of the switch from which the configuration has been uploaded.	int	
CORE_SWITCH_ID		int	
BACKUP_DATE_TIME	The date/time stamp at which the configuration has been uploaded.	timestamp	
CONFIG_DATA	The actual switch configuration data.	text	
CEE_CONFIG_DATA	Switch configuration data for CEE	text	
KEEP_COPY	The column value (1) helps to preserve the configuration even after the expiration of its age.	smallint	
CREATED_BY	The column value helps to figure out who triggered the configuration upload operation.	varchar	64
CONFIG_TYPE	Configuration Type • FC=0 • CEE_RUNNING=1 • CEE_STARTUP=2 • INVALID=-1 Default value is 0.	smallint	
COMMENTS	Brief comments about this configuration.	varchar	256

TABLE 564 SWITCH_CONFIG_DETAIL

Field	Definition	Format	Size
SWITCH_CONFIG_ID		int	
IP_ADDRESS	IP Address of the switch for which the configuration was uploaded either on demand or schedule.	varchar	128
WWN	WWN of the switch for which the configuration was uploaded either on demand or schedule.	char	23
PHYSICAL_SWITCH_WWN	CORE WWN of the switch for which the configuration was uploaded either on demand or schedule.	char	23
MODEL_NUMBER	Model Number of the switch for which the configuration was uploaded either on demand or schedule.	varchar	32

TABLE 564 SWITCH_CONFIG_DETAIL (Continued)

Field	Definition	Format	Size
IS_BASELINE	Indicates if this is a baseline configuration or not. 0 - NOT A BASELINE 1 - BASELINE	smallin	
BACKUP_TYPE	The operation based on which this configuration was retrieved -1 - NOT AVAILABLE 0 - IMPORTED 1 - DISCOVERY 2 - RESYNC 3 - MANUAL 4 - SCHEDULE	int	
DRIFT_STATUS	Indicates if the current switch configuration has deviated from the baseline configuration. -1 - NO_BASELINE 0 - NO_DEVIATION 1 - DEVIATED	int	

TABLE 565 SPX_PORT_DETAILS

Field	Definition	Format	Size
ID*	Primary key and it's serial number.	int	
SPX_PORT	Connected spx_port of both CB and PE units.	varchar	128
PE_GROUP_NAME	PE Group Name configured by user.	varchar	255
UNIT_NUMBER	The spx_port belonging unit. (Both CB and PE units).	int	
DEVICE_ID	The foreign reference of Device table.	int	
INTERFACE_ID	The foreign key reference of Interface table.	int	
CONNECTED_PE_UNIT_NUMBER	The Connected PE Units to the CB SPX port.	int	

TABLE 566 SWITCH_LICENSE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
CORE_SWITCH_ID	Refers to the entry in the CORE_SWITCH table.	int	
LICENSE_KEY	Stores the license key obtained from the switch.	varchar	256

TABLE 567 SWITCH_MODEL

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
SWBD_TYPE	Switch type number, universally used by all the Management application module implementation.	smallint	

TABLE 567 SWITCH_MODEL (Continued)

Field	Definition	Format	Size
SUBTYPE	Switch subtype. At present no subtypes for existing model records are defined. Default value is 0.	smallint	
DESCRIPTION	Model description, such as FC link speed, port count and whether multi-card (director) class switch or other type of switch. Default is 'Not Available'.	varchar	256
MODEL	Switch model string.	varchar	32
REMARK	Remarks, such as an internal project name.	varchar	64
SYS_OID	This will represent the sys_oid for each product type.	varchar	255
PRODUCT_FAMILY	This represents the product family that each OID belongs to.	varchar	128
BRIEF_PRODUCT_FAMILY	Shorter name for the product family.	varchar	32
SPEED	Switch max speed. Value 0 represents NotAvailable.	smallint	
MULTI_CP_CAPABLE	Switch is multi cp cabable or not. 0 means single CP and 1 means multi.	smallint	
MIN_IMAGE_VERSION	Supported min firmware version.	varchar	64
MAX_IMAGE_VERSION	Supported max firmware version.	varchar	64

TABLE 568 SWITCH_PORT

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
VIRTUAL_SWITCH_ID	DB ID of virtual_switch to which this port belongs.	int	
WWN	WWN of the port.	char	23
NAME	User friendly name of the port.	char	32
SLOT_NUMBER	Slot number. Default value is 0.	int	
PORT_NUMBER	The logical port number of the user port. There is no assumption of any relation to the physical location of a port within a chassis.	smallint	
USER_PORT_NUMBER	User port number. Unique port number in a chassis.	smallint	
PORT_ID	Port ID of this port.	varchar	8
PORT_INDEX	Number used for identifying port in zoning.	smallint	
AREA_ID	Area number the port is assigned to.	smallint	
MAC_ADDRESS	MAC address of this port.	varchar	64
PORT_MOD	Stores the port module type. Not applicable if port doesn't have a GBIC installed.	varchar	64
TYPE	Port type. The specific mode currently enabled for the port. The port type could be U-Port, F-Port, E-Port etc.	varchar	16
FULL_TYPE	Refers to the full type of the port, U-Port, F-Port etc.	varchar	128

TABLE 568 SWITCH_PORT (Continued)

Field	Definition	Format	Size
STATUS	Refers to the Status of the port. Eg. No Light, No Module, Mod_inv, Online etc.	varchar	64
HEALTH	Refers to the Health of the port. Eg. Unmonitored, Healthy, Offline , Error etc.	varchar	16
STATUS_MESSAGE	Any additional port level status similar to what is seen in CLI, like Segmented, Speed Mismatch, Trunk master etc are stored here.	varchar	255
PHYSICAL_PORT	Indicates if the port is physical port. It will be 0 for virtual ports Eg. Ports created by LISLs.	smallint	
LOCKED_PORT_TYPE	Indicates the locked port type of the port. Ports can be locked down so that they can come up only in that mode.	varchar	16
CATEGORY	Denotes the category of the switch. 1 denotes FC port and 2 denotes gige port.	smallint	
PROTOCOL	The protocol used by the port. FC, FCIP etc.	varchar	16
SPEED	Actual speed at which the port is currently operating.	varchar	64
SPEEDS_SUPPORTED	The supported port speed as a list of comma separated values.	varchar	32
MAX_PORT_SPEED	The maximum speed the port is capable of supporting, in bits per second.	int	
DESIRED_CREDITS	How many BB credits are desired for the port.	int	
BUFFER_ALLOCATED	How many BB credits are allocated for the port.	int	
ESTIMATED_DISTANCE	The estimated physical distance of the connection between ports.	int	
ACTUAL_DISTANCE	The physical distance of the connection on the port in relation to the other port.	int	
LONG_DISTANCE_SETTING	Whether long distance enabled.	int	
DEGRADED_PORT	Denotes if the port is in a degraded state. Has value as N/A for ports that are not online.	varchar	16
REMOTE_NODE_WWN	Node WWN of the attached port.	varchar	255
REMOTE_PORT_WWN	WWN of the attached port.	varchar	255
LICENSED	1 = the port is licensed; otherwise, 0.	smallint	
SWAPPED	1 = port is swapped; otherwise, 0.	smallint	
TRUNKED	1 = port is trunked; otherwise, 0.	smallint	
TRUNK_MASTER	1 = the port is trunk master; otherwise, 0.	smallint	
PERSISTENT_DISABLE	1 = port is persistently disabled.	smallint	
FICON_SUPPORTED	1 = FICON is supported; otherwise, 0.	smallint	
BLOCKED	1 = port is blocked; otherwise, 0.	smallint	
PROHIBIT_PORT_NUMBERS	Indicates the ports prohibited with the current port as configured in the allow prohibit matric (PDCM).	varchar	1024
PROHIBIT_PORT_COUNT	The count of prohibited ports.	smallint	

TABLE 568 SWITCH_PORT (Continued)

Field	Definition	Format	Size
NPIV	Whether NPIV mode is enabled.	smallint	
NPIV_CAPABLE	Instance NPIV mode capability: 1 = indicates port has NPIV capability 2 = NPIV license is enabled	smallint	
NPIV_ENABLED	Whether NPIV mode is enabled.	smallint	
FC_FAST_WRITE_ENABLED	1 = FC fast write is enabled.	smallint	
ISL_RRDY_ENABLED	Denotes if ISL receiver ready is enabled.	smallint	
RATE_LIMIT_CAPABLE	Denotes if the port is capable of Rate Limiting.	smallint	
RATE_LIMITED	Denotes if the port has Rate Limiting Enabled.	smallint	
QOS_CAPABLE	Indicates if the port is QOS capable.	smallint	
QOS_ENABLED	Indicates if the port is QOS enabled.	smallint	
TUNNEL_CONFIGURED	Denotes if the port has a fcip tunnel configured.	smallint	
FCIP_TUNNEL_UP	Denotes if the fcip tunnel that is configured is up.	smallint	
FCR_FABRIC_ID	Stores the FCR fabric ID. Applicable if the port is configured as an EX port.	smallint	
FCR_INTEROP_MODE	The interop mode of the FCR fabric. Applicable if the port is an EX port.	smallint	
CALCULATED_STATUS	The calculated status of the port. Eg. Healthy, Down, Marginal etc.	varchar	64
USER_DEFINED_VALUE1	User defined value used for annotation.	varchar	256
USER_DEFINED_VALUE2	User defined value used for annotation.	varchar	256
USER_DEFINED_VALUE3	User defined value used for annotation.	varchar	256
KIND	Stores the port kind from the NVP portKind.	varchar	32
STATE	The state of the port whether it is online or offline	varchar	64
PREVIOUS_STATUS	This table can hold the same values as STATUS column. But this will be holding the previous status of the PORT. These values to be populated by switch asset collector.	varchar	64
AUTO_DISABLE_CONFIGURED	To represent auto disable configuration state (set by user). Default value is 0.	smallint	
AUTO_DISABLED	To represent auto disabled status (set by switch). Default value is 0.	smallint	
OCCUPIED	Default value is 0.	smallint	
LAST_UPDATE	Last update time stored as long value. Elapsed time from 1970 in milliseconds.	bigint	
PORT_BIT_MASK	F-Port trunk bit mask value. Default value is 0.	int	
LOGICAL_PORT_NUMBER	F-Port trunk logical port number. Default value is -1.	smallint	

TABLE 568 SWITCH_PORT (Continued)

Field	Definition	Format	Size
DEFAULT_AREA_ID	Default Area id of F-Port trunk port. Default value is -1.	smallint	
LOGICAL_PORT_WWN	Logical port WWN of F-Port trunk group.	char	23
PREVIOUS_TYPE	This fields copies the old state of the port type. The field could be used to track the state change information for the switch port type. SwitchAssetCollector sets this field during the collection time. SMIA requested this information but could be used by any module which needs to track the type state change.	varchar	16
LATENCY_DETECT_SUPPORTED	Whether the port supports latency detection. 1 means true and 0 means false	smallint	
PREVIOUS_STATE	Fields copies the old state of the port . The field could be used to track the state change information for the switch port . SwitchAssetCollector sets this field during the collection time. SMIA requested this information but could be used by any module which needs to track the state change.	varchar	64
EPORT_DISABLED	Represents the eportDisabled field from switch.html. Values populated by SwitchAssetcollector during the collection time. Possible values includes 0 and 1. Default value is -1.	smallint	
SPEED_NEGOTIATED	This column indicates if the port speed is negotiated or not. If port speed is negotiated then value is 1 else it will be 0. Default value is -1.	smallint	
MAX_FRAME_MONITOR	Maximum frame monitor supported for switch port.	int	
MAX_FRAME_MONITOR_OFFSET	Maximum offset supported in fame monitor for switch port.	int	
	Contains the features supported as a bit mask at port level.	int	
IDENTIFIER	Switch port identifier extracted from interface name	char	80
PORT_CAPABILITIES	'List of capabilities of this port specified as bit mask. Each bit would represent capability like FEC, Encryption and compression, NPIV etc.';	int	
XISL_PORT_LIST	This field is applicable only for logical ports created for LISLs. It denotes the list of XISL ports associated with the current logical port. Will be blank for non-logical ports.	varchar	256

TABLE 568 SWITCH_PORT (Continued)

Field	Definition	Format	Size
PORT_COMMISSION_STATE	Indicates whether port decommission/recommission was in progress or completed, based on this status we will show the decommission/recommission icon on ports and Indicates the Decommissioned/Recommissioned status of the ports which was performed from the Management application. None - 0, Decommission In Progress - 1 , Decommissioned - 2, Recommission In Progress - 3, Recommissioned - 4. If the decommission is performed through CLI or other Management application server then the state would be None (0).	int	
FEATURES_ENABLED	Holds as a bit mask the features that are enabled . Refer FEATURES_ACTIVE for the active/inactive status of a feature. Each bit would represent features like Encryption, compression etc.' The bit mask and their corresponding Features are defined as an enum in the domain model class - SwitchPort.java.	int	
FEATURES_ACTIVE	Holds as a bit mask the features that are active. Please note that this is different from the enabled value which is found in the FEATURES_ENABLED column. Each bit would represent features like Encryption, compression etc. The bit mask and their corresponding Features are defined as an enum in the domain model class - SwitchPort.java.	int	
DISABLED_REASON	The Switch Port disabled reason.	varchar	1024
FENCED	1 means port is fenced.	smallint	
MASTER_PORT_NUMBER	This column will have the trunk master port number for the trunk members. For trunk master, it will have its own port number. For non-trunk ports, it will have the default value -1.	int	
SPEED_TYPE	Stores the speed type of the port. It contains one of the following values: • 1 - Indicates speed is in Mbps. • 2 - Indicates speed is in Gbps.	int	
EXT_TYPE	Refers to the extended type of the port . Eg Mirror-Port.	varchar	128

TABLE 569 SWITCH_PORT_PERFORMANCE

Field	Definition	Format	Size
PORT_ID	Primary key of the Switch Port.	int	
SWITCH_ID	Primary key of Virtual Switch which this port is present	int	
TX	The number of octets or bytes that have been transmitted by this port. One second periodic polling of the port. This value is saved and compared with the next polled value to compute net throughput. Note, for Fibre Channel, ordered sets are not included in the count	double precision	

TABLE 569 SWITCH_PORT_PERFORMANCE (Continued)

Field	Definition	Format	Size
RX	The number of octets or bytes that have been received by this port. One second periodic polling of the port. This value is saved and compared with the next polled value to compute net throughput. Note, for Fibre Channel, ordered sets are not included in the count	double precision	
TX_UTILIZATION	The computed value of TX based on speed of port	double precision	
RX_UTILIZATION	'The computed value of RX based on speed of port	double precision	
LINK_FAILURE	Count of link failures. This count is part of the Link Error Status Block (LESB). (FC-PH 29.8). Note, this is a Fibre Channel only stat	double precision	
TX_LINK_RESETS	Count of Link resets. This is the number of LRs received. Note, this is a Fibre Channel only stat	double precision	
RX_LINK_RESETS	Count of Link resets. This is the number LRs transmitted. Note, this is a Fibre Channel only stat	double precision	
SYNC_LOSSES	Count of instances of synchronization loss detected at port. This count is part of the Link Error Status Block (LESB). (FC-PH 29.8). Note, this is a Fibre Channel only stat.	double precision	
SIGNAL_LOSSES	Count of instances of signal loss detected at port. This count is part of the Link Error Status Block (LESB). (FC-PH 29.8). Note, this is a Fibre Channel only stat	double precision	
SEQUENCE_ERRORS	Count of primitive sequence protocol errors detected at this port. This count is part of the Link Error Status Block (LESB). (FC-PH 29.8). Note, this is a Fibre Channel only stat	double precision	
INVALID_TRANSMISSIONS	Count of invalid transmission words received at this port. This count is part of the Link Error Status Block (LESB). FC-PH 29.8). Note, this is a Fibre Channel only stat	double precision	
CRC_ERRORS	Count of frames received with invalid CRC. This count is part of the Link Error Status Block (LESB). (FC-PH 29.8). Loop ports should not count CRC errors passing through when monitoring. Note, this is a Fibre Channel only stat.'	double precision	
LAST_UPDATE_TIME	Time when this stats record was updated	timestamp	
FAKE_PORT	Denotes whether the port is a fake port created by the Management application for drawing connections or a real one obtained during collection from the switch. 1 denotes a fake port and 0 denotes a real port.	smallint	
SPEED_TYPE	Stores the speed type of the port. It will contain one of the following values: 1 - indicates speed is in Mbps 2 - indicates speed is in Gbps	int	

TABLE 570 SWITCH_THRESHOLD_SETTING

Field	Definition	Format	Size
SWITCH_ID*	References the ID in CORE_SWITCH table.	int	
POLICY_ID*	References the ID in THRESHOLD_POLICY table.	int	
STATUS	The status of applied to the switch.	smallint	
OVERRIDDEN	Policy is overridden or not overridden.	smallint	
DESCRIPTION	Description about the status of policy applied to the switch.	varchar	100

TABLE 571 SYSTEM_CARD_ENGINE_MAPPING

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
ENCRYPTION_ENGINE_ID	Foreign key reference to the ENCRYPTION_ENGINE for which a system card is registered	int	
SMART_CARD_ID	Foreign key reference to the SMART_CARD that is registered as a system card for the encryption engine.	int	

TABLE 572 SYSTEM_PROPERTY

Field	Definition	Format	Size
NAME*	The name of the property.	char	64
VALUE	The value for the property.	varchar	2048

TABLE 573 TARGET_TYPE

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
TYPE	Type of the target device. Some possible values are • Switch • Device • Port • Host • Port Group • Product Group • VLAN • Fabric	varchar	64

TABLE 574 THIRD_PARTY_DEVICE

Field	Definition	Format	Size
DEVICE_ID	Primary key for this table.	int	
DEVICE_TYPE	Type of the third party device. As of now, we have two types Wireless Location Manager and LANcope device.	varchar	64

TABLE 575 THRESHOLD_MEASURE

Field	Definition	Format	Size
MEASURE_ID*	References the ID In PM_MEASURE table, where all measures are defined.	int	
HIGH_BOUNDARY	Configured high boundary threshold value for measure ID.	int	
LOW_BOUNDARY	Configured low boundary threshold value for measure ID.	int	
BUFFER_SIZE	Configured buffer size for measure ID.	int	
POLICY_ID*	References the ID in THRESHOLD_POLICY table.	int	

TABLE 576 THRESHOLD_POLICY

Field	Definition	Format	Size
ID*		int	
NAME	Name of the policy.	varchar	100
TYPE	Type of the policy.	varchar	20
DESCRIPTION	Description about the policy.	varchar	100

TABLE 577 TIME_SERIES_DATA

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For device level collector the target type is 0, for port level it is 1.	smallint	
MEASURE_ID	ID of the measure.	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId, for port level it will use interfaceId.	int	
COLLECTOR_ID	ID of the data_collector.	int	
MEASURE_INDEX	'Stores the index_map value in case of anexpression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	30 mins aggregated data.	double precision	

TABLE 578 TIME_SERIES_DATA_1DAY

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when the record is inserted.	int	
TARGET_TYPE	Target type of the PM collector data. For device level collector the target type is 0, for port level it is 1.	smallint	
MEASURE_ID	ID of the measure.	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId, for port level it will use interfaceId.	int	
COLLECTOR_ID	ID of the data_collector.	int	

TABLE 578 TIME_SERIES_DATA_1DAY (Continued)

Field	Definition	Format	Size
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	30 mins aggregated data.	double precision	

TABLE 579 TIME_SERIES_DATA_2HOUR

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when the record is inserted.	int	
TARGET_TYPE	Target type of the PM collector data. For device level collector the target type is 0, for port level it is 1.	smallint	
MEASURE_ID	ID of the measure.	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId, for port level it will use interfaceId.	int	
COLLECTOR_ID	ID of the data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	30 mins aggregated data.	double precision	

TABLE 580 TIME_SERIES_DATA_30MIN

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when the record is inserted.	int	
TARGET_TYPE	Target type of the PM collector data. For device level collector the target type is 0, for port level it is 1.	smallint	
MEASURE_ID	ID of the measure.	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId, for port level it will use interfaceId.	int	
COLLECTOR_ID	ID of the data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	30 mins aggregated data.	double precision	

TABLE 581 TIME_SERIES_DATA_1

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables. Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 582 TIME_SERIES_DATA_1_1DAY

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores One day aggregated data.	double precision	

TABLE 582 TIME_SERIES_DATA_1_1DAY (Continued)

Field	Definition	Format	Size
MIN_VALUE	Minimum value in 2 hour table while aggregating 1 day data.	double precision	
MAX_VALUE	Maximum value in 2 hour table while aggregating 1 day data.	double precision	

TABLE 583 TIME_SERIES_DATA_1_2HOUR

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the 2 hours aggregated data.	double precision	
MIN_VALUE	Minimum value in 30 min table while aggregating 2 hours of data.	double precision	
MAX_VALUE	Maximum value in 30 min table while aggregating 2 hours of data.	double precision	

TABLE 584 TIME_SERIES_DATA_1_30MIN

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	

TABLE 584 TIME_SERIES_DATA_1_30MIN (Continued)

Field	Definition	Format	Size
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the 30 minutes aggregated data.	double precision	
MIN_VALUE	Minimum value in raw performance statistics table while aggregating 30 minutes of data.	double precision	
MAX_VALUE	Maximum value in raw performance statistics table while aggregating 30 minutes of data.	double precision	

TABLE 585 TIME_SERIES_DATA_2

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables. Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 586 TIME_SERIES_DATA_2_1DAY

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores One day aggregated data.	double precision	
MIN_VALUE	Minimum value in 2 hour table while aggregating 1 day data.	double precision	
MAX_VALUE	Maximum value in 2 hour table while aggregating 1 day data.	double precision	

TABLE 587 TIME_SERIES_DATA_2_2HOUR

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the 2 hours aggregated data.	double precision	

TABLE 587 TIME_SERIES_DATA_2_2HOUR (Continued)

Field	Definition	Format	Size
MIN_VALUE	Minimum value in 30 min table while aggregating 2 hours of data.	double precision	
MAX_VALUE	Maximum value in 30 min table while aggregating 2 hours of data.	double precision	

TABLE 588 TIME_SERIES_DATA_2_30MIN

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the 30 minutes aggregated data.	double precision	
MIN_VALUE	Minimum value in raw performance statistics table while aggregating 30 minutes of data.	double precision	
MAX_VALUE	Maximum value in raw performance statistics table while aggregating 30 minutes of data.	double precision	

TABLE 589 TIME_SERIES_DATA_3

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data.	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	

TABLE 589 TIME_SERIES_DATA_3 (Continued)

Field	Definition	Format	Size
VALUE	Stores the raw data received from the device.	double precision	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables.Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 590 TIME_SERIES_DATA_3_1DAY

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data.	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores One day aggregated data.	double precision	
MIN_VALUE	Minimum value in 2 hour table while aggregating 1 day data.	double precision	
MAX_VALUE	Maximum value in 2 hour table while aggregating 1 day data.	double precision	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables.Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 591 TIME_SERIES_DATA_3_2HOUR

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data.	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256

TABLE 591 TIME_SERIES_DATA_3_2HOUR (Continued)

Field	Definition	Format	Size
ME_ID	ME_ID of the target.	int	
VALUE	Stores the 2 hours aggregated data.	double precision	
MIN_VALUE	Minimum value in 30 min table while aggregating 2 hours of data.	double precision	
MAX_VALUE	Maximum value in 30 min table while aggregating 2 hours of data.	double precision	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables.Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 592 TIME_SERIES_DATA_3_30MIN

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data.	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the 30 minutes aggregated data.	double precision	
MIN_VALUE	Minimum value in raw performance statistics table while aggregating 30 minutes of data.	double precision	
MAX_VALUE	Maximum value in raw performance statistics table while aggregating 30 minutes of data.	double precision	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables.Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 593 TIME_SERIES_DATA_4

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data.	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	

TABLE 593 TIME_SERIES_DATA_4 (Continued)

Field	Definition	Format	Size
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the raw data received from the device.	double precision	
MIN_VALUE	Stores the minimum value received from the device.	double precision,	
MAX_VALUE	Stores the maximum value received from the device.	double precision	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables.Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 594 TIME_SERIES_DATA_4_1DAY

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data.	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores One day aggregated data.	double precision	
MIN_VALUE	Minimum value in 2 hour table while aggregating 1 day data.	double precision	
MAX_VALUE	Maximum value in 2 hour table while aggregating 1 day data.	double precision	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables.Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 595 TIME_SERIES_DATA_4_2HOUR

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data.	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the 2 hours aggregated data.	double precision	
MIN_VALUE	Minimum value in 30 min table while aggregating 2 hours of data.	double precision	
MAX_VALUE	Maximum value in 30 min table while aggregating 2 hours of data.	double precision	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables.Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 596 TIME_SERIES_DATA_4_30MIN

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data.	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the 30 minutes aggregated data.	double precision	
MIN_VALUE	Minimum value in raw performance statistics table while aggregating 30 minutes of data.	double precision	

TABLE 596 TIME_SERIES_DATA_4_30MIN (Continued)

Field	Definition	Format	Size
MAX_VALUE	Maximum value in raw performance statistics table while aggregating 30 minutes of data.	double precision	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables.Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TABLE 597 TIME_SERIES_DATA_5

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the raw data received from the device.	double precision	
SUM_VALUE	Named after SUM_VALUE to be consistent with column names in aggregated data tables. Stores the delta changes for counter values between two samples, only used for counter values, 0 for all other types of measures.	double precision	

TIME_SERIES_DATA_5_30MIN

Field	Definition	Format	Size
TIME_IN_SECONDS	Time when value of the measure retrieved from the corresponding target.	int	
TARGET_TYPE	Target type of the PM collector data. For IP_DEVICE(0), IP_PORT(1), IP_TRUNK(2), FOS_DEVICE(3), FC_PORT(4), GE_PORT(5), TE_PORT(6), HBA_PORT(7), CNA_PORT(8), VIRTUAL_FCOE_PORT(9), FCIP_TUNNEL(10), EE_MONITOR(11), IP_DEVICE_GROUP(12), IP_PORT_GROUP(13), VIRTUAL_GROUP(14), TRILL_TRUNK(15), ALL_SAN_PRODUCTS(16).	smallint	
MEASURE_ID	ID of the measure (MIB/Expression).	int	

TIME_SERIES_DATA_5_30MIN (Continued)

Field	Definition	Format	Size
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId/virtualSwitchId, for port level it will use interfaceId/switchPortId/ fcipTunnelId/devicePortId.	int	
COLLECTOR_ID	DB ID of the pm_data_collector.	int	
MEASURE_INDEX	Stores the index_map value in case of an expression.	varchar	256
ME_ID	ME_ID of the target.	int	
VALUE	Stores the 30 minutes aggregated data.	double precision	
MIN_VALUE	Minimum value in raw performance statistics table while aggregating 30 minutes of data.		
MAX_VALUE	Maximum value in raw performance statistics table while aggregating 30 minutes of data.		
SUM_VALUE	Summation of SUM_VALUE in raw stats table while aggregating to 30 minutes data.	double precision	

TABLE 598 TOOL_APP

Field	Definition	Format	Size
TOOL_MENU_TEXT*	Text to be displayed for the Tool Menu.	varchar	256
TOOL_ID	A Tool in the TOOL_PATH table where the tools are defined.	int	
PARAMETERS	Default path for launching the tool.	varchar	256
KEY_STROKE	Short cut key stroke to the application.	varchar	30

TABLE 599 TOOL_PATH

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
TOOL_NAME	Name of the tool.	varchar	256
PATH	Path of the tool where installed or available.	varchar	1057
WORKING_FOLDER	Working folder for that application.	varchar	512

TABLE 600 TOPO_MAP_IMAGE

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
NAME	Image name in the foo.png format	varchar	256
IMAGE_OBJECT	'Image Object BLOB	bytea	

TABLE 601 TRILL

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
CLUSTER_ME_ID	The Management Element ID of the VCS Cluster in the VirtualSwitch	int	

TABLE 601 TRILL (Continued)

Field	Definition	Format	Size
SOURCE_ME_ID	The Management Element ID of the source VirtualSwitch.	int	
SOURCE_DOMAIN_ID	The source vcs member id	int	
SOURCE_PORT	The source port number as retrieved from the switch.	int	
SOURCE_PORT_NUMBER	The source port represented as a tuple of member/slot/port	char	30
DEST_ME_ID	The Management Element ID of the destination VirtualSwitch.	int	
DEST_DOMAIN_ID	The destination vcs member id	int	
DEST_PORT	The dest port number as retrieved from the switch.	int	
DEST_PORT_NUMBER	The source port represented as a tuple of member/slot/port	char	30
COST	Cost for the given trill link	int	
TYPE	Type of the given trill link	smallint	
TRUSTED	Is this trill link trusted	smallint	
TRUNKED	Is this trill link part of a trunk	smallint	
CREATION_TIME	Time when the TRILL link record is created between source and destination.	timestamp	
MISSING	Is this trill link was discovered and is now missing	smallint	
MISSING_TIME	Time when the TRILL link is missing from the switch.	timestamp	
SOURCE_PORT_NAME	Switch port name for the source	char	30
DEST_PORT_NAME	Switch port name for the destination	char	30

TABLE 602 TRILL_TRUNK_GROUP

Field	Definition	Format	Size
ID	Primary key for this table. Serial number which is uniquely generated by DB.	int	
ME_ID	The Management Element ID of the VCS member in the VirtualSwitch	int	
MASTER_PORT_NUMBER	The master port represented as a tuple of member/slot/port	varchar	30

TABLE 603 TRILL_TRUNK_MEMBER

Field	Definition	Format	Size
GROUP_ID		int	
PORT_NUMBER	The source port represented as a tuple of member/slot/port	varchar	30

TABLE 604 TRUNK_GROUP_INTERFACE

Field	Definition	Format	Size
INTERFACE_ID		int	
VLAG	Specifies whether the lag is a vlag or not	smallint	

TABLE 605 TRUNK_GROUP_MEMBER

Field	Definition	Format	Size
TRUNK_GROUP_MEMBER_ID	Primary key for this table.	int	
INTERFACE_ID	Foreign key which refers INTERACE table.	int	
TRUNK_INTERFACE_ID	Foreign key which refers TRUNK_GROUP_INTERACE table.	int	
LAG_NAME	Lag name of the trunk.	varchar	64

TABLE 606 TYPE_AHEAD_DEFINITION

Field	Definition	Format	Size
KEY	Unique key to identify specific Query definition for type ahead text field. Module should define the unique key which would be used by framework to execute the defined query.	varchar	256
QUERY_STRING	Actual Query string to be executed to get the data for the type ahead text field. All the query string should have dynamic placeholder {SEARCH_TEXT} and it can also have other dynamic parameters which can be changed dynamically in the definition. Framework would replace these parameters with actual text strings provided by the user or module.	text	

TABLE 607 USER_

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
NAME	User name.	varchar	128
DESCRIPTION	User description.	varchar	512
PASSWORD	User password.	varchar	512
EMAIL	User e-mail ID.	varchar	1024
NOTIFICATION_ENABLED	Flag for e-mail notification. Default value is 0.	smallint	
FULL_NAME	User''s Full Name.	varchar	512
PHONE_NUMBER	User''s Phone number.	varchar	32
INVALID_LOGIN_COUNT	This is a counter filed to identify the number of invalid login attempts. **NOTE:** After successful login this filed will be set to NULL. Default value is 0.	smallint	
LOCKED_OUT_DATETIME	The date time stamp when a user got locked out because of exceeding max number of invalid login attempts.	timestamp	
STATUS	User's account status: • 0=Disabled • 1=Enabled Default value is 1.	smallint	

TABLE 607 USER_ (Continued)

Field	Definition	Format	Size
SOURCE_OF_CREATION	To identify the source for creating the user account. • 0= User created through Management application Client • 1= User created when authenticated through external server. **NOTE:** At present there is no direct use of this field however this can be referred in future to build certain reports. Default value is 0.	smallint	
IP_PRODUCT_LOGIN_NAME	User CLI credential login user name.	varchar	256
IP_PRODUCT_LOGIN_PASSWORD	User CLI credential login password.	varchar	768
IP_PRODUCT_ENABLE_USER_NAME	User CLI credential enable user name.	varchar	256
IP_PRODUCT_ENABLE_PASSWORD	User CLI credential enable password.	varchar	768

TABLE 608 USER_AOR_MAP

Field	Definition	Format	Size
USER_NAME		varchar	128
AOR_ID	AOR ID where user has membership.	smallint	

TABLE 609 USER_DEFINED_DEVICE_DETAIL

Field	Definition	Format	Size
WWN	WWN of the device.	char	23
NAME	'Name of the device which is updated by the user.	varchar	256
TYPE	Type of the device (Initiator or Target.	varchar	32
IP_ADDRESS	IP address of the device which is updated by the user.	varchar	63
CONTACT	Contact detail of the device which is updated by the user.	varchar	256
LOCATION	Location of the device which is updated by the user.	varchar	256
DESCRIPTION	Description of the device which is updated by user.	varchar	256
USER_DEFINED_VALUE1	Value of the user defined property 1.	varchar	256
USER_DEFINED_VALUE2	Value of the user defined property 2.	varchar	256
USER_DEFINED_VALUE3	Value of the user defined property 3.	varchar	256

TABLE 610 USERDEFINED_NETWORK_SCOPE

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
NAME	Name of the Scope	varchar	128

TABLE 610 USERDEFINED_NETWORK_SCOPE (Continued)

Field	Definition	Format	Size
USER_ID	Foreign Key USER_.ID. ID of the user who created the Custom Dashboard.	int	
OWNED_BY	Holds the feature identifier using which scope is created. 0 - Dashboard, 1 - FCP.	int	

TABLE 611 USERDEFINED_NETWORK_SCOPE_MEMBERSHIP

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SCOPE_ID	Foreign Key USERDEFINED_NETWORK_SCOPE.ID. The ID of the user defined network scope to which this membership belongs.	int	
FABRIC_ID	Foreign Key FABRIC.ID. The ID of the fabric in the membership. This can be null if user does not include Fabric in his custom membership.	int	
PRODUCT_ME_ID	Foreign Key MANAGED_ELEMENT.ID. The ME ID of the device in the membership. This can be null if user does not include Switch in his custom membership.	int	
SWITCH_PORT_ID	Foreign Key SWITCH_PORT.ID. The ID of the switch Port in the membership. This can be null if user does not include Switch Port in his custom membership.	int	
INTERFACE_ID	Foreign Key INTERFACE. INTERFACE_ID. The ID of the Interface in the membership. This can be null if user does not include Interface in his custom membership.	int	
DEVICE_PORT_ID	Foreign Key DEVICE_PORT.ID. The ID of the Device Port in the membership. This can be null if user does not include Device Port in his custom membership.	int	
ZONE_NAME	Holds the zone name. This can be null if the membership is created using other scopes.	varchar	64
ZONE_ALIAS	Holds the zone alias name. This can be null if the membership is created using other scopes.	varchar	64

TABLE 612 USER_PREFERENCE

Field	Definition	Format	Size
USER_NAME *	User name whose preferences are saved. It corresponds to user_name in USER_table.	varchar	128
CATEGORY *	The name for a set of related preferences.	varchar	128
CONTENT	The set of preferences saved as name-value pairs.	text	

TABLE 613 USER_REALTIME_MEASURE_MAPPING

Field	Definition	Format	Size
ID	Primary Key.	int	
USER_ID	Foreign key reference to the user_ Table.	int	
EXPRESSION_ID	Foreign key reference to Measure Table.	int	

TABLE 613 USER_REALTIME_MEASURE_MAPPING (Continued)

Field	Definition	Format	Size
MIB_OBJECT_ID	Foreign key reference to Measure Table.	int	
MEASURE_TYPE	This identifies the collectible type. 0 for MIBs and 1 for Expressions.	int	

TABLE 614 USER_REALTIME_MEASURE_SETTING

Field	Definition	Format	Size
ID	Primary Key field for the user_realtime_measure_setting table	int	
USER_ID	This is the foreign key reference key to the user_ Table	int	
EXPRESSION_ID	This is the foreign key reference key to the snmp_expression Table	int	
MIB_OBJECT_ID	This is the foreign key reference key to the mib_object Table	int	
TYPE	This identifies the collectible type. 0 for MIBs, 1 for Expressions	int	

TABLE 615 USER_RESOURCE_MAP

Field	Definition	Format	Size
USER_NAME*	User name.	varchar	128
RESOURCE_GROUP_ID*	Resource group name, which is mapped for the user.	int	

TABLE 616 USER_ROLE_MAP

Field	Definition	Format	Size
USER_NAME*	User name.	varchar	128
ROLE_ID*	Role ID, which is mapped for the user.	int	

TABLE 617 USER_STATE_MAP

Field	Definition	Format	Size
USER_NAME		varchar	128
STATE	Current user state. The possible values are: • 0 - Locked out by user manager • 1 - Locked Out Threshold Reached • 2 - Password Expired • 3 - Password History Policy Violated • 4 - Password Format Policy Violated **NOTE:** This numeric state value will be mapped to associated ENUM at DTO side	smallint	

TABLE 618 V_PORT_DETAIL

Field	Definition	Format	Size
DEVICE_PORT_ID	Primary key from the owner device port table.	int	
STATE	Flag to indicate whether port is online or offline	varchar	32

TABLE 618 V_PORT_DETAIL (Continued)

Field	Definition	Format	Size
FCP_INITIATOR	The role of the virtual port; for example, FCP Initiator	varchar	256
SWITCH_IP	IP of the switch, the V port is connected to	varchar	128
VF_ID	VF ID for the V port	smallint	

TABLE 619 VCN_ICL

Field	Definition	Format	Size
VCN_ICL_ID	Virtual Cluster Node ICL DB ID.	int	
ICL_NAME	ICL name.	varchar	100
ICL_PORT_ID	ICL port foreign key.	int	
VCN_MEMBER_ID	Virtual Cluster Node member id foreign key.	int	

TABLE 620 VCN_MEMBER

Field	Definition	Format	Size
VCN_MEMBER_ID	Virtual Cluster Node member db id.	int	
CLUSTER_ID	Cluster id.	int	
CLUSTER_NAME	Cluster name.	varchar	100
CLUSTER_RBRIDGE_ID	Cluster rbridge id.	int	
SESSION_VLAN	Session VLAN id.	smallint	
KEEP_ALIVE_VLAN	Keep alive VLAN id.	smallint	
CLIENT_ISOLATION_MODE	Cluster isolation state: • Loose(0) • Strict(1).	smallint	
IS_CLIENTS_SHUTDOWN	Is MCT Client interfaces shutdown?	numeric	(1,0)
MEMBER_VLAN_RANGE	Configured member VLAN range.	varchar	256
ACTIVE_MEMBER_VLAN_RANGE	Active member VLAN range.	varchar	256
CLUSTER_DEPLOY_STATE	Cluster deployment state: • Deployed(0) • Undeployed(1).	smallint	
DEVICE_ID	Device id foreign key.	int	

TABLE 621 VCN_PEER

Field	Definition	Format	Size
VCN_PEER_ID	Virtual Cluster Node Peer db id.	int	
IP_ADDRESS	Peer ip address.	varchar	100
RBRIDGE_ID	Peer rbridge id.	int	
ICL_NAME	Cluster ICL name used for this peer.	varchar	100
FAST_FAILOVER_STATE	Cluster Peer fast failover state: • Disabled(0) • Enabled(1).	smallint	

TABLE 621 VCN_PEER (Continued)

Field	Definition	Format	Size
KEEP_ALIVE_INTERVAL	Cluster Peer keep alive interval in seconds.	INET	
HOLD_TIME	Cluster Peer hold time in seconds.	int	
ACTIVE_MEMBER_VLAN_RANGE	Cluster Peer Active member VLAN range.	varchar	256
PEER_OPER_STATE	Cluster Peer operational state.	smallint	
PEER_DOWN_REASON	Cluster Peer down reason.	int	
PEER_UP_TIME	Cluster Peer up time.	int	
VCN_MEMBER_ID	Virtual Cluster Node member id foreign key.	int	
PEER_DEVICE_ID	Peer device id.	int	

TABLE 622 VCS_CLUSTER_MEMBER

Field	Definition	Format	Size
CLUSTER_ME_ID	The Management Element ID of the VCS Cluster in the VirtualSwitch.	int	
MEMBER_ME_ID	The Management Element ID of the cluster member in the VirtualSwitch.	int	
CREATION_TIME	Creation time of the record	timestamp	
TRUSTED	Describes whether the member is trusted. Possible values are 1 and 0. 1 means trusted and 0 means untrusted.	smallint	
MISSING	Describes whether the member is missing or not. Possible values are 1 and 0. 1 means missing and 0 means not missing	smallint	
MISSING_TIME	Time when the member gone missing.	timestamp	
STATE	Indicates the state of the member with respect to cluster. States can be Online, Offline, Rejoining etc.	varchar	64
FABRIC_STATUS	Stores the fabric level status of the node like Unknown and Online. Status is unknown when: • A node is going through a reboot or ISLs have not formed yet. • A node is not part of a cluster yet. Status is Online when: • A node is waiting to rejoin a cluster. • A node joins a cluster and all the ports are up and ISLs are formed.	varchar	64

TABLE 623 VCEM_PROFILE

Field	Definition	Format	Size
ID			
VERSION	Version of the VCEM server.	varchar	32
NETWORK_ADDRESS	The version number of the VCEM server.	varchar	128
PORT_NUMBER	The SOAP API port number on the VCEM server.	int	

TABLE 623 VCEM_PROFILE (Continued)

Field	Definition	Format	Size
USERNAME	The username to be used to logon to the VCEM server.	varchar	512
PASSWORD	The password to be used to logon to the VCEM server. Will be store encrypted.	varchar	512
DISCOVERY_STATUS	The current discovery status of the VCEM server (discovery pending, ,active, failed, deleted pending etc.).	smallint	
LAST_DISCOVERY_STATUS	The discovery status of the VCEM server in the previous discovery cycle.	smallint	
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	
LAST_FAILURE_TIMESTAMP	The time of the last failed collection.	timestamp	
LAST_SUCCESSFUL_TIMESTAMP	The time of the last successful collection.	timestamp	

TABLE 624 VIP_SERVER

Field	Definition	Format	Size
ID	Primary Key field for the VIP_SERVER	int	
TYPE	Even Policy Type • 0? Virtual Server • 1 ? Real Server	smallint	
DEVICE_ID	This is the foreign key reference key to the Device Table	int	
IP_ADDRESS	The IP Address for the Virtual Server or Real Server	varchar	128
NAME	The Name of Virtual Server or Real Server	varchar	256

TABLE 625 VIP_SERVER_BINDING

Field	Definition	Format	Size
ID	Primary Key field for the VIP_SERVER_BINDING	int	
DEVICE_ID	This is the foreign key reference key to the Device Table	int	
VIRTUAL_SERVER_IP_ADDRESS	The IP Address for the Virtual Server	varchar	128
VIRUTAL_SERVER_PORT	The Port number of the Virtual Server	int	
REAL_SERVER_IP_ADDRESS	The IP Address for the Real Server	varchar	128
REAL_SERVER_PORT	The Port Number for the Real Server	int	

TABLE 626 VIRTUAL_CIRCUIT_INTERFACE

Field	Definition	Format	Size
INTERFACE_ID		int	

TABLE 627 VIRTUAL_FCOE_PORT

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
VIRTUAL_SWITCH_ID	The unique id of switch the virtual fcoe port belongs to.	int	
PORT_WWN	WWN of port	varchar	64
PORT_SPEED	Will be 10G.	varchar	32
PORT_TYPE	Will be Virtual-FCoE-Port	varchar	16
ENABLED	Enabled/disabled	smallint	
STATUS	Status	varchar	64
TRUNK_INDEX	Trunk index	smallint	
PORT_NUMBER	Port number	smallint	
NAME	Name	varchar	64
SLOT_NUMBER	The Slot number in the switch to which this Virtual FCoE Port belongs	int	
VLAN_ID	Comma Separated values of the VLANs associated with this Virtual FCoE Port	varchar	64
DEVICE_COUNT	The number of devices associated with this Virtual FCoE Port. The default value is 0.	smallint	
PEER_MAC	The Peer FCF MAC if this Virtual FCoE Port is a FCoE VE-port	varchar	

TABLE 628 VIRTUAL_FCOE_PORT_MAC_MEMBER

Field	Definition	Format	Size
VIRTUAL_FCOE_PORT_ID	The unique id of virtual fcoe port the member belongs to	int	
MAC_ADDRESS	Mac address of member.	varchar	64

TABLE 629 VIRTUAL_FCOE_PORT_STAT

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SWITCH_ID		int	
PORT_ID		int	
TX	The number of valid frames sent from the port	double precision	
RX	The number of valid frames received at this port	double precision	

TABLE 629 VIRTUAL_FCOE_PORT_STAT (Continued)

Field	Definition	Format	Size
TX_UTILIZATION	The computed value of TX based on speed of port (for MarchingAnts)	double precision	
RX_UTILIZATION	The computed value of RX based on speed of port (for MarchingAnts)	double precision	
CREATION_TIME	The time this stats record was created	timestamp	
ACTIVE_STATE	Used for error scenario	smallint	
LINK_FAILURES	Link failures	double precision	
TX_LINK_RESETS	TX Link resets	double precision	
RX_LINK_RESETS	RX link resets	double precision	
SYNC_LOSSES	Synchronization losses	double precision	
SIGNAL_LOSSES	Signal losses	double precision	
SEQUENCE_ERRORS	Sequence Errors	double precision	
INVALID_TX	Invalid transmissions	double precision	
CRC_ERRORS	Cyclic Redundancy check error	double precision	

TABLE 630 VIRTUAL_PORT_WWN_DETAILS

Field	Definition	Format	Size
ID	Unique generated database identifier.		
SWITCH_ID	If the VPWWN is constructed based on AG Node WWN and AG_Port_Index then this is id of connected switch.	int	
SWITCH_PORT_NUMBER	If the VPWWN is configured for AG , this value will have the default value(-1).	smallint	
AG_NODE_WWN	If the VPWWN is configured for Switch Port , this value will have the default value.	char	23
AG_PORT_NUMBER	If the VPWWN is configured for Switch Port , this value will have the default value.	smallint	
TYPE	Active WWN 0-Auto is the switch created VPWWN and User is user defined VPWWN'; 1-User	smallint	
STATUS	Enable or disable the VPWWN feature on switch port or AG-port. • 1-Enabled • 0-disabled	smallint	
USER_VPWWN	User created VPWWN.	char	23

TABLE 630 VIRTUAL_PORT_WWN_DETAILS (Continued)

Field	Definition	Format	Size
AUTO_VPWWN	VPWWN created by Switch.	char	23
DEVICE_PORT_WWN	Physical port WWN of the device for which VPWWN is assigned.	char	23
SLOT_NUMBER	Slot number of the switch, This will be -1 for AG.	smallint	

TABLE 631 VIRTUAL_SWITCH

Field	Definition	Format	Size
ID*	Primary key for the table.	int	
NAME	Stores the switch name.	varchar	64
WWN	WWN of the Switch.	char	23
VIRTUAL_FABRIC_ID	Virtaul fabric ID of the switch. A positive value will be stored if VF is enabled else -1.	smallint	
DOMAIN_ID	Domain ID of the switch.	smallint	
BASE_SWITCH	Indidates whether its a base switch. 1 is base switch and 0 is not.	smallint	
SWITCH_MODE	Stores the switch mode. • 0 is switch mode • 2 is ag mode.	smallint	
ROLE	Stores the role of the switch like Primary, Subordinate, Cluster etc.	varchar	32
FCS_ROLE	FCS role for the Switch . This is used only when FCS policy is turned on.	varchar	16
AD_CAPABLE	Stores the switch capability for Admin domain. • 1 is capable • 0 is not capable.	smallint	
FABRIC_IDID_MODE	Denotes if Insistent Domain ID mode is enabled.	smallint	
OPERATIONAL_STATUS	Stores the operational status of the switch.	varchar	128
MAX_ZONE_CONFIG_SIZE	Denotes the maximum supported zone DB size in bytes.	int	
CREATION_TIME	Creation time of the record.	timestamp	
LAST_UPDATE_TIME	Stores the timestamp of the last database update.	timestamp	
USER_NAME	Stores the telnet user name used to login to switch.	varchar	128
PASSWORD	Password used to login to the switch.	varchar	128
MANAGEMENT_STATE	Management state of the switch. This is a bit mask that indicates the switches manageability state, like switch not reachable, credentials invalid, not ready for manangement etc . • 0 means management state is Ok • non zero value will indicate manageability issues.	bigint	
STATE	Stores the switch state like Online, offline etc.	varchar	32
STATUS	Stores the status value here : UNKNOWN(0), MARGINAL(2),DOWN(3),HEALTHY(1).	varchar	32

TABLE 631 VIRTUAL_SWITCH (Continued)

Field	Definition	Format	Size
STATUS_REASON	Stores the status reason, which states the contributors for the status.	varchar	2048
USER_DEFINED_VALUE_1	User defined value used for annotation.	varchar	256
USER_DEFINED_VALUE_2	User defined value used for annotation.	varchar	256
USER_DEFINED_VALUE_3	User defined value used for annotation.	varchar	256
CORE_SWITCH_ID	Reference to Core Switch record.	int	
INTEROP_MODE	Interop mode for the switch. • 0 is native • 2 is McData • 3 is open fabric.	smallint	
CRYPTO_CAPABLE	Stores the switch capability for crypto support . • 1 is capable • 0 is not capable.	smallint	
FCR_CAPABLE	Stores the switch capability for FCR support . • 1 is capable • 0 is not capable.	smallint	
FCIP_CAPABLE	Stores the switch capability for FCIP support . • 1 is capable • 0 is not capable.	smallint	
FCOE_CAPABLE	If the switch supports FCoE. Default value is 0.	smallint	
L2_CAPABLE	If the switch supports L2.	smallint	
L3_CAPABLE	If the switch supports L3.	smallint	
LF_ENABLED	Logical Fabric Enabled/Disabled for a Virtual Switch. Default value is 0.	smallint	
DEFAULT_LOGICAL_SWITCH	Check to see whether virtual switch is a default logical switch or not. 1 is true and 0 is false. Default value is 0.	smallint	
FEATURES_SUPPORTED	Contains the features supported as a bit mask. Default value is 0.	int	
FMS_MODE	Stores FMS mode in FICON environment.	smallint	
DYNAMIC_LOAD_SHARING	Stores the switch capability for dynamic load sharing, • 1 is capable • 0 is not capable.	smallint	
PORT_BASED_ROUTING	Indicates whether the port based routing is present. • 1 is present • 0 is absent.	smallint	
IN_ORDER_DELIVERY	Indicates whether in order delivery is enabled or disabled. • 1 is enabled • 0 is disabled.	smallint	

TABLE 631 VIRTUAL_SWITCH (Continued)

Field	Definition	Format	Size
INSISTENT_DID_MODE	Indicates whether persistent domain ID is enabled on the switch. • 1 is enabled • 0 is disabled.	smallint	
LAST_SCAN_TIME	Stores the timestamp of the last scan time, the time which the switch was contacted for update.	timestamp	
DOMAIN_MODE_239	Stores the domain mode offset. Its only used in the mixed fabric (FOS+EOS).	smallint	
DOMAIN_ID_OFFSET	Stores the domain id offset value. Its only used in the mixed fabric (FOS+EOS).	smallint	
PREVIOUS_OPERATIONAL_STATUS	This table can hold the same values as OEPRATION_STATUS column. But this will be holding the previous OPERATIONAL_STATUS of the Virtual switch. These values to be populated by FCS during Fabric Refresh task	varchar	128
FCOE_LOGIN_ENABLED	The FCoE Login Management Status of the switch. Default value is 0.	smallint	
FCIP_CIRCUIT_CAPABLE	Whether the switch can create FCIP Circuits. 1 means true and 0 means false. Default value is 0.	smallint	
DISCOVERED_PORT_COUNT	Reflects the number of managed ports in the discovered switch. Default value is 0.	smallint	
LAST_PORT_MEMBERSHIP_CHANGE	Stores the timestamp of the last port member ship update.	bigint	
MAX_FCIP_TUNNELS	The maximun number of tunnels that can be created in this switch,-1 means not supported. Default value is -1.	int	
MAX_FCIP_CIRCUITS	The maximun number of circuits that can be created in this switch, -1 means not supported. Default value is -1.	int	
FCIP_LICENSED	FCIP Advanced Extension Licensing is available. 1 means licensed and 0 means not licensed, -1 means not supported. Default value is -1.	smallint	
ADDRESSING_MODE	This column to represent the logical switch addressing modes to assign Port Addresses, There are three different addressing modes supported. Fixed (0), Flat or 10 bit (1), Dynamic (2). Default value is -1.	smallint	
PREVIOUS_STATE	This fields copies the old state of the switch . The field could be used to track the state change information for the switch.These values to be populated by FCS during Fabric Refresh task.SMIA requested this information but could be used by any module which needs to track the state change	varchar	32

TABLE 631 VIRTUAL_SWITCH (Continued)

Field	Definition	Format	Size
MANAGED_ELEMENT_ID	A unique managed element ID for this virtual switch. Also a foreign key reference to the MANAGED_ELEMENT table.	int	
HIF_ENABLED	The HIF Enabled bit on the switch. Values are 1 for enabled and 0 for not enabled. -1 the default, stands for not supported and will be used for older firmwares. Default value is -1.	smallint	
CLUSTER_MODE	This column is used to determine whether VCS Cluster is in Standalone mode or Cluster mode. The values will be populated by the VCS collector during the discovery of the VCS switch. The default value of -1 means that its a non VCS device. Following Enum will be defined as NON_VCS(-1), STANDALONE(0), CLUSTER(1).	smallint	
VCS_ID	This column is used to store the VCS ID of the device. The value will be populated by the NosSwitchAssetCollector during the discovery of the VCS Cluster. The non zero value will be stored as VCS ID. Default value is -1.	smallint	
CLUSTER_TYPE	This column is used to determine whether VCS is in Fabric Cluster or Logical Chassis. The values are populated by the VCS collector during the discovery of the VCS switch. The default value -1 means that its a non-VCS device. Following are the values and their types: • 0 - Unknown • 1 - Standalone • 2 - Fabric Cluster • 3 - Logical Chassis	smallint	
SWITCH_ID	Represents the Switch embedded port destination identifier.	int	
MONITORED	To identify whether the switch is monitored or unmonitored. 0 is Unmonitored and 1 is Monitored.	int	
FEATURES_ENABLED	Holds as a bit mask the features that are active / enabled. Each bit would represent features like Lossless etc.	int	
MAPS_ENABLED_ACTIONS	Bitmask of Maps actions enabled on the switch. 0-None, 1-Raslog, 2-SNMP, 4-Email, 8-Fence Port, 16-SW Down, 32-SW Marginal	int	
FABRIC_STATUS	Stores the fabric level status of the node like Unknown and Online. Status is Unknown when: • A node is going through a reboot or ISLs have not formed yet. • A node is not part of a cluster yet. Status is Online when: • A node is waiting to rejoin a cluster. • A node joins a cluster and all the ports are up and ISLs are formed.	varchar	64

TABLE 631 VIRTUAL_SWITCH (Continued)

Field	Definition	Format	Size
ROUTING_POLICY	This column represents the routing policy (APT-Advanced Performance Tuning) configured in the switch. 1 = Not Applicable. Routing policy is not applicable for FOS AG switches/NOS/mEOS switches. 0 = Unknown. For non-AG FOS switches unknown value will be set as the initial value. Once asset collection is successful, the corresponding value will be fetched from switch. 1 = Port Based Routing 2 = Device Based Routing 3 = Exchange Based Routing	int	
PROTOCOL	Stores the communication protocol used to communicate with this device. Following are the values this column can hold: 0 - NA, for NOS devices 1 - Communication using HTTP 2 - Communication using HTTPS	int	
BOUND	Indicates whether this BNA Server is bound to Logical AMP or not. 1- If this BNA server bound to Logical AMP. 0 â€“ If bind operation is not initiated or Logical AMP is bound to some other BNA server.	smallint	
BOUND_BNA_IP_ADDRESS	The IP Address of BNA server bound to Logical AMP. Empty or 0.0.0.0 if no BNA server bound to Logical AMP.	varchar	128

TABLE 632 VIRTUAL_SWITCH_CAPABILITY

Field	Definition	Format	Size
VIRTUAL-SWITCH_ID *	DB ID of virtual switch.	int	
CAPABILITY_ *	Name of capability detected on virtual switch.	varchar	256
ENABLED	1 = the capability is enabled on the virtual switch.	int	

TABLE 633 VIRTUAL_SWITCH_CHECKSUM

Field	Definition	Format	Size
VIRTUAL_SWITCH_ID *	DB ID of virtual switch.	int	
CHECKSUM_KEY *	Checksum key.	varchar	32
CHECKSUM	Checksum value.	varchar	16

TABLE 634 VIRTUAL_SWITCH_COLLECTION

Field	Definition	Format	Size
VIRTUAL_SWITCH_ID *	DB ID of virtual switch.	int	
COLLECTOR_NAME *	Collector name.	varchar	256
LAST_VIRTUAL_SW_MODIFICATION	Last modified time on switch.	timestamp	

TABLE 635 VLAN

Field	Definition	Format	Size
VLAN_DB_ID	Unique database generated identifier.	int	
DEVICE_ID	Database ID of the DEVICE instance which is associated with the vlan.	int	
NAME	Name for vlan.	varchar	128
TABLE_SUBTYPE	Table subtype possible value is VLAN.	varchar	32

TABLE 636 VLAN_DYNAMIC_INTERFACE_MEMBER

Field	Definition	Format	Size
VLAN_INTERFACE_RELATIO N_ID	Database ID of the VLAN_INTERFACE_RELATION instance which is associated with the dynamic interface member.	int	

TABLE 637 VLAN_EXCLUDED_INTERFACE

Field	Definition	Format	Size
VLAN_INTERFACE_RELATIO N_ID	Database ID of the VLAN_INTERFACE_RELATION instance which is associated with the excluded interface member.	int	

TABLE 638 VLAN_INTERFACE_MEMBER

Field	Definition	Format	Size
VLAN_INTERFACE_RELATIO N_ID	Database ID of the VLAN_INTERFACE_RELATION instance which is associated with the interface member.	int	

TABLE 639 VLAN_INTERFACE_RELATION

Field	Definition	Format	Size
VLAN_INTERFACE_RELATIO N_ID	Unique database generated identifier.	int	
VLAN_DB_ID	Database ID of the VLAN instance which is associated with the interface.	int	
INTERFACE_ID	Database ID of the INTERFACE instance which is associated with the vlan.	int	
TABLE_SUBTYPE	Table subtype possible value is VLAN_INTERFACE_RELATION.		
DEVICE_ID	VLAN associated device id.	int	

TABLE 640 VLAN_INT_C_TAG_RELATION

Field	Definition	Format	Size
VLAN_INTERFACE_RELATIO N_ID	Foreign Key Reference to VLAN_INTERFACE_RELATION table.	int	
C_TAG_ID	This as an Incoming customer tag (c-tag) associated with a GVLAN and its applicable only for trunk mode. If the TLS (Transparent LAN Service) is enabled in the device, a pre-defined range of values are used for C-Tag.	text	

TABLE 641 VLAN_INT_MAC_GROUP_RELATION

Field	Definition	Format	Size
VLAN_INTERFACE_RELATIO N_ID	Foreign Key Reference to VLAN_INTERFACE_RELATION table.	int	
MAC_GROUP_DB_ID	Foreign key Reference to ID field of MAC_GROUP table.	text	

TABLE 642 VLAN_STATIC_INTERFACE_MEMBER

Field	Definition	Format	Size
VLAN_INTERFACE_RELATIO N_ID	Database ID of the VLAN_INTERFACE_RELATION instance which is associated with the static interface member.	int	

TABLE 643 VLL_DEVICE_RELATION

Field	Definition	Format	Size
MPLS_SERVICE_DEVICE_REL ATION_DB_ID	Database ID inherited from MPLS_SERVICE_DEVICE_RELATION.	int	
VLL_DEVICE_RELATION.VLL_ MODE	Represents the VLL mode. Possible values are Unknown-0, Raw-1 and Taggged-2.	int	

TABLE 644 VLL_ENDPOINT_RELATION

Field	Definition	Format	Size
MPLS_SERVICE_ENDPOINT_ RELATION_DB_ID	Database ID inherited from MPLS_SERVICE_ENDPOINT_RELATION.	int	
PW_ENET_PW_INSTANCE	Represents the Index of Ethernet tables associated with this endpoint Instance.	int	
COS	This value indicates the Class Of Service for this endpoint. For VLL, this value is used to select the appropriate tunnel whose COS value is either same, or almost approaching this value. For VLL-local, this value is applied to the ingress packet of an endpoint. Allowed range is 0-7 and 255. 255 means COS is not explicitly configured.	smallint	

TABLE 645 VMOTION_EVENT

Field	Definition	Format	Size
ID	Uniquely identifies the vmotion event.	int	
SOURCE_HOST_NAME	The name of the source host at the time of the vmotion.	varchar	256
SOURCE_IP_ADDRESS	IP address of the source host at the time of the vmotion.	varchar	128
SOURCE_HOST_UUID	The uuid assigned by the hypervisor to the source host.	varchar	64
DEST_HOST_NAME	The name of the destination host at the time of the vmotion.	varchar	256
DEST_IP_ADDRESS	IP address of the destination host at the time of the vmotion.	varchar	128

TABLE 645 VMOTION_EVENT (Continued)

Field	Definition	Format	Size
DEST_HOST_UUID	The uuid assigned by the hypervisor to the destination host. This can be null in case of a failed vmotion.	varchar	64
SOURCE_DATACENTER_NAME	Source Datacenter name.	varchar	256
DEST_DATACENTER_NAME	Destination Dataccenter name. Can be null in case of a failed vmotion.	varchar	256
VM_UUID	Unique identifier for the VM to identify that VM across vmotions.	varchar	64
VM_NAME	User-assigned name for the VM.	varchar	80
VM_IP_ADDRESS	The primary IPv4 or IPv6 address used by the VM on the management LAN, if any.	varchar	32
VCENTER_HOST	The FQDN or the ip address of the vcenter.	varchar	256
VNIC_MACS	Comma separated vnic mac addresses.	varchar	256
START_TIME	Start time of the vmotion event.	timestamp	
END_TIME	End time of the vmotion event, could be null cause of a failed vmotion.	timestamp	
STATUS	VMotion event status. 0 = info, 1 = warning, 2 = failed.	smallint	
DRS_TRIGGERED	Identifies whether the events was due to DRS. 0 = No, 1 = Yes.	smallint	
USER_NAME	Identifies that user who initiated the vmotion.	varchar	80
DESCRIPTION	Event message that is received.	varchar	256

TABLE 646 VMOTION_PNIC_DETAILS

Field	Definition	Format	Size
ID	Identifies an entry for the source or destination pnic and the connected switch details.	int	
VMOTION_EVENT_ID	Foreign key to the vmotion_event table.	int	
PNIC_TYPE	Pnic type. 0 = source, 1 = destination, identifies if the pNIC is from the source or the destination host.	smallint	
PNIC_MAC	Physical Nic mac addresse of the connected Pnic on the host.	varchar	256
PORT_PROFILES	Comma separated PP Name-SwitchName for all the port profiles associated with the vNics that are being migrated as a result of the vmotion.	varchar	256
SWITCH_NAME	Switch names entry for connected switch to the pNic.	varchar	256
SWITCH_IP_ADDRESS	Switch ip addresses entries for connected switch to the pNic.	varchar	256

TABLE 647 VM_APPLICATION_DETAILS

Field	Definition	Format	Size
VM_INSTANCE_UUID	identifies the Instance UUID of the Virtual Machine.	varchar	256
APPLICATIONS_NAME	Names of the applications running in the VM. More than one application can be specified as comma separated values.	varchar	256

TABLE 648 VM_DATA_CENTER

Field	Definition	Format	Size
ID	Unique generated database identifier.		
NAME	Data center name.	varchar	256
VCENTER_ID	Id of the vCenter server managing this Data center.	int	
MOR_ID	The managed object reference number assigned by the hypervisor.	int	

TABLE 649 VM_DATASTORE_DETAILS

Field	Definition	Format	Size
ID	Primary key.	int	
DATACENTER_ID	Foreign to vm_data_center.	int	
NAME	Name of the datastore.	varchar	256
ACCESSIBLE	The connectivity status of this datastore. If this is set to false, meaning the datastore is not accessible, this datastores capacity and freespace properties cannot be validated. 0 = no 1 = yes.	smallint	
STATUS	Status of the datastore could be normal, enteringMaintenance, inMaintenance.	varchar	20
FILE_SYSTEM_TYPE	Type of file system volume, such as VMFS or NFS.	varchar	20
TOTAL_CAPACITY	Maximum capacity of this datastore, in bytes. This value is updated periodically by the server.	bigint	
FREE_SPACE	Available space of this datastore, in bytes. The server periodically updates this value.	bigint	
LAST_UPDATE_TIME	Time when the free-space and capacity values in DatastoreInfo and DatastoreSummary were updated.	timestamp	
RDM_SUPPORTED	Flag Indicates whether or not raw disk mappings can be created on this datastore. 0 = no 1 = yes.	smallint	
PERFILE_THIN_PROVISIONING_SUPPORTED	Flag indicating whether or not the per file thin provisioning is supported or not. 0 = no 1 = yes. When thin provisioning is used, backing storage is lazily allocated.	smallint	
STORAGE_IORM_SUPPORTED	Indicates whether the datastore supports Storage I/O Resource Management. 0 = no 1 = yes.	smallint	

TABLE 649 VM_DATASTORE_DETAILS (Continued)

Field	Definition	Format	Size
DIRECTORY_HIERARCHY_SUPPORTED	Indicates whether or not directories can be created on this datastore. 0 = no 1 = yes.	smallint	
LOCATION	The unique locator for the datastore.	varchar	256
MOR_ID	The managed object reference number assigned by the hypervisor.	int	

TABLE 650 VM_DV_PORT

Field	Definition	Format	Size
VM_DV_SWITCH_ID	Foreign key to the VM_DVSWITCH table. The dvSwitch on which this port group exists.	int	
VM_DV_PORT_GROUP_ID	Foreign Key to the VM_DV_PORTGROUP table. The dvPortgroup in which this dvPort instance may exist (in case it's not a standalone port)	int	
NAME	The name of the port	varchar	256
DESCRIPTION	A description string of the port.	varchar	256
CONFLICT	Whether the port is a conflict port. A port could be marked as conflict if an entity is discovered connecting to a port that is already occupied, or if the port is created by the host without conferring with Virtual Center Server. A conflict port will not have its runtime state persisted and the port can''t move away from the host, i.e no vMotion if a Virtual Machine is using a conflict port	smallint	
CONNECTEE_TYPE	The type of the connectee. One of: hostConsoleVnic hostVmkVnic pnic vmVnic	smallint	
CONNECTEE_ADDRESS_HINT	A hint on address info of the nic that connects to this port	varchar	256
MTU	The MTU of the port. Currently, this property can only be set at the switch level. Attempt to change it at the portgroup or port level will raise exception	int	
MAC_ADDRESS	The mac address that is used at this port	varchar	64
RUNTIME_LINK_UP_STATE	Whether the port is in linkUp status	varchar	128
RUNTIME_LINK_PEER	The name of the connected entity	varchar	128
RUNTIME_BLOCKED	Whether the port is blocked by switch implementation	smallint	
TRUNKING_MODE	True if the port VLAN tagging/stripping is disabled	smallint	
VLAN_IDS	The VLAN id of the port	varchar	256
PROXY_HOST_NAME	The host that services this port	varchar	256

TABLE 650 VM_DV_PORT (Continued)

Field	Definition	Format	Size
KEY	The key for the port	varchar	64
MOR_ID	The managed object reference number assigned by the hypervisor	int	

TABLE 651 VM_DV_PORT_GROUP

Field	Definition	Format	Size
VM_DV_SWITCH_ID	Foreign Key to the vm_dvswitch table. The dvSwitch on which this port group exists	int	
NAME	The name of the portgroup.	varchar	256
NUM_PORTS	Number of ports in the portgroup	int	
TYPE	The type of portgroup. One of: earlyBinding ephemeral lateBinding	smallint	
DESCRIPTION	A description string of the portgroup	varchar	256
UPLINK_PORT_GROUP	Whether this portgroup is an uplink portgroup	smallint	
KEY	The key for the port group	varchar	64
MOR_ID	The managed object reference number assigned by the hypervisor	int	

TABLE 652 VM_DV_SWITCH

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
UUID	The generated UUID of the switch. Unique across VC inventory and instances	varchar	256
NAME	The name of the switch	varchar	256
MAX_PORTS	The maximum number of ports allowed in the switch, not including conflict ports	int	
DESCRIPTION	A description string of the switch	varchar	1024
PORT_COUNT	Current number of ports, not including conflict ports	int	
STANDALONE_PORT_COUNT	The number of standalone ports in the switch. Standalone ports are ports that don''t belong to any portgroup	int	
ADMIN_NAME	The name of the person that is responsible for the switch	varchar	256
ADMIN_CONTACT	The contact information for the person	varchar	256
BUILD	Build string for the server on which this call is made. For example, x.y.z-num. This string does not apply to the API	varchar	256
PRODUCT_NAME	Short form of the product name	varchar	256
VENDOR_NAME	Name of the vendor of this product	varchar	256

TABLE 652 VM_DV_SWITCH (Continued)

Field	Definition	Format	Size
VERSION	Dot-separated version string. For example, "1.2"	varchar	256
FORWARDING_CLASS	Forwarding class of the distributed virtual switch	varchar	256
DV_PORT_GROUP_OPER_SUPPORTED	Whether this switch allow Virtual Center users to modify DVS configuration at portgroup level, except for host memeber, policy and scope operations	smallint	
DV_PORT_OPER_SUPPORTED	Whether this switch allow Virtual Center users to modify DVS configuration at port level, except for host memeber, policy and scope operations	smallint	
DVS_OPER_SUPPORTED	Whether this switch allow Virtual Center users to modify DVS configuration at switch level, except for host memeber, policy and scope operations	smallint	
CREATION_TIME	The create time of the switch	timestamp	
UPLINK_PORT_NAME	The uniform name of uplink ports on each host	varchar	256
VM_DATA_CENTER_ID	A foreign key referencing VM_DATACENTER table instance to which this host is associated with	int	
MOR_ID	The managed object reference number assigned by the hypervisor	int	
IP_ADDRESS	The IP address currently assigned to the DV switch.	varchar	64
IPFIX_ENABLED	Whether netflow is enabled on the dvswitch, 0 implies false and 1 implies that its enabled.	smallint	
DISCOVERY_PROTOCOL	Neighbor discovery protocol 0 = CDP else 1 which is LLDP.	smallint	
DISCOVERY_OPERATION	Discovery operation default is 0 = listen, 1= advertise, 2 = both, 3 = none.	smallint	
CDP_ENABLED	Whether CDP is enabled on the dvswitch, 0 implies false and 1 implies that its enabled.	smallint	
VSPAN_ENABLED	Whether Port Mirroring is enabled on the dvswitch, 0 implies false and 1 implies that its enabled.	smallint	
MAXIMUM_MTU	The maximum transmission unit (MTU) associated with this distributed virtual switch in bytes.	int	

TABLE 653 VM_DV_SWITCH_HOST_MEMBER

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VM_DV_SWITCH_ID	A foreign key referencing VM_DV_SWITCH (ID)	int	
VM_HOST_ID	A foreign key referencing VM_HOST (ID)	int	

TABLE 654 VM_FC_HBA

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
NODE_WWN	The world wide node name for the adapter	varchar	23
PORT_WWN	The world wide port name for the adapter	varchar	23
PORT_TYPE	The type of the fiber channel port. One of : • Fabric • Loop • Point to point • Unknown	smallint	
SPEED	The current operating speed of the adapter in bits per second.	varchar	64
BUS	The host bus number	int	
DEVICE_NAME	The device name of host bus adapter	varchar	256
DRIVER	The name of the driver	varchar	256
MODEL	The model name of the host bus adapter	varchar	256
PCI	The Peripheral Connect Interface (PCI) ID of the device representing the host bus adapter	varchar	256
STATUS	The operational status of the adapter. Valid values include : online offline fault	smallint	
VM_HOST_ID	A foreign key referencing VM_HOST table instance to which this host is associated with	int	
MOR_ID	The managed object reference number assigned by the hypervisor	int	

TABLE 655 VM_FC_HBA_DEVICE_PORT_MAP

Field	Definition	Format	Size
DEVICE_PORT_ID	A foreign key referencing DEVICE_PORT table instance to which this host is associated with	int	
VM_FC_HBA_ID	A foreign key referencing VM_FC_HBA table instance to which this host is associated with	int	

TABLE 656 VM_HOST

Field	Definition	Format	Size
DEVIE_ENCLOSURE_ID	Identifies a server running a supported hypervisor. The ID value is the same as the ID of the corresponding DEVICE_ENCLOSURE record.	int	
NODE_WWN	The Node WWN for this host.	char	23
HYPERVISOR_NAME	Hypervisor name and version, such as VMware ESX Server v3.5.0	varchar	64

TABLE 656 VM_HOST (Continued)

Field	Definition	Format	Size
HYPERVISOR_TYPE	Numeric hypervisor type ID. 1 = VMware, 2 = Hyper-V. The default value is 0.	smallint	
CPU_COUNT	Number of CPUs in the server. The default value is 0.	int	
CPU_TYPE	Text summary of CPU hardware, such as: Intel(R) Xeon(TM) CPU 2.6 GHz	varchar	64
CPU_RESOURCES	Text summary of CPU resources, such as "20 GHz total, 15 GHz reserved". May be a different format for different VM vendors	varchar	64
MEM_RESOURCES	Text summary of memory resources, such as "7 GB total, 5 GB reserved". May be a different format for different VM vendors	varchar	64
LICENSE_SERVER	IP address or hostname of VM Hypervisor''s license server.	varchar	128
BOOT_TIME	Date and time that the host was last started	timestamp	
VM_DATACENTER_ID	A foreign key referencing VM_DATACENTER table instance to which this host is associated with.	int	
DVS_HOSTMEMBER_STATUS	• 1 - disconnected The host is in disconnected or not responding state. • 2 - down The host component is down. • 3 - outOfSync The switch configuration in the host component is not the same as the configuration in VirtualCenter server. • 4 - pending The host component is waiting to be initialized. • 5 - up The host component is up and running. • 6 - warning The host requires attention.	smallint	
DVS_PRODUCT_NAME	Short form of the product name of proxy switch module of a dvSwitch.	varchar	256
DVS_PRODUCT_VENDOR	Name of the vendor of this product.	varchar	256
DVS_PRODUCT_VERSION	Dot-separated version string. For example, "1.2".	varchar	256
CLUSTER_NAME	The name of the cluster of which this host is a member of.	varchar	256
MOR_ID	The managed object reference number assigned by the hypervisor.	int	
UUID	UUID to uniquely identify the host.	varchar	64

TABLE 657 VM_HOST_END_DEV_CONNECTIVITY

Field	Definition	Format	Size
VM_PHYSICAL_NIC_ID	A foreign key referencing VM_PHYSICAL_NIC (ID)	int	
INTERFACE_ID	A foreign key referencing INTERFACE (ID)	int	

TABLE 658 VM_HOST_PROXY_SWITCH

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VM_HOST_ID	Foreign Key to the vm_host table	int	
DVS_NAME	The name of the DistributedVirtualSwitch that the HostProxySwitch is part of	varchar	256
DVS_UUID	The uuid of the DistributedVirtualSwitch that the HostProxySwitch is a part of	varchar	256
KEY_	The proxy switch key	varchar	256
NUM_PORTS	The number of ports that this switch currently has	int	
NUM_PORTS_AVAILABLE	The number of ports that are available on this virtual switch	int	
UPLINK_PORT_NAMES	The list of ports that can be potentially used by physical nics. This property contains the names of such ports	varchar	256

TABLE 659 VM_HOST_PROXY_SWITCH_PNIC_SPEC

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VM_HOST_PROXY_SWITCH_ID	Foreign Key to the vm_host_proxy_switch table	int	
PNIC_DEVICE	The physical NIC to be added in the switch	varchar	256
UPLINK_PORT_GROUP_KEY	The key of the portgroup to be connected to the physical NIC	varchar	256
UPLINK_PORT_KEY	The key of the port to be connected to the physical NICs	varchar	256
UPLINK_PORT_NAME	The name of the port to be connected to the physical NICs	varchar	256

TABLE 660 VM_HOST_VIRTUAL_NIC

Field	Definition	Format	Size
ID	Unique Auto Generated DB ID.	serial	
DEVICE_NAME	Device Name for the virtual NIC.	varchar	256
MAC	The media access control (MAC) address of the virtual network adapter	varchar	64
DHCP_ENABLED	The flag to indicate whether or not DHCP (dynamic host control protocol) is enabled. If this property is set to true, the ipAddress and the subnetMask strings cannot be set explicitly	smallint	
IP_ADDRESS	The IP address currently used by the network adapter. All IP addresses are specified using IPv4 dot notation	varchar	128
SUBNET_MASK	Subnet mask for the virtual NIC.	varchar	64

TABLE 660 VM_HOST_VIRTUAL_NIC (Continued)

Field	Definition	Format	Size
VM_STD_VSWITCH_PORT_GROUP_ID	Foreign Key to the VM_STD_VSWITCH_PORT_GROUP table. Port group with which this vmknic is associated	int	
VM_DV_PORT_ID	Foreign key to the vm_dv_port table. DV Port with which this vmknic is associated	int	
MTU	The MTU of the port	int	
VM_HOST_ID	FOREIGN KEY to the vm_host table	int	
MOR_ID	The managed object reference number assigned by the hypervisor	int	
PORT_GROUP_KEY	The key for the port group	varchar	256
BINARY_MAC	MAC address in binary format.	bytea	
BINARY_IP	IP address in binary format.	bytea	

TABLE 661 VM_NETWORK_SETTINGS

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VLAN_TYPE	One of: Private VLAN Trunk VLAN Access VLAN	smallint	
VLAN_IDS	Single or range of VLANs configured on the port	varchar	256
BLOCKED	Whether this port is blocked, i.e. packet forwarding is stopped	int	
VM_STD_VSWITCH_PORT_GROUP_ID	ID of standard vSwitch port group	int	
VM_STANDARD_VIRTUAL_SWITCH_ID	ID of standard vSwitch	int	
VM_DV_SWITCH_ID	ID of distributed vSwitch	int	
VM_DV_PORT_GROUP_ID	ID of distributed vSwitch port group	int	
VM_DV_PORT_ID	ID of distributed vSwitch port	int	

TABLE 662 VM_NIC_TEAMING_POLICY

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
NOTIFY_SWITCHES	Flag to specify whether or not to notify the physical switch if a link fails. If this property is true, ESX Server will respond to the failure by sending a RARP packet from a different physical adapter, causing the switch to update its cache.	smallint	

TABLE 662 VM_NIC_TEAMING_POLICY (Continued)

Field	Definition	Format	Size
POLICY	Network adapter teaming policy includes failover and load balancing, It can be one of the following: • loadbalance_ip: route based on ip hash. • loadbalance_srcmac: route based on source MAC hash. • loadbalance_srcid: route based on the source of the port ID. • failover_explicit: use explicit failover order.	smallint	
REVERSE_POLICY	The flag to indicate whether or not the teaming policy is applied to inbound frames as well. For example, if the policy is explicit failover, a broadcast request goes through uplink1 and comes back through uplink2. Then if the reverse policy is set, the frame is dropped when it is received from uplink2. This reverse policy is useful to prevent the virtual machine from getting reflections.	smallint	
ROLLING_ORDER	The flag to indicate whether or not to use a rolling policy when restoring links. For example, assume the explicit link order is (vmnic9, vmnic0), therefore vmnic9 goes down, vmnic0 comes up. However, when vmnic9 comes backup, if rollingOrder is set to be true, vmnic0 continues to be used, otherwise, vmnic9 is restored as specified in the explicitly order.	smallint	
ACTIVE_NICS_ORDER	Comma separated list of active network adapters used for load balancing.	varchar	1056
STANDBY_NICS_ORDER	Standby network adapters used for failover.	varchar	1056
NIC_FAIL_CRITERIA_CHK_BEACON	Failover detection policy for this network adapter team. The bridge must be BondBridge for this property to be valid. The flag to indicate whether or not to enable this property to enable beacon probing as a method to validate the link status of a physical network adapter. checkBeacon can be enabled only if the VirtualSwitch has been configured to use the beacon. Attempting to set checkBeacon on a PortGroup or VirtualSwitch that does not have beacon probing configured for the applicable VirtualSwitch results in an error.	smallint	
VM_NETWORK_SETTINGS_ID	ID of network settings table.	int	
UNUSED_NICS_ORDER	Comma separated list of unused network adapters.	varchar	1056

TABLE 663 VM_PATH

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
HOST_ID	Identifies the host containing this path. This is a foreign key reference to VM_HOST.ID	int	
VM_ID	Identifies the VM using this path to a LUN. If the path is used by the host hypervisor instead of a VM, VM_ID is 0. When non-zero, this value matches VIRTUAL_MACHINE.ID	int	

TABLE 663 VM_PATH (Continued)

Field	Definition	Format	Size
STORAGE_ID	Identifies the LUN that is assigned to the VM. Not a foreign key, but the value matches VM_LUN.ID	int	
NAME	The VM-assigned name for this path. For VMware, this is the device name, such as vmhba0:0:1.	varchar	128
FABRIC_ID	Identifies the fabric that contains this path. Not a foreign key reference. Copied here for convenience. Determined by locating the HBA port WWN or target port WWN in the DEVICE_PORT table. Zero if the fabric is not managed. The default value is 0.	int	
HBA_PORT	The HBAs physical port WWN for this path	char	23
VM_PORT_WWN	The initiator port WWN used by the VM. If NPIV is used, this is a virtual port WWN assigned by the VM to this HBA port. If NPIV is not used, this WWN is the same as the HBA Port WWN	char	23
TARGET_PORT	The port WWN of the destination target.	char	23
ENABLED	'0 = path disabled, 1 = path enabled. The default value is 0.	smallint	
ACTIVE	0 = path inactive, 1 = path active. The default value is 0.	smallint	
PREFERRED	0 = not preferred, 1 = preferred path. The preferred path is used whenever available when the path policy is Fixed.The default value is 0.	smallint	
USAGE	Identifies how a VMware VM uses this LUN. 0 = NA (used for Hyper-V), 1 = VMFS (datastores), 2 = RDM (Raw Device Mapping). The default value is 0.	smallint	
HBA_NODE	The HBA physical node WWN for this path	char	23
VM_NODE_WWN	The initiator node WWN used by the VM. If NPIV is used, this is a virtual node WWN assigned to the VM. If NPIV is not used, this WWN is the same as the node WWN of one of the HBAs in the host.	char	23
TARGET_NODE	The node WWN of the destination target	char	23
HBA_NAME	The hypervisor device name of the HBA used in this path, such as vmhba1	varchar	64
FS_TYPE	This field will identify the filesystem type to be either: VMFS, NFS or RDM.	varchar	32

TABLE 664 VM_PHYSICAL_NIC

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
DEVICE_NAME	The device name of the physical network adapter.	varchar	256
DRIVER	The name of the driver	varchar	256

TABLE 664 VM_PHYSICAL_NIC (Continued)

Field	Definition	Format	Size
LINK_SPEED_MBPS	The bit rate on the link	int	
DUPLEX	The flag to indicate whether or not the link is capable of full-duplex ("true") or only half-duplex ("false").	smallint	
MAC_ADDRESS	The media access control (MAC) address of the physical network adapter.	varchar	17
PCI	Device hash of the PCI device corresponding to this physical network adapter.	varchar	256
WAKE_ON_LAN_SUPPORTED	Flag indicating whether the NIC is wake-on-LAN capable. 0 - false, 1 - true.	smallint	
DHCP_ENABLED	The flag to indicate whether or not DHCP (dynamic host control protocol) is enabled. If this property is set to true, the ipAddress and the subnetMask strings cannot be set explicitly. 0 - false, 1 - true.	smallint	
IP_ADDRESS	The IP address currently used by the network adapter. All IP addresses are specified using IPv4 dot notation. For example, "192.168.0.1". Subnet addresses and netmasks are specified using the same notation.	varchar	64
SUBNET_MASK	Subnet mask for the Physical NIC.	varchar	64
VM_HOST_ID	A foreign key referencing VM_HOST(ID).	int	
VM_STANDARD_VIRTUAL_SWITCH_ID	A foreign key referencing VM_STANDARD_VIRTUAL_SWITCH(ID).	int	
VM_DV_PORT_ID	A foreign key referencing VM_DV_PORT(ID).	int	
MOR_ID	The managed object reference number assigned by the hypervisor.	int	
BINARY_MAC	MAC address in binary format.	bytea	

TABLE 665 VM_SECURITY_POLICY

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
ALLOW_PROMISCUOUS	The flag to indicate whether or not all traffic is seen on the port. 0 - false, 1 - true	smallint	
FORGED_TRANSMITS	The flag to indicate whether or not the virtual network adapter should be allowed to send network traffic with a different MAC address than that of the virtual network adapter. 0 - false, 1 - true	smallint	
MAC_CHANGES	The flag to indicate whether or not the Media Access Control (MAC) address can be changed. 0 - false, 1 - true	smallint	
VM_NETWORK_SETTINGS_ID	ID of network settings table.	int	

TABLE 666 VM_STANDARD_VIRTUAL_SWITCH

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
NAME	The name of the virtual switch.	varchar	32
PORTS_COUNT	The number of ports that this virtual switch currently has.	int	
PORTS_AVAILABLE	The number of ports that are available on this virtual switch.	int	
MTU	The maximum transmission unit (MTU) associated with this virtual switch in bytes.	int	
BRIDGE_TYPE	The bridge specification describes how physical network adapters can be bridged to a virtual switch. One of: Auto Bridge - 0, Bond Bridge - 1, Simple Bridge - 2.	smallint	
VM_HOST_ID	References the ESX host in which this switch exists.	int	
MOR_ID	The managed object reference number assigned by the hypervisor.	int	

TABLE 667 VM_STANDARD_VSWITCH_PORT

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
MAC	The Media Access Control (MAC) address of network service of the virtual machine connected on this port.	varchar	64
TYPE	The type of component connected on this port. One of: • VMKernel • Service Console • Unknown • VM	smallint	
VM_STD_VSWITCH_PORT_GROUP_ID	Foreign Key to the VM_STD_VSWITCH_PORT_GROUP table. Port group in which this port exists.	int	
MOR_ID	The managed object reference number assigned by the hypervisor.	int	

TABLE 668 VM_STD_VSWITCH_PORT_GROUP

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
NAME	The name of the port group.	varchar	256
VM_STANDARD_VIRTUAL__SWITCH_ID	Foreign Key to the vm_standard_virtual_switch table. The standard virtual swtich on which this port group exists.	int	
MOR_ID	The managed object reference number assigned by the hypervisor.	int	

TABLE 669 VM_STORAGE

Field	Definition	Format	Size
ID	Uniquely identifies this LUN.	serial	
HOST_ID	Identifies the server that accesses this LUN.	int	
NAME	The VM-assigned device name for this LUN, such as vmhba1:0:0. For VMware, this is the canonical name.	varchar	512
TARGET_NODE	The Node WWN or iSCSI target name for the storage device (target) that contains this LUN.	char	256
VENDOR	Vendor name, such as Seagate.	varchar	64
MODEL	Target model name, such as ST581.	varchar	64
SERIAL_NUMBER	The device''s serial number.	varchar	64
TYPE	0 = disk, 1 = tape.	smallint	
CAPACITY	For disks, the disk capacity in GB.	double precision	
STATUS	The status reported by the host. 0 = offline, 1 = online.	smallint	
PATH_POLICY	Determines how multiple paths to this LUN are used. 0 = fixed, 1 = Most Recently Used, 2 = Round Robin.	smallint	
UUID	Universal unique ID	varchar	
DATASTORE_URL	The unique locator for the datastore.	varchar	256
DATASTORE_NAME	Name of the datastore in case this LUN/NAS volume is exposed as an extent of a VMFS/NFS datastore.	varchar	256
ISCSI_TARGET_ADDRESS	IP address or host name of the iSCSI target.	varchar	256
ISCSI_TARGET_PORT	The TCP port of the storage device. If not specified, the standard default of 3260 is used.	varchar	10
NAS_REMOTE_HOST	The host that runs the NFS/CIFS server.	varchar	64
NAS_REMOTE_PATH	The remote path of NFS/CIFS mount point.	varchar	256
NAS_REMOTE_USER	In case of CIFS, the user name used while connecting to the server.	varchar	256
TARGET_PORT	Target Port WWN that the storage is connected to or the iSCSI target address.	varchar	256)

TABLE 670 VM_STORAGE_HBA_REMOTE_PORT_MAP

Field	Definition	Format	Size
VM_STORAGE_ID	A foreign key referencing VM_STORAGE (ID).	int	
HBA_REMOTE_PORT_ID	A foreign key referencing HBA_REMOTE_PORT (ID).	int	

TABLE 671 VM_TRAFFIC_SHAPING_POLICY

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
ENABLED	'The flag to indicate whether or not traffic shaper is enabled on the port. 0 - false, 1 - true	smallint	
AVERAGE_BANDWIDTH	The average bandwidth in bits per second if shaping is enabled on the port.	bigint	
BURST_SIZE	The maximum burst size allowed in bytes if shaping is enabled on the port.	bigint	
PEAK_BANDWIDTH	The peak bandwidth during bursts in bits per second if traffic shaping is enabled on the port.	bigint	
VM_NETWORK_SETTINGS_ID	ID of network settings table.	int	
TYPE	Type of traffic shaping policy, whether ingress or egress. 0 is ingress, 1 is egress traffic shaping policy.	smallint	

TABLE 672 VM_VCENTER

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
HOST	The FQDN or the ip address of the host.	varchar	256
PORT	The port of the VCENTER server on the host.	int	
USER_NAME	The username to login into the VCENTER.	varchar	64
PASSWORD	The password to login into the VCENTER.	varchar	512
VERSION	The version of VCENTER.	varchar	10
TOKEN_ID	The id to map the each VCENTER on the host.	varchar	64
PLUGIN_STATUS	Status of Plug-in registration to the vCenter server.	varchar	32
PLUGIN_ENABLED	Whether plug-in enabled or disabled.	smallint	
PLUGIN_FORWARD_EVENTS	Whether to forward events from Management application to the vCenter server or not	smallint	
DISCOVERY_STATUS	vCenter server discovery status. Can be one of the below values: 1. Active 2. Failed - Authentication Failure 3. Failed - Not reachable	smallint	
DELETED_DISCOVERY	The vCenter server discovery has been deleted. Such a deleted vCenter server entry will not be discovered.	smallint	
MANAGED_ELEMENT_ID	A foreign key referencing MANAGED_ELEMENT(ID).	int	
FAULT_MONITORING_STATE	Flag to indicate whether fault monitoring is registered or not for a VM Host. Possible values are: 1.Not registered 2.Registered (Default)	smallint	

TABLE 672 VM_VCENTER (Continued)

Field	Definition	Format	Size
NAME	The name of the VCenter.	varchar	64
UUID	Unique identifier for vCenter server instance.	varchar	64

TABLE 673 VM_VCENTER_MEMBER

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
HOST_NAME	Hostname of VM host.	varchar	256
IP_ADDRESS	IP address of VM host.	varchar	128
STATUS	Discovery status of VM host. This can be one of the following: 1. Discovery Pending 2. Excluded 3. Conflict - Existing Host 4. Disconnected 5.Not responding.	smallint	
REASON	In case the status is 3 (Conflict - Existing host) then this field will be used to persist the hostname for conflicting user defined host.	varchar	1024
VM_VCENTER_ID	Id of the vCenter server managing this host.	int	
VM_HOST_ID	Foreign Key to the vm_host table.	int	

TABLE 674 VM_VIRTUAL_ETHERNET_ADAPTER

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
DISPLAY_LABEL	Display label for the virtual ethernet adapter.	varchar	256
DISPLAY_SUMMARY	Summary description.	varchar	256
KEY	This property is a unique key that distinguishes this device from other devices in the same virtual machine. Keys are immutable but may be recycled; that is, a key does not change as long as the device is associated with a particular virtual machine. However, once a device is removed, its key may be used when another device is added.	int	
ADDRESS_TYPE	MAC address type. Valid values for address type are: • Manual • Statically assigned MAC address. • Generated • Automatically generated MAC address. • Assigned • MAC address assigned by VirtualCenter.	smallint	
MAC_ADDRESS	MAC address assigned to the virtual network adapter. Clients can set this property to any of the allowed address types. The server might override the specified value for "Generated" or "Assigned" if it does not fall in the right ranges or is determined to be a duplicate.	varchar	64

TABLE 674 VM_VIRTUAL_ETHERNET_ADAPTER (Continued)

Field	Definition	Format	Size
WAKE_ON_LAN_ENABLED	Indicates whether wake-on-LAN is enabled on this virtual network adapter. Clients can set this property to selectively enable or disable wake-on-LAN.	smallint	
VIRTUAL_MACHINE_ID	Foreign Key to the vm_virtual_machine table. References the VM to which this vnic is attached.	int	
ADAPTER_TYPE	One of: • E1000 • Vmxnet • Pcnet32	smallint	
VM_STD_VSWITCH_PORT_GROUP_ID	Foreign Key to the VM_STD_VSWITCH_PORT_GROUP table. References the vSS port group to which the vnic may be associated with.	int	
VM_DV_PORT_ID	Foreign key to the vm_dv_port table. References dvPort to which this vnic is attached to.	int	
DV_PORT_KEY	The key of the port.	varchar	64
DV_PORT_GROUP_KEY	The key of portgroup.	varchar	64
DV_SWITCH_UUID	The UUID of the switch.	varchar	64
PORT_GROUP_NAME	The port group name.	varchar	256
MOR_ID	The managed object reference number assigned by the hypervisor.	int	
BINARY_MAC	MAC address in binary format.	bytea	
IP_ADDRESS	IPv4 address of VNIC.	varchar	32
BINARY_IP	IP address in binary format.	bytea	

TABLE 675 VM_VIRTUAL_MACHINE

Field	Definition	Format	Size
ID	Uniquely identifies the virtual machine	serial	
HOST_ID	Identifies the server that contains this VM	int	
HYPERVISOR_VM_ID	The VM number assigned by the hypervisor. Some hypervisors identify VMs by number as well as by name	int	
NAME	User-assigned name for the VM	varchar	80
DESCRIPTION	Optional user-entered notes describing the VM. (Annotation in VMware terminology.)	varchar	256
OS	Operating system name and version.	varchar	64
STATUS	VM status. 0 = stopped, 1 = running, 2 = suspended.	smallint	
VCPU_COUNT	Number of virtual CPUs used by the VM.	int	
CPU_RESOURCES	Summary of CPU resource configuration. Format may depend on VM vendor.	varchar	64

TABLE 675 VM_VIRTUAL_MACHINE (Continued)

Field	Definition	Format	Size
MEM_RESOURCES	Summary of memory resource configuration. Format may depend on VM vendor.	varchar	64
IP_ADDRESS	The primary IPv4 or IPv6 IP address used by the VM on the management LAN, if any. Primary is defined by the VM vendor.	varchar	32
HOSTNAME	The primary hostname assigned to this VM.	varchar	128
BOOT_TIME	The date and time the VM was last started.	timestamp	
DATASTORE_NAME	The user-assigned name for the VMs datastore. The datastore holds the VMs virtual disks, swap file, and configuration data.	varchar	80
DATASTORE_LOCATION	The location of the VMs datastore May be a SAN target disk or a locally-attached host disk folder. For VMware, this is a target LUN name.	varchar	64
NODE_WWN	The Node WWN for this VM. If NPIV is not being used, this will be the same as the Node WWN in the host''s DEVICE_ENCLOSURE record. If NPIV is being used, each VM has a unique Node WWN.	char	23
UUID		varchar	64
BINARY_IP	IP address in binary format.	bytea	
CONNECTION_STATE	The connectivity state of a virtual machine. • 0 = not available • 1 = connected • 2 = disconnected • 3 = inaccessible • 4 = invalid • 5 =orphaned	smallint	
COMMITTED_STORAGE	Used storage by a particular virtual machine.	varchar	64
UNCOMMITTED_STORAGE	Additional Provisioned storage for a particular virtual machine.	varchar	64
UNSHARED_STORAGE	Exclusive storage for a particular virtual machine.	varchar	64
INSTANCE_UUID	Instance Id that uniquely Identifies the Virtual Machine.	varchar	256

TABLE 676 VM_VIRTUAL_MACHINE_DATASTORE_MAP

Field	Definition	Format	Size
VM_DATASTORE_DETAILS_ID	A foreign key referencing VM_DATASTORE_DETAILS(ID).	int	
VIRTUAL_MACHINE_ID	A foreign key referencing VM_VIRTUAL_MACHINE(ID).	int	

TABLE 676 VM_VIRTUAL_MACHINE_DATASTORE_MAP (Continued)

Field	Definition	Format	Size
PROVISIONED_STORAGE	Additional storage space, in bytes, potentially used by the virtual machine on this datastore. Additional space may be needed for example when lazily allocated disks grow, or storage for swap is allocated when powering on the virtual machine.	bigint	
NOT_SHARED_STORAGE	Storage space, in bytes, occupied by the virtual machine on this datastore that is not shared with any other virtual machine.	bigint	
USED_STORAGE	Storage space, in bytes, on this datastore that is actually being used by the virtual machine.It includes space actually occupied by disks, logs, snapshots, configuration files etc. Files of the virtual machine which are present on a different datastore (e.g. a virtual disk on another datastore) are not included here.	bigint	

TABLE 677 VPLS_DEVICE_RELATION

Field	Definition	Format	Size
MPLS_SERVICE_DEVICE_RELATION_DB_ID	Database ID inherited from MPLS_SERVICE_DEVICE_RELATION.	int	
VPLS_CONFIG_INDEX	Represents the unique config index of VPLS endpoint.	int	
MAC_LIMIT	The maximum number of MAC address entries that can be learned for this VPLS Instance.	int	

TABLE 678 VPLS_ENDPOINT_RELATION

Field	Definition	Format	Size
MPLS_SERVICE_ENDPOINT_RELATION_DB_ID	Database ID inherited from MPLS_SERVICE_ENDPOINT_RELATION.	int	
ISID	The ISID value for that endpoint. Valid ISID value is between 256 (0x100) and 16777214 (0xFFFFFE). Default is 0 which indicates the endpoint is not configured with ISID.	int	

TABLE 679 VR_CONN_DOMAIN

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VCEM_PROFILE_ID	Foreign key references the ID of the VCEM server that the domain belongs to.	int	
VR_CONN_DOMAIN_GROUP_ID	Nullable foreign key references the ID of the domain group that the domain may belong to.	int	
VCEM_ASSIGNED_ID	The ID assigned by the VCEM server.	bigint	
GUID		varchar	512
NAME		varchar	256

TABLE 679 VR_CONN_DOMAIN (Continued)

Field	Definition	Format	Size
IP_ADDRESS		varchar	128
STATUS		varchar	256
FIRMWARE_VERSION		varchar	128
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	

TABLE 680 VR_CONN_DOMAIN_GROUP

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VCEM_PROFILE_ID	Foreign key references the ID of the VCEM server that the domain group belongs to.	int	
VCEM_ASSIGNED_ID	The ID assigned by the VCEM server.	bigint	
NAME		varchar	256
STATUS		varchar	256
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	

TABLE 681 VR_CONN_FC_CONNECTION

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VR_CONN_SERVER_PROFILE_ID	Foreign key references the ID of the server profile that the FC connection belongs to.	int	
PORT_NUMBER		smallint	
CONNECTION_BAY		smallint	
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	

TABLE 682 VR_CONN_MODULE

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VR_CONN_DOMAIN_ID	Foreign key references the domain ID that the module belongs to.	int	
VCEM_ASSIGNED_ID	The ID assigned by VCEM.	varchar	256
WWN	The WWN of the module.	char	23
PRODUCT_NAME	The product name of the module.	varchar	256
SERIAL_NUMBER	The serial number of the module.	varchar	32
STATUS	The current status of the module.	varchar	256
LAST_STATUS	The previous status of the module.	varchar	256
IO_BAY	The bay number of the module.	int	

TABLE 682 VR_CONN_MODULE (Continued)

Field	Definition	Format	Size
VENDOR	Subject to chnage. May not be able to differentiate module maker. Maker of the module. 0: unknown 1: Brocade 2: QLogic	int	
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	

TABLE 683 VR_CONN_MODULE_PORT

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VR_CONN_MODULE_ID	The ID of the module that the port belongs to.	int	
WWN	The WWN of the Virtual Connect port.	char	23
POSITION_	The port number of the port within the module.	smallint	
FABRIC_NAME	The fabric name of the VCEM.	varchar	256
SPEED		varchar	64
STATUS		varchar	64
LAST_STATUS		varchar	64
REMOTE_NODE_WWN	The WWN of the connected remote switch.	char	23
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	

TABLE 684 VR_CONN_SERVER_PROFILE

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
VCEM_PROFILE_ID	Foreign key references the ID of the VCEM server that the server profile belongs to.	int	
VR_CONN_DOMAIN_GROUP_ID	Nullable foreign key references the ID of the domain group that the server profile may belong to.	int	
VCEM_ASSIGNED_ID	The ID assigned by the VCEM server.	bigint	
NAME		varchar	256
BAY_NAME		varchar	256
BAY_NUMBER		smallint	
VIRTUAL_SERIAL_NUMBER		varchar	32
CREATION_TIME		timestamp	

TABLE 684 VR_CONN_SERVER_PROFILE (Continued)

Field	Definition	Format	Size
LAST_UPDATE_TIME		timestamp	
BAY_ENCLOSURE_UUID	The UUID extracted from the enclosure object inside the bay object inside the server profile. The value matches the domain GUID.	varchar	512

TABLE 685 VR_CONN_WWN

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
VR_CONN_FC_CONNECTION_ID	Foreign key references the ID of the FC connection that the WWN belongs to	int	
PORT_ADDRESS	Port WWN	char	23
NODE_ADDRESS	Node WWN	char	23
SAN_NAME		varchar	256
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	

TABLE 686 WIRELESS_PRODUCT_DETAILS

Field	Definition	Format	Size
ID	The primary key of the table.	int	
DEVICE_ID	Each AP or Controller has an entry in the table. No other device will have entries here. The foreign key reference to device table.	int	
CONTROLLER_DEVICE_ID	The reference to the APs controller in device table. If APs controller gets deleted, the value sets to null. If the entry is controller the value is null.	int	
PROFILE_NAME	Profile name that the AP is using.	varchar	64
RF_DOMAIN_NAME	RF domain name set for the AP.	varchar	64
TIME_ZONE	Time zone set for the AP.	varchar	80
COUNTRY	Country set for the AP.	varchar	32
VLAN_FOR_CONTROL_TRAFFIC	VLAN for control traffic set for the AP.	varchar	512
CLIENT_COUNT	Number of wireless clients or stations that connected or associated to the AP.	int	

TABLE 687 WIRELESS_PRODUCT_RELATION

Field	Definition	Format	Size
ID	The primary key of the table.	int	
AP_DEVICE_ID	The foreign key reference to device table for AP.	int	

TABLE 687 WIRELESS_PRODUCT_RELATION (Continued)

Field	Definition	Format	Size
AP_INTERFACE_ID	The reference to the AP interface in interface table. In case the AP interface is not found or discovered, the value is null.	int	
CONNECTED_SWITCH_INTERFACE_ID	The reference to the switch interface in interface table which connected to the AP.	int	

TABLE 688 WT_ARCHIVE

Field	Definition	Format	Size
FIRMWARE_VERSION	Firmware version for which jar files are downloaded	varchar	128
JAR_LIST	List of jar files as comma separated string	varchar	256

TABLE 689 ZONE_DB

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
FABRIC_ID	PK of the owning fabric.	int	
NAME	Zone DB name for offline Zone DBs.	varchar	256
OFFLINE	Offline Zone DB (1 = offline).	smallint	
CREATED	Created timestamp.	timestamp	
LAST_MODIFIED	Last modified timestamp.	timestamp	
LAST_APPLIED	Last saved to switch timestamp.	timestamp	
CREATED_BY	Created by user name.	varchar	128
LAST_MODIFIED_BY	Last modified by user name.	varchar	128
LAST_APPLIED_BY	Last saved to switch user name.	varchar	128
DEFAULT_ZONE_STATUS	All access or no access when no active zone configuration.	smallint	
ZONE_TXN_SUPPORTED	Zoning commands support transaction.	smallint	
MCDATA_DEFAULT_ZONE	McData switch default zoning mode.	smallint	
MCDATA_SAFE_ZONE	McData switch safe zoning mode.	smallint	
ZONE_CONFIG_SIZE	Zone configuration string length.	int	
ZONE_AVAILABLE_SIZE	Available zone DB size in the switch. Default value is -1.	int	

TABLE 690 ZONE_DB_CONFIG

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
ZONE_DB_ID	PK of the owning zone DB	int	
DEFINED_CONTENT	Defined zone raw config string, wrapped with $ to prevent special char trimming	text	

TABLE 690 ZONE_DB_CONFIG (Continued)

Field	Definition	Format	Size
ACTIVE_CONTENT	Active zone raw config string	text	
TI_ZONE_CONTENT	TI zone raw config string	text	

TABLE 691 ZONE_DB_CONTENT

Field	Definition	Format	Size
ID*		int	
ZONE_DB_ID	PK of the owning offline zone DB.	int	
CONTENT	Saved online content before offline was saved to switch.	text	
TI_CONTENT	TI_CONTENT saved online TI zone content before offline was saved to switch.	text	
DEFINED		text	
ACTIVE		text	

TABLE 692 ZONE_DB_USERS

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
ZONE_DB_ID	PK of the owning zone DB.	int	
USER_NAME	List of users currently editing this zone DB.	varchar	128

TABLE 693 ZONE_PID_DETAILS

Field	Definition	Format	Size
FABRIC_ID	Fabric which the zone alias/active zone present.	int	
CFG_TYPE	Zone configuration type possible values are ACTIVEZONE â€“ 1, ZONEALIAS â€“2.	int	
CFG_NAME	Name of active zone or zone alias based on CFG_TYPE.	varchar	128
DEVICE_PORT_LIST	List of device port identifiers which are members of this zone alias or active zone. If the member is D,P then device ports connected to this switch port identified by D,P is added to this list.	int	

TABLE 694 ZONE_TRANSACTION

Field	Definition	Format	Size
FABRIC_ID	The id of the fabric on which the zoning transaction is open. This is the primary key for this table and is a foreign key from the FABRIC table where ZONE_TR.FABRIC_ID == FABRIC.ID.	int	
USER_NAME	The dcfm username of the cimclient who has opened the zoning transaction on the fabric. This is a valid dcfm username in the dcfm db.	varchar	128
LAST_TIME_USED	The last time this transaction was used for.	timestamp	
LSAN_ZONING	1 if transaction is to carry out lsan zoning, 0 otherwise.	smallint	

Views

ADAPTER_PORT_CONFIG_INFO

```
create or replace view ADAPTER_PORT_CONFIG_INFO as
select
    ADAPTER_PORT_CONFIG.ID,
    ADAPTER_PORT_CONFIG.NAME as CONFIG_NAME,
    ADAPTER_PORT_CONFIG.TYPE as TYPE,
    ADAPTER_PORT_CONFIG_PROPERTY.NAME as PROPERTY_NAME,
    ADAPTER_PORT_CONFIG_DETAILS.VALUE as PROPERTY_VALUE
from
    ADAPTER_PORT_CONFIG,
    ADAPTER_PORT_CONFIG_DETAILS,
    ADAPTER_PORT_CONFIG_PROPERTY
where
    ADAPTER_PORT_CONFIG.ID = ADAPTER_PORT_CONFIG_DETAILS.CONFIG_ID
    and ADAPTER_PORT_CONFIG_PROPERTY.ID= ADAPTER_PORT_CONFIG_DETAILS.PROPERTY_ID;
```

AG_CONNECTION_INFO

```
create or replace view AG_CONNECTION_INFO as
select
    AG_N_PORT.VIRTUAL_SWITCH_ID as SOURCE_SWITCH_ID,
    AG_N_PORT.ID as SOURCE_PORT_ID,
    AG_N_PORT.WWN as SOURCE_PORT_WWN,
    AG_N_PORT.TYPE as SOURCE_PORT_TYPE,
    AG_N_PORT.USER_PORT_NUMBER as SOURCE_USER_PORT_NUMBER,
    EDGE_F_PORT.VIRTUAL_SWITCH_ID as DESTINATION_SWITCH_ID,
    EDGE_F_PORT.ID as DESTINATION_PORT_ID,
    EDGE_F_PORT.WWN as DESTINATION_PORT_WWN,
    EDGE_F_PORT.TYPE as DESTINATION_PORT_TYPE,
    EDGE_F_PORT.USER_PORT_NUMBER as DESTINATION_USER_PORT_NUMBER
from
    SWITCH_PORT AG_N_PORT,
    SWITCH_PORT EDGE_F_PORT
where
    ((AG_N_PORT.REMOTE_PORT_WWN = EDGE_F_PORT.WWN)
        or (AG_N_PORT.REMOTE_PORT_WWN = EDGE_F_PORT.LOGICAL_PORT_WWN
            and EDGE_F_PORT.TRUNK_MASTER = 1))
    and AG_N_PORT.TYPE = 'N-Port';
```

BIRTREPORT_SCHEDULE_INFO

```
CREATE OR REPLACE VIEW birtreport_schedule_info AS

 SELECT birtreport_schedule_config.id AS schedule_id,
    birtreport_schedule_config.name AS schedule_name,
    birtreport_schedule_config.report_store_location,
    birtreport_schedule_config.overwrite,
    birtreport_schedule_config.format_type,
    birtreport_schedule_config.created_by AS birtreport_schedule_config_created_by,
    birtreport_schedule_config.email_delivery,
    birtreport_schedule_config.folder_delivery,
    birtreport_schedule_config.email_recipients,
    birtreport_schedule_config.email_from,
```

```
    birtreport_schedule_config.email_replyto,
    birtreport_schedule_config.email_subject,
    birtreport_schedule_config.email_prologue,
    birtreport_schedule_config.email_epilogue,
    birtreport_schedule_config.created_time AS birtreport_schedule_config_created_time,
    birtreport_schedule_config.last_modified_time AS

birtreport_schedule_config_last_modified_time,
    birtreport_schedule_config.deployment_id,
    deployment_configuration.name AS deploment_configuration_name,
    deployment_configuration.configuration_type,
    deployment_configuration.deploy_option,
    deployment_configuration.deployment_handler_id,
    deployment_configuration.schedule_enabled,
    deployment_configuration.snapshot_enabled,
    deployment_configuration.cli_template_id,
    deployment_configuration.snapshot_setting,
    deployment_configuration.post_deployment_delay,
    deployment_configuration.created_by AS user_id_who_created_schedule,
    deployment_configuration.last_modified_by AS deplyment_config_last_modified_by,
    deployment_configuration.management_flag,
    deployment_configuration.description AS deployment_config_description,
    deployment_status.id AS deploment_status_id,
    deployment_status.deployment_time, deployment_status.status AS

deployment_status_status,
    deployment_status.status_message AS deployment_status_status_message,

deployment_status.trigger_source,
    schedule_entry.schedule_entry_id, schedule_entry.module,
    schedule_entry.user_id, schedule_entry.minutes, schedule_entry.hours,
    schedule_entry.week_days, schedule_entry.days, schedule_entry.months,
    schedule_entry.years, schedule_entry.type AS schedule_entry_type,
    schedule_entry.status AS status_from_schedule_entry,
    schedule_entry.duration, schedule_entry.table_name
   FROM birtreport_schedule_config
   JOIN deployment_configuration ON birtreport_schedule_config.deployment_id =

deployment_configuration.id
   LEFT JOIN deployment_status ON deployment_configuration.id =

deployment_status.deployment_configuration_id AND deployment_status.deployment_time =

(SELECT max(deployment_status.deployment_time) FROM deployment_status WHERE

deployment_status.deployment_configuration_id = deployment_configuration.id)
   LEFT JOIN schedule_entry ON schedule_entry.identity::text =

deployment_configuration.id::character varying(16)::text AND schedule_entry.table_name::text =

'DEPLOYMENT_CONFIGURATION'::text;
```

BOOT_IMAGE_FILE_DETAILS_INFO

```
create or replace view BOOT_IMAGE_FILE_DETAILS_INFO as
select
    BOOT_IMAGE_FILE_DETAILS.BOOT_IMAGE_NAME,
    BOOT_IMAGE_FILE_DETAILS.MAJOR_VERSION,
    BOOT_IMAGE_FILE_DETAILS.MINOR_VERSION,
    BOOT_IMAGE_FILE_DETAILS.MAINTENANCE,
    BOOT_IMAGE_FILE_DETAILS.PATCH,
    BOOT_IMAGE_FILE_DETAILS.IMPORTED_DATE,
```

```
    BOOT_IMAGE_FILE_DETAILS.RELEASE_DATE,
    BOOT_IMAGE_FILE_DETAILS.RELEASE_NOTES_LOCATION,
    BOOT_IMAGE_FILE_DETAILS.LOCATION,
    BOOT_IMAGE_DRIVER_MAP.SUPPORTED_DRIVERS
from
    BOOT_IMAGE_FILE_DETAILS,
    BOOT_IMAGE_DRIVER_MAP
where
    BOOT_IMAGE_FILE_DETAILS.DRIVER_MAPPING_ID= BOOT_IMAGE_DRIVER_MAP.ID;
```

CNA_ETH_PORT_CONFIG_INFO

```
create or replace view CNA_ETH_PORT_CONFIG_INFO as
select
    CNA_PORT.ID,
    CNA_PORT.PORT_NUMBER,
    CNA_PORT.PORT_WWN,
    CNA_PORT.NODE_WWN,
    CNA_PORT.PHYSICAL_PORT_TYPE,
    CNA_PORT.NAME,
    CNA_PORT.MAC_ADDRESS,
    CNA_PORT.MEDIA,
    CNA_PORT.CEE_STATE,
    CNA_PORT.HBA_ID,
    CNA_ETH_PORT_CONFIG.CNA_ETH_PORT_ID as CNA_ETH_PORT_ID,
    CNA_ETH_PORT_CONFIG.ID as CNA_ETH_PORT_CONFIG_ID,
    CNA_ETH_PORT_CONFIG.CURRENT_MAC_ADDRESS,
    CNA_ETH_PORT_CONFIG.MAX_BANDWIDTH,
    CNA_ETH_PORT_CONFIG.PCIF_INDEX,
    CNA_ETH_PORT_CONFIG.MIN_BANDWIDTH,
    CNA_ETH_PORT_CONFIG.PORT_NUMBER as ETH_PORT_CONFIG_PORT_NUMBER,
    CNA_ETH_PORT_CONFIG.PORT_TYPE,
    CNA_ETH_PORT_CONFIG.CONFIGURATION_STATUS
from
    CNA_PORT
        left outer join CNA_ETH_PORT_CONFIG on CNA_PORT.ID = CNA_ETH_PORT_CONFIG.CNA_PORT_ID;
```

CNA_PORT_DETAILS_INFO

```
create or replace view CNA_PORT_DETAILS_INFO as
select
    CNA_PORT.ID,
    CNA_PORT.PORT_NUMBER,
    CNA_PORT.PORT_WWN,
    CNA_PORT.NODE_WWN,
    CNA_PORT.PHYSICAL_PORT_TYPE,
    CNA_PORT.NAME,
    CNA_PORT.MAC_ADDRESS,
    CNA_PORT.MEDIA,
    CNA_PORT.CEE_STATE,
    CNA_PORT.HBA_ID,
    CNA_PORT.CREATION_TIME as CNA_PORT_CREATION_TIME,
    CNA_ETH_PORT.ID as ETH_PORT_ID,
    CNA_ETH_PORT.ETH_DEV,
    CNA_ETH_PORT.ETH_LOG_LEVEL,
    CNA_ETH_PORT.NAME as ETH_PORT_NAME,
    CNA_ETH_PORT.MAC_ADDRESS as ETH_MAC_ADDRESS,
    CNA_ETH_PORT.IOC_ID,
    CNA_ETH_PORT.HARDWARE_PATH,
    CNA_ETH_PORT.STATUS,
    CNA_ETH_PORT.CREATION_TIME as ETH_PORT_CREATION_TIME,
```

```
    CNA_ETH_PORT.CURRENT_MAC_ADDRESS as CURRENT_MAC_ADDRESS,
    CNA_ETH_PORT.MAX_BANDWIDTH,
    CNA_ETH_PORT.PCIF_INDEX,
    CNA_ETH_PORT.MAX_PCIF,
    CNA_ETH_PORT.MIN_BANDWIDTH,
    CNA_ETH_PORT.MTU,
    CNA_PORT.ALARM_WARNING
from
    CNA_PORT
        left outer join CNA_ETH_PORT on CNA_PORT.ID = CNA_ETH_PORT.CNA_PORT_ID;
```

CNA_PORT_INFO

```
create or replace view CNA_PORT_INFO as
select
    CNA_PORT.ID,
    CNA_PORT.PORT_NUMBER,
    CNA_PORT.PORT_WWN,
    CNA_PORT.NODE_WWN,
    CNA_PORT.PHYSICAL_PORT_TYPE,
    CNA_PORT.NAME,
    CNA_PORT.MAC_ADDRESS,
    CNA_PORT.MEDIA,
    CNA_PORT.CEE_STATE,
    CNA_PORT.HBA_ID,
    CNA_PORT.CREATION_TIME as CNA_PORT_CREATION_TIME,
    CNA_ETH_PORT.ID as ETH_PORT_ID,
    CNA_ETH_PORT.ETH_DEV,
    CNA_ETH_PORT.ETH_LOG_LEVEL,
    CNA_ETH_PORT.NAME as ETH_PORT_NAME,
    CNA_ETH_PORT.MAC_ADDRESS as ETH_MAC_ADDRESS,
    CNA_ETH_PORT.IOC_ID,
    CNA_ETH_PORT.HARDWARE_PATH,
    CNA_ETH_PORT.STATUS,
    CNA_ETH_PORT.CREATION_TIME as ETH_PORT_CREATION_TIME,
    HBA_PORT.DEVICE_PORT_ID,
    CNA_ETH_PORT.MTU,
    CNA_PORT.ALARM_WARNING
from
    CNA_PORT
        left outer join HBA_PORT on CNA_PORT.ID = HBA_PORT.CNA_PORT_ID
        left outer join CNA_ETH_PORT on CNA_PORT.ID = CNA_ETH_PORT.CNA_PORT_ID;
```

CORE_SWITCH_DETAILS_INFO

```
create or replace view CORE_SWITCH_DETAILS_INFO as
select
   CORE_SWITCH.ID,
   CORE_SWITCH.IP_ADDRESS,
   CORE_SWITCH.WWN,
   CORE_SWITCH.NAME,
   CORE_SWITCH.TYPE,
   CORE_SWITCH.MODEL,
   CORE_SWITCH.FIRMWARE_VERSION,
   CORE_SWITCH.VENDOR,
   CORE_SWITCH.MAX_VIRTUAL_SWITCHES,
   CORE_SWITCH.NUM_VIRTUAL_SWITCHES,
   CORE_SWITCH.REACHABLE,
   CORE_SWITCH.UNREACHABLE_TIME,
   CORE_SWITCH.OPERATIONAL_STATUS,
   CORE_SWITCH.CREATION_TIME,
```

```
    CORE_SWITCH.LAST_SCAN_TIME,
    CORE_SWITCH.LAST_UPDATE_TIME,
    CORE_SWITCH.SYSLOG_REGISTERED,
    CORE_SWITCH.CALL_HOME_ENABLED,
    CORE_SWITCH.SNMP_REGISTERED,
    CORE_SWITCH.USER_IP_ADDRESS,
    CORE_SWITCH.NIC_PROFILE_ID,
    CORE_SWITCH.MANAGING_SERVER_IP_ADDRESS,
    CORE_SWITCH.VF_ENABLED,
    CORE_SWITCH.VF_SUPPORTED,
    CORE_SWITCH.MANAGED_ELEMENT_ID as CORE_MANAGED_ELEMENT_ID,
    CORE_SWITCH_DETAILS.ETHERNET_MASK,
    CORE_SWITCH_DETAILS.FC_MASK,
    CORE_SWITCH_DETAILS.FC_IP,
    CORE_SWITCH_DETAILS.FC_CERTIFICATE,
    CORE_SWITCH_DETAILS.SW_LICENSE_ID,
    CORE_SWITCH_DETAILS.SUPPLIER_SERIAL_NUMBER,
    CORE_SWITCH_DETAILS.PART_NUMBER,
    CORE_SWITCH_DETAILS.CHECK_BEACON,
    CORE_SWITCH_DETAILS.TIMEZONE,
    CORE_SWITCH_DETAILS.MAX_PORT,
    CORE_SWITCH_DETAILS.CHASSIS_SERVICE_TAG,
    CORE_SWITCH_DETAILS.BAY_ID,
    CORE_SWITCH_DETAILS.TYPE_NUMBER,
    CORE_SWITCH_DETAILS.MODEL_NUMBER,
    CORE_SWITCH_DETAILS.MANUFACTURER,
    CORE_SWITCH_DETAILS.PLANT_OF_MANUFACTURER,
    CORE_SWITCH_DETAILS.SWITCH_SERIAL_NUMBER,
    CORE_SWITCH_DETAILS.ACT_CP_PRI_FW_VERSION,
    CORE_SWITCH_DETAILS.ACT_CP_SEC_FW_VERSION,
    CORE_SWITCH_DETAILS.STBY_CP_PRI_FW_VERSION,
    CORE_SWITCH_DETAILS.STBY_CP_SEC_FW_VERSION,
    CORE_SWITCH_DETAILS.EGM_CAPABLE,
    CORE_SWITCH_DETAILS.SUB_TYPE,
    CORE_SWITCH_DETAILS.PARTITION,
    CORE_SWITCH_DETAILS.MAX_NUM_OF_BLADES,
    CORE_SWITCH_DETAILS.VENDOR_VERSION,
    CORE_SWITCH_DETAILS.VENDOR_PART_NUMBER,
    CORE_SWITCH_DETAILS.RNID_SEQUENCE_NUMBER,
    CORE_SWITCH_DETAILS.CONTACT,
    CORE_SWITCH_DETAILS.LOCATION,
    CORE_SWITCH_DETAILS.DESCRIPTION,
    CORE_SWITCH_DETAILS.IP_ADDRESS_PREFIX,
    CORE_SWITCH_DETAILS.DOMAIN_NAME,
    CORE_SWITCH_DETAILS.FRAME_LOG_SIZE,
    CORE_SWITCH_DETAILS.FRAME_LOG_ENABLED,
    CORE_SWITCH_DETAILS.MAPS_ENABLED
from
    CORE_SWITCH LEFT OUTER JOIN CORE_SWITCH_DETAILS
    on CORE_SWITCH_DETAILS.CORE_SWITCH_ID = CORE_SWITCH.ID;
```

CRYPTO_HOST_LUN_INFO

```
create or replace view CRYPTO_HOST_LUN_INFO as
select
    LUN.CRYPTO_HOST_ID,
    LUN.ID CRYPTO_LUN_ID,
    LUN.LUN_NUMBER,
    LUN.CRYPTO_TARGET_CONTAINER_ID,
    LUN.SERIAL_NUMBER,
    LUN.ENCRYPTION_STATE,
    LUN.STATUS,
```

```
    LUN.REKEY_INTERVAL,
    LUN.VOLUME_LABEL_PREFIX,
    LUN.LAST_REKEY_DATE,
    LUN.LAST_REKEY_STATUS,
    LUN.LAST_REKEY_PROGRESS,
    LUN.CURRENT_VOLUME_LABEL,
    LUN.PRIOR_ENCRYPTION_STATE,
    LUN.ENCRYPTION_FORMAT,
    LUN.ENCRYPT_EXISTING_DATA,
    LUN.DECRYPT_EXISTING_DATA,
    LUN.KEY_ID,
    LUN.BLOCK_SIZE,
    LUN.TOTAL_BLOCKS,
    LUN.LUN_STATE,
    LUN.LUN_FLAGS,
    LUN.ENCRYPTION_ALGORITHM,
    LUN.KEY_ID_STATE,
    LUN.REKEY_SESSION_NUMBER,
    LUN.PERCENTAGE_COMPLETE,
    LUN.REKEY_ROLE,
    LUN.CURRENT_LBA,
    LUN.LUN_STATE_STRING,
    LUN.NEW_LUN,
    LUN.NEW_LUN_TYPE,
    LUN.DISABLE_WRITE_EARLY_ACK,
    LUN.DISABLE_READ_AHEAD,
    LUN.TIME_LEFT_FOR_AUTO_REKEY,
    CRYPTO_HOST.HOST_PORT_WWN,
    CRYPTO_HOST.HOST_NODE_WWN
     LUN.THIN_PROVISION_LUN
from
    CRYPTO_LUN LUN,
    CRYPTO_HOST
where
    LUN.CRYPTO_HOST_ID = CRYPTO_HOST.ID;
```

CRYPTO_TARGET_ENGINE_INFO

```
create or replace view CRYPTO_TARGET_ENGINE_INFO as
select
    CRYPTO_TARGET_CONTAINER.ID TARGET_CONTAINER_ID,
    CRYPTO_TARGET_CONTAINER.NAME,
    CRYPTO_TARGET_CONTAINER.VT_NODE_WWN,
    CRYPTO_TARGET_CONTAINER.VT_PORT_WWN,
    CRYPTO_TARGET_CONTAINER.FAILOVER_STATUS,
    CRYPTO_TARGET_CONTAINER.FAILOVER_STATUS_2,
    CRYPTO_TARGET_CONTAINER.DEVICE_STATUS,
    CRYPTO_TARGET_CONTAINER.DEVICE_TYPE,
    CRYPTO_TARGET_CONTAINER.TARGET_PORT_WWN,
    CRYPTO_TARGET_CONTAINER.TARGET_NODE_WWN,
    CRYPTO_TARGET_CONTAINER.CONTAINER_FIELD_DATA,
    CRYPTO_TARGET_CONTAINER.CONFIGURATION_STATUS,
    CRYPTO_TARGET_CONTAINER.FRONT_END_N_PORT_NUMBER,
    ENCRYPTION_ENGINE.STATUS ENCRYPTION_ENGINE_STATUS,
    ENCRYPTION_ENGINE.HA_CLUSTER_ID,
    ENCRYPTION_ENGINE.SYSTEM_CARD_STATUS,
    ENCRYPTION_ENGINE.WWN_POOLS_AVAILABLE,
    ENCRYPTION_ENGINE.STATE ENCRYPTION_ENGINE_STATE,
    ENCRYPTION_ENGINE.ID ENCRYPTION_ENGINE_ID,
    CRYPTO_SWITCH.SWITCH_ID SWITCH_ID,
    CRYPTO_SWITCH.ENCRYPTION_GROUP_ID ENCRYPTION_GROUP_ID
from
```

```
   CRYPTO_TARGET_CONTAINER,
   ENCRYPTION_ENGINE,
   CRYPTO_SWITCH
where
   CRYPTO_TARGET_CONTAINER.ENCRYPTION_ENGINE_ID = ENCRYPTION_ENGINE.ID
   and CRYPTO_SWITCH.SWITCH_ID = ENCRYPTION_ENGINE.SWITCH_ID;
```

DASHBOARD_PREFERENCES_INFO

```
create or replace view DASHBOARD_PREFERENCES_INFO as
select
   DASHBOARD.NAME as DASHBOARD_NAME,
   DASHBOARD.DESCRIPTION as DASHBOARD_DESC,
   DASHBOARD.CREATED_BY,
   DASHBOARD.INSTALLATION_TYPE,
   DASHBOARD.SHARED,
   DASHBOARD_CANVAS.NAME as CANVAS_NAME,
   DASHBOARD_CANVAS.DESCRIPTION as CANVAS_DESC,
   DASHBOARD_CANVAS_PREFERENCE.SCOPE_ID,
   DASHBOARD_CANVAS_PREFERENCE.SCOPE_TYPE,
   DASHBOARD_CANVAS_PREFERENCE.DASHBOARD_ID,
   DASHBOARD_CANVAS_PREFERENCE.DASHBOARD_CANVAS_ID,
   DASHBOARD_CANVAS_PREFERENCE.VISIBLE,
   DASHBOARD_CANVAS_PREFERENCE.TIME_SCOPE,
   DASHBOARD_CANVAS_PREFERENCE.USER_ID
from
  DASHBOARD,
  DASHBOARD_CANVAS,
  DASHBOARD_CANVAS_PREFERENCE
where
  DASHBOARD.ID = DASHBOARD_CANVAS_PREFERENCE.DASHBOARD_ID
  and DASHBOARD_CANVAS.ID = DASHBOARD_CANVAS_PREFERENCE.DASHBOARD_CANVAS_ID;
```

DEPLOYMENT_INFO

```
create or replace view DEPLOYMENT_INFO as
select DEPLOYMENT_CONFIGURATION.ID, DEPLOYMENT_CONFIGURATION.NAME,
    DEPLOYMENT_CONFIGURATION.DESCRIPTION, DEPLOYMENT_HANDLER.MODULE,
    DEPLOYMENT_HANDLER.SUB_MODULE, DEPLOYMENT_STATUS.DEPLOYMENT_TIME,
    DEPLOYMENT_CONFIGURATION.DEPLOY_OPTION as DEPLOYMENT_OPTION,
    DEPLOYMENT_STATUS.STATUS, DEPLOYMENT_STATUS.DEPLOYED_BY,
    DEPLOYMENT_CONFIGURATION.CREATED_BY as CREATOR,
    DEPLOYMENT_CONFIGURATION.SCHEDULE_ENABLED,
    DEPLOYMENT_CONFIGURATION.SNAPSHOT_ENABLED, SCHEDULE_ENTRY.TYPE as FREQUENCY,
    DEPLOYMENT_CONFIGURATION.MANAGEMENT_FLAG, DEPLOYMENT_HANDLER.PRIVILEGE_ID,
    DEPLOYMENT_HANDLER.HANDLER_CLASS,
    DEPLOYMENT_HANDLER.CLIENT_ACTION_HANDLER_CLASS,
    DEPLOYMENT_STATUS.ID as STATUS_ID, DEPLOYMENT_HANDLER.MODULE_DISPLAYNAME,
    DEPLOYMENT_REPORT_TEMPLATE.HEADER, DEPLOYMENT_REPORT_TEMPLATE.FOOTER,
    DEPLOYMENT_CONFIGURATION. IS_FCP_POLICY
   from DEPLOYMENT_CONFIGURATION
   join DEPLOYMENT_HANDLER on DEPLOYMENT_CONFIGURATION.DEPLOYMENT_HANDLER_ID = DEPLOYMENT_HANDLER.ID
   left join DEPLOYMENT_STATUS on DEPLOYMENT_STATUS.DEPLOYMENT_TIME = (( select
max(DEPLOYMENT_STATUS.DEPLOYMENT_TIME) as max
   from DEPLOYMENT_STATUS
  where DEPLOYMENT_STATUS.DEPLOYMENT_CONFIGURATION_ID = DEPLOYMENT_CONFIGURATION.ID))
   left join SCHEDULE_ENTRY on SCHEDULE_ENTRY.IDENTITY::TEXT = DEPLOYMENT_CONFIGURATION.ID::CHARACTER
VARYING(16)::TEXT and SCHEDULE_ENTRY.TABLE_NAME::TEXT = 'deployment_configuration'::TEXT
   left join DEPLOYMENT_REPORT_TEMPLATE on DEPLOYMENT_REPORT_TEMPLATE.DEPLOYMENT_HANDLER_ID =
DEPLOYMENT_HANDLER.ID;
```

DEPLOYMENT_LOG

```
create or replace view DEPLOYMENT_LOG as
select
    DEPLOYMENT_CONFIGURATION.ID,
    DEPLOYMENT_CONFIGURATION.NAME,
    DEPLOYMENT_CONFIGURATION.DESCRIPTION,
    DEPLOYMENT_HANDLER.MODULE,
    DEPLOYMENT_HANDLER.SUB_MODULE,
    DEPLOYMENT_STATUS.DEPLOYMENT_TIME,
    DEPLOYMENT_CONFIGURATION.DEPLOY_OPTION as DEPLOYMENT_OPTION,
    DEPLOYMENT_STATUS.STATUS, DEPLOYMENT_STATUS.DEPLOYED_BY,
    DEPLOYMENT_CONFIGURATION.CREATED_BY as CREATOR,
    DEPLOYMENT_CONFIGURATION.SCHEDULE_ENABLED,
    DEPLOYMENT_CONFIGURATION.SNAPSHOT_ENABLED,
    DEPLOYMENT_CONFIGURATION.MANAGEMENT_FLAG,
    DEPLOYMENT_HANDLER.PRIVILEGE_ID,
    DEPLOYMENT_HANDLER.HANDLER_CLASS,
    DEPLOYMENT_HANDLER.CLIENT_ACTION_HANDLER_CLASS,
    DEPLOYMENT_STATUS.ID as STATUS_ID,
    DEPLOYMENT_HANDLER.MODULE_DISPLAYNAME,
    DEPLOYMENT_STATUS.TRIGGER_SOURCE as TRIGGER_SOURCE,
    DEPLOYMENT_REPORT_TEMPLATE.HEADER,
    DEPLOYMENT_REPORT_TEMPLATE.FOOTER
from
    DEPLOYMENT_CONFIGURATION
        inner join DEPLOYMENT_HANDLER
            on DEPLOYMENT_CONFIGURATION.DEPLOYMENT_HANDLER_ID = DEPLOYMENT_HANDLER.ID
        inner join DEPLOYMENT_STATUS
            on DEPLOYMENT_CONFIGURATION.ID = DEPLOYMENT_STATUS.DEPLOYMENT_CONFIGURATION_ID
        left outer join DEPLOYMENT_REPORT_TEMPLATE on DEPLOYMENT_REPORT_TEMPLATE.DEPLOYMENT_HANDLER_ID =
DEPLOYMENT_HANDLER.ID;
```

DEVICE_CONNECTION_INFO

```
create or replace view DEVICE_CONNECTION_INFO as
select
     DEVICE_CONNECTION.ID,
          DEVICE_CONNECTION.FABRIC_ID,
    DEVICE_CONNECTION.DEVICE_PORT_ID,
          DEVICE_CONNECTION.SWITCH_PORT_ID,
    DEVICE_CONNECTION.AG_PORT_ID,
    COALESCE(DEVICE_ENCLOSURE_MEMBER.ENCLOSURE_ID, HBA.HOST_ID, VM_HOST.DEVICE_ENCLOSURE_ID) as
DEVICE_ENCLOSURE_ID,
    DEVICE_CONNECTION.CREATION_TIME,
          DEVICE_CONNECTION.LAST_UPDATED_TIME,
    DEVICE_PORT.NODE_ID,
          DEVICE_CONNECTION.MISSING,
    DEVICE_CONNECTION.MISSING_TIME,
          SWPORT.VIRTUAL_SWITCH_ID,
    DEVICE_CONNECTION.TRUSTED,
          AGPORT.VIRTUAL_SWITCH_ID as AG_SWITCH_ID,
    DEVICE_PORT.WWN as DEVICE_PORT_WWN,
    COALESCE(USERDEFINEDDETAILS.TYPE, DN.TYPE) as DEVICE_TYPE
from DEVICE_CONNECTION
    left join DEVICE_PORT on DEVICE_CONNECTION.DEVICE_PORT_ID = DEVICE_PORT.ID
    left join SWITCH_PORT SWPORT on DEVICE_CONNECTION.SWITCH_PORT_ID = SWPORT.ID
    left join SWITCH_PORT AGPORT on DEVICE_CONNECTION.AG_PORT_ID = AGPORT.ID
    left join HBA_PORT_DEVICE_PORT_MAP on DEVICE_PORT.ID = HBA_PORT_DEVICE_PORT_MAP.DEVICE_PORT_ID
    left join HBA_PORT on HBA_PORT_DEVICE_PORT_MAP.HBA_PORT_ID = HBA_PORT.DEVICE_PORT_ID
    left join HBA on HBA_PORT.HBA_ID = HBA.ID
    left join VM_FC_HBA_DEVICE_PORT_MAP ON VM_FC_HBA_DEVICE_PORT_MAP.DEVICE_PORT_ID = DEVICE_PORT.ID
```

```
    left join VM_FC_HBA ON VM_FC_HBA.ID = VM_FC_HBA_DEVICE_PORT_MAP.VM_FC_HBA_ID
    left join VM_HOST ON VM_HOST.DEVICE_ENCLOSURE_ID = VM_FC_HBA.VM_HOST_ID
    left join DEVICE_ENCLOSURE_MEMBER on DEVICE_PORT.ID = DEVICE_ENCLOSURE_MEMBER.DEVICE_PORT_ID
    left join DEVICE_NODE DN on DEVICE_PORT.NODE_ID = DN.ID
    left join USER_DEFINED_DEVICE_DETAIL USERDEFINEDDETAILS on DN.WWN = USERDEFINEDDETAILS.WWN;
```

EE_MONITOR_STATS_5MIN_INFO

```
create or replace view EE_MONITOR_STATS_5MIN_INFO as
select VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
  ME_ID,
  TARGET_ID as EE_MONITOR_ID,
  timestamp with time zone 'epoch' + TIME_IN_SECONDS * interval '1 second' as CREATION_TIME,
  sum(case when MEASURE_ID = 208 then value else 0 end) as TX_UTILIZATION,
  sum(case when MEASURE_ID = 209 then value else 0 end) as RX_UTILIZATION,
  sum(case when MEASURE_ID = 210 then value else 0 end) as CRC_ERRORS
  from TIME_SERIES_DATA_1, VIRTUAL_SWITCH
  where ME_ID = MANAGED_ELEMENT_ID and COLLECTOR_ID = 16
  group by ME_ID, TARGET_TYPE, TARGET_ID, TIME_IN_SECONDS,VIRTUAL_SWITCH_ID order by TIME_IN_SECONDS desc;
```

EE_MONITOR_STATS_30MIN_INFO

```
create or replace view EE_MONITOR_STATS_30MIN_INFO as
select VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
  ME_ID,
  TARGET_ID as EE_MONITOR_ID,
  timestamp with time zone 'epoch' + TIME_IN_SECONDS * interval '1 second' as CREATION_TIME,
  sum(case when MEASURE_ID = 208 then value else 0 end) as TX_UTILIZATION,
  sum(case when MEASURE_ID = 209 then value else 0 end) as RX_UTILIZATION,
  sum(case when MEASURE_ID = 210 then value else 0 end) as CRC_ERRORS
  from TIME_SERIES_DATA_1_30MIN, VIRTUAL_SWITCH
  where ME_ID = MANAGED_ELEMENT_ID and COLLECTOR_ID = 16
  group by ME_ID, TARGET_TYPE, TARGET_ID, TIME_IN_SECONDS,VIRTUAL_SWITCH_ID order by TIME_IN_SECONDS desc;
```

EE_MONITOR_STATS_2HOUR_INFO

```
create or replace view EE_MONITOR_STATS_2HOUR_INFO as
select VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
  ME_ID,
  TARGET_ID as EE_MONITOR_ID,
  timestamp with time zone 'epoch' + TIME_IN_SECONDS * interval '1 second' as CREATION_TIME,
  sum(case when MEASURE_ID = 208 then value else 0 end) as TX_UTILIZATION,
  sum(case when MEASURE_ID = 209 then value else 0 end) as RX_UTILIZATION,
  sum(case when MEASURE_ID = 210 then value else 0 end) as CRC_ERRORS
  from TIME_SERIES_DATA_1_2HOUR, VIRTUAL_SWITCH
  where ME_ID = MANAGED_ELEMENT_ID and COLLECTOR_ID = 16
  group by ME_ID, TARGET_TYPE, TARGET_ID, TIME_IN_SECONDS,VIRTUAL_SWITCH_ID order by TIME_IN_SECONDS desc;
```

EE_MONITOR_STATS_1DAY_INFO

```
create or replace view EE_MONITOR_STATS_1DAY_INFO as
select VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
  ME_ID,
  TARGET_ID as EE_MONITOR_ID,
  timestamp with time zone 'epoch' + TIME_IN_SECONDS * interval '1 second' as CREATION_TIME,
  sum(case when MEASURE_ID = 208 then value else 0 end) as TX_UTILIZATION,
  sum(case when MEASURE_ID = 209 then value else 0 end) as RX_UTILIZATION,
```

```
  sum(case when MEASURE_ID = 210 then value else 0 end) as CRC_ERRORS
  from TIME_SERIES_DATA_1_1DAY, VIRTUAL_SWITCH
  where ME_ID = MANAGED_ELEMENT_ID and COLLECTOR_ID = 16
  group by ME_ID, TARGET_TYPE, TARGET_ID, TIME_IN_SECONDS,VIRTUAL_SWITCH_ID order by TIME_IN_SECONDS desc;
```

TE_PORT_STATS_5MIN_INFO

```
create or replace view TE_PORT_STATS_5MIN_INFO as
select VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
  ME_ID,
  TARGET_ID as PORT_ID,
  timestamp with time zone 'epoch' + TIME_IN_SECONDS * interval '1 second' as CREATION_TIME,
  sum(case when MEASURE_ID = 193 then value else 0 end) as RECEIVE_OK_PERCENT_UTIL,
  sum(case when MEASURE_ID = 194 then value else 0 end) as TRANSMIT_OK_PERCENT_UTIL,
  sum(case when MEASURE_ID = 196 then value else 0 end) as RECEIVE_OK,
  sum(case when MEASURE_ID = 195 then value else 0 end) as TRANSMIT_OK,
  sum(case when MEASURE_ID = 36 then value else 0 end) as RECEIVE_EOF,
  sum(case when MEASURE_ID = 40 then value else 0 end) as UNDERFLOW_ERRORS,
  sum(case when MEASURE_ID = 41 then value else 0 end) as OVERFLOW_ERRORS,
  sum(case when MEASURE_ID = 43 then value else 0 end) as ALIGNMENT_ERRORS,
  sum(case when MEASURE_ID = 42 then value else 0 end) as RUNT_ERRORS,
  sum(case when MEASURE_ID = 43 then value else 0 end) as TOO_LONG_ERRORS,
  sum(case when MEASURE_ID = 39 then value else 0 end) as CRC_ERRORS
  from TIME_SERIES_DATA_1, VIRTUAL_SWITCH
  where ME_ID = MANAGED_ELEMENT_ID and COLLECTOR_ID = 12
  group by ME_ID, TARGET_TYPE, TARGET_ID, TIME_IN_SECONDS,VIRTUAL_SWITCH_ID order by TIME_IN_SECONDS desc;
```

TE_PORT_STATS_30MIN_INFO

```
create or replace view TE_PORT_STATS_30MIN_INFO as
select VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
  ME_ID,
  TARGET_ID as PORT_ID,
  timestamp with time zone 'epoch' + TIME_IN_SECONDS * interval '1 second' as CREATION_TIME,
  sum(case when MEASURE_ID = 193 then value else 0 end) as RECEIVE_OK_PERCENT_UTIL,
  sum(case when MEASURE_ID = 194 then value else 0 end) as TRANSMIT_OK_PERCENT_UTIL,
  sum(case when MEASURE_ID = 196 then value else 0 end) as RECEIVE_OK,
  sum(case when MEASURE_ID = 195 then value else 0 end) as TRANSMIT_OK,
  sum(case when MEASURE_ID = 36 then value else 0 end) as RECEIVE_EOF,
  sum(case when MEASURE_ID = 40 then value else 0 end) as UNDERFLOW_ERRORS,
  sum(case when MEASURE_ID = 41 then value else 0 end) as OVERFLOW_ERRORS,
  sum(case when MEASURE_ID = 43 then value else 0 end) as ALIGNMENT_ERRORS,
  sum(case when MEASURE_ID = 42 then value else 0 end) as RUNT_ERRORS,
  sum(case when MEASURE_ID = 43 then value else 0 end) as TOO_LONG_ERRORS,
  sum(case when MEASURE_ID = 39 then value else 0 end) as CRC_ERRORS
  from TIME_SERIES_DATA_1_30MIN, VIRTUAL_SWITCH
  where ME_ID = MANAGED_ELEMENT_ID and COLLECTOR_ID = 12
  group by ME_ID, TARGET_TYPE, TARGET_ID, TIME_IN_SECONDS,VIRTUAL_SWITCH_ID order by TIME_IN_SECONDS desc;
```

TE_PORT_STATS_2HOUR_INFO

```
create or replace view TE_PORT_STATS_2HOUR_INFO as
select VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
  ME_ID,
  TARGET_ID as PORT_ID,
  timestamp with time zone 'epoch' + TIME_IN_SECONDS * interval '1 second' as CREATION_TIME,
  sum(case when MEASURE_ID = 193 then value else 0 end) as RECEIVE_OK_PERCENT_UTIL,
  sum(case when MEASURE_ID = 194 then value else 0 end) as TRANSMIT_OK_PERCENT_UTIL,
  sum(case when MEASURE_ID = 196 then value else 0 end) as RECEIVE_OK,
```

```
  sum(case when MEASURE_ID = 195 then value else 0 end) as TRANSMIT_OK,
  sum(case when MEASURE_ID = 36 then value else 0 end) as RECEIVE_EOF,
  sum(case when MEASURE_ID = 40 then value else 0 end) as UNDERFLOW_ERRORS,
  sum(case when MEASURE_ID = 41 then value else 0 end) as OVERFLOW_ERRORS,
  sum(case when MEASURE_ID = 43 then value else 0 end) as ALIGNMENT_ERRORS,
  sum(case when MEASURE_ID = 42 then value else 0 end) as RUNT_ERRORS,
  sum(case when MEASURE_ID = 43 then value else 0 end) as TOO_LONG_ERRORS,
  sum(case when MEASURE_ID = 39 then value else 0 end) as CRC_ERRORS
  from TIME_SERIES_DATA_1_2HOUR, VIRTUAL_SWITCH
  where ME_ID = MANAGED_ELEMENT_ID and COLLECTOR_ID = 12
  group by ME_ID, TARGET_TYPE, TARGET_ID, TIME_IN_SECONDS,VIRTUAL_SWITCH_ID order by TIME_IN_SECONDS desc;
```

TE_PORT_STATS_1DAY_INFO

```
create or replace view TE_PORT_STATS_1DAY_INFO as
select VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
  ME_ID,
  TARGET_ID as PORT_ID,
  timestamp with time zone 'epoch' + TIME_IN_SECONDS * interval '1 second' as CREATION_TIME,
  sum(case when MEASURE_ID = 193 then value else 0 end) as RECEIVE_OK_PERCENT_UTIL,
  sum(case when MEASURE_ID = 194 then value else 0 end) as TRANSMIT_OK_PERCENT_UTIL,
  sum(case when MEASURE_ID = 196 then value else 0 end) as RECEIVE_OK,
  sum(case when MEASURE_ID = 195 then value else 0 end) as TRANSMIT_OK,
  sum(case when MEASURE_ID = 36 then value else 0 end) as RECEIVE_EOF,
  sum(case when MEASURE_ID = 40 then value else 0 end) as UNDERFLOW_ERRORS,
  sum(case when MEASURE_ID = 41 then value else 0 end) as OVERFLOW_ERRORS,
  sum(case when MEASURE_ID = 43 then value else 0 end) as ALIGNMENT_ERRORS,
  sum(case when MEASURE_ID = 42 then value else 0 end) as RUNT_ERRORS,
  sum(case when MEASURE_ID = 43 then value else 0 end) as TOO_LONG_ERRORS,
  sum(case when MEASURE_ID = 39 then value else 0 end) as CRC_ERRORS
  from TIME_SERIES_DATA_1_1DAY, VIRTUAL_SWITCH
  where ME_ID = MANAGED_ELEMENT_ID and COLLECTOR_ID = 12
  group by ME_ID, TARGET_TYPE, TARGET_ID, TIME_IN_SECONDS,VIRTUAL_SWITCH_ID order by TIME_IN_SECONDS desc;
```

SPX_PORT_DETAILS_INFO

```
create or replace view SPX_PORT_DETAILS_INFO as
select
PHYSICAL_INTERFACE.INTERFACE_ID,
SPX_PORT_DETAILS.PE_GROUP_NAME, PHYSICAL_INTERFACE.UNIT_NUMBER, INTERFACE.IDENTIFIER,
PHYSICAL_INTERFACE.IS_STACKING_INTERFACE, PHYSICAL_DEVICE.UNIT_ROLE, DEVICE.DEVICE_ID,
PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID, INTERFACE. IF_INDEX, SPX_PORT_DETAILS.CONNECTED_PE_UNIT_NUMBER
from PHYSICAL_DEVICE
right join PHYSICAL_INTERFACE on PHYSICAL_INTERFACE.PHYSICAL_DEVICE_ID  = PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID
right join DEVICE on DEVICE.DEVICE_ID = PHYSICAL_DEVICE.DEVICE_ID
right join INTERFACE on INTERFACE.INTERFACE_ID  = PHYSICAL_INTERFACE.INTERFACE_ID
left join SPX_PORT_DETAILS on SPX_PORT_DETAILS.DEVICE_ID  = PHYSICAL_DEVICE.DEVICE_ID and
SPX_PORT_DETAILS.INTERFACE_ID = PHYSICAL_INTERFACE.INTERFACE_ID
where PHYSICAL_INTERFACE.IS_STACKING_INTERFACE = 3
and PHYSICAL_DEVICE.DEVICE_ID = DEVICE.DEVICE_ID
and PHYSICAL_INTERFACE.PHYSICAL_DEVICE_ID = PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID;
```

SWITCH_INFO

```
create or replace view SWITCH_INFO as
select
   CORE_SWITCH.ID as PHYSICAL_SWITCH_ID,
   CORE_SWITCH.NAME as PHYSICAL_SWITCH_NAME,
```

```
CORE_SWITCH.IP_ADDRESS,
CORE_SWITCH.WWN as PHYSICAL_SWITCH_WWN,
CORE_SWITCH.OPERATIONAL_STATUS as PHYSICAL_OPERATIONAL_STATUS,
CORE_SWITCH.TYPE,
CORE_SWITCH.MAX_VIRTUAL_SWITCHES,
CORE_SWITCH.NUM_VIRTUAL_SWITCHES,
CORE_SWITCH.FIRMWARE_VERSION,
CORE_SWITCH.VENDOR,
CORE_SWITCH.REACHABLE,
CORE_SWITCH.UNREACHABLE_TIME,
CORE_SWITCH.MODEL,
CORE_SWITCH.SYSLOG_REGISTERED,
CORE_SWITCH.SNMP_REGISTERED,
CORE_SWITCH.CALL_HOME_ENABLED,
CORE_SWITCH.USER_IP_ADDRESS,
CORE_SWITCH.NIC_PROFILE_ID,
CORE_SWITCH.MANAGING_SERVER_IP_ADDRESS,
CORE_SWITCH.VF_ENABLED,
CORE_SWITCH.VF_SUPPORTED,
CORE_SWITCH.MANAGED_ELEMENT_ID as CORE_MANAGED_ELEMENT_ID,
CORE_SWITCH.NAT_PRIVATE_IP_ADDRESS,
CORE_SWITCH.ALTERNATE_IP_ADDRESS,
CORE_SWITCH.MAC_ADDRESS,
VIRTUAL_SWITCH.ID,
VIRTUAL_SWITCH.NAME,
VIRTUAL_SWITCH.OPERATIONAL_STATUS,
VIRTUAL_SWITCH.SWITCH_MODE,
VIRTUAL_SWITCH.AD_CAPABLE,
VIRTUAL_SWITCH.WWN,
VIRTUAL_SWITCH.ROLE,
VIRTUAL_SWITCH.FCS_ROLE,
VIRTUAL_SWITCH.DOMAIN_ID,
VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID,
VIRTUAL_SWITCH.BASE_SWITCH,
VIRTUAL_SWITCH.MAX_ZONE_CONFIG_SIZE,
VIRTUAL_SWITCH.CREATION_TIME,
VIRTUAL_SWITCH.LAST_UPDATE_TIME,
VIRTUAL_SWITCH.USER_NAME,
VIRTUAL_SWITCH.PASSWORD,
VIRTUAL_SWITCH.MANAGEMENT_STATE,
VIRTUAL_SWITCH.STATE,
VIRTUAL_SWITCH.STATUS,
VIRTUAL_SWITCH.STATUS_REASON,
VIRTUAL_SWITCH.FABRIC_IDID_MODE,
VIRTUAL_SWITCH.LOGICAL_ID,
VIRTUAL_SWITCH.USER_DEFINED_VALUE_1,
VIRTUAL_SWITCH.USER_DEFINED_VALUE_2,
VIRTUAL_SWITCH.USER_DEFINED_VALUE_3,
VIRTUAL_SWITCH.INTEROP_MODE,
VIRTUAL_SWITCH.CRYPTO_CAPABLE,
VIRTUAL_SWITCH.FCR_CAPABLE,
VIRTUAL_SWITCH.FCIP_CAPABLE,
VIRTUAL_SWITCH.LF_ENABLED,
VIRTUAL_SWITCH.FCOE_CAPABLE,
VIRTUAL_SWITCH.L2_CAPABLE,
VIRTUAL_SWITCH.L3_CAPABLE,
VIRTUAL_SWITCH.DEFAULT_LOGICAL_SWITCH,
VIRTUAL_SWITCH.FEATURES_SUPPORTED,
VIRTUAL_SWITCH.FMS_MODE,
VIRTUAL_SWITCH.DYNAMIC_LOAD_SHARING,
VIRTUAL_SWITCH.PORT_BASED_ROUTING,
VIRTUAL_SWITCH.IN_ORDER_DELIVERY,
VIRTUAL_SWITCH.INSISTENT_DID_MODE,
```

```
   VIRTUAL_SWITCH.PREVIOUS_OPERATIONAL_STATUS,
   VIRTUAL_SWITCH.LAST_SCAN_TIME,
   VIRTUAL_SWITCH.DOMAIN_MODE_239,
   VIRTUAL_SWITCH.DOMAIN_ID_OFFSET,
   VIRTUAL_SWITCH.DISCOVERED_PORT_COUNT,
   VIRTUAL_SWITCH.FCOE_LOGIN_ENABLED,
   VIRTUAL_SWITCH.LAST_PORT_MEMBERSHIP_CHANGE,
   VIRTUAL_SWITCH.FCIP_CIRCUIT_CAPABLE,
   VIRTUAL_SWITCH.MAX_FCIP_TUNNELS,
   VIRTUAL_SWITCH.MAX_FCIP_CIRCUITS,
   VIRTUAL_SWITCH.FCIP_LICENSED,
   VIRTUAL_SWITCH.ADDRESSING_MODE,
   VIRTUAL_SWITCH.PREVIOUS_STATE,
   VIRTUAL_SWITCH.MANAGED_ELEMENT_ID,
   VIRTUAL_SWITCH.HIF_ENABLED,
   VIRTUAL_SWITCH.AUTO_SNMP,
   VIRTUAL_SWITCH.RNID_SEQUENCE_NUMBER,
   VIRTUAL_SWITCH.VCS_ID,
   VIRTUAL_SWITCH.CLUSTER_TYPE,
   VIRTUAL_SWITCH.CLUSTER_MODE,
   VIRTUAL_SWITCH.RNID_TAG,
   VIRTUAL_SWITCH.SWITCH_ID,
   VIRTUAL_SWITCH.MONITORED,
   VIRTUAL_SWITCH.FEATURES_ENABLED,
   VIRTUAL_SWITCH.MAPS_ENABLED_ACTIONS,
   VIRTUAL_SWITCH.ROUTING_POLICY,
   VIRTUAL_SWITCH.FABRIC_STATUS,
   VIRTUAL_SWITCH.PROTOCOL,
   FABRIC_MEMBER.FABRIC_ID,
   FABRIC_MEMBER.TRUSTED,
   FABRIC_MEMBER.MISSING,
   FABRIC_MEMBER.MISSING_TIME,
   FABRIC.MANAGED as FABRIC_MANAGED,
  FABRIC.PRINCIPAL_SWITCH_WWN,
   FABRIC.SEED_SWITCH_WWN,
   FABRIC.TYPE as FABRIC_TYPE
VIRTUAL_SWITCH.BOUND,
   VIRTUAL_SWITCH.BOUND_BNA_IP_ADDRESS
from
   CORE_SWITCH,
   VIRTUAL_SWITCH,
   FABRIC_MEMBER,
   FABRIC
where
   VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
   and FABRIC_MEMBER.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
   and FABRIC_MEMBER.FABRIC_ID = FABRIC.ID;
```

SWITCH_REPORT_INFO

```
create or replace view SWITCH_REPORT_INFO as
SELECT core_switch.id AS physical_switch_id,
    core_switch.name AS physical_switch_name, core_switch.ip_address,
    core_switch.wwn AS physical_switch_wwn,
    core_switch.operational_status AS physical_operational_status,
    core_switch.type, core_switch.max_virtual_switches,
    core_switch.num_virtual_switches, core_switch.firmware_version,
    core_switch.vendor, core_switch.reachable, core_switch.unreachable_time,
    core_switch.model, core_switch.syslog_registered,
    core_switch.snmp_registered, core_switch.call_home_enabled,
    core_switch.user_ip_address, core_switch.nic_profile_id,
    core_switch.managing_server_ip_address, core_switch.vf_enabled,
```

```
    core_switch.vf_supported,
    core_switch.managed_element_id AS core_managed_element_id,
    core_switch.nat_private_ip_address, core_switch.alternate_ip_address,
    core_switch.mac_address, virtual_switch.id, virtual_switch.name,
    virtual_switch.operational_status, virtual_switch.switch_mode,
    virtual_switch.ad_capable, virtual_switch.wwn, virtual_switch.role,
    virtual_switch.fcs_role, virtual_switch.domain_id,
    virtual_switch.virtual_fabric_id, virtual_switch.base_switch,
    virtual_switch.max_zone_config_size, virtual_switch.creation_time,
    virtual_switch.last_update_time, virtual_switch.user_name,
    virtual_switch.password, virtual_switch.management_state,
    virtual_switch.state, virtual_switch.status, virtual_switch.status_reason,
    virtual_switch.fabric_idid_mode, virtual_switch.logical_id,
    virtual_switch.user_defined_value_1, virtual_switch.user_defined_value_2,
    virtual_switch.user_defined_value_3, virtual_switch.interop_mode,
    virtual_switch.crypto_capable, virtual_switch.fcr_capable,
    virtual_switch.fcip_capable, virtual_switch.lf_enabled,
    virtual_switch.fcoe_capable, virtual_switch.l2_capable,
    virtual_switch.l3_capable, virtual_switch.default_logical_switch,
    virtual_switch.features_supported, virtual_switch.fms_mode,
    virtual_switch.dynamic_load_sharing, virtual_switch.port_based_routing,
    virtual_switch.in_order_delivery, virtual_switch.insistent_did_mode,
    virtual_switch.previous_operational_status, virtual_switch.last_scan_time,
    virtual_switch.domain_mode_239, virtual_switch.domain_id_offset,
    virtual_switch.discovered_port_count, virtual_switch.fcoe_login_enabled,
    virtual_switch.last_port_membership_change,
    virtual_switch.fcip_circuit_capable, virtual_switch.max_fcip_tunnels,
    virtual_switch.max_fcip_circuits, virtual_switch.fcip_licensed,
    virtual_switch.addressing_mode, virtual_switch.previous_state,
    virtual_switch.managed_element_id, virtual_switch.hif_enabled,
    virtual_switch.auto_snmp, virtual_switch.rnid_sequence_number,
    virtual_switch.vcs_id, virtual_switch.cluster_type,
    virtual_switch.cluster_mode, virtual_switch.rnid_tag,
    virtual_switch.switch_id, virtual_switch.monitored,
    virtual_switch.features_enabled, virtual_switch.maps_enabled_actions,
    virtual_switch.routing_policy, fabric_member.fabric_id,
        CASE
            WHEN virtual_switch.switch_mode = 2 THEN 'Yes'::text
            ELSE 'No'::text
        END AS ag_mode,
    fabric_member.trusted, fabric_member.missing, fabric_member.missing_time,
    fabric.san_id AS fabric_san_id, fabric.name AS fabric_name,
    fabric.contact AS fabric_contact, fabric.location AS fabric_location, fabric.description AS
fabric_description,
    fabric.fabric_name AS fabric_fabric_name, fabric.active_zoneset_name AS fabric_active_zoneset_name,
    fabric.status AS fabric_status, fabric.bottleneck_status AS fabric_bottleneck_status,
    fabric.managed AS fabric_managed, fabric.principal_switch_wwn,
    fabric.seed_switch_wwn, fabric.type AS fabric_type,
    core_switch_details.model_number AS switchtype,
    core_switch_details.part_number, core_switch_details.max_port AS max_port_supported,
core_switch_details.type_number,
    core_switch_details.switch_serial_number, core_switch_details.type AS blade_type,
core_switch_details.sub_type,
    core_switch_details.partition AS partitions_supported, core_switch_details.max_num_of_blades,
core_switch_details.contact,
    core_switch_details.location, core_switch_details.description, core_switch_details.firmware_version AS
core_switch_details_firmware_version,
    core_switch_details.vendor_version, core_switch_details.vendor_part_number,
    ( SELECT count(switch_port.id) AS count
           FROM switch_port
```

```
         WHERE switch_port.virtual_switch_id = virtual_switch.id AND switch_port.licensed = 1 and
SWITCH_PORT.KIND::TEXT <> 'ICL'::text and SWITCH_PORT.PHYSICAL_PORT = 1 AND ((switch_port.type::text = ANY
(ARRAY['E-Port'::character varying, 'G-Port'::character varying, 'U-Port'::character varying,
'F-Port'::character varying, 'L-Port'::character varying, 'EX-Port'::character varying, 'FL-Port'::character
varying, 'SIM-Port'::character varying, 'N-Port'::character varying]::text[])) OR switch_port.type
like'LB-Port%')) AS switch_port_count,
    ( SELECT count(device_node.id) AS count
           FROM device_node, device_port, switch_port
          WHERE device_node.id = device_port.node_id AND device_port.switch_port_wwn = switch_port.wwn AND
switch_port.licensed = 1 AND switch_port.virtual_switch_id = virtual_switch.id AND device_node.type::text ~~
'Initiator'::text) AS initiatorcount,
    ( SELECT count(device_node.id) AS count
           FROM device_node, device_port, switch_port
          WHERE device_node.id = device_port.node_id AND device_port.switch_port_wwn = switch_port.wwn AND
switch_port.licensed = 1 AND switch_port.virtual_switch_id = virtual_switch.id AND device_node.type::text ~~
'Target'::text) AS tagetcount,
    ( SELECT count(device_node.id) AS count
           FROM device_node, device_port, switch_port
          WHERE device_node.id = device_port.node_id AND device_port.switch_port_wwn = switch_port.wwn AND
switch_port.licensed = 1 AND switch_port.virtual_switch_id = virtual_switch.id AND device_node.type::text ~~
'Initiator+Target'::text) AS initiatortargetcount,
    ( SELECT count(device_node.id) AS count
           FROM device_node, device_port, switch_port
          WHERE device_node.id = device_port.node_id AND device_port.switch_port_wwn = switch_port.wwn AND
switch_port.licensed = 1 AND switch_port.virtual_switch_id = virtual_switch.id AND device_node.type::text ~~
'Unknown'::text) AS unknowncount,
     ( SELECT count(switch_port.id) AS count
           FROM switch_port
          WHERE switch_port.virtual_switch_id = virtual_switch.id AND switch_port.licensed = 1 and
SWITCH_PORT.KIND::text <> 'ICL'::text and SWITCH_PORT.PHYSICAL_PORT = 1 AND ((switch_port.type::text = ANY
(ARRAY['E-Port'::character varying, 'G-Port'::character varying, 'U-Port'::character varying,
'F-Port'::character varying, 'L-Port'::character varying, 'EX-Port'::character varying, 'FL-Port'::character
varying, 'SIM-Port'::character varying, 'N-Port'::character varying]::text[]))OR switch_port.type
like'LB-Port%') AND switch_port.occupied = 1 ) AS occupied_switch_port_count,
    ( SELECT count(switch_port.id) AS count
           FROM switch_port
          WHERE switch_port.virtual_switch_id = virtual_switch.id AND switch_port.licensed = 1 and
SWITCH_PORT.KIND::text <> 'ICL'::text and SWITCH_PORT.PHYSICAL_PORT = 1 AND ((switch_port.type::text = ANY
(ARRAY['E-Port'::character varying, 'G-Port'::character varying, 'U-Port'::character varying,
'F-Port'::character varying, 'L-Port'::character varying, 'EX-Port'::character varying, 'FL-Port'::character
varying, 'SIM-Port'::character varying, 'N-Port'::character varying]::text[]))OR switch_port.type
like'LB-Port%') AND switch_port.occupied = 0 ) AS free_switch_port_count,
    ( SELECT count(switch_port.id) AS count
           FROM switch_port
          WHERE switch_port.virtual_switch_id = virtual_switch.id AND switch_port.licensed = 0 AND
((switch_port.type::text = ANY (ARRAY['E-Port'::character varying, 'G-Port'::character varying,
'U-Port'::character varying, 'F-Port'::character varying, 'L-Port'::character varying, 'EX-Port'::character
varying, 'FL-Port'::character varying, 'SIM-Port'::character varying, 'N-Port'::character
varying]::text[]))OR switch_port.type like'LB-Port%')) AS unlicensed_switch_port_count
   FROM fabric, fabric_member, virtual_switch,
    core_switch
   LEFT JOIN core_switch_details ON core_switch_details.core_switch_id = core_switch.id
  WHERE virtual_switch.core_switch_id = core_switch.id AND core_switch.type NOT IN(40,41) AND
virtual_switch.monitored = 1 AND (virtual_switch.switch_mode = ANY (ARRAY[0, 2])) AND
fabric_member.virtual_switch_id = virtual_switch.id AND fabric_member.fabric_id = fabric.id AND
fabric.managed = 1 AND (fabric.type <> ALL (ARRAY[65, 66, 4]));
```

SWITCH_PORT_DETAILS_INVENTORY_INFO

```
create or replace view SWITCH_PORT_DETAILS_INVENTORY_INFO as
select
```

```
  local_switch_port.ID as SWITCH_PORT_DB_ID,
  local_switch_port.VIRTUAL_SWITCH_ID,
  local_switch_port.CATEGORY,
  (case when local_switch_port.CATEGORY=2 then local_gige_port.PORT_NAME
        else local_switch_port.NAME end) as SWITCH_PORT_NAME,
  local_switch_port.WWN,
  local_gige_port.MAC_ADDRESS,
  local_switch_port.TYPE as PORT_TYPE,
  local_switch_port.SPEED as PORT_SPEED,
  local_switch_port.STATUS as PORT_STATUS,
  local_switch_port.STATUS_MESSAGE as PORT_STATUS_MESSAGE,
  local_switch_port.DISABLED_REASON as PORT_DISABLED_REASON,
  local_switch_port.STATE as PORT_STATE,
  local_switch_port.PHYSICAL_PORT AS PHYSICAL_PORT,
  local_switch_port.EXT_TYPE as SWITCH_PORT_EXT_TYPE,
  coalesce(nullif(remote_switch_isl_port.NAME::text, ''::text),
                  remote_switch_isl_port.WWN::text,
           nullif(remote_switch_san_conn_port.NAME::text, ''::text),
                 remote_switch_san_conn_port.WWN::text,
           nullif(remote_ag_switch_conn_port.NAME::text, ''::text),
                  remote_ag_switch_conn_port.WWN::text,
           nullif(udd_port.NAME::text, ''::text),
                  remote_device_port.WWN::text, ''::character varying::text) AS REMOTE_DEVICE_PORT_NAME,
  coalesce(nullif(remote_switch.NAME::text, ''::text),
                  remote_switch.WWN::text,
           nullif(udd_node.NAME::text, ''::text),
                  remote_device_node.wwn::text, ''::character varying::text) AS REMOTE_DEVICE_NAME,
  (case when (local_virtual_switch.SWITCH_MODE != 4 and local_switch_port.type='AE-Port') then 'Analytics
Monitoring Platform'
       when coalesce(remote_switch_isl_port.ID,remote_switch_san_conn_port.ID) IS NOT NULL THEN
'Switch'::character varying
       when nullif(switch_ag_connection.DESTINATION_PORT_ID,-1) is not null then 'Access Gateway'::character
varying
      else coalesce(udd_node.TYPE, udd_port.TYPE, remote_device_node.TYPE,''::character varying) end) as
CONNECTED_DEVICE_TYPE,
  remote_device_port.ID as REMOTE_DEVICE_PORT_ID,
  coalesce(remote_switch_isl_port.ID,remote_switch_san_conn_port.ID) AS REMOTE_SWITCH_PORT_ID,
  remote_ag_switch_conn_port.ID AS REMOTE_AG_PORT_ID,
  coalesce(remote_device_port.NPV_PHYSICAL,0) as NPIV_PHYSICAL,
  coalesce(local_switch_port.PROTOCOL, ''::character varying) as PROTOCOL
from SWITCH_PORT local_switch_port
  join VIRTUAL_SWITCH local_virtual_switch
               on local_switch_port.VIRTUAL_SWITCH_ID = local_virtual_switch.ID
  join FABRIC_MEMBER local_fabric_member
               on local_virtual_switch.ID = local_fabric_member.VIRTUAL_SWITCH_ID
  join FABRIC local_fabric on local_fabric.ID = local_fabric_member.FABRIC_ID
  left join GIGE_PORT local_gige_port on local_switch_port.ID = local_gige_port.SWITCH_PORT_ID
  left join ISL_CONNECTION on local_switch_port.ID = ISL_CONNECTION.SOURCE_SWITCH_PORT_ID
  left join SAN_CONNECTION on local_switch_port.ID = SAN_CONNECTION.SOURCE_PORT_ID
  left join SAN_CONNECTION switch_ag_connection on local_switch_port.ID =
switch_ag_connection.DESTINATION_PORT_ID
  left join DEVICE_CONNECTION on (local_switch_port.ID = DEVICE_CONNECTION.SWITCH_PORT_ID and
DEVICE_CONNECTION.AG_PORT_ID = -1)
  left join DEVICE_CONNECTION device_ag_connection on local_switch_port.ID = device_ag_connection.AG_PORT_ID
  left join SWITCH_PORT remote_switch_isl_port on ISL_CONNECTION.TARGET_SWITCH_PORT_ID =
remote_switch_isl_port.ID
  left join SWITCH_PORT remote_switch_san_conn_port on SAN_CONNECTION.DESTINATION_PORT_ID =
remote_switch_san_conn_port.ID
  left join SWITCH_PORT remote_ag_switch_conn_port on switch_ag_connection.SOURCE_PORT_ID =
remote_ag_switch_conn_port.ID
  left join VIRTUAL_SWITCH remote_switch on (remote_switch_isl_port.VIRTUAL_SWITCH_ID = remote_switch.ID
                                        or remote_switch_san_conn_port.VIRTUAL_SWITCH_ID = remote_switch.ID
                                        or remote_ag_switch_conn_port.VIRTUAL_SWITCH_ID = remote_switch.ID)
```

```
  left join DEVICE_PORT remote_device_port on (DEVICE_CONNECTION.DEVICE_PORT_ID = remote_device_port.ID
                                               or device_ag_connection.DEVICE_PORT_ID = remote_device_port.ID)
  left join DEVICE_NODE remote_device_node  on remote_device_port.NODE_ID = remote_device_node.ID
  left join USER_DEFINED_DEVICE_DETAIL udd_port on remote_device_port.WWN = udd_port.WWN
  left join USER_DEFINED_DEVICE_DETAIL udd_node  on remote_device_node.WWN = udd_node.WWN
where local_switch_port.FAKE_PORT = 0 and local_virtual_switch.MONITORED = 1 and local_fabric.MANAGED = 1 and
(local_virtual_switch.SWITCH_MODE = any (array[0, 2,4]));
```

DEVICE_INFO

```
create or replace view DEVICE_INFO as
select distinct
  DEVICE_NODE.ID as DEVICE_NODE_ID,
  DEVICE_NODE.WWN as DEVICE_NODE_WWN,
  DEVICE_NODE.TYPE as DEVICE_NODE_TYPE,
  DEVICE_NODE.SYMBOLIC_NAME as DEVICE_NODE_SYMBOLIC_NAME,
  DEVICE_NODE.DEVICE_TYPE,
  DEVICE_NODE.FDMI_HOST_NAME,
  DEVICE_NODE.VENDOR,
  DEVICE_NODE.CAPABILITY_,
  DEVICE_NODE.AG,
  DEVICE_NODE.SIMULATED,
  DEVICE_PORT.ID as DEVICE_PORT_ID,
  DEVICE_PORT.DOMAIN_ID as DEVICE_PORT_DOMAIN_ID,
  DEVICE_PORT.WWN as DEVICE_PORT_WWN,
  DEVICE_PORT.NUMBER,
  DEVICE_PORT.PORT_ID,
  DEVICE_PORT.TYPE as DEVICE_PORT_TYPE,
  DEVICE_PORT.SYMBOLIC_NAME as DEVICE_PORT_SYMBOLIC_NAME,
  DEVICE_PORT.FC4_TYPE,
  DEVICE_PORT.IP_PORT,
  DEVICE_PORT.HARDWARE_ADDRESS,
  DEVICE_PORT.TRUSTED as DEVICE_PORT_TRUSTED,
  DEVICE_PORT.MISSING as DEVICE_PORT_MISSING,
  DEVICE_PORT.COS,
  DEVICE_PORT.NPV_PHYSICAL,
  SWITCH_PORT.ID as SWITCH_PORT_ID,
  SWITCH_PORT.WWN as SWITCH_PORT_WWN,
  SWITCH_PORT.NAME as SWITCH_PORT_NAME,
  SWITCH_PORT.SLOT_NUMBER,
  SWITCH_PORT.PORT_NUMBER,
  SWITCH_PORT.PORT_INDEX,
  SWITCH_PORT.TYPE as SWITCH_PORT_TYPE,
  SWITCH_PORT.FULL_TYPE as SWITCH_PORTFULL_TYPE,
  SWITCH_PORT.EXT_TYPE as SWITCH_PORT_EXT_TYPE,
  SWITCH_PORT.STATUS as SWITCH_PORT_STATUS,
  SWITCH_PORT.HEALTH as SWITCH_PORT_HEALTH,
  SWITCH_PORT.SPEED,
  SWITCH_PORT.MAX_PORT_SPEED,
  SWITCH_PORT.NPIV,
  SWITCH_PORT.NPIV_CAPABLE,
  SWITCH_PORT.CALCULATED_STATUS,
  SWITCH_PORT.AREA_ID,
  SWITCH_PORT.PHYSICAL_PORT,
  SWITCH_PORT.CATEGORY,
  SWITCH_PORT.PERSISTENT_DISABLE,
  SWITCH_PORT.BLOCKED,
  SWITCH_PORT.FCR_INTEROP_MODE,
  SWITCH_PORT.SPEED_TYPE,
  SWITCH_INFO.IP_ADDRESS,
```

```
  SWITCH_INFO.PHYSICAL_SWITCH_WWN,
  SWITCH_INFO.FIRMWARE_VERSION,
  SWITCH_INFO.REACHABLE,
  SWITCH_INFO.SYSLOG_REGISTERED,
  SWITCH_INFO.SNMP_REGISTERED,
  SWITCH_INFO.ID as VIRTUAL_SWITCH_ID,
  SWITCH_INFO.NAME as VIRTUAL_SWITCH_NAME,
  SWITCH_INFO.OPERATIONAL_STATUS,
  SWITCH_INFO.SWITCH_MODE,
  SWITCH_INFO.WWN as VIRTUAL_SWITCH_WWN,
  SWITCH_INFO.DOMAIN_ID as VIRTUAL_SWITCH_DOMAIN_ID,
  SWITCH_INFO.VIRTUAL_FABRIC_ID,
  SWITCH_INFO.BASE_SWITCH,
  SWITCH_INFO.STATE as VIRTUAL_SWITCH_STATE,
  SWITCH_INFO.STATUS as VIRTUAL_SWITCH_STATUS,
  SWITCH_INFO.FABRIC_ID,
  SWITCH_INFO.MONITORED,
  SWITCH_INFO.CRYPTO_CAPABLE
from
  DEVICE_NODE, DEVICE_PORT, SWITCH_PORT, SWITCH_INFO
where
  DEVICE_PORT.NODE_ID = DEVICE_NODE.ID and
  DEVICE_PORT.SWITCH_PORT_WWN = SWITCH_PORT.WWN and
  SWITCH_PORT.VIRTUAL_SWITCH_ID = SWITCH_INFO.ID and
  DEVICE_NODE.FABRIC_ID = SWITCH_INFO.FABRIC_ID;
```

N2F_PORT_MAP_INFO

```
create or replace view N2F_PORT_MAP_INFO as
select
    N2F_PORT_MAP.VIRTUAL_SWITCH_ID,
    N2F_PORT_MAP.N_PORT,
    N2F_PORT_MAP.F_PORT,
    AG_N_PORT.REMOTE_PORT_WWN as EDGE_SWITCH_PORT_WWN,
    AG_N_PORT.WWN as AG_N_PORT_WWN,
    AG_F_PORT.WWN as AG_F_PORT_WWN,
    AG_F_PORT.REMOTE_NODE_WWN,
    AG_F_PORT.REMOTE_PORT_WWN as DEVICE_PORT_WWN
from
    N2F_PORT_MAP,
    SWITCH_PORT AG_N_PORT,
    SWITCH_PORT AG_F_PORT,
    VIRTUAL_SWITCH AG_SWITCH
where
    N2F_PORT_MAP.VIRTUAL_SWITCH_ID = AG_N_PORT.VIRTUAL_SWITCH_ID
    and N2F_PORT_MAP.N_PORT = AG_N_PORT.USER_PORT_NUMBER
    and N2F_PORT_MAP.VIRTUAL_SWITCH_ID = AG_F_PORT.VIRTUAL_SWITCH_ID
    and N2F_PORT_MAP.F_PORT = AG_F_PORT.USER_PORT_NUMBER
    and AG_N_PORT.VIRTUAL_SWITCH_ID = AG_SWITCH.ID
    and AG_SWITCH.MONITORED = 1;
```

DEVICE_NODE_INFO

```
create or replace view DEVICE_NODE_INFO as
select
    DEVICE_NODE.ID,
    DEVICE_NODE.FABRIC_ID,
    DEVICE_NODE.WWN,
    DEVICE_NODE.TYPE,
    DEVICE_NODE.DEVICE_TYPE,
    DEVICE_NODE.SYMBOLIC_NAME,
```

```
    DEVICE_NODE.FDMI_HOST_NAME,
    DEVICE_NODE.VENDOR,
    DEVICE_NODE.CAPABILITY_,
    DEVICE_NODE.TRUSTED,
    DEVICE_NODE.CREATION_TIME,
    DEVICE_NODE.MISSING,
    DEVICE_NODE.MISSING_TIME,
    DEVICE_NODE.PROXY_DEVICE,
    DEVICE_NODE.AG,
    DEVICE_NODE.PREVIOUS_MISSING_STATE,
    USER_DEFINED_DEVICE_DETAIL.NAME,
    USER_DEFINED_DEVICE_DETAIL.TYPE as USER_DEFINED_TYPE,
    USER_DEFINED_DEVICE_DETAIL.IP_ADDRESS,
    USER_DEFINED_DEVICE_DETAIL.CONTACT,
    USER_DEFINED_DEVICE_DETAIL.LOCATION,
    USER_DEFINED_DEVICE_DETAIL.DESCRIPTION,
    USER_DEFINED_DEVICE_DETAIL.USER_DEFINED_VALUE1,
    USER_DEFINED_DEVICE_DETAIL.USER_DEFINED_VALUE2,
    USER_DEFINED_DEVICE_DETAIL.USER_DEFINED_VALUE3,
    FABRIC.PRINCIPAL_SWITCH_WWN as PRINCIPAL_WWN,
    DEVICE_FDMI_DETAILS.SERIAL_NUMBER AS FDMI_SERIAL_NUMBER,
    DEVICE_FDMI_DETAILS.FIRMWARE_VERSION AS FDMI_FIRMWARE_VERSION,
    DEVICE_FDMI_DETAILS.DRIVER_VERSION AS FDMI_DRIVER_VERSION,
    DEVICE_FDMI_DETAILS.MANUFACTURER AS FDMI_MANUFACTURER,
    DEVICE_FDMI_DETAILS.MODEL AS FDMI_MODEL,
    DEVICE_FDMI_DETAILS.HARDWARE_VERSION AS FDMI_HARDWARE_VERSION,
    DEVICE_FDMI_DETAILS.MODEL_DESCRIPTION AS FDMI_MODEL_DESCRIPTION,
    DEVICE_FDMI_DETAILS.NODE_NAME AS FDMI_NODE_NAME
from
    DEVICE_NODE
        left outer join USER_DEFINED_DEVICE_DETAIL
            on DEVICE_NODE.WWN = USER_DEFINED_DEVICE_DETAIL.WWN
        left outer join FABRIC
            on DEVICE_NODE.FABRIC_ID = FABRIC.ID
        left outer join DEVICE_FDMI_DETAILS
            on DEVICE_NODE.ID = DEVICE_FDMI_DETAILS.DEVICE_NODE_ID;
```

DEVICE_PORT_INFO

```
CREATE VIEW device_port_info AS
    select
    DEVICE_PORT.ID,
    DEVICE_PORT.NODE_ID,
    DEVICE_PORT.DOMAIN_ID,
    DEVICE_PORT.WWN,
    DEVICE_PORT.SWITCH_PORT_WWN,
    DEVICE_PORT.NUMBER,
    DEVICE_PORT.PORT_ID,
    DEVICE_PORT.TYPE,
    DEVICE_PORT.SYMBOLIC_NAME,
    DEVICE_PORT.FC4_TYPE,
    DEVICE_PORT.COS,
    DEVICE_PORT.IP_PORT,
    DEVICE_PORT.HARDWARE_ADDRESS,
    DEVICE_PORT.TRUSTED,
    DEVICE_PORT.CREATION_TIME,
    DEVICE_PORT.MISSING,
    DEVICE_PORT.MISSING_TIME,
    DEVICE_PORT.NPV_PHYSICAL,
    DEVICE_PORT.EDGE_SWITCH_PORT_WWN,
    DEVICE_PORT.LOGGED_TO_AG,
    DEVICE_PORT.AG_NODE_WWN,
```

```
    DEVICE_PORT.AG_N_PORT_WWN,
    DEVICE_PORT.MISSING_REASON,
    FICON_DEVICE_PORT.TYPE_NUMBER,
    FICON_DEVICE_PORT.MODEL_NUMBER,
    FICON_DEVICE_PORT.MANUFACTURER,
    FICON_DEVICE_PORT.MANUFACTURER_PLANT,
    FICON_DEVICE_PORT.SEQUENCE_NUMBER,
    FICON_DEVICE_PORT.TAG,
    FICON_DEVICE_PORT.FLAG,
    FICON_DEVICE_PORT.PARAMS,
    USER_DEFINED_DEVICE_DETAIL.NAME,
    USER_DEFINED_DEVICE_DETAIL.TYPE as USER_DEFINED_TYPE,
    USER_DEFINED_DEVICE_DETAIL.IP_ADDRESS,
    USER_DEFINED_DEVICE_DETAIL.CONTACT,
    USER_DEFINED_DEVICE_DETAIL.LOCATION,
    USER_DEFINED_DEVICE_DETAIL.DESCRIPTION,
    USER_DEFINED_DEVICE_DETAIL.USER_DEFINED_VALUE1,
    USER_DEFINED_DEVICE_DETAIL.USER_DEFINED_VALUE2,
    USER_DEFINED_DEVICE_DETAIL.USER_DEFINED_VALUE3,
    DEVICE_NODE.WWN as DEVICE_NODE_WWN,
    DEVICE_NODE.FDMI_HOST_NAME,
    DEVICE_NODE.SYMBOLIC_NAME as DEVICE_SYMBOLIC_NAME,
    DEVICE_NODE.AG as AG_PORT,
    coalesce(SWITCH_PORT.NAME, VIRTUAL_FCOE_PORT.NAME) as SWITCH_PORT_NAME,
    coalesce (SWITCH_PORT.TYPE, VIRTUAL_FCOE_PORT.PORT_TYPE) as SWITCH_PORT_TYPE,
    SWITCH_PORT.LOGICAL_PORT_WWN,
    coalesce(VS1.WWN, VS2.WWN) as SWITCH_WWN,
    coalesce(VS1.MANAGEMENT_STATE, VS2.MANAGEMENT_STATE) as MANAGEMENT_STATE,
    coalesce(VS1.MONITORED, VS2.MONITORED) as MONITORED,
    FABRIC.PRINCIPAL_SWITCH_WWN as PRINCIPAL_WWN,
    FABRIC.ID as FABRIC_ID
from
    DEVICE_PORT
        left outer join USER_DEFINED_DEVICE_DETAIL
            on DEVICE_PORT.WWN = USER_DEFINED_DEVICE_DETAIL.WWN
        left outer join FICON_DEVICE_PORT
            on DEVICE_PORT.ID = FICON_DEVICE_PORT.DEVICE_PORT_ID
         left outer join DEVICE_NODE
             on DEVICE_PORT.NODE_ID = DEVICE_NODE.ID
         left outer join SWITCH_PORT
             on DEVICE_PORT.SWITCH_PORT_WWN = SWITCH_PORT.WWN
         left outer join VIRTUAL_FCOE_PORT
             on DEVICE_PORT.SWITCH_PORT_WWN = VIRTUAL_FCOE_PORT.PORT_WWN
         left outer join VIRTUAL_SWITCH VS1
             on SWITCH_PORT.VIRTUAL_SWITCH_ID = VS1.ID
         left outer join VIRTUAL_SWITCH VS2
             on VIRTUAL_FCOE_PORT.VIRTUAL_SWITCH_ID = VS2.ID
         left outer join FABRIC
             on DEVICE_NODE.FABRIC_ID = FABRIC.ID;
```

DEVICE_REPORT_INFO

```
create or replace view DEVICE_REPORT_INFO as
select VS.NAME as VIRTUAL_SWITCH_NAME, VS.WWN as VIRTUAL_SWITCH_WWN,
VS.MANAGEMENT_STATE, VS.MONITORED, VS.SWITCH_MODE,
coalesce(F.NAME, F.FABRIC_NAME) as FABRIC_NAME, F.MANAGED as FABRIC_MANAGED,
F.TYPE as FABRIC_TYPE, F.SEED_SWITCH_WWN,
F.PRINCIPAL_SWITCH_WWN as PRINCIPAL_WWN, DN.ID as DEVICE_NODE_ID,
DN.FABRIC_ID, DN.SYMBOLIC_NAME as DEVICE_NODE_SYMBOLIC_NAME,
DN.TYPE as DEVICE_NODE_TYPE, DN.FDMI_HOST_NAME, DN.VENDOR, DN.CAPABILITY_,
DN.TRUSTED as DEVICE_NODE_TRUSTED,
DN.CREATION_TIME as DEVICE_NODE_CREATION_TIME,
```

```
DN.MISSING as DEVICE_NODE_MISSING,
DN.MISSING_TIME as DEVICE_NODE_MISSING_TIME, DN.PROXY_DEVICE, DN.AG,
DN.PREVIOUS_MISSING_STATE, DN.SIMULATED, DN.WWN as DEVICE_NODE_WWN,
DP.WWN as DEVICE_PORT_WWN, DP.PORT_ID as DEVICE_PORT_FC_ADDRESS,
DP.NUMBER as DEVICE_PORT_NUMBER, DP.ID as DEVICE_PORT_ID, DP.DOMAIN_ID,
DP.TYPE as DEVICE_PORT_TYPE, DP.SYMBOLIC_NAME as DEVICE_PORT_SYMBOLIC_NAME,
DP.FC4_TYPE, DP.COS, DP.IP_PORT, DP.NPV_PHYSICAL, DP.HARDWARE_ADDRESS,
DP.TRUSTED as DEVICE_PORT_TRUSTED, DP.EDGE_SWITCH_PORT_WWN,
DP.CREATION_TIME as DEVICE_PORT_CREATION_TIME,
DP.MISSING as DEVICE_PORT_MISSING,
DP.MISSING_TIME as DEVICE_PORT_MISSING_TIME, DP.LOGGED_TO_AG,
DP.AG_NODE_WWN, DP.AG_N_PORT_WWN, DP.MISSING_REASON,
case
when DP.NPV_PHYSICAL = 0 then 'PHYSICAL'::text
when DP.NPV_PHYSICAL = 1 then 'VIRTUAL'::text
when DP.NPV_PHYSICAL = 2 then 'NPIV'::text
when DP.NPV_PHYSICAL = 3 then 'ISCSI'::text
when DP.NPV_PHYSICAL = 4 then 'PHY+VIR'::text
else ''::text
end as DEVICE_TYPE,
DP.SWITCH_PORT_WWN, SP.ID as SWITCH_PORT_ID,
SP.STATUS as SWITCH_PORT_STATUS, SP.SLOT_NUMBER as SWITCH_PORT_SLOT_NUMBER,
SP.PORT_NUMBER as SWITCH_PORT_PORT_NUMBER,
SP.PORT_ID as SWITCH_PORT_FC_ADDRESS, SP.PORT_INDEX,
SP.SPEED as SWITCH_PORT_SPEED, SP.TYPE as SWITCH_PORT_TYPE, SP.LICENSED,
SP.NAME as SWITCH_PORT_NAME, SP.LOGICAL_PORT_WWN, SP.VIRTUAL_SWITCH_ID,
CS.IP_ADDRESS as SWITCH_IP_ADDRESS,
coalesce(USER_DEFINED_DEVICE_DETAIL.TYPE, DN.TYPE, ''::character varying) as USER_DEFINED_DEVICE_TYPE,
USER_DEFINED_DEVICE_DETAIL.NAME as USER_DEFINED_NAME,
USER_DEFINED_DEVICE_DETAIL.IP_ADDRESS as USER_DEFINED_IP_ADDRESS,
USER_DEFINED_DEVICE_DETAIL.CONTACT,
USER_DEFINED_DEVICE_DETAIL.LOCATION,
USER_DEFINED_DEVICE_DETAIL.DESCRIPTION
from DEVICE_PORT DP
join DEVICE_NODE DN on DP.NODE_ID = DN.ID
left join USER_DEFINED_DEVICE_DETAIL on DN.WWN = USER_DEFINED_DEVICE_DETAIL.WWN
left join SWITCH_PORT SP on DP.SWITCH_PORT_WWN = SP.WWN
left join VIRTUAL_SWITCH VS on SP.VIRTUAL_SWITCH_ID = VS.ID
left join CORE_SWITCH CS on VS.CORE_SWITCH_ID = CS.ID
left join FABRIC F on DN.FABRIC_ID = F.ID
where DN.AG <> 1 and F.MANAGED = 1 and (F.TYPE <> all (array[65, 66, 4])) and (SP.LICENSED is null or
SP.LICENSED = 1) and (VS.MONITORED is null or  VS.MONITORED = 1) and (VS.SWITCH_MODE is null or
(VS.SWITCH_MODE = any (array[0, 2])));
```

DEV_PORT_GIGE_PORT_LINK_INFO

```
create or replace view DEV_PORT_GIGE_PORT_LINK_INFO as
select
    DEVICE_PORT_GIGE_PORT_LINK.DEVICE_PORT_ID,
    DEVICE_PORT_GIGE_PORT_LINK.GIGE_PORT_ID,
    DEVICE_PORT_GIGE_PORT_LINK.DIRECT_ATTACH,
    DEVICE_PORT_GIGE_PORT_LINK.VIRTUAL_FCOE_PORT_ID,
    DEVICE_PORT.TRUSTED,
    DEVICE_PORT.CREATION_TIME,
    DEVICE_PORT.MISSING,
    DEVICE_PORT.MISSING_TIME,
    DEVICE_PORT_GIGE_PORT_LINK.LAG_ID
from
    DEVICE_PORT_GIGE_PORT_LINK,
    DEVICE_PORT
where
    DEVICE_PORT_GIGE_PORT_LINK.DEVICE_PORT_ID = DEVICE_PORT.ID;
```

DEV_PORT_MAC_ADDR_MAP_INFO

```
create or replace view DEV_PORT_MAC_ADDR_MAP_INFO as
select
    DEVICE_PORT_MAC_ADDRESS_MAP.DEVICE_PORT_ID,
    DEVICE_PORT_MAC_ADDRESS_MAP.MAC_ADDRESS,
    DEVICE_NODE.ID as DEVICE_NODE_ID,
    DEVICE_NODE.FABRIC_ID,
    DEVICE_PORT.TRUSTED,
    DEVICE_PORT.CREATION_TIME,
    DEVICE_PORT.MISSING,
    DEVICE_PORT.MISSING_TIME
from
    DEVICE_PORT_MAC_ADDRESS_MAP,
    DEVICE_PORT,
    DEVICE_NODE
where
    DEVICE_PORT_MAC_ADDRESS_MAP.DEVICE_PORT_ID = DEVICE_PORT.ID
    and DEVICE_PORT.NODE_ID = DEVICE_NODE.ID;
```

ISL_CONNECTION_INFO

```
create or replace view ISL_CONNECTION_INFO as
select
 distinct ISL_CONNECTION.ID,
 ISL_CONNECTION.FABRIC_ID,
 ISL_CONNECTION.SOURCE_SWITCH_PORT_ID,
 ISL_CONNECTION.TARGET_SWITCH_PORT_ID,
 ISL_CONNECTION.COST,
 ISL_CONNECTION.TYPE,
 ISL_CONNECTION.TRUSTED,
 ISL_CONNECTION.MISSING,
 ISL_CONNECTION.MISSING_TIME,
 ISL_CONNECTION.CREATION_TIME,
 ISL_CONNECTION.TRUNKED,
 ISL_CONNECTION.MISSING_REASON,
 ISL_CONNECTION.MASTER_CONNECTION_ID,
 ISL_CONNECTION.SOURCE_MASTER_PORT AS SOURCE_MASTER_PORT_NUMBER,
 ISL_CONNECTION.TARGET_MASTER_PORT AS DEST_MASTER_PORT_NUMBER,
 ISL_CONNECTION.SOURCE_PORT_SPEED AS SOURCE_PORT_SPEED,
 ISL_CONNECTION.TARGET_PORT_SPEED AS DEST_PORT_SPEED,
 SOURCE_SWITCH_PORT.VIRTUAL_SWITCH_ID as SOURCE_SWITCH_ID,
 SOURCE_SWITCH_PORT.USER_PORT_NUMBER as SOURCE_SWITCH_PORT_NUMBER,
 DEST_SWITCH_PORT.VIRTUAL_SWITCH_ID as DEST_SWITCH_ID,
 DEST_SWITCH_PORT.USER_PORT_NUMBER as DEST_SWITCH_PORT_NUMBER,
 COALESCE(ISL_TRUNK_GROUP.MEMBER_TRACKING_STATUS, -1) AS MEMBER_TRACKING_STATUS
 COALESCE(ISL_TRUNK_GROUP.ID, (-1)) AS TRUNK_ID

from
 ISL_CONNECTION
LEFT JOIN SWITCH_PORT SOURCE_SWITCH_PORT ON (ISL_CONNECTION.SOURCE_SWITCH_PORT_ID = SOURCE_SWITCH_PORT.ID)
LEFT JOIN SWITCH_PORT DEST_SWITCH_PORT ON (ISL_CONNECTION.TARGET_SWITCH_PORT_ID = DEST_SWITCH_PORT.ID)
LEFT JOIN ISL_TRUNK_GROUP ON (ISL_CONNECTION.TRUNKED = 1
        AND ISL_TRUNK_GROUP.MASTER_USER_PORT = ISL_CONNECTION.SOURCE_MASTER_PORT
        AND ISL_TRUNK_GROUP.VIRTUAL_SWITCH_ID =  SOURCE_SWITCH_PORT.VIRTUAL_SWITCH_ID );
```

ISL_INFO

```
create or replace view ISL_INFO as
```

```
select distinct
    ISL.ID,
    ISL.FABRIC_ID,
    ISL.COST,
    ISL.TYPE,
    ISL.SOURCE_DOMAIN_ID,
    ISL.SOURCE_PORT_NUMBER,
    ISL.MISSING,
    ISL.MISSING_TIME,
    ISL.TRUSTED,
    ISL.CREATION_TIME,
    ISL.TRUNKED,
    SOURCE_VIRTUAL_SWITCH.ID as SOURCE_SWITCH_ID,
    SOURCE_VIRTUAL_SWITCH.NAME as SOURCE_SWITCH_NAME,
    SOURCE_VIRTUAL_SWITCH.WWN as SOURCE_SWITCH_WWN,
    SOURCE_VIRTUAL_SWITCH.CORE_SWITCH_ID as SOURCE_CORE_SWITCH_ID,
    SOURCE_VIRTUAL_SWITCH.BASE_SWITCH as SOURCE_BASE_SWITCH,
    SOURCE_VIRTUAL_SWITCH.MANAGEMENT_STATE as SOURCE_VIRTUAL_SWITCH_MANAGEMENT_STATE,
    SOURCE_VIRTUAL_SWITCH.MONITORED as SOURCE_VIRTUAL_SWITCH_MONITORED,
    SOURCE_SWITCH_PORT.ID as SOURCE_SWITCH_PORT_ID,
    SOURCE_SWITCH_PORT.WWN as SOURCE_SWITCH_PORT_WWN,
    SOURCE_SWITCH_PORT.NAME as SOURCE_SWITCH_PORT_NAME,
    SOURCE_SWITCH_PORT.TYPE as PORT_TYPE,
    SOURCE_SWITCH_PORT.KIND as SOURCE_SWITCH_PORT_KIND,
    SOURCE_SWITCH_PORT.PHYSICAL_PORT as SOURCE_PHYSICAL_PORT,
    SOURCE_SWITCH_PORT.TRUNKED as SOURCE_SWITCH_PORT_TRUNKED,
    ISL.DEST_DOMAIN_ID,
    ISL.DEST_PORT_NUMBER,
    DEST_VIRTUAL_SWITCH.ID as DEST_SWITCH_ID,
    DEST_VIRTUAL_SWITCH.NAME as DEST_SWITCH_NAME,
    DEST_VIRTUAL_SWITCH.WWN as DEST_SWITCH_WWN,
    DEST_VIRTUAL_SWITCH.CORE_SWITCH_ID as DEST_CORE_SWITCH_ID,
    DEST_VIRTUAL_SWITCH.BASE_SWITCH as DEST_BASE_SWITCH,
    DEST_VIRTUAL_SWITCH.MANAGEMENT_STATE as DEST_VIRTUAL_SWITCH_MANAGEMENT_STATE,
    DEST_VIRTUAL_SWITCH.MONITORED as DEST_VIRTUAL_SWITCH_MONITORED,
    DEST_SWITCH_PORT.ID as DEST_SWITCH_PORT_ID,
    DEST_SWITCH_PORT.WWN as DEST_SWITCH_PORT_WWN,
    DEST_SWITCH_PORT.NAME as DEST_SWITCH_PORT_NAME,
    DEST_SWITCH_PORT.KIND as DEST_SWITCH_PORT_KIND,
    DEST_SWITCH_PORT.PHYSICAL_PORT as DEST_PHYSICAL_PORT,
    DEST_SWITCH_PORT.TRUNKED as DEST_SWITCH_PORT_TRUNKED,
    FABRIC.PRINCIPAL_SWITCH_WWN as PRINCIPAL_SWITCH_WWN
from
    ISL,
    FABRIC_MEMBER SOURCE_FABRIC_MEMBER,
    VIRTUAL_SWITCH         SOURCE_VIRTUAL_SWITCH,
    SWITCH_PORT   SOURCE_SWITCH_PORT,
    FABRIC_MEMBER DEST_FABRIC_MEMBER,
    VIRTUAL_SWITCH         DEST_VIRTUAL_SWITCH,
    SWITCH_PORT   DEST_SWITCH_PORT,
    FABRIC
where
    SOURCE_FABRIC_MEMBER.FABRIC_ID = ISL.FABRIC_ID and
    SOURCE_VIRTUAL_SWITCH.ID = SOURCE_FABRIC_MEMBER.VIRTUAL_SWITCH_ID and
    SOURCE_VIRTUAL_SWITCH.DOMAIN_ID = ISL.SOURCE_DOMAIN_ID and
    SOURCE_SWITCH_PORT.VIRTUAL_SWITCH_ID = SOURCE_VIRTUAL_SWITCH.ID and
    SOURCE_SWITCH_PORT.CATEGORY = 1 and
    SOURCE_SWITCH_PORT.USER_PORT_NUMBER = ISL.SOURCE_PORT_NUMBER and
    DEST_FABRIC_MEMBER.FABRIC_ID = ISL.FABRIC_ID and
    DEST_VIRTUAL_SWITCH.ID = DEST_FABRIC_MEMBER.VIRTUAL_SWITCH_ID and
    DEST_VIRTUAL_SWITCH.DOMAIN_ID = ISL.DEST_DOMAIN_ID and
    DEST_SWITCH_PORT.VIRTUAL_SWITCH_ID = DEST_VIRTUAL_SWITCH.ID and
    DEST_SWITCH_PORT.CATEGORY = 1 and
```

```
    DEST_SWITCH_PORT.USER_PORT_NUMBER = ISL.DEST_PORT_NUMBER and
    FABRIC.ID = ISL.FABRIC_ID;
```

ETHERNET_ISL_INFO

```
create or replace view ETHERNET_ISL_INFO as
select
    ETHERNET_ISL.ID as ETHERNET_ISL_ID,
    ETHERNET_ISL.SOURCE_PORT_ID,
    ETHERNET_ISL.DEST_PORT_ID,
    ETHERNET_ISL.TRUSTED,
    ETHERNET_ISL.CREATION_TIME,
    ETHERNET_ISL.MISSING,
    ETHERNET_ISL.MISSING_TIME,
    SOURCE_SWITCH_PORT.VIRTUAL_SWITCH_ID as SOURCE_SWITCH_ID,
    SOURCE_SWITCH_PORT.USER_PORT_NUMBER as SOURCE_PORT_NUMBER,
    SOURCE_SWITCH_PORT.TYPE as SOURCE_PORT_TYPE,
    SOURCE_VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID as SOURCE_VIRTUAL_FABRIC_ID,
    DEST_SWITCH_PORT.VIRTUAL_SWITCH_ID as DEST_SWITCH_ID,
    DEST_SWITCH_PORT.USER_PORT_NUMBER as DEST_PORT_NUMBER,
    DEST_SWITCH_PORT.TYPE as DEST_PORT_TYPE,
    DEST_VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID as DEST_VIRTUAL_FABRIC_ID
from
    ETHERNET_ISL,
    GIGE_PORT       SOURCE_GIGE_PORT,
    VIRTUAL_SWITCH  SOURCE_VIRTUAL_SWITCH,
    SWITCH_PORT     SOURCE_SWITCH_PORT,
    GIGE_PORT       DEST_GIGE_PORT,
    VIRTUAL_SWITCH  DEST_VIRTUAL_SWITCH,
    SWITCH_PORT     DEST_SWITCH_PORT
where
    SOURCE_GIGE_PORT.ID  = ETHERNET_ISL.SOURCE_PORT_ID and
    SOURCE_GIGE_PORT.SWITCH_PORT_ID = SOURCE_SWITCH_PORT.ID and
    SOURCE_SWITCH_PORT.VIRTUAL_SWITCH_ID = SOURCE_VIRTUAL_SWITCH.ID and
    DEST_GIGE_PORT.ID  = ETHERNET_ISL.DEST_PORT_ID and
    DEST_GIGE_PORT.SWITCH_PORT_ID = DEST_SWITCH_PORT.ID and
    DEST_SWITCH_PORT.VIRTUAL_SWITCH_ID = DEST_VIRTUAL_SWITCH.ID;
```

EVENT_DETAILS_INFO

```
create or replace view EVENT_DETAILS_INFO (ID, ME_ID, SEVERITY, AREA, ACKNOWLEDGED, SOURCE_NAME, SOURCE_ADDR,
LAST_OCCURRENCE_HOST_TIME, FIRST_OCCURRENCE_HOST_TIME, EVENT_COUNT, EVENT_KEY, AUDIT, RESOLVED, ACKED_TIME,
EVENT_ACTION_ID, DEVICE_GROUP_ID, PORT_GROUP_ID, SPECIAL_EVENT, CALLHOME_EVENT, ORIGIN, EVENT_CATEGORY,
DESCRIPTION, MODULE, RAS_LOG_ID, PRODUCT_ADDRESS, CONTRIBUTORS, NODE_WWN, PORT_WWN, OPERATIONAL_STATUS,
FIRST_OCCURRENCE_SWITCH_TIME, LAST_OCCURRENCE_SWITCH_TIME, VIRTUAL_FABRIC_ID, UNIT, SLOT, PORT, OID,
USER_NAME, EVENT_NUMBER, FRU_CODE, REASON_CODE, FRU_POSITION, INTERFACE_TYPE, PORT_NAME, MAC_ADDRESS,
ANNOTATED_BY, ANNOTATIONS) as
select
    EVENT.ID as ID,
    EVENT.ME_ID as ME_ID,
    EVENT.SEVERITY as SEVERITY,
    EVENT.AREA as AREA,
    EVENT.ACKNOWLEDGED as ACKNOWLEDGED,
    EVENT.SOURCE_NAME as SOURCE_NAME,
    EVENT.SOURCE_ADDR as SOURCE_ADDR,
    EVENT.LAST_OCCURRENCE_HOST_TIME as LAST_OCCURRENCE_HOST_TIME,
    EVENT.FIRST_OCCURRENCE_HOST_TIME as FIRST_OCCURRENCE_HOST_TIME,
    EVENT.EVENT_COUNT as EVENT_COUNT,
    EVENT.EVENT_KEY as EVENT_KEY,
    EVENT.EVENT_AUDIT as AUDIT,
    EVENT.RESOLVED as RESOLVED,
```

```
    EVENT.ACKED_TIME as ACKED_TIME,
    EVENT.EVENT_ACTION_ID as EVENT_ACTION_ID,
    EVENT.DEVICE_GROUP_ID as DEVICE_GROUP_ID,
    EVENT.PORT_GROUP_ID as PORT_GROUP_ID,
    EVENT.SPECIAL_EVENT,
    EVENT.CALLHOME_EVENT,
    EVENT_ORIGIN.ID as ORIGIN,
    EVENT_CATEGORY.ID as EVENT_CATEGORY,
EVENT_CATEGORY. DESCRIPTION as EVENT_CATEGORY_DESCRIPTION,
    EVENT_DESCRIPTION.DESCRIPTION as DESCRIPTION,
    EVENT_MODULE.ID as MODULE,
    EVENT_DETAILS.RAS_LOG_ID as RAS_LOG_ID,
    EVENT_DETAILS.PRODUCT_ADDRESS as PRODUCT_ADDRESS,
    EVENT_DETAILS.CONTRIBUTORS as CONTRIBUTORS,
    EVENT_DETAILS.NODE_WWN as NODE_WWN,
    EVENT_DETAILS.PORT_WWN as PORT_WWN,
    EVENT_DETAILS.OPERATIONAL_STATUS as OPERATIONAL_STATUS,
    EVENT_DETAILS.FIRST_OCCURRENCE_SWITCH_TIME as FIRST_OCCURRENCE_SWITCH_TIME,
    EVENT_DETAILS.LAST_OCCURRENCE_SWITCH_TIME as LAST_OCCURRENCE_SWITCH_TIME,
    EVENT_DETAILS.VIRTUAL_FABRIC_ID as VIRTUAL_FABRIC_ID,
    EVENT_DETAILS.UNIT as UNIT,
    EVENT_DETAILS.SLOT as SLOT,
    EVENT_DETAILS.PORT as PORT,
    EVENT_DETAILS.OID,
    EVENT_DETAILS.USER_NAME as USER_NAME,
    EVENT_CALL_HOME.EVENT_NUMBER as EVENT_NUMBER,
    EVENT_CALL_HOME.FRU_CODE as FRU_CODE,
    EVENT_CALL_HOME.REASON_CODE as REASON_CODE,
    EVENT_CALL_HOME.FRU_POSITION as FRU_POSITION,
    EVENT_DETAILS.INTERFACE_TYPE as INTERFACE_TYPE,
    EVENT_DETAILS.PORT_NAME as PORT_NAME,
    EVENT_DETAILS.MAC_ADDRESS,
    EVENT_DETAILS.ANNOTATED_BY ,
    EVENT_DETAILS.ANNOTATIONS

from
    EVENT
        left outer join EVENT_ORIGIN on EVENT.EVENT_ORIGIN_ID = EVENT_ORIGIN.ID
        left outer join EVENT_CATEGORY on EVENT.EVENT_CATEGORY_ID = EVENT_CATEGORY.ID
        left outer join EVENT_MODULE on EVENT.EVENT_MODULE_ID = EVENT_MODULE.ID
        left outer join EVENT_DESCRIPTION on EVENT.EVENT_DESCRIPTION_ID = EVENT_DESCRIPTION.ID
        left outer join EVENT_DETAILS on EVENT.ID = EVENT_DETAILS.EVENT_ID
                    left outer join EVENT_CALL_HOME on EVENT.ID = EVENT_CALL_HOME.EVENT_ID;
```

EVENT_INFO

```
create or replace view EVENT_INFO as
select
    EVENT.ID as ID,
    EVENT.ME_ID as ME_ID,
    EVENT.SEVERITY as SEVERITY,
    EVENT.AREA as AREA,
    EVENT.ACKNOWLEDGED as ACKNOWLEDGED,
    EVENT.SOURCE_NAME as SOURCE_NAME,
    EVENT.SOURCE_ADDR as SOURCE_ADDR,
    EVENT.LAST_OCCURRENCE_HOST_TIME as LAST_OCCURRENCE_HOST_TIME,
    EVENT.FIRST_OCCURRENCE_HOST_TIME as FIRST_OCCURRENCE_HOST_TIME,
    EVENT.EVENT_COUNT as EVENT_COUNT,
    EVENT.EVENT_AUDIT as AUDIT,
    EVENT.EVENT_ACTION_ID,
    EVENT.SPECIAL_EVENT,
    EVENT.CALLHOME_EVENT,
```

```
    EVENT_ORIGIN.ID as ORIGIN,
    EVENT_CATEGORY.ID as EVENT_CATEGORY,
    EVENT_DESCRIPTION.DESCRIPTION as DESCRIPTION,
    EVENT_MODULE.ID as MODULE,
    EVENT_DETAILS.RAS_LOG_ID as RAS_LOG_ID,
    EVENT_DETAILS.PRODUCT_ADDRESS as PRODUCT_ADDRESS,
    EVENT_DETAILS.CONTRIBUTORS as CONTRIBUTORS,
    EVENT_DETAILS.NODE_WWN as NODE_WWN,
    EVENT_DETAILS.OPERATIONAL_STATUS as OPERATIONAL_STATUS,
    EVENT_DETAILS.FIRST_OCCURRENCE_SWITCH_TIME as FIRST_OCCURRENCE_SWITCH_TIME,
    EVENT_DETAILS.LAST_OCCURRENCE_SWITCH_TIME as LAST_OCCURRENCE_SWITCH_TIME,
    EVENT_DETAILS.VIRTUAL_FABRIC_ID as VIRTUAL_FABRIC_ID,
    EVENT_DETAILS.USER_NAME as USER_NAME,
    EVENT_DETAILS.PORT_NAME as PORT_NAME,
    EVENT_DETAILS.MAC_ADDRESS
from
    EVENT
        left join EVENT_DETAILS on EVENT.ID = EVENT_DETAILS.EVENT_ID, EVENT_ORIGIN, EVENT_CATEGORY,
EVENT_MODULE, EVENT_DESCRIPTION
        where EVENT.EVENT_ORIGIN_ID = EVENT_ORIGIN.ID and EVENT.EVENT_CATEGORY_ID = EVENT_CATEGORY.ID and
EVENT.EVENT_MODULE_ID = EVENT_MODULE.ID
        and   EVENT.EVENT_DESCRIPTION_ID = EVENT_DESCRIPTION.ID;
```

FABRIC_INFO

```
create or replace view FABRIC_INFO as
select
        FABRIC.ID,
        FABRIC.SAN_ID,
        FABRIC.SEED_SWITCH_WWN,
        FABRIC.NAME,
        FABRIC.ACTIVE_ZONESET_NAME,
        FABRIC.MANAGEMENT_STATE,
        FABRIC.LAST_FABRIC_CHANGED,
        FABRIC.SECURE,
        FABRIC.AD_ENVIRONMENT,
        FABRIC.MANAGED,
        FABRIC.CONTACT,
        FABRIC.LOCATION,
        FABRIC.DESCRIPTION,
        FABRIC.CREATION_TIME,
        FABRIC.LAST_SCAN_TIME,
        FABRIC.LAST_UPDATE_TIME,
        FABRIC.TRACK_CHANGES,
        FABRIC.TYPE,
       FABRIC.HAS_NOS_AG,
    FABRIC.USER_DEFINED_VALUE_1,
    FABRIC.USER_DEFINED_VALUE_2,
    FABRIC.USER_DEFINED_VALUE_3,
                FABRIC.PRINCIPAL_SWITCH_WWN,
                FABRIC.ZONE_TRANSACTION_TIMEOUT,
    FABRIC.FABRIC_MODEL,
    FABRIC.ENHANCED_TI_ZONE_SUPPORT,
    FABRIC.FABRIC_NAME,
    VIRTUAL_SWITCH.ID as SEED_SWITCH_ID,
                VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID,
    VIRTUAL_SWITCH.INTEROP_MODE,
                CORE_SWITCH.IP_ADDRESS as SEED_SWITCH_IP_ADDRESS,
                (select count(*) from FABRIC_MEMBER
        where FABRIC_MEMBER.FABRIC_ID = FABRIC.ID) as SWITCH_COUNT
from
                FABRIC, CORE_SWITCH, VIRTUAL_SWITCH
```

```
where
                FABRIC.SEED_SWITCH_WWN = VIRTUAL_SWITCH.WWN and
                VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID;
```

FCIP_TUNNEL_CIRCUIT_INFO

```
CREATE VIEW fcip_tunnel_circuit_info AS
    select
    FCIP_TUNNEL_CIRCUIT.ID,
    FCIP_TUNNEL_CIRCUIT.TUNNEL_ID,
    FCIP_TUNNEL_CIRCUIT.CIRCUIT_NUMBER,
    FCIP_TUNNEL_CIRCUIT.COMPRESSION_ENABLED,
    FCIP_TUNNEL_CIRCUIT.TURBO_WRITE_ENABLED,
    FCIP_TUNNEL_CIRCUIT.TAPE_ACCELERATION_ENABLED,
    FCIP_TUNNEL_CIRCUIT.IKE_POLICY_NUM,
    FCIP_TUNNEL_CIRCUIT.IPSEC_POLICY_NUM,
    FCIP_TUNNEL_CIRCUIT.PRESHARED_KEY,
    FCIP_TUNNEL_CIRCUIT.SOURCE_IP,
    FCIP_TUNNEL_CIRCUIT.DEST_IP,
    FCIP_TUNNEL_CIRCUIT.VLAN_TAG,
    FCIP_TUNNEL_CIRCUIT.SELECTIVE_ACK,
    FCIP_TUNNEL_CIRCUIT.QOS_MAPPING,
    FCIP_TUNNEL_CIRCUIT.PATH_MTU_DISCOVERY,
    FCIP_TUNNEL_CIRCUIT.MIN_COMM_RATE,
    FCIP_TUNNEL_CIRCUIT.MAX_COMM_RATE,
    FCIP_TUNNEL_CIRCUIT.MIN_RETRANSMIT_TIME,
    FCIP_TUNNEL_CIRCUIT.MAX_RETRANSMIT_TIME,
    FCIP_TUNNEL_CIRCUIT.KEEP_ALIVE_TIMEOUT,
    FCIP_TUNNEL_CIRCUIT.ADMIN_STATUS,
    FCIP_TUNNEL_CIRCUIT.METRIC,
    FCIP_TUNNEL_CIRCUIT.DATA_L2_COS,
    FCIP_TUNNEL_CIRCUIT.DSCP_DATA,
    FCIP_TUNNEL_CIRCUIT.MAX_RETRANSMISSIONS,
    FCIP_TUNNEL_CIRCUIT.SLOT_NUMBER,
    FCIP_TUNNEL_CIRCUIT.VE_PORT_NUMBER,
    FCIP_TUNNEL_CIRCUIT.SECURITY_FLAG,
    FCIP_TUNNEL_CIRCUIT.DSCP_CONTROL,
    FCIP_TUNNEL_CIRCUIT.CIRCUIT_STATUS,
    FCIP_TUNNEL_CIRCUIT.ENABLED,
    FCIP_TUNNEL_CIRCUIT.MISMATCHED_CONFIGURATIONS,
    FCIP_TUNNEL_CIRCUIT.CIRCUIT_STATUS_STRING,
    FCIP_TUNNEL_CIRCUIT.L2COS_F_CLASS,
    FCIP_TUNNEL_CIRCUIT.L2_COS_HIGH,
    FCIP_TUNNEL_CIRCUIT.L2_COS_MEDIUM,
    FCIP_TUNNEL_CIRCUIT.L2_COS_LOW,
    FCIP_TUNNEL_CIRCUIT.DSCP_F_CLASS,
    FCIP_TUNNEL_CIRCUIT.DSCP_HIGH,
    FCIP_TUNNEL_CIRCUIT.DSCP_MEDIUM,
    FCIP_TUNNEL_CIRCUIT.DSCP_LOW,
    FCIP_TUNNEL_CIRCUIT.FAILOVER_CIRCUIT,
    FCIP_TUNNEL_CIRCUIT.FAILOVER_GROUP_ID,
    GIGE_PORT.PORT_NUMBER GIGE_PORT_NUMBER,
    GIGE_PORT.SLOT_NUMBER GIGE_PORT_SLOT_NUMBER,
    FCIP_CIRCUIT_PORT_MAP.SWITCH_PORT_ID GIGE_PORT_ID,
    SWITCH_PORT.VIRTUAL_SWITCH_ID,
    SWITCH_PORT.USER_PORT_NUMBER
from
    FCIP_TUNNEL_CIRCUIT
        left outer join FCIP_CIRCUIT_PORT_MAP on
            FCIP_CIRCUIT_PORT_MAP.CIRCUIT_ID = FCIP_TUNNEL_CIRCUIT.ID
        left outer join GIGE_PORT
            on FCIP_CIRCUIT_PORT_MAP.SWITCH_PORT_ID = GIGE_PORT.ID
```

```
        left outer join SWITCH_PORT
            on GIGE_PORT.SWITCH_PORT_ID = SWITCH_PORT.ID;
```

FCIP_TUNNEL_REPORT_INFO

```
CREATE OR REPLACE VIEW fcip_tunnel_report_info AS
 SELECT fcip_tunnel.id, fcip_tunnel.tunnel_id, fcip_tunnel.vlan_tag,
    fcip_tunnel.source_ip, fcip_tunnel.dest_ip, fcip_tunnel.local_wwn,
    fcip_tunnel.remote_wwn_restrict, fcip_tunnel.communication_rate,
    fcip_tunnel.min_retransmit_time, fcip_tunnel.selective_ack_enabled,
    fcip_tunnel.keep_alive_timeout, fcip_tunnel.max_retransmission,
    fcip_tunnel.wan_tov_enabled, fcip_tunnel.tunnel_status,
    fcip_tunnel.ha_status, fcip_tunnel.description,
    fcip_tunnel.ficon_trb_id_enabled, fcip_tunnel.ficon_tt_emul_enabled,
    fcip_tunnel.ficon_dla_emul_enabled, fcip_tunnel.ficon_tape_write_max_pipe,
    fcip_tunnel.ficon_tape_read_max_pipe, fcip_tunnel.ficon_tape_write_max_ops,
    fcip_tunnel.ficon_tape_read_max_ops, fcip_tunnel.ficon_tape_write_timer,
    fcip_tunnel.ficon_tape_max_write_chain, fcip_tunnel.ficon_oxid_base,
    fcip_tunnel.ficon_xrc_emulation_enabled, fcip_tunnel.ficon_tw_emul_enabled,
    fcip_tunnel.ficon_tr_emul_enabled, fcip_tunnel.ficon_debug_flags,
    fcip_tunnel.remote_wwn, fcip_tunnel.cdc, fcip_tunnel.admin_status,
    fcip_tunnel.control_l2_cos, fcip_tunnel.dscp_control,
    fcip_tunnel.trunking_algorithm, fcip_tunnel.extended_tunnel,
    fcip_tunnel.virtual_switch_id, fcip_tunnel.circuit_count,
    fcip_tunnel.mismatched_config_details, fcip_tunnel.last_update, fcip_tunnel.slot_number,
    fcip_tunnel.ficon_enabled, fcip_tunnel.tperf_enabled, fcip_tunnel.auth_key,
    fcip_tunnel.connected_count, fcip_tunnel.tunnel_status_string,
    fcip_tunnel.compression_mode, fcip_tunnel.turbo_write_enabled,
    fcip_tunnel.tape_acceleration_enabled, fcip_tunnel.ipsec_enabled,
    fcip_tunnel.preshared_key, fcip_tunnel.ipsec_policy_name,
    fcip_tunnel.qos_high, fcip_tunnel.qos_medium, fcip_tunnel.qos_low,
    fcip_tunnel.backward_compatible, fcip_tunnel.ficon_teradata_read_enabled,
    fcip_tunnel.ficon_teradata_write_enabled,
    local_virtual_switch.id AS local_virtual_switch_id, local_virtual_switch.managed_element_id AS
local_virtual_switch_managed_element_id, local_virtual_switch.name AS local_virtual_switch_name,
    local_virtual_switch.switch_mode AS local_virtual_switch_switch_mode, local_virtual_switch.domain_id AS
local_virtual_switch_domain_id,
    local_virtual_switch.wwn AS local_virtual_switch_wwn, local_virtual_switch.operational_status AS
local_virtual_switch_operational_status,
    local_virtual_switch.management_state AS local_virtual_switch_management_state,
local_virtual_switch.state AS local_virtual_switch_state,
    local_virtual_switch.status AS local_virtual_switch_status, local_virtual_switch.status_reason AS
local_virtual_switch_status_reason,
    local_core_switch.id AS local_core_switch_id, local_core_switch.ip_address AS
local_core_switch_ip_address,
    local_core_switch.wwn AS local_core_switch_wwn, local_core_switch.name AS local_core_switch_name,
    local_core_switch.type AS local_core_switch_type, local_core_switch.model AS local_core_switch_model,
    local_core_switch.vendor AS local_core_switch_vendor, local_core_switch.reachable AS
local_core_switch_reachable,
    local_core_switch.operational_status AS local_core_switch_operational_status,
    local_fabric_member.fabric_id AS local_fabric_id, local_fabric.seed_switch_wwn AS
local_fabric_seed_switch_wwn, local_fabric.name AS local_fabric_name,
    local_fabric.management_state AS local_fabric_management_state, local_fabric.principal_switch_wwn AS
local_fabric_principal_switch_wwn,
    local_fabric.fabric_name AS local_fabric_switch_persist_fabric_name, local_fabric.status AS
local_fabric_status, local_fabric.bottleneck_status AS local_fabric_bottleneck_status,
    local_switch_port.id AS local_switch_port_id, local_switch_port.wwn AS local_switch_port_wwn,
local_switch_port.name AS local_switch_port_name,
    local_switch_port.remote_port_wwn AS remote_switch_port_wwn, local_switch_port.remote_node_wwn,
local_switch_port.slot_number AS local_switch_port_slot_number,
    local_switch_port.port_number AS local_switch_port_port_number, local_switch_port.user_port_number AS
local_switch_port_user_port_number,
```

```
    local_switch_port.port_index AS local_switch_port_port_index, local_switch_port.port_id AS
local_switch_port_port_id, local_switch_port.status_message AS local_switch_port_status_message,
    remote_switch_port.id AS remote_switch_port_id, remote_switch_port.name AS remote_switch_port_name,
    remote_switch_port.slot_number AS remote_switch_port_slot_number, remote_switch_port.port_number AS
remote_switch_port_port_number,
   remote_switch_port.user_port_number AS remote_switch_port_user_port_number, remote_switch_port.port_index
AS remote_switch_port_port_index,
    remote_switch_port.port_id AS remote_switch_port_port_id, remote_switch_port.status_message AS
remote_switch_port_status_message,
    remote_virtual_switch.id AS remote_virtual_switch_id, remote_virtual_switch.managed_element_id AS
remote_virtual_switch_managed_element_id, remote_virtual_switch.name AS remote_virtual_switch_name,
    remote_virtual_switch.switch_mode AS remote_virtual_switch_switch_mode, remote_virtual_switch.domain_id
AS remote_virtual_switch_domain_id,
    remote_virtual_switch.wwn AS remote_virtual_switch_wwn, remote_virtual_switch.operational_status AS
remote_virtual_switch_operational_status,
    remote_virtual_switch.management_state AS remote_virtual_switch_management_state,
remote_virtual_switch.state AS remote_virtual_switch_state,
    remote_virtual_switch.status AS remote_virtual_switch_status, remote_virtual_switch.status_reason AS
remote_virtual_switch_status_reason,
    remote_core_switch.id AS remote_core_switch_id, remote_core_switch.ip_address AS
remote_core_switch_ip_address,
    remote_core_switch.wwn AS remote_core_switch_wwn, remote_core_switch.name AS remote_core_switch_name,
   remote_core_switch.type AS remote_core_switch_type, remote_core_switch.model AS remote_core_switch_model,
    remote_core_switch.vendor AS remote_core_switch_vendor, remote_core_switch.reachable AS
remote_core_switch_reachable,
    remote_core_switch.operational_status AS remote_core_switch_operational_status,
    remote_fabric_member.fabric_id AS remote_fabric_id, remote_fabric.seed_switch_wwn AS
remote_fabric_seed_switch_wwn, remote_fabric.name AS remote_fabric_name,
    remote_fabric.management_state AS remote_fabric_management_state, remote_fabric.principal_switch_wwn AS
remote_fabric_principal_switch_wwn,
    remote_fabric.fabric_name AS remote_fabric_switch_persist_fabric_name, remote_fabric.status AS
remote_fabric_status, remote_fabric.bottleneck_status AS remote_fabric_bottleneck_status
   FROM fcip_tunnel
                JOIN virtual_switch AS local_virtual_switch ON fcip_tunnel.virtual_switch_id =
local_virtual_switch.id
                JOIN core_switch AS local_core_switch ON local_virtual_switch.core_switch_id =
local_core_switch.id
                JOIN fabric_member AS local_fabric_member ON local_virtual_switch.id =
local_fabric_member.virtual_switch_id
                JOIN fabric AS local_fabric ON local_fabric_member.fabric_id = local_fabric.id
                LEFT JOIN fcip_port_tunnel_map ON fcip_port_tunnel_map.tunnel_id = fcip_tunnel.id
                LEFT JOIN switch_port AS local_switch_port ON fcip_port_tunnel_map.switchport_id =
local_switch_port.id
               LEFT JOIN switch_port AS remote_switch_port ON local_switch_port.remote_port_wwn IS NOT NULL
AND local_switch_port.remote_port_wwn != '' AND local_switch_port.wwn IS NOT NULL AND local_switch_port.wwn
!= ''
                                  AND local_switch_port.remote_port_wwn = remote_switch_port.wwn AND
local_switch_port.wwn = remote_switch_port.remote_port_wwn
                LEFT JOIN virtual_switch AS remote_virtual_switch ON remote_switch_port.virtual_switch_id =
remote_virtual_switch.id
                LEFT JOIN core_switch AS remote_core_switch ON remote_virtual_switch.core_switch_id =
remote_core_switch.id
                LEFT JOIN fabric_member AS remote_fabric_member ON remote_virtual_switch.id =
remote_fabric_member.virtual_switch_id
                LEFT JOIN fabric AS remote_fabric ON remote_fabric_member.fabric_id = remote_fabric.id
  WHERE local_fabric.managed = 1 AND local_virtual_switch.monitored = 1 AND (remote_fabric.managed IS NULL OR
remote_fabric.managed = 1) AND (remote_virtual_switch.monitored IS NULL OR remote_virtual_switch.monitored =
1);
```

FCIP_TUNNEL_INFO

```
create or replace view FCIP_TUNNEL_INFO as
```

```
select FCIP_TUNNEL.ID,
   FCIP_TUNNEL.TUNNEL_ID,
   FCIP_TUNNEL.VLAN_TAG,
   FCIP_TUNNEL.SOURCE_IP,
   FCIP_TUNNEL.DEST_IP,
   FCIP_TUNNEL.LOCAL_WWN,
   FCIP_TUNNEL.REMOTE_WWN_RESTRICT,
   FCIP_TUNNEL.COMMUNICATION_RATE,
   FCIP_TUNNEL.MIN_RETRANSMIT_TIME,
   FCIP_TUNNEL.SELECTIVE_ACK_ENABLED,
   FCIP_TUNNEL.KEEP_ALIVE_TIMEOUT,
   FCIP_TUNNEL.MAX_RETRANSMISSION,
   FCIP_TUNNEL.WAN_TOV_ENABLED,
   FCIP_TUNNEL.TUNNEL_STATUS,
   FCIP_TUNNEL.HA_STATUS,
   FCIP_TUNNEL.DESCRIPTION,
   FCIP_TUNNEL.FICON_TRB_ID_ENABLED,
   FCIP_TUNNEL.FICON_TT_EMUL_ENABLED,
   FCIP_TUNNEL.FICON_DLA_EMUL_ENABLED,
   FCIP_TUNNEL.FICON_TAPE_WRITE_MAX_PIPE,
   FCIP_TUNNEL.FICON_TAPE_READ_MAX_PIPE,
   FCIP_TUNNEL.FICON_TAPE_WRITE_MAX_OPS,
   FCIP_TUNNEL.FICON_TAPE_READ_MAX_OPS,
   FCIP_TUNNEL.FICON_TAPE_WRITE_TIMER,
   FCIP_TUNNEL.FICON_TAPE_MAX_WRITE_CHAIN,
   FCIP_TUNNEL.FICON_OXID_BASE,
   FCIP_TUNNEL.FICON_XRC_EMULATION_ENABLED,
   FCIP_TUNNEL.FICON_TW_EMUL_ENABLED,
   FCIP_TUNNEL.FICON_TR_EMUL_ENABLED,
   FCIP_TUNNEL.FICON_DEBUG_FLAGS,
   FCIP_TUNNEL.REMOTE_WWN,
   FCIP_TUNNEL.CDC,
   FCIP_TUNNEL.ADMIN_STATUS,
   FCIP_TUNNEL.CONTROL_L2_COS,
   FCIP_TUNNEL.DSCP_CONTROL,
   FCIP_TUNNEL.TRUNKING_ALGORITHM,
   FCIP_TUNNEL.EXTENDED_TUNNEL,
   FCIP_TUNNEL.VIRTUAL_SWITCH_ID,
   FCIP_TUNNEL.CIRCUIT_COUNT,
   FCIP_TUNNEL.MISMATCHED_CONFIG_DETAILS,
   FCIP_TUNNEL.SLOT_NUMBER,
   FCIP_TUNNEL.FICON_ENABLED,
   FCIP_TUNNEL.TPERF_ENABLED,
   FCIP_TUNNEL.AUTH_KEY,
   FCIP_TUNNEL.CONNECTED_COUNT,
   FCIP_TUNNEL.TUNNEL_STATUS_STRING,
   FCIP_TUNNEL.IP_EXTN_MODE,
   FCIP_TUNNEL.COMPRESSION_MODE,
   FCIP_TUNNEL.IP_COMPRESSION_MODE,
   FCIP_TUNNEL.TURBO_WRITE_ENABLED,
   FCIP_TUNNEL.TAPE_ACCELERATION_ENABLED,
   FCIP_TUNNEL.IPSEC_ENABLED,
   FCIP_TUNNEL.PRESHARED_KEY,
   FCIP_TUNNEL.IPSEC_POLICY_NAME,
   FCIP_TUNNEL.QOS_DISTRIBUTION_MODE,
   FCIP_TUNNEL.QOS_DISTRIBUTION_VALUE,
   FCIP_TUNNEL.QOS_HIGH,
   FCIP_TUNNEL.QOS_MEDIUM,
   FCIP_TUNNEL.QOS_LOW,
   FCIP_TUNNEL.IP_QOS,
   FCIP_TUNNEL.BACKWARD_COMPATIBLE,
   FCIP_TUNNEL.FICON_TERADATA_READ_ENABLED,
   FCIP_TUNNEL.FICON_TERADATA_WRITE_ENABLED,
```

```
    FCIP_TUNNEL.LOAD_LEVEL,
    PORT.WWN as VIRTUAL_PORT_WWN,
    PORT.REMOTE_PORT_WWN,
    PORT.REMOTE_NODE_WWN,
    PORT.ID as SWITCH_PORT_ID,
    PORT.PORT_NUMBER as SWITCH_PORT_NUMBER,
    PORT.USER_PORT_NUMBER,
    PORT.PORT_INDEX,
    PORT.STATUS_MESSAGE
   from FCIP_TUNNEL
     left join FCIP_PORT_TUNNEL_MAP on FCIP_PORT_TUNNEL_MAP.TUNNEL_ID = FCIP_TUNNEL.ID
     left join SWITCH_PORT PORT on FCIP_PORT_TUNNEL_MAP.SWITCHPORT_ID = PORT.ID;
```

FCOE_DEVICE_INFO

```
create or replace view FCOE_DEVICE_INFO as
select
    FCOE_DEVICE.DEVICE_NODE_ID,
    FCOE_DEVICE.DIRECT_ATTACH,
    FCOE_DEVICE.ATTACH_ID,
    FCOE_DEVICE.MAC_ADDRESS,
    DEVICE_NODE.TRUSTED,
    DEVICE_NODE.CREATION_TIME,
    DEVICE_NODE.MISSING,
    DEVICE_NODE.MISSING_TIME
from
    FCOE_DEVICE,
    DEVICE_NODE
where
    FCOE_DEVICE.DEVICE_NODE_ID = DEVICE_NODE.ID;
```

FRU_INFO

```
create or replace view FRU_INFO as
select
    FRU.ID,
    FRU.CORE_SWITCH_ID,
    FRU.TAG,
    FRU.PART_NUMBER,
    FRU.SERIAL_NUMBER,
    FRU.VENDOR_PART_NUMBER,
    FRU.VENDOR_SERIAL_NUMBER,
    FRU.CAN_BE_FRUED,
    FRU.SLOT_NUMBER,
    FRU.MANUFACTURER_DATE,
    FRU.UPDATE_DATE,
    FRU.VERSION,
    FRU.MANUFACTURER,
    FRU.VENDOR_EQUIPMENT_TYPE,
    FRU.OPERATIONAL_STATUS,
    FRU.TOTAL_OUTPUT_POWER,
    FRU.SPEED,
    FRU.CREATION_TIME,
    FRU.LAST_UPDATE_TIME,
    FRU.PREVIOUS_OP_STATUS,
    FRU.VENDOR,
    CORE_SWITCH.WWN as PHYSICAL_SWITCH_WWN,
    VIRTUAL_SWITCH.SWITCH_MODE as VIRTUAL_SWITCH_MODE,
    VIRTUAL_SWITCH.MANAGEMENT_STATE,
    VIRTUAL_SWITCH.MONITORED
from
```

```
    FRU,
    CORE_SWITCH,
    VIRTUAL_SWITCH
where
    FRU.CORE_SWITCH_ID = CORE_SWITCH.ID and
    FRU.CORE_SWITCH_ID = VIRTUAL_SWITCH.CORE_SWITCH_ID;
```

GIGE_PORT_ECLOUD_LINK_INFO

```
create or replace view GIGE_PORT_ECLOUD_LINK_INFO as
select
    GIGE_PORT_ETHERNET_CLOUD_LINK.ID,
    GIGE_PORT_ETHERNET_CLOUD_LINK.SWITCH_PORT_ID as GIGE_PORT_ID,
    GIGE_PORT_ETHERNET_CLOUD_LINK.CLOUD_ID,
    GIGE_PORT_ETHERNET_CLOUD_LINK.TRUSTED,
    GIGE_PORT_ETHERNET_CLOUD_LINK.CREATION_TIME,
    GIGE_PORT_ETHERNET_CLOUD_LINK.MISSING,
    GIGE_PORT_ETHERNET_CLOUD_LINK.MISSING_TIME,
    GIGE_PORT.SWITCH_PORT_ID,
    GIGE_PORT.PORT_TYPE,
    SWITCH_PORT.VIRTUAL_SWITCH_ID,
    SWITCH_PORT.USER_PORT_NUMBER,
    VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID
from
    GIGE_PORT_ETHERNET_CLOUD_LINK,
    GIGE_PORT,
    SWITCH_PORT,
    VIRTUAL_SWITCH
where
    GIGE_PORT_ETHERNET_CLOUD_LINK.SWITCH_PORT_ID = GIGE_PORT.ID and
    GIGE_PORT.SWITCH_PORT_ID = SWITCH_PORT.ID and
    SWITCH_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID;
```

GIGE_PORT_INFO

```
create or replace view GIGE_PORT_INFO as
select
    GIGE_PORT.ID,
    GIGE_PORT.SWITCH_PORT_ID,
    GIGE_PORT.PORT_NUMBER,
    GIGE_PORT.SLOT_NUMBER,
    GIGE_PORT.ENABLED,
    GIGE_PORT.SPEED,
    GIGE_PORT.MAX_SPEED,
    GIGE_PORT.MAC_ADDRESS,
    GIGE_PORT.PORT_NAME,
    GIGE_PORT.OPERATIONAL_STATUS,
    GIGE_PORT.LED_STATE,
    GIGE_PORT.SPEED_LED_STATE,
    GIGE_PORT.PORT_TYPE,
    GIGE_PORT.PERSISTENTLY_DISABLED,
    GIGE_PORT.INTERFACE_TYPE,
    GIGE_PORT.CHECKSUM,
    GIGE_PORT.FCIP_CAPABLE,
    coalesce(CARD.FCIP_CIRCUIT_CAPABLE, VIRTUAL_SWITCH.FCIP_CIRCUIT_CAPABLE) as FCIP_CIRCUIT_CAPABLE,
    GIGE_PORT.ISCSI_CAPABLE,
    GIGE_PORT.REMOTE_MAC_ADDRESS,
    GIGE_PORT.INBAND_MANAGEMENT_STATUS,
    GIGE_PORT.LAST_UPDATE,
    SWITCH_PORT.VIRTUAL_SWITCH_ID,
    SWITCH_PORT.USER_PORT_NUMBER,
```

```
    SWITCH_PORT.PORT_INDEX,
    SWITCH_PORT.SPEED_TYPE,
    VIRTUAL_SWITCH.WWN as VIRTUAL_SWITCH_WWN,
    VIRTUAL_SWITCH.MANAGEMENT_STATE,
    VIRTUAL_SWITCH.MONITORED,
    CORE_SWITCH.WWN as PHYSICAL_SWITCH_WWN
from
    GIGE_PORT,
    SWITCH_PORT,
    CORE_SWITCH
        left outer join CARD on CORE_SWITCH.ID = CARD.CORE_SWITCH_ID,
    VIRTUAL_SWITCH
where
    GIGE_PORT.SWITCH_PORT_ID = SWITCH_PORT.ID and
    SWITCH_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID and
    VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID and
    GIGE_PORT.SLOT_NUMBER in (0, CARD.SLOT_NUMBER);
```

GENERATED_BIRTREPORT_INFO

```
CREATE OR REPLACE VIEW generated_birtreport_info AS
SELECT generated_birtreport.id AS generated_report_id,
    generated_birtreport.file_name, generated_birtreport.rptdoc_store_location,
    generated_birtreport.generate_time, generated_birtreport.format,
    generated_birtreport.flagged, generated_birtreport.shared,
    generated_birtreport.report_template_id,
    report_template.name AS report_template_name,
    report_template.title AS report_template_title,
    report_template.created_time AS report_template_creation_time,
    report_template.created_by AS template_created_by,
    report_template.report_type AS template_type,
    report_template.parameterized AS template_is_parameterized,
    report_template.category AS template_category,
    report_template.shared AS template_is_shared,
    generated_birtreport.generated_by AS report_generated_by_user_id,
    user_.name AS report_generated_by_user_name,
    user_.email AS report_generated_by_user_email,
    generated_birtreport.scheduled, generated_birtreport.schedule_id,
    birtreport_schedule_config.name AS schedule_name,
    birtreport_schedule_config.report_store_location,
    birtreport_schedule_config.format_type AS scheduled_report_format,
    birtreport_schedule_config.deployment_id,
    birtreport_schedule_config.email_delivery,
    birtreport_schedule_config.folder_delivery,
CREATE OR REPLACE VIEW generated_birtreport_info AS
 SELECT generated_birtreport.id AS generated_report_id,
    generated_birtreport.file_name, generated_birtreport.rptdoc_store_location,
    generated_birtreport.generate_time, generated_birtreport.format,
    generated_birtreport.flagged, generated_birtreport.shared,
    generated_birtreport.report_template_id,
    generated_birtreport.report_generation_status,
    generated_birtreport.report_generation_status_description,
    report_template.name AS report_template_name,
    report_template.title AS report_template_title,
    report_template.created_time AS report_template_creation_time,
    report_template.created_by AS template_created_by,
    report_template.report_type AS template_type,
    report_template.parameterized AS template_is_parameterized,
    report_template.category AS template_category,
    report_template.shared AS template_is_shared,
    generated_birtreport.generated_by AS report_generated_by_user_id,
```

```
    user_.name AS report_generated_by_user_name,
    user_.email AS report_generated_by_user_email,
    generated_birtreport.scheduled, generated_birtreport.schedule_id,
    birtreport_schedule_config.name AS schedule_name,
    birtreport_schedule_config.report_store_location,
    birtreport_schedule_config.format_type AS scheduled_report_format,
    birtreport_schedule_config.deployment_id,
    birtreport_schedule_config.email_delivery,
    birtreport_schedule_config.folder_delivery,
    birtreport_schedule_config.email_recipients,
    birtreport_schedule_config.email_from,
    birtreport_schedule_config.email_replyto,
    birtreport_schedule_config.email_subject,
    birtreport_schedule_config.email_prologue,
    birtreport_schedule_config.email_epilogue,
    birtreport_schedule_config.last_modified_time AS last_modified_time_for_schedule,
    birtreport_schedule_config.created_time AS created_time_for_schedule
   FROM generated_birtreport
   LEFT JOIN report_template ON generated_birtreport.report_template_id = report_template.id
   LEFT JOIN user_ ON generated_birtreport.generated_by = user_.id
   LEFT JOIN birtreport_schedule_config ON generated_birtreport.schedule_id = birtreport_schedule_config.id;
```

HBA_PORT_DETAILS_INFO

```
create or replace view HBA_PORT_DETAILS_INFO as
select
    HBA_PORT.DEVICE_PORT_ID,
    HBA_PORT.CONFIGURED_STATE,
    HBA_PORT.CONFIGURED_SPEED,
    HBA_PORT.CONFIGURED_TOPOLOGY,
    HBA_PORT.MAX_SPEED_SUPPORTED,
    HBA_PORT.OPERATING_STATE,
    HBA_PORT.OPERATING_TOPOLOGY,
    HBA_PORT.SUPPORTED_FC4_TYPES,
    HBA_PORT.SUPPORTED_COS,
    HBA_PORT.TRUSTED as HBA_PORT_TRUSTED,
    HBA_PORT.CREATION_TIME as HBA_PORT_CREATION_TIME,
    HBA_PORT.MISSING as HBA_PORT_MISSING,
    HBA_PORT.MISSING_TIME as HBA_PORT_MISSING_TIME,
    HBA_PORT.OPERATING_SPEED,
    HBA_PORT.CNA_PORT_ID,
    HBA_PORT.PORT_NWWN,
    HBA_PORT.PHYSICAL_PORT_WWN,
    HBA_PORT.SWITCH_IP,
    HBA_PORT.PRINCIPAL_SWITCH_WWN,
    HBA_PORT.HBA_ID,
    HBA_PORT.PORT_NUMBER,
    HBA_PORT.NAME,
    HBA_PORT.FACTORY_PORT_WWN,
    HBA_PORT.FACTORY_NODE_WWN,
    HBA_PORT.PREBOOT_CREATED,
    HBA_PORT.MAX_BANDWIDTH,
    HBA_PORT.PCIF_INDEX,
    HBA_PORT.MAX_PCIF,
    HBA_PORT_DETAIL.PERSISTENT_BINDING,
    HBA_PORT_DETAIL.FABRIC_NAME,
    HBA_PORT_DETAIL.BOOT_OVER_SAN,
    HBA_PORT_DETAIL.BOOT_OPTION,
    HBA_PORT_DETAIL.BOOT_SPEED,
    HBA_PORT_DETAIL.BOOT_TOPOLOGY,
    HBA_PORT_DETAIL.BB_CREDIT,
    HBA_PORT_DETAIL.FRAME_DATA_FIELD_SIZE,
```

```
    HBA_PORT_DETAIL.HARDWARE_PATH,
    HBA_PORT_DETAIL.V_PORT_COUNT,
    HBA_PORT_DETAIL.QUEUE_DEPTH,
    HBA_PORT_DETAIL.INTERRUPT_CONTROL_COALESCE,
    HBA_PORT_DETAIL.INTERRUPT_CONTROL_LATENCY,
    HBA_PORT_DETAIL.INTERRUPT_CONTROL_DELAY,
    HBA_PORT_DETAIL.BEACON_STATE,
    HBA_PORT_DETAIL.LINK_BEACON_STATE,
    HBA_PORT_DETAIL.MPIO_MODE_STATE,
    HBA_PORT_DETAIL.PATH_TIME_OUT,
    HBA_PORT_DETAIL.LOGGING_LEVEL,
    HBA_PORT_DETAIL.TARGET_RATE_LIMIT,
    HBA_PORT_DETAIL.DEFAULT_RATE_LIMIT,
    HBA_PORT_DETAIL.VF_MODE,
    HBA_PORT_DETAIL.RECIEVE_BUFFER_CREDIT,
    HBA_PORT_DETAIL.TRANSMIT_BUFFER_CREDIT,
    HBA_PORT_DETAIL.FCSP_AUTH_STATE,
    HBA_PORT_DETAIL.FCSP_STATUS,
    HBA_PORT_DETAIL.FCSP_ALGORITHM,
    HBA_PORT_DETAIL.FCSP_GROUP,
    HBA_PORT_DETAIL.FCSP_ERROR_STATUS,
    HBA_PORT_DETAIL.QOS_CONFIGURED_STATE,
    HBA_PORT_DETAIL.QOS_OPERATING_STATE,
    HBA_PORT_DETAIL.QOS_TOTAL_BB_CREDIT,
    HBA_PORT_DETAIL.QOS_PRIORITY_LEVEL,
    HBA_PORT_DETAIL.QOS_HIGH_BW_ALLOCATION,
    HBA_PORT_DETAIL.QOS_MEDIUM_BW_ALLOCATION,
    HBA_PORT_DETAIL.QOS_LOW_BW_ALLOCATION,
    HBA_PORT_DETAIL.MEDIA as MEDIA,
    HBA_PORT_DETAIL.IOC_ID as IOC_ID,
    HBA_PORT_DETAIL.PREBOOT_DISABLED,
    HBA_PORT_FCOE_DETAILS.BANDWIDTH as FCOE_BANDWIDTH,
    HBA_PORT_FCOE_DETAILS.FIP_STATE,
    HBA_PORT_FCOE_DETAILS.DISCOVERY_PRIORITY,
    HBA_PORT_FCOE_DETAILS.FCF_FCMAP,
    HBA_PORT_FCOE_DETAILS.FCF_FPMA_MAC,
    HBA_PORT_FCOE_DETAILS.FCF_MAC,
    HBA_PORT_FCOE_DETAILS.FCF_MODE,
    HBA_PORT_FCOE_DETAILS.FCF_NAMEID,
    HBA_PORT_FCOE_DETAILS.FCPIM_MPIO_MODE,
    HBA_PORT_FCOE_DETAILS.PORT_LOG_ENABLED,
    HBA_PORT_FCOE_DETAILS.MAX_FRAME_SIZE as FCOE_MAX_FRAME_SIZE,
    HBA_PORT_FCOE_DETAILS.MTU as FCOE_MTU,
    HBA_PORT_FCOE_DETAILS.PATH_TOV as FCOE_PATH_TOV,
    HBA_PORT_FCOE_DETAILS.SCSI_QUEUE_DEPTH as FCOE_SCSI_QUEUE_DEPTH,
    HBA_PORT_FCOE_DETAILS.STATE as FCOE_STATE,
    HBA_PORT_FCOE_DETAILS.SUPPORTED_CLASS as FCOE_SUPPORTED_CLASS,
    HBA_PORT_FCOE_DETAILS.TRL_SPEED as FCOE_TRL_SPEED,
    HBA_PORT_FCOE_DETAILS.TRL_STATE as FCOE_TRL_STATE,
    HBA_PORT_FCOE_DETAILS.PG_ID as FCOE_PG_ID,
    HBA_PORT_FCOE_DETAILS.PRIORITIES as FCOE_PRIORITIES,
    HBA_PORT_FCOE_DETAILS.FCOE_MAC,
    HBA_PORT.SYNTHETIC_FC,
    HBA_PORT_DETAIL.ALARM_WARNING,
    HBA_PORT_DETAIL.IO_EXEC_THROTTLE_MAX,
    HBA_PORT_DETAIL.IO_EXEC_THROTTLE_OPERATIONAL,
    HBA_PORT_DETAIL.IO_EXEC_THROTTLE_CONFIGURED,
    HBA_PORT_DETAIL.BOOTUP_DELAY,
    HBA_PORT_DETAIL.FEC_STATE,
    HBA_PORT_DETAIL.BB_CREDIT_RECOVERY_STATUS,
    HBA_PORT_DETAIL.CONFIGURED_BB_SCN_COUNT,
    HBA_PORT_DETAIL.NEGOTIATED_BB_SCN_COUNT
from
```

```
    HBA_PORT
        left outer join HBA_PORT_DETAIL
            on HBA_PORT.DEVICE_PORT_ID = HBA_PORT_DETAIL.DEVICE_PORT_ID
        left outer join HBA_PORT_FCOE_DETAILS
            on HBA_PORT.DEVICE_PORT_ID = HBA_PORT_FCOE_DETAILS.DEVICE_PORT_ID;
```

HBA_TARGET_INFO

```
create or replace view HBA_TARGET_INFO as
select
    HBA_TARGET.DEVICE_PORT_ID,
    HBA_TARGET.HBA_REMOTE_PORT_LUN_ID,
    HBA_TARGET.BOOT_LUN,
    HBA_TARGET.TRUSTED,
    HBA_TARGET.CREATION_TIME,
    HBA_TARGET.MISSING,
    HBA_TARGET.MISSING_TIME,
    HBA_TARGET.TARGET_ID as HBA_PORT_TARGET_ID,
    HBA_REMOTE_PORT.ID as HBA_REMOTE_PORT_ID,
    HBA_REMOTE_PORT.SYMBOLIC_NAME,
    HBA_REMOTE_PORT.PORT_WWN,
    HBA_REMOTE_PORT.NODE_WWN,
    HBA_REMOTE_PORT.NAME,
    HBA_REMOTE_PORT.FC_ADDRESS,
    HBA_REMOTE_PORT.FRAME_DATA_SIZE,
    HBA_REMOTE_PORT.SPEED,
    HBA_REMOTE_PORT.STATE,
    HBA_REMOTE_PORT.SUPPORTED_COS,
    HBA_REMOTE_PORT.DEVICE_TYPE,
    HBA_REMOTE_PORT.BIND_TYPE,
    HBA_REMOTE_PORT.TARGET_ID,
    HBA_REMOTE_PORT.ROLE,
    HBA_REMOTE_PORT.VENDOR,
    HBA_REMOTE_PORT.PRODUCT_ID,
    HBA_REMOTE_PORT.PRODUCT_VERSION,
    HBA_REMOTE_PORT.QOS_PRIORITY,
    HBA_REMOTE_PORT.QOS_FLOW_ID,
    HBA_REMOTE_PORT.CURRENT_SPEED,
    HBA_REMOTE_PORT.TRL_ENFORCED,
    HBA_REMOTE_PORT.BUS_NO,
    HBA_REMOTE_PORT_LUN.FCP_LUN,
    HBA_REMOTE_PORT_LUN.CAPACITY,
    HBA_REMOTE_PORT_LUN.BLOCK_SIZE,
    HBA_REMOTE_PORT_LUN.VENDOR as LUN_VENDOR,
    HBA_REMOTE_PORT_LUN.PRODUCT_ID as LUN_PRODUCT_ID,
    HBA_REMOTE_PORT_LUN.PRODUCT_VERSION as LUN_PRODUCT_VERSION,
    HBA_REMOTE_PORT_LUN.PRODUCT_SERIAL_NO,
    HBA_REMOTE_PORT_LUN.TARGET_WWN,
    HBA_REMOTE_PORT_LUN.PHYSICAL_LUN,
    HBA_REMOTE_PORT_LUN.LUN_ID,
    HBA_REMOTE_PORT.FCP_IM_STATE,
    HBA_REMOTE_PORT.IO_LATENCY_MIN,
    HBA_REMOTE_PORT.IO_LATENCY_MAX,
    HBA_REMOTE_PORT.IO_LATENCY_AVERAGE,
    HBA_REMOTE_PORT.DATA_RETRANSMISSION_SUPPORT,
    HBA_REMOTE_PORT.REC_SUPPORT,
    HBA_REMOTE_PORT.TASK_RENTRY_IDENT_SUPPORT,
    HBA_REMOTE_PORT.CONFIRMED_COMPLETIONS_SUPPORT
from
    HBA_TARGET, HBA_REMOTE_PORT, HBA_REMOTE_PORT_LUN
where
    HBA_TARGET.HBA_REMOTE_PORT_LUN_ID = HBA_REMOTE_PORT_LUN.ID and
```

```
    HBA_REMOTE_PORT.ID = HBA_REMOTE_PORT_LUN.HBA_REMOTE_PORT_ID;
```

HEALTH_STATUS_INFO

```
create or replace view HEALTH_STATUS_INFO as
select
    DEPLOYMENT_CONFIGURATION.ID as CONFIGURATION_ID,
    DEPLOYMENT_CONFIGURATION.NAME,
    DEPLOYMENT_STATUS.ID as STATUS_ID,
    DEPLOYMENT_STATUS.DEPLOYMENT_TIME,
    DEPLOYMENT_STATUS.DEPLOYED_BY,
    HEALTH_STATUS.RULE_ID,
    HEALTH_STATUS.RULE_DESCRIPTION,
    HEALTH_TARGET_STATUS.TARGET_ID,
    HEALTH_TARGET_STATUS.TARGET_TYPE,
    HEALTH_TARGET_STATUS.STATUS,
    HEALTH_TARGET_STATUS.MESSAGE,
    HEALTH_TARGET_STATUS.LEGACY_NAME
from
    DEPLOYMENT_CONFIGURATION,
    DEPLOYMENT_STATUS,
    HEALTH_STATUS,
    HEALTH_TARGET_STATUS
where
    DEPLOYMENT_STATUS.DEPLOYMENT_CONFIGURATION_ID = DEPLOYMENT_CONFIGURATION.ID
    and HEALTH_STATUS.DEPLOYMENT_STATUS_ID = DEPLOYMENT_STATUS.ID
    and HEALTH_TARGET_STATUS.HEALTH_STATUS_ID = HEALTH_STATUS.ID;
```

HOST_INVENTORY_REPORT_INFO

```
CREATE OR REPLACE VIEW host_inventory_report_info AS
select
DEVICE_ENCLOSURE.ID as DEVICE_ENCLOSURE_ID,
DEVICE_ENCLOSURE.NAME as HOST_NAME,
DEVICE_ENCLOSURE.IP_ADDRESS as HOST_IP,
case when (DEVICE_ENCLOSURE.VENDOR is null or DEVICE_ENCLOSURE.VENDOR = '') then 'NA' else

DEVICE_ENCLOSURE.VENDOR end as HOST_VENDOR,
DEVICE_ENCLOSURE.MODEL as HOST_MODEL,
DEVICE_ENCLOSURE.OS as HOST_OS,
DEVICE_ENCLOSURE.LOCATION as HOST_LOCATION,
DEVICE_ENCLOSURE.CONTACT as HOST_CONTACT,
DEVICE_ENCLOSURE.DESCRIPTION as HOST_DESC,
DEVICE_ENCLOSURE.HCM_AGENT_VERSION as AGENT_VERSION,
USER_DEFINED_DEVICE_DETAIL.NAME as ADAPTER_NAME,
HBA.NAME as ADAPTER_NAME_HCM,
HBA.WWN as ADAPTER_WWN,
HBA.OPERATING_STATUS as ADAPTER_STATUS,
HBA.MODEL as ADAPTER_MODEL,
HBA.VENDOR as ADPATER_VENDOR,
HBA.SERIAL_NUMBER as ADAPTER_SERIAL_NO,
HBA.FIRMWARE_VERSION as ADAPTER_FIRMWARE,
HBA.BIOS_VERSION as ADAPTER_BIOS_VERSION,
HBA.DRIVER_VERSION as ADAPTER_DRIVER_VERSION,
case when (ADAPTER_PORT.CNA_PORT_ID is not null or HBA.VPD_OEM_INFO = '') then 'N/A' else

HBA.VPD_OEM_INFO end as ADAPTER_OEM_INFO,
case when ADAPTER_PORT.PCIF_INDEX != '' then ADAPTER_PORT.PCIF_INDEX else

ADAPTER_PORT.PORT_NUMBER :: varchar end as ADAPTER_PORT_ID,
case when ADAPTER_PORT.TYPE = 'IP-Port' then ADAPTER_PORT.ETH_DEV else
```

```
DEVICE_PORT_INFO.NAME end as ADAPTER_PORT_NAME,
ADAPTER_PORT.NAME as ADAPTER_PORT_NAME_HCM,
DEVICE_PORT_INFO.SYMBOLIC_NAME as ADAPTER_SYMBOLIC_NAME,
case when ADAPTER_PORT.TYPE != 'IP-Port' then ADAPTER_PORT.PORT_NWWN else

ADAPTER_PORT.MAC_ADDRESS end as ADAPTER_NODE_WWN,
case when ADAPTER_PORT.type != 'IP-Port' then DEVICE_PORT_INFO.WWN else

ADAPTER_PORT.MAC_ADDRESS end as ADAPTER_PORT_WWN,
HBA_PORT_DETAIL.QOS_CONFIGURED_STATE as ADAPTER_PORT_QOS_CONFIGURED_STATE,
HBA_PORT_DETAIL.FRAME_DATA_FIELD_SIZE as ADAPTER_PORT_FRAME_DATA_FIELD_SIZE,
HBA_PORT_DETAIL.MEDIA as ADAPTER_PORT_MEDIA,
HBA_PORT_DETAIL.ALARM_WARNING as ADAPTER_ALARM_WARNING,
DEVICE_PORT_INFO.port_id as adapter_port_FCaadress,
case when ADAPTER_PORT.type = 'FC' then DEVICE_PORT_INFO.type else ADAPTER_PORT.type

end as adapter_port_type,
HBA_PORT_DETAIL.FABRIC_NAME as ADAPTER_PORT_FABRIC_NAME,
case when (HBA_PORT_DETAIL.FAA_STATUS = 0) then 'NA' else (case when

(HBA_PORT_DETAIL.FAA_STATUS = 1) then 'Disabled' else 'Enabled' end) end as

ADAPTER_PORT_FAA_STATUS,
case when (HBA_PORT_DETAIL.WWN_SOURCE is null or HBA_PORT_DETAIL.WWN_SOURCE = '')

then 'NA' else HBA_PORT_DETAIL.WWN_SOURCE end as ADAPTER_PORT_WWN_SOURCE,
HBA_PORT_DETAIL.BOOT_OVER_SAN as ADAPTER_PORT_BOOT_OVER_SAN,
ADAPTER_PORT.MAX_SPEED_SUPPORTED as ADAPTER_PORT_MAX_SPEED_SUPPORTED,
ADAPTER_PORT.OPERATING_STATE as PORT_OPERATING_STATE,
ADAPTER_PORT_COUNT.PORT_COUNT as HBA_PORT_COUNT,
HBA.ID as HBA_PORT_ID,
ADAPTER_PORT_FABRIC_MAP.FABRIC_NAME,
ADAPTER_PORT_FABRIC_MAP.FABRIC_ID,
ADAPTER_PORT_FABRIC_MAP.FABRIC_PRINCIPAL_SWITCH_WWN,
ADAPTER_PORT_FABRIC_MAP.FABRIC_SEED_SWITCH_WWN,
ADAPTER_PORT_FABRIC_MAP. device_node_id,
ADAPTER_PORT_FABRIC_MAP. virtual_swcith_id,
ADAPTER_PORT_FABRIC_MAP.edge_virtual_switch_wwn,
ADAPTER_PORT_FABRIC_MAP.edge_switch_name,
ADAPTER_PORT_FABRIC_MAP.edge_switch_virtual_fabric_id,
ADAPTER_PORT_FABRIC_MAP.edge_virtual_switch_monitored,
ADAPTER_PORT_FABRIC_MAP.edge_virtual_switch_domain_id,
ADAPTER_PORT_FABRIC_MAP.sp_category,
ADAPTER_PORT_FABRIC_MAP.sp_licensed,
ADAPTER_PORT_FABRIC_MAP.sp_name,
ADAPTER_PORT_FABRIC_MAP.sp_slot_number,
ADAPTER_PORT_FABRIC_MAP.sp_port_number,
ADAPTER_PORT_FABRIC_MAP.sp_port_id,
ADAPTER_PORT_FABRIC_MAP.sp_port_index,
ADAPTER_PORT_FABRIC_MAP.sp_area_id,
ADAPTER_PORT_FABRIC_MAP.sp_mac_address,
ADAPTER_PORT_FABRIC_MAP.sp_status,
ADAPTER_PORT_FABRIC_MAP.sp_state,
ADAPTER_PORT_FABRIC_MAP.edge_core_switch_smodel,
ADAPTER_PORT_FABRIC_MAP.edge_core_switch_ip_address,
ADAPTER_PORT_FABRIC_MAP.edge_core_switch_physical_switch_wwn,
ADAPTER_PORT_FABRIC_MAP.edge_core_switch_operational_status,
ADAPTER_PORT_FABRIC_MAP.edge_core_switch_name,
ADAPTER_PORT_FABRIC_MAP.edge_core_switch_type,
ADAPTER_PORT_FABRIC_MAP.edge_core_switch_model,
ADAPTER_PORT_FABRIC_MAP.edge_core_switch_vendor
```

```
from
DEVICE_ENCLOSURE,
HBA left join USER_DEFINED_DEVICE_DETAIL on (HBA.WWN =

USER_DEFINED_DEVICE_DETAIL.WWN),
HBA_PORT_DETAIL,
DEVICE_PORT_INFO left join
(select FABRIC.ID as FABRIC_ID, FABRIC.NAME as FABRIC_NAME, DEVICE_PORT_INFO.WWN as

ADAPTER_PORT_WWN, FABRIC.SEED_SWITCH_WWN AS

FABRIC_SEED_SWITCH_WWN,FABRIC.PRINCIPAL_SWITCH_WWN AS

FABRIC_PRINCIPAL_SWITCH_WWN , dn.id As device_node_id, vs.id As virtual_swcith_id,
vs.wwn AS edge_virtual_switch_wwn,
vs.name AS edge_switch_name,
vs.virtual_fabric_id AS edge_switch_virtual_fabric_id,
vs.monitored AS  edge_virtual_switch_monitored,
vs.domain_id As edge_virtual_switch_domain_id,
sp.category AS sp_category,
sp.licensed AS sp_licensed,
sp.name AS sp_name,
sp.slot_number AS sp_slot_number,
sp.port_number AS sp_port_number,
sp.port_id AS sp_port_id,
sp.port_index AS sp_port_index,
sp.area_id AS sp_area_id,
sp.mac_address AS sp_mac_address,
sp.status AS sp_status,
sp.state AS sp_state,
cs.model AS edge_core_switch_smodel,
cs.ip_address AS edge_core_switch_ip_address,
cs.wwn AS  edge_core_switch_physical_switch_wwn,
cs.operational_status AS edge_core_switch_operational_status,
cs.name AS edge_core_switch_name,
cs.type AS edge_core_switch_type,
cs.model AS edge_core_switch_model,
cs.vendor AS edge_core_switch_vendor
from FABRIC, DEVICE_PORT_INFO
LEFT JOIN device_node dn ON DEVICE_PORT_INFO.node_id = dn.id
LEFT JOIN switch_port sp ON DEVICE_PORT_INFO.switch_port_wwn = sp.wwn
LEFT JOIN virtual_switch vs ON sp.virtual_switch_id = vs.id
LEFT JOIN core_switch cs ON vs.core_switch_id = cs.id
where FABRIC.ID = DEVICE_PORT_INFO.FABRIC_ID and FABRIC.TYPE not in (65, 66)
)as ADAPTER_PORT_FABRIC_MAP on ADAPTER_PORT_FABRIC_MAP.ADAPTER_PORT_WWN =

DEVICE_PORT_INFO.WWN,

(select DEVICE_PORT_ID, CONFIGURED_STATE, CONFIGURED_SPEED, CONFIGURED_TOPOLOGY,
MAX_SPEED_SUPPORTED, OPERATING_STATE, OPERATING_TOPOLOGY, SUPPORTED_FC4_TYPES,
SUPPORTED_COS, TRUSTED, CREATION_TIME, MISSING, MISSING_TIME,
OPERATING_SPEED, CNA_PORT_ID, PORT_NWWN, PHYSICAL_PORT_WWN, SWITCH_IP,
PRINCIPAL_SWITCH_WWN, HBA_ID, PORT_NUMBER, NAME, FACTORY_PORT_WWN,
FACTORY_NODE_WWN, PREBOOT_CREATED, MAX_BANDWIDTH, PCIF_INDEX,
MAX_PCIF, SYNTHETIC_FC, 'FC' as TYPE, 0 as ETH_PORT_ID, '' as ETH_DEV, '' as MAC_ADDRESS
from HBA_PORT

union

select DEVICE_PORT_ID, CONFIGURED_STATE, CONFIGURED_SPEED, CONFIGURED_TOPOLOGY,
MAX_SPEED_SUPPORTED, OPERATING_STATE, OPERATING_TOPOLOGY, SUPPORTED_FC4_TYPES,
SUPPORTED_COS, HBA_PORT.TRUSTED, CNA_ETH_PORT.CREATION_TIME, HBA_PORT.MISSING,
```

```
HBA_PORT.MISSING_TIME,
OPERATING_SPEED, HBA_PORT.CNA_PORT_ID, PORT_NWWN, PHYSICAL_PORT_WWN, SWITCH_IP,
PRINCIPAL_SWITCH_WWN, HBA_PORT.HBA_ID, HBA_PORT.PORT_NUMBER, CNA_PORT.NAME,

HBA_PORT.FACTORY_PORT_WWN,
HBA_PORT.FACTORY_NODE_WWN, HBA_PORT.PREBOOT_CREATED,

CNA_ETH_PORT.MAX_BANDWIDTH, CNA_ETH_PORT.PCIF_INDEX,
CNA_ETH_PORT.MAX_PCIF, SYNTHETIC_FC, 'IP-Port' as TYPE,
CNA_ETH_PORT.ID as ETH_PORT_ID, CNA_ETH_PORT.ETH_DEV, CNA_ETH_PORT.MAC_ADDRESS

from DEVICE_PORT, HBA_PORT, CNA_PORT, CNA_ETH_PORT
where DEVICE_PORT.ID = HBA_PORT.DEVICE_PORT_ID and HBA_PORT.CNA_PORT_ID =

CNA_ETH_PORT.CNA_PORT_ID and CNA_PORT.ID = CNA_ETH_PORT.CNA_PORT_ID and

DEVICE_PORT.NPV_PHYSICAL=0
) as ADAPTER_PORT,

(select HBA_PORT.HBA_ID, count(HBA_PORT.DEVICE_PORT_ID) as PORT_COUNT from HBA_PORT,

DEVICE_PORT_INFO where HBA_PORT.DEVICE_PORT_ID = DEVICE_PORT_INFO.ID and

DEVICE_PORT_INFO.NPV_PHYSICAL = 0 group by (HBA_PORT.HBA_ID)) as ADAPTER_PORT_COUNT

where
DEVICE_ENCLOSURE.MANAGED_BY in (2,4)
and DEVICE_ENCLOSURE.ID = HBA.HOST_ID
and ADAPTER_PORT.HBA_ID = HBA.ID
and ADAPTER_PORT.DEVICE_PORT_ID = DEVICE_PORT_INFO.ID
and HBA_PORT_DETAIL.DEVICE_PORT_ID = DEVICE_PORT_INFO.ID
and ADAPTER_PORT_COUNT.HBA_ID = HBA.ID
order by
DEVICE_ENCLOSURE.IP_ADDRESS, HBA.WWN, ADAPTER_PORT.PHYSICAL_PORT_WWN;
```

HOST_DISCOVERY_REQUEST_INFO

```
create or replace view HOST_DISCOVERY_REQUEST_INFO as
select
    HOST_DISCOVERY_REQUEST.ID,
    HOST_DISCOVERY_REQUEST.HOST_NAME AS REQUEST_HOST_NAME,
    HOST_DISCOVERY_REQUEST.DEVICE_ENCLOSURE_ID,
    HOST_DISCOVERY_REQUEST.REQUEST_GROUP_ID,
    HOST_DISCOVERY_REQUEST.HOST_DISCOVERY_OPTION_ID,
    HOST_DISCOVERY_REQUEST.VM_MANAGEMENT_STATE,
    HOST_DISCOVERY_REQUEST.JSON_MANAGEMENT_STATE,
    HOST_DISCOVERY_REQUEST.CIM_MANAGEMENT_STATE,
    HOST_DISCOVERY_REQUEST.MANAGEMENT_STATE,
    HOST_DISCOVERY_OPTION.DISCOVER_JSON,
    HOST_DISCOVERY_OPTION.JSON_USERNAME,
    HOST_DISCOVERY_OPTION.JSON_PASSWD,
    HOST_DISCOVERY_OPTION.DISCOVER_CIM,
    HOST_DISCOVERY_OPTION.CIM_IMPL,
    HOST_DISCOVERY_OPTION.CIM_USERNAME,
    HOST_DISCOVERY_OPTION.CIM_PASSWORD,
    HOST_DISCOVERY_OPTION.CIM_NAMESPACE,
    HOST_DISCOVERY_OPTION.CIM_PORT,
    HOST_DISCOVERY_OPTION.DISCOVER_VM,
    HOST_DISCOVERY_OPTION.VM_USERNAME,
    HOST_DISCOVERY_OPTION.VM_PASSWORD,
    HOST_DISCOVERY_OPTION.JSON_PORT,
    HOST_DISCOVERY_OPTION.VM_PORT,
```

```
    HOST_DISCOVERY_OPTION.Application_Name_USER_NAME,
    HOST_DISCOVERY_OPTION.Application_Name_SERVER_ADDRESS,
    DEVICE_ENCLOSURE.NAME,
    DEVICE_ENCLOSURE.TYPE,
    DEVICE_ENCLOSURE.ICON,
    DEVICE_ENCLOSURE.OS,
    DEVICE_ENCLOSURE.APPLICATIONS,
    DEVICE_ENCLOSURE.DEPARTMENT,
    DEVICE_ENCLOSURE.CONTACT,
    DEVICE_ENCLOSURE.LOCATION,
    DEVICE_ENCLOSURE.DESCRIPTION,
    DEVICE_ENCLOSURE.COMMENT_,
    DEVICE_ENCLOSURE.IP_ADDRESS,
    DEVICE_ENCLOSURE.VENDOR,
    DEVICE_ENCLOSURE.MODEL,
    DEVICE_ENCLOSURE.SERIAL_NUMBER,
    DEVICE_ENCLOSURE.FIRMWARE,
    DEVICE_ENCLOSURE.USER_DEFINED_VALUE1,
    DEVICE_ENCLOSURE.USER_DEFINED_VALUE2,
    DEVICE_ENCLOSURE.USER_DEFINED_VALUE3,
    DEVICE_ENCLOSURE.HCM_AGENT_VERSION,
    DEVICE_ENCLOSURE.OS_VERSION,
    DEVICE_ENCLOSURE.CREATED_BY,
    DEVICE_ENCLOSURE.TRACK_CHANGES,
    DEVICE_ENCLOSURE.LAST_UPDATE_TIME,
    DEVICE_ENCLOSURE.LAST_UPDATE_MODULE,
    DEVICE_ENCLOSURE.TRUSTED,
    DEVICE_ENCLOSURE.CREATION_TIME,
    DEVICE_ENCLOSURE.MISSING,
    DEVICE_ENCLOSURE.MISSING_TIME,
    DEVICE_ENCLOSURE.HOST_NAME,
    DEVICE_ENCLOSURE.SYSLOG_REGISTERED,
    DEVICE_ENCLOSURE.VIRTUALIZATION,
    DEVICE_ENCLOSURE.MANAGED_ELEMENT_ID,
    HOST_DISCOVERY_REQUEST.MANAGEMENT_STATE_DETAILS
from
    HOST_DISCOVERY_REQUEST
        join HOST_DISCOVERY_OPTION on HOST_DISCOVERY_REQUEST.HOST_DISCOVERY_OPTION_ID =
HOST_DISCOVERY_OPTION.ID
        left outer join DEVICE_ENCLOSURE on HOST_DISCOVERY_REQUEST.DEVICE_ENCLOSURE_ID = DEVICE_ENCLOSURE.ID;
```

HOST_DISCOVERY_REQUESTS_INFO

```
create or replace view HOST_DISCOVERY_REQUESTS_INFO as
 select HOST_DISCOVERY_REQUESTS.ID,
    HOST_DISCOVERY_REQUESTS.HOST_NAME as REQUEST_HOST_NAME,
    HOST_DISCOVERY_REQUESTS.DEVICE_ENCLOSURE_ID,
    HOST_DISCOVERY_REQUESTS.HOST_DISCOVERY_REQ_GROUP_ID,
    HOST_DISCOVERY_REQUESTS.HOST_DISCOVERY_OPTIONS_ID,
    HOST_DISCOVERY_REQUESTS.DISCOVERY_STATE,
    HOST_DISCOVERY_REQUESTS.REQUEST_STATE,
    HOST_DISCOVERY_REQUESTS.REQUEST_STATE_DETAILS,
    HOST_DISCOVERY_OPTIONS.DISCOVERY_TYPE,
    HOST_DISCOVERY_OPTIONS.USERNAME,
    HOST_DISCOVERY_OPTIONS.PASSWORD,
    HOST_DISCOVERY_OPTIONS.PORT,
    HOST_DISCOVERY_OPTIONS.SSL_ENABLED,
    HOST_DISCOVERY_OPTIONS.NAMESPACE,
    DEVICE_ENCLOSURE.NAME,
    DEVICE_ENCLOSURE.TYPE,
    DEVICE_ENCLOSURE.ICON,
    DEVICE_ENCLOSURE.OS,
```

```
     DEVICE_ENCLOSURE.APPLICATIONS,
     DEVICE_ENCLOSURE.DEPARTMENT,
     DEVICE_ENCLOSURE.CONTACT,
     DEVICE_ENCLOSURE.LOCATION,
     DEVICE_ENCLOSURE.DESCRIPTION,
     DEVICE_ENCLOSURE.COMMENT_,
     DEVICE_ENCLOSURE.IP_ADDRESS,
     DEVICE_ENCLOSURE.VENDOR,
     DEVICE_ENCLOSURE.MODEL,
     DEVICE_ENCLOSURE.SERIAL_NUMBER,
     DEVICE_ENCLOSURE.FIRMWARE,
     DEVICE_ENCLOSURE.USER_DEFINED_VALUE1,
     DEVICE_ENCLOSURE.USER_DEFINED_VALUE2,
     DEVICE_ENCLOSURE.USER_DEFINED_VALUE3,
     DEVICE_ENCLOSURE.HCM_AGENT_VERSION,
     DEVICE_ENCLOSURE.OS_VERSION,
     DEVICE_ENCLOSURE.CREATED_BY,
     DEVICE_ENCLOSURE.TRACK_CHANGES,
     DEVICE_ENCLOSURE.LAST_UPDATE_TIME,
     DEVICE_ENCLOSURE.LAST_UPDATE_MODULE,
     DEVICE_ENCLOSURE.TRUSTED,
     DEVICE_ENCLOSURE.CREATION_TIME,
     DEVICE_ENCLOSURE.MISSING,
     DEVICE_ENCLOSURE.MISSING_TIME,
     DEVICE_ENCLOSURE.HOST_NAME,
     DEVICE_ENCLOSURE.SYSLOG_REGISTERED,
     DEVICE_ENCLOSURE.VIRTUALIZATION,
     DEVICE_ENCLOSURE.MANAGED_ELEMENT_ID
    from HOST_DISCOVERY_REQUESTS
    join HOST_DISCOVERY_OPTIONS on HOST_DISCOVERY_REQUESTS.HOST_DISCOVERY_OPTIONS_ID =
HOST_DISCOVERY_OPTIONS.ID
    left join DEVICE_ENCLOSURE on HOST_DISCOVERY_REQUESTS.DEVICE_ENCLOSURE_ID = DEVICE_ENCLOSURE.ID;

    create or replace view FCIP_TUNNEL_CIRCUIT_INFO as
 select FCIP_TUNNEL_CIRCUIT.ID, FCIP_TUNNEL_CIRCUIT.TUNNEL_ID,
     FCIP_TUNNEL_CIRCUIT.CIRCUIT_NUMBER, FCIP_TUNNEL_CIRCUIT.COMPRESSION_ENABLED,
     FCIP_TUNNEL_CIRCUIT.TURBO_WRITE_ENABLED,
     FCIP_TUNNEL_CIRCUIT.TAPE_ACCELERATION_ENABLED,
     FCIP_TUNNEL_CIRCUIT.IKE_POLICY_NUM, FCIP_TUNNEL_CIRCUIT.IPSEC_POLICY_NUM,
     FCIP_TUNNEL_CIRCUIT.PRESHARED_KEY, FCIP_TUNNEL_CIRCUIT.SOURCE_IP,
     FCIP_TUNNEL_CIRCUIT.DEST_IP, FCIP_TUNNEL_CIRCUIT.VLAN_TAG,
     FCIP_TUNNEL_CIRCUIT.DP1_SOURCE_IP,FCIP_TUNNEL_CIRCUIT.DP1_DEST_IP,FCIP_TUNNEL_CIRCUIT.DP1_VLAN_TAG,
    FCIP_TUNNEL_CIRCUIT.SELECTIVE_ACK, FCIP_TUNNEL_CIRCUIT.QOS_MAPPING,
     FCIP_TUNNEL_CIRCUIT.PATH_MTU_DISCOVERY, FCIP_TUNNEL_CIRCUIT.MIN_COMM_RATE,
     FCIP_TUNNEL_CIRCUIT.MAX_COMM_RATE, FCIP_TUNNEL_CIRCUIT.MIN_RETRANSMIT_TIME,
     FCIP_TUNNEL_CIRCUIT.MAX_RETRANSMIT_TIME,
     FCIP_TUNNEL_CIRCUIT.KEEP_ALIVE_TIMEOUT, FCIP_TUNNEL_CIRCUIT.ADMIN_STATUS,
     FCIP_TUNNEL_CIRCUIT.METRIC, FCIP_TUNNEL_CIRCUIT.DATA_L2_COS,
     FCIP_TUNNEL_CIRCUIT.DSCP_DATA, FCIP_TUNNEL_CIRCUIT.MAX_RETRANSMISSIONS,
     FCIP_TUNNEL_CIRCUIT.SLOT_NUMBER, FCIP_TUNNEL_CIRCUIT.VE_PORT_NUMBER,
     FCIP_TUNNEL_CIRCUIT.SECURITY_FLAG, FCIP_TUNNEL_CIRCUIT.DSCP_CONTROL,
     FCIP_TUNNEL_CIRCUIT.CIRCUIT_STATUS, FCIP_TUNNEL_CIRCUIT.ENABLED,
     FCIP_TUNNEL_CIRCUIT.MISMATCHED_CONFIGURATIONS,
     FCIP_TUNNEL_CIRCUIT.CIRCUIT_STATUS_STRING,
     FCIP_TUNNEL_CIRCUIT.L2COS_F_CLASS, FCIP_TUNNEL_CIRCUIT.L2_COS_HIGH,
     FCIP_TUNNEL_CIRCUIT.L2_COS_MEDIUM, FCIP_TUNNEL_CIRCUIT.L2_COS_LOW,
     FCIP_TUNNEL_CIRCUIT.DSCP_F_CLASS, FCIP_TUNNEL_CIRCUIT.DSCP_HIGH,
     FCIP_TUNNEL_CIRCUIT.DSCP_MEDIUM, FCIP_TUNNEL_CIRCUIT.DSCP_LOW,
     FCIP_TUNNEL_CIRCUIT.FAILOVER_CIRCUIT, FCIP_TUNNEL_CIRCUIT.FAILOVER_GROUP_ID,
     DP0_GIGE_PORT.PORT_NUMBER as GIGE_PORT_NUMBER,
     DP0_GIGE_PORT.SLOT_NUMBER as GIGE_PORT_SLOT_NUMBER,
     FCIP_CIRCUIT_PORT_MAP.SWITCH_PORT_ID as GIGE_PORT_ID,
     FCIP_TUNNEL.VIRTUAL_SWITCH_ID, DP0_SWITCH_PORT.USER_PORT_NUMBER,
```

```
    DP1_GIGE_PORT.PORT_NUMBER as DP1_GIGE_PORT_NUMBER,
    DP1_GIGE_PORT.SLOT_NUMBER as DP1_GIGE_PORT_SLOT_NUMBER,
    FCIP_CIRCUIT_PORT_MAP.DP1_SWITCH_PORT_ID as DP1_GIGE_PORT_ID,
    DP1_SWITCH_PORT.USER_PORT_NUMBER as DP1_USER_PORT_NUMBER
   from FCIP_TUNNEL,
    FCIP_TUNNEL_CIRCUIT
   left join FCIP_CIRCUIT_PORT_MAP on FCIP_CIRCUIT_PORT_MAP.CIRCUIT_ID = FCIP_TUNNEL_CIRCUIT.ID
   left join GIGE_PORT as DP0_GIGE_PORT on FCIP_CIRCUIT_PORT_MAP.SWITCH_PORT_ID = DP0_GIGE_PORT.ID
   left join SWITCH_PORT as DP0_SWITCH_PORT on DP0_GIGE_PORT.SWITCH_PORT_ID = DP0_SWITCH_PORT.ID
   left join GIGE_PORT as DP1_GIGE_PORT ON FCIP_CIRCUIT_PORT_MAP.DP1_SWITCH_PORT_ID = DP1_GIGE_PORT.ID
   left join SWITCH_PORT as DP1_SWITCH_PORT ON DP1_GIGE_PORT.SWITCH_PORT_ID = DP1_SWITCH_PORT.ID
  where FCIP_TUNNEL_CIRCUIT.TUNNEL_ID = FCIP_TUNNEL.ID;
```

IFL_INFO

```
create or replace view IFL_INFO as
select
    IFL.ID as IFL_ID,
    IFL.EDGE_FABRIC_ID,
    (select distinct FCR_PORT.VIRTUAL_SWITCH_ID
        from SWITCH_PORT FCR_PORT
        where FCR_PORT.WWN = IFL.BB_PORT_WWN)
        as FCR_SWITCH_ID,
    IFL.EDGE_PORT_WWN,
    IFL.BB_FABRIC_ID,
    IFL.BB_PORT_WWN ,
    IFL.BB_RA_TOV,
    IFL.BB_ED_TOV,
    IFL.BB_PID_FORMAT,
    SWITCH_PORT.VIRTUAL_SWITCH_ID as EDGE_SWITCH_ID,
    SWITCH_PORT.ID as EDGE_PORT_ID,
    SWITCH_PORT.USER_PORT_NUMBER as EDGE_PORT_NUMBER,
    SWITCH_PORT.TYPE as EDGE_PORT_TYPE
from IFL
    left outer join SWITCH_PORT
        on IFL.EDGE_PORT_WWN = SWITCH_PORT.WWN;
```

IFL_REPORT_INFO

```
CREATE OR REPLACE VIEW ifl_report_info AS
SELECT
ifl.id AS ifl_id,
ifl.edge_fabric_id,
edge_fabric.id AS edge_fabric_db_id,
ifl.bb_fabric_id,
ifl.edge_port_wwn,
ifl.bb_port_wwn,
ifl.bb_ra_tov,
ifl.bb_ed_tov,
ifl.bb_pid_format,
bb_switch_port.virtual_switch_id AS bb_virtual_switch_id,
bb_switch_port.id AS bb_switch_port_id,
bb_switch_port.wwn AS bb_switch_port_wwn,
bb_switch_port.name AS bb_switch_port_name,
bb_switch_port.slot_number AS bb_switch_port_slot_number,
bb_switch_port.port_number AS bb_switch_port_port_number,
bb_switch_port.port_id AS bb_switch_port_port_id,
bb_switch_port.port_index AS bb_switch_port_port_index,
bb_switch_port.area_id AS bb_switch_port_area_id,
bb_switch_port.mac_address AS bb_switch_port_mac_address,
bb_switch_port.status AS bb_switch_port_status,
```

```
bb_switch_port.state AS bb_switch_port_state,
bb_switch_port.health AS bb_switch_port_health,
bb_switch_port.status_message AS bb_switch_port_status_message,
bb_switch_port.category AS bb_switch_port_category,
bb_switch_port.licensed AS bb_switch_port_licensed,
bb_switch_port.type AS bb_switch_port_type,
bb_switch_port.kind AS bb_switch_port_kind,
bb_switch_port.physical_port AS bb_switch_port_physical_port,
bb_switch_port.trunked AS bb_switch_port_trunked,
bb_switch_port.trunk_master AS bb_switch_port_trunk_master,
bb_switch_port.master_port_number AS bb_switch_port_master_port_number,
bb_switch_port.identifier AS bb_port_identifier,
bb_vs.id AS bb_vs_id,
bb_vs.wwn AS bb_virtual_switch_wwn,
bb_vs.name AS bb_switch_name,
bb_vs.virtual_fabric_id AS bb_switch_virtual_fabric_id,
bb_vs.operational_status AS bb_switch_operational_status,
bb_vs.state AS bb_switch_state,
bb_vs.status AS bb_switch_status,
bb_vs.status_reason AS bb_switch_status_reason,
bb_vs.core_switch_id AS bb_core_switch_id,
bb_vs.base_switch AS bb_base_switch,
bb_vs.management_state AS bb_virtual_switch_management_state,
bb_vs.monitored AS bb_virtual_switch_monitored,
bb_vs.domain_id As bb_virtual_switch_domain_id,
bb_cs.model AS bb_core_switch_smodel,
bb_cs.ip_address AS bb_core_switch_ip_address,
bb_cs.wwn AS  bb_core_switch_physical_switch_wwn,
bb_cs.operational_status AS bb_core_switch_operational_status,
bb_cs.name AS bb_core_switch_name,
bb_cs.type AS bb_core_switch_type,
bb_cs.model AS bb_core_switch_model,
bb_cs.vendor AS bb_core_switch_vendor,
bb_cs.reachable AS bb_core_switch_reachable,
bb_fabric.seed_switch_wwn as bb_fabric_seed_switch_wwn,
bb_fabric.name AS bb_fabric_name,
bb_fabric.fabric_name AS bb_fabric_fabric_name,
bb_fabric.principal_switch_wwn AS bb_fabric_principal_switch_wwn,
bb_fabric.management_state AS bb_fabric_management_state,
edge_switch_port.virtual_switch_id AS edge_virtual_switch_id,
edge_switch_port.id AS edge_switch_port_id,
edge_switch_port.wwn AS edge_switch_port_wwn,
edge_switch_port.name AS edge_switch_port_name,
edge_switch_port.slot_number AS edge_switch_port_slot_number,
edge_switch_port.port_number AS edge_switch_port_port_number,
edge_switch_port.port_id AS edge_switch_port_port_id,
edge_switch_port.port_index AS edge_switch_port_port_index,
edge_switch_port.area_id AS edge_switch_port_area_id,
edge_switch_port.mac_address AS edge_switch_port_mac_address,
edge_switch_port.status AS edge_switch_port_status,
edge_switch_port.state AS edge_switch_port_state,
edge_switch_port.health AS edge_switch_port_health,
edge_switch_port.status_message AS edge_switch_port_status_message,
edge_switch_port.category AS edge_switch_port_category,
edge_switch_port.licensed AS edge_switch_port_licensed,
edge_switch_port.type AS edge_switch_port_type,
edge_switch_port.kind AS edge_switch_port_kind,
edge_switch_port.physical_port AS edge_switch_port_physical_port,
edge_switch_port.trunked AS edge_switch_port_trunked,
edge_switch_port.trunk_master AS edge_switch_port_trunk_master,
edge_switch_port.master_port_number AS edge_switch_port_master_port_number,
edge_switch_port.identifier AS edge_port_identifier,
edge_vs.id AS edge_vs_id,
```

```
edge_vs.wwn AS edge_virtual_switch_wwn,
edge_vs.name AS edge_switch_name,
edge_vs.virtual_fabric_id AS edge_switch_virtual_fabric_id,
edge_vs.operational_status AS edge_switch_operational_status,
edge_vs.state AS edge_switch_state,
edge_vs.status AS edge_switch_status,
edge_vs.status_reason AS edge_switch_status_reason,
edge_vs.core_switch_id AS edge_core_switch_id,
edge_vs.base_switch AS edge_base_switch,
edge_vs.management_state AS  edge_virtual_switch_management_state,
edge_vs.monitored AS  edge_virtual_switch_monitored,
edge_vs.domain_id As edge_virtual_switch_domain_id,
edge_cs.model AS edge_core_switch_smodel,
edge_cs.ip_address AS edge_core_switch_ip_address,
edge_cs.wwn AS  edge_core_switch_physical_switch_wwn,
edge_cs.operational_status AS edge_core_switch_operational_status,
edge_cs.name AS edge_core_switch_name,
edge_cs.type AS edge_core_switch_type,
edge_cs.model AS edge_core_switch_model,
edge_cs.vendor AS edge_core_switch_vendor,
edge_cs.reachable AS edge_core_switch_reachable,
edge_fabric.seed_switch_wwn as edge_fabric_seed_switch_wwn,
edge_fabric.name AS edge_fabric_name,
edge_fabric.fabric_name AS edge_fabric_fabric_name,
edge_fabric.principal_switch_wwn AS edge_fabric_principal_switch_wwn,
edge_fabric.management_state AS edge_fabric_management_state,
CASE
WHEN bb_switch_port.trunk_master = 1 AND (edge_switch_port.trunk_master IS NULL OR

edge_switch_port.trunk_master = 1) THEN 'IFL Trunk'
ELSE 'IFL'
END AS connection_type

FROM ifl
JOIN switch_port bb_switch_port ON ifl.bb_port_wwn = bb_switch_port.wwn
JOIN virtual_switch bb_vs ON bb_switch_port.virtual_switch_id = bb_vs.id
JOIN core_switch bb_cs ON bb_vs.core_switch_id = bb_cs.id
JOIN fabric_member ON fabric_member.virtual_switch_id = bb_vs.id
JOIN fabric bb_fabric ON bb_fabric.id = fabric_member.fabric_id

LEFT JOIN switch_port edge_switch_port ON ifl.edge_port_wwn = edge_switch_port.wwn
LEFT JOIN virtual_switch edge_vs ON edge_switch_port.virtual_switch_id = edge_vs.id
LEFT JOIN core_switch edge_cs ON edge_vs.core_switch_id = edge_cs.id
LEFT JOIN fabric_member fs ON fs.virtual_switch_id = edge_vs.id
LEFT JOIN fabric edge_fabric ON edge_fabric.id = fs.fabric_id

WHERE bb_fabric.managed = 1 AND bb_vs.monitored = 1 AND (edge_fabric.managed IS NULL OR

edge_fabric.managed = 1) AND (edge_vs.monitored IS NULL OR edge_vs.monitored = 1);
```

ISL_INFO

```
create or replace view ISL_INFO as
select distinct
                ISL.ID,
                ISL.FABRIC_ID,
                ISL.COST,
                ISL.TYPE,
                ISL.SOURCE_DOMAIN_ID,
                ISL.SOURCE_PORT_NUMBER,
                ISL.MISSING,
                ISL.MISSING_TIME,
```

```
ISL.TRUSTED,
ISL.CREATION_TIME,
ISL.TRUNKED,
ISL.MISSING_REASON,
                SOURCE_VIRTUAL_SWITCH.ID as SOURCE_SWITCH_ID,
                SOURCE_VIRTUAL_SWITCH.NAME as SOURCE_SWITCH_NAME,
                SOURCE_VIRTUAL_SWITCH.WWN as SOURCE_SWITCH_WWN,
SOURCE_VIRTUAL_SWITCH.CORE_SWITCH_ID as SOURCE_CORE_SWITCH_ID,
SOURCE_VIRTUAL_SWITCH.BASE_SWITCH as SOURCE_BASE_SWITCH,
SOURCE_VIRTUAL_SWITCH.MANAGEMENT_STATE as SOURCE_VIRTUAL_SWITCH_MANAGEMENT_STATE,
SOURCE_VIRTUAL_SWITCH.MONITORED as SOURCE_VIRTUAL_SWITCH_MONITORED,
SOURCE_SWITCH_PORT.ID as SOURCE_SWITCH_PORT_ID,
                SOURCE_SWITCH_PORT.WWN as SOURCE_SWITCH_PORT_WWN,
                SOURCE_SWITCH_PORT.NAME as SOURCE_SWITCH_PORT_NAME,
SOURCE_SWITCH_PORT.TYPE as PORT_TYPE,
SOURCE_SWITCH_PORT.KIND as SOURCE_SWITCH_PORT_KIND,
SOURCE_SWITCH_PORT.PHYSICAL_PORT as SOURCE_PHYSICAL_PORT,
SOURCE_SWITCH_PORT.TRUNKED as SOURCE_SWITCH_PORT_TRUNKED,
SOURCE_SWITCH_PORT.SPEED as SOURCE_PORT_SPEED,
ISL.DEST_DOMAIN_ID,
                ISL.DEST_PORT_NUMBER,
                DEST_VIRTUAL_SWITCH.ID as DEST_SWITCH_ID,
                DEST_VIRTUAL_SWITCH.NAME as DEST_SWITCH_NAME,
DEST_VIRTUAL_SWITCH.WWN as DEST_SWITCH_WWN,
DEST_VIRTUAL_SWITCH.CORE_SWITCH_ID as DEST_CORE_SWITCH_ID,
DEST_VIRTUAL_SWITCH.BASE_SWITCH as DEST_BASE_SWITCH,
DEST_VIRTUAL_SWITCH.MANAGEMENT_STATE as DEST_VIRTUAL_SWITCH_MANAGEMENT_STATE,
DEST_VIRTUAL_SWITCH.MONITORED as DEST_VIRTUAL_SWITCH_MONITORED,
                DEST_SWITCH_PORT.ID as DEST_SWITCH_PORT_ID,
                DEST_SWITCH_PORT.WWN as DEST_SWITCH_PORT_WWN,
DEST_SWITCH_PORT.NAME as DEST_SWITCH_PORT_NAME,
DEST_SWITCH_PORT.KIND as DEST_SWITCH_PORT_KIND,
DEST_SWITCH_PORT.PHYSICAL_PORT as DEST_PHYSICAL_PORT,
DEST_SWITCH_PORT.TRUNKED as DEST_SWITCH_PORT_TRUNKED,
DEST_SWITCH_PORT.SPEED as DEST_PORT_SPEED,
FABRIC.PRINCIPAL_SWITCH_WWN as PRINCIPAL_SWITCH_WWN
from
                ISL,
FABRIC_MEMBER                                           SOURCE_FABRIC_MEMBER,
VIRTUAL_SWITCH            SOURCE_VIRTUAL_SWITCH,
SWITCH_PORT                                                 SOURCE_SWITCH_PORT,
FABRIC_MEMBER                                           DEST_FABRIC_MEMBER,
VIRTUAL_SWITCH            DEST_VIRTUAL_SWITCH,
SWITCH_PORT                                                 DEST_SWITCH_PORT,
FABRIC
where
                SOURCE_FABRIC_MEMBER.FABRIC_ID = ISL.FABRIC_ID and
                SOURCE_VIRTUAL_SWITCH.ID = SOURCE_FABRIC_MEMBER.VIRTUAL_SWITCH_ID and
                SOURCE_VIRTUAL_SWITCH.DOMAIN_ID = ISL.SOURCE_DOMAIN_ID and
                SOURCE_SWITCH_PORT.VIRTUAL_SWITCH_ID = SOURCE_VIRTUAL_SWITCH.ID and
SOURCE_SWITCH_PORT.CATEGORY = 1 and
                SOURCE_SWITCH_PORT.USER_PORT_NUMBER = ISL.SOURCE_PORT_NUMBER and
                DEST_FABRIC_MEMBER.FABRIC_ID = ISL.FABRIC_ID and
                DEST_VIRTUAL_SWITCH.ID = DEST_FABRIC_MEMBER.VIRTUAL_SWITCH_ID and
                DEST_VIRTUAL_SWITCH.DOMAIN_ID = ISL.DEST_DOMAIN_ID and
                DEST_SWITCH_PORT.VIRTUAL_SWITCH_ID = DEST_VIRTUAL_SWITCH.ID and
DEST_SWITCH_PORT.CATEGORY = 1 and
                DEST_SWITCH_PORT.USER_PORT_NUMBER = ISL.DEST_PORT_NUMBER and
FABRIC.ID = ISL.FABRIC_ID;
```

ISL_REPORT_INFO

```
CREATE OR REPLACE VIEW ISL_REPORT_INFO AS
SELECT ISL.ID, ISL.FABRIC_ID, ISL.COST, ISL.TYPE, ISL.TRUSTED,
    ISL.CREATION_TIME, ISL.MISSING, ISL.MISSING_TIME, ISL.TRUNKED, ISL.MISSING_REASON,
    ISL.SOURCE_DOMAIN_ID, ISL.SOURCE_PORT_NUMBER,
    SOURCE_VIRTUAL_SWITCH.ID AS SOURCE_SWITCH_ID,
    SOURCE_VIRTUAL_SWITCH.NAME AS SOURCE_SWITCH_NAME,
    SOURCE_VIRTUAL_SWITCH.WWN AS SOURCE_SWITCH_WWN,
    SOURCE_VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID AS SOURCE_SWITCH_VIRTUAL_FABRIC_ID,
    SOURCE_VIRTUAL_SWITCH.OPERATIONAL_STATUS AS SOURCE_SWITCH_OPERATIONAL_STATUS,
    SOURCE_VIRTUAL_SWITCH.STATE AS SOURCE_SWITCH_STATE,
    SOURCE_VIRTUAL_SWITCH.STATUS AS SOURCE_SWITCH_STATUS,
    SOURCE_VIRTUAL_SWITCH.STATUS_REASON AS SOURCE_SWITCH_STATUS_REASON,
    SOURCE_VIRTUAL_SWITCH.CORE_SWITCH_ID AS SOURCE_CORE_SWITCH_ID,
    SOURCE_VIRTUAL_SWITCH.BASE_SWITCH AS SOURCE_BASE_SWITCH,
    SOURCE_VIRTUAL_SWITCH.MANAGEMENT_STATE AS SOURCE_VIRTUAL_SWITCH_MANAGEMENT_STATE,
    SOURCE_VIRTUAL_SWITCH.MONITORED AS SOURCE_VIRTUAL_SWITCH_MONITORED,
    SOURCE_CORE_SWITCH.IP_ADDRESS AS SOURCE_CORE_SWITCH_IP_ADDRESS,
    SOURCE_CORE_SWITCH.WWN AS SOURCE_CORE_SWITCH_WWN,
    SOURCE_CORE_SWITCH.NAME AS SOURCE_CORE_SWITCH_NAME,
    SOURCE_CORE_SWITCH.TYPE AS SOURCE_CORE_SWITCH_TYPE,
    SOURCE_CORE_SWITCH.OPERATIONAL_STATUS AS SOURCE_CORE_SWITCH_OPERATIONAL_STATUS,
    SOURCE_CORE_SWITCH.MODEL AS SOURCE_CORE_SWITCH_MODEL,
    SOURCE_CORE_SWITCH.VENDOR AS SOURCE_CORE_SWITCH_VENDOR,
    SOURCE_CORE_SWITCH.REACHABLE AS SOURCE_CORE_SWITCH_REACHABLE,
    SOURCE_SWITCH_PORT.ID AS SOURCE_SWITCH_PORT_ID,
    SOURCE_SWITCH_PORT.WWN AS SOURCE_SWITCH_PORT_WWN,
    SOURCE_SWITCH_PORT.NAME AS SOURCE_SWITCH_PORT_NAME,
    SOURCE_SWITCH_PORT.SLOT_NUMBER AS SOURCE_SWITCH_PORT_SLOT_NUMBER,
    SOURCE_SWITCH_PORT.PORT_NUMBER AS SOURCE_SWITCH_PORT_PORT_NUMBER,
    SOURCE_SWITCH_PORT.PORT_ID AS SOURCE_SWITCH_PORT_PORT_ID,
    SOURCE_SWITCH_PORT.PORT_INDEX AS SOURCE_SWITCH_PORT_PORT_INDEX,
    SOURCE_SWITCH_PORT.AREA_ID AS SOURCE_SWITCH_PORT_AREA_ID,
    SOURCE_SWITCH_PORT.MAC_ADDRESS AS SOURCE_SWITCH_PORT_MAC_ADDRESS,
    SOURCE_SWITCH_PORT.STATUS AS SOURCE_SWITCH_PORT_STATUS,
    SOURCE_SWITCH_PORT.STATE AS SOURCE_SWITCH_PORT_STATE,
    SOURCE_SWITCH_PORT.HEALTH AS SOURCE_SWITCH_PORT_HEALTH,
    SOURCE_SWITCH_PORT.STATUS_MESSAGE AS SOURCE_SWITCH_PORT_STATUS_MESSAGE,
    SOURCE_SWITCH_PORT.CATEGORY AS SOURCE_SWITCH_PORT_CATEGORY,
    SOURCE_SWITCH_PORT.LICENSED AS SOURCE_SWITCH_PORT_LICENSED,
    SOURCE_SWITCH_PORT.TYPE AS SOURCE_SWITCH_PORT_TYPE,
    SOURCE_SWITCH_PORT.KIND AS SOURCE_SWITCH_PORT_KIND,
    SOURCE_SWITCH_PORT.PHYSICAL_PORT AS SOURCE_PHYSICAL_PORT,
    SOURCE_SWITCH_PORT.TRUNKED AS SOURCE_SWITCH_PORT_TRUNKED,
    SOURCE_SWITCH_PORT.TRUNK_MASTER AS SOURCE_SWITCH_PORT_TRUNK_MASTER,
    SOURCE_SWITCH_PORT.MASTER_PORT_NUMBER AS SOURCE_SWITCH_PORT_MASTER_PORT_NUMBER,
    SOURCE_SWITCH_PORT.SPEED as SOURCE_SWITCH_PORT_SPEED,
    ISL.DEST_DOMAIN_ID, ISL.DEST_PORT_NUMBER,
    DEST_VIRTUAL_SWITCH.ID AS DEST_SWITCH_ID,
    DEST_VIRTUAL_SWITCH.NAME AS DEST_SWITCH_NAME,
    DEST_VIRTUAL_SWITCH.WWN AS DEST_SWITCH_WWN,
    DEST_VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID AS DEST_SWITCH_VIRTUAL_FABRIC_ID,
    DEST_VIRTUAL_SWITCH.OPERATIONAL_STATUS AS DEST_SWITCH_OPERATIONAL_STATUS,
    DEST_VIRTUAL_SWITCH.STATE AS DEST_SWITCH_STATE,
    DEST_VIRTUAL_SWITCH.STATUS AS DEST_SWITCH_STATUS,
    DEST_VIRTUAL_SWITCH.STATUS_REASON AS DEST_SWITCH_STATUS_REASON,
    DEST_VIRTUAL_SWITCH.CORE_SWITCH_ID AS DEST_CORE_SWITCH_ID,
    DEST_VIRTUAL_SWITCH.BASE_SWITCH AS DEST_BASE_SWITCH,
    DEST_VIRTUAL_SWITCH.MANAGEMENT_STATE AS DEST_VIRTUAL_SWITCH_MANAGEMENT_STATE,
    DEST_VIRTUAL_SWITCH.MONITORED AS DEST_VIRTUAL_SWITCH_MONITORED,
    DEST_CORE_SWITCH.IP_ADDRESS AS DEST_CORE_SWITCH_IP_ADDRESS,
    DEST_CORE_SWITCH.WWN AS DEST_CORE_SWITCH_WWN,
```

```
    DEST_CORE_SWITCH.NAME AS DEST_CORE_SWITCH_NAME,
    DEST_CORE_SWITCH.TYPE AS DEST_CORE_SWITCH_TYPE,
    DEST_CORE_SWITCH.OPERATIONAL_STATUS AS DEST_CORE_SWITCH_OPERATIONAL_STATUS,
    DEST_CORE_SWITCH.MODEL AS DEST_CORE_SWITCH_MODEL,
    DEST_CORE_SWITCH.VENDOR AS DEST_CORE_SWITCH_VENDOR,
    DEST_CORE_SWITCH.REACHABLE AS DEST_CORE_SWITCH_REACHABLE,
    DEST_SWITCH_PORT.ID AS DEST_SWITCH_PORT_ID,
    DEST_SWITCH_PORT.WWN AS DEST_SWITCH_PORT_WWN,
    DEST_SWITCH_PORT.NAME AS DEST_SWITCH_PORT_NAME,
    DEST_SWITCH_PORT.SLOT_NUMBER AS DEST_SWITCH_PORT_SLOT_NUMBER,
    DEST_SWITCH_PORT.PORT_NUMBER AS DEST_SWITCH_PORT_PORT_NUMBER,
    DEST_SWITCH_PORT.PORT_ID AS DEST_SWITCH_PORT_PORT_ID,
    DEST_SWITCH_PORT.PORT_INDEX AS DEST_SWITCH_PORT_PORT_INDEX,
    DEST_SWITCH_PORT.AREA_ID AS DEST_SWITCH_PORT_AREA_ID,
    DEST_SWITCH_PORT.MAC_ADDRESS AS DEST_SWITCH_PORT_MAC_ADDRESS,
    DEST_SWITCH_PORT.STATUS AS DEST_SWITCH_PORT_STATUS,
    DEST_SWITCH_PORT.STATE AS DEST_SWITCH_PORT_STATE,
    DEST_SWITCH_PORT.HEALTH AS DEST_SWITCH_PORT_HEALTH,
    DEST_SWITCH_PORT.STATUS_MESSAGE AS DEST_SWITCH_PORT_STATUS_MESSAGE,
    DEST_SWITCH_PORT.CATEGORY AS DEST_SWITCH_PORT_CATEGORY,
    DEST_SWITCH_PORT.LICENSED AS DEST_SWITCH_PORT_LICENSED,
    DEST_SWITCH_PORT.TYPE AS DEST_SWITCH_PORT_TYPE,
    DEST_SWITCH_PORT.KIND AS DEST_SWITCH_PORT_KIND,
    DEST_SWITCH_PORT.PHYSICAL_PORT AS DEST_PHYSICAL_PORT,
    DEST_SWITCH_PORT.TRUNKED AS DEST_SWITCH_PORT_TRUNKED,
    DEST_SWITCH_PORT.TRUNK_MASTER AS DEST_SWITCH_PORT_TRUNK_MASTER,
    DEST_SWITCH_PORT.MASTER_PORT_NUMBER AS DEST_SWITCH_PORT_MASTER_PORT_NUMBER,
    DEST_SWITCH_PORT.SPEED as DEST_SWITCH_PORT_SPEED,
CASE WHEN SOURCE_SWITCH_PORT.TRUNK_MASTER = 1 AND DEST_SWITCH_PORT.TRUNK_MASTER = 1 THEN 'TRUNK' ELSE 'ISL'
END AS CONNECTION_TYPE,
    FABRIC.SEED_SWITCH_WWN AS FABRIC_SEED_SWITCH_WWN, FABRIC.NAME AS FABRIC_NAME, FABRIC.FABRIC_NAME AS
FABRIC_FABRIC_NAME,
    FABRIC.PRINCIPAL_SWITCH_WWN AS FABRIC_PRINCIPAL_SWITCH_WWN, FABRIC.MANAGEMENT_STATE AS
FABRIC_MANAGEMENT_STATE, FABRIC.TYPE AS FABRIC_TYPE, FABRIC.MANAGED AS FABRIC_MANAGED
FROM ISL
LEFT JOIN FABRIC_MEMBER SOURCE_FABRIC_MEMBER ON SOURCE_FABRIC_MEMBER.FABRIC_ID = ISL.FABRIC_ID JOIN
VIRTUAL_SWITCH SOURCE_VIRTUAL_SWITCH ON SOURCE_VIRTUAL_SWITCH.ID = SOURCE_FABRIC_MEMBER.VIRTUAL_SWITCH_ID AND
SOURCE_VIRTUAL_SWITCH.DOMAIN_ID = ISL.SOURCE_DOMAIN_ID
JOIN CORE_SWITCH SOURCE_CORE_SWITCH ON SOURCE_CORE_SWITCH.ID = SOURCE_VIRTUAL_SWITCH.CORE_SWITCH_ID
LEFT JOIN SWITCH_PORT SOURCE_SWITCH_PORT ON SOURCE_SWITCH_PORT.VIRTUAL_SWITCH_ID = SOURCE_VIRTUAL_SWITCH.ID
AND SOURCE_SWITCH_PORT.USER_PORT_NUMBER = ISL.SOURCE_PORT_NUMBER
LEFT JOIN FABRIC_MEMBER DEST_FABRIC_MEMBER ON DEST_FABRIC_MEMBER.FABRIC_ID = ISL.FABRIC_ID JOIN
VIRTUAL_SWITCH DEST_VIRTUAL_SWITCH ON DEST_VIRTUAL_SWITCH.ID = DEST_FABRIC_MEMBER.VIRTUAL_SWITCH_ID AND
DEST_VIRTUAL_SWITCH.DOMAIN_ID = ISL.DEST_DOMAIN_ID
JOIN CORE_SWITCH DEST_CORE_SWITCH ON DEST_CORE_SWITCH.ID = DEST_VIRTUAL_SWITCH.CORE_SWITCH_ID
LEFT JOIN SWITCH_PORT DEST_SWITCH_PORT ON DEST_SWITCH_PORT.VIRTUAL_SWITCH_ID = DEST_VIRTUAL_SWITCH.ID AND
DEST_SWITCH_PORT.USER_PORT_NUMBER = ISL.DEST_PORT_NUMBER
JOIN FABRIC ON FABRIC.ID = ISL.FABRIC_ID
WHERE SOURCE_VIRTUAL_SWITCH.MONITORED = 1 AND SOURCE_CORE_SWITCH.TYPE NOT IN (40,41) AND
(SOURCE_SWITCH_PORT.CATEGORY IS NULL OR SOURCE_SWITCH_PORT.CATEGORY = 1)
AND DEST_VIRTUAL_SWITCH.MONITORED = 1 AND DEST_CORE_SWITCH.TYPE NOT IN (40,41) AND (DEST_SWITCH_PORT.CATEGORY
IS NULL OR DEST_SWITCH_PORT.CATEGORY = 1)
AND FABRIC.MANAGED = 1 AND FABRIC.TYPE NOT IN (65,66,4);
```

ISL_TRILL_INFO

```
create or replace view ISL_TRILL_INFO as
select distinct
    VCS_DEVICE.DEVICE_ID as VCS_DEVICE_ID,
    SOURCE_CLUSTER_MEMBER.CLUSTER_ME_ID,
    ISL.ID,
    ISL.FABRIC_ID,
```

```
    ISL.COST,
    ISL.MISSING,
    ISL.SOURCE_DOMAIN_ID,
    ISL.SOURCE_PORT_NUMBER,
    SOURCE_DEVICE.MANAGED_ELEMENT_ID as SOURCE_ME_ID,
    SOURCE_DEVICE.DEVICE_ID as SOURCE_DEVICE_ID,
    SOURCE_DEVICE.SYS_NAME as SOURCE_DEVICE_NAME,
    SOURCE_SWITCH_PORT.ID as SOURCE_SWITCH_PORT_ID,
    SOURCE_SWITCH_PORT.NAME as SOURCE_SWITCH_PORT_NAME,
    SOURCE_SWITCH_PORT.IDENTIFIER as SOURCE_SWITCH_PORT_IDENTIFIER,
    SOURCE_SWITCH_PORT.TYPE as PORT_TYPE,
    SOURCE_SWITCH_PORT.KIND as SOURCE_SWITCH_PORT_KIND,
    SOURCE_SWITCH_PORT.PHYSICAL_PORT as SOURCE_PHYSICAL_PORT,
    SOURCE_SWITCH_PORT.TRUNKED as SOURCE_SWITCH_PORT_TRUNKED,
    ISL.DEST_DOMAIN_ID,
    ISL.DEST_PORT_NUMBER,
    DEST_DEVICE.DEVICE_ID as DEST_DEVICE_ID,
    DEST_DEVICE.MANAGED_ELEMENT_ID AS DEST_ME_ID,
    DEST_DEVICE.SYS_NAME as DEST_DEVICE_NAME,
    DEST_SWITCH_PORT.ID as DEST_SWITCH_PORT_ID,
    DEST_SWITCH_PORT.NAME as DEST_SWITCH_PORT_NAME,
    DEST_SWITCH_PORT.IDENTIFIER as DEST_SWITCH_PORT_IDENTIFIER,
    DEST_SWITCH_PORT.KIND as DEST_SWITCH_PORT_KIND,
    DEST_SWITCH_PORT.PHYSICAL_PORT as DEST_PHYSICAL_PORT,
    DEST_SWITCH_PORT.TRUNKED as DEST_SWITCH_PORT_TRUNKED

from
    ISL,
    DEVICE VCS_DEVICE,
    VCS_CLUSTER_MEMBER SOURCE_CLUSTER_MEMBER,
    VCS_CLUSTER_MEMBER DEST_CLUSTER_MEMBER,
    DEVICE SOURCE_DEVICE,
    SWITCH_PORT SOURCE_SWITCH_PORT,
    FABRIC_MEMBER SOURCE_FABRIC_MEMBER,
    VIRTUAL_SWITCH SOURCE_VIRTUAL_SWITCH,
    DEVICE DEST_DEVICE,
    SWITCH_PORT DEST_SWITCH_PORT,
    FABRIC_MEMBER DEST_FABRIC_MEMBER,
    VIRTUAL_SWITCH DEST_VIRTUAL_SWITCH,
    FABRIC

where
    SOURCE_FABRIC_MEMBER.FABRIC_ID = ISL.FABRIC_ID and
    SOURCE_VIRTUAL_SWITCH.ID = SOURCE_FABRIC_MEMBER.VIRTUAL_SWITCH_ID and
    SOURCE_VIRTUAL_SWITCH.DOMAIN_ID = ISL.SOURCE_DOMAIN_ID and
    SOURCE_SWITCH_PORT.VIRTUAL_SWITCH_ID = SOURCE_VIRTUAL_SWITCH.ID and
    SOURCE_SWITCH_PORT.USER_PORT_NUMBER = ISL.SOURCE_PORT_NUMBER and
    DEST_FABRIC_MEMBER.FABRIC_ID = ISL.FABRIC_ID and
    DEST_VIRTUAL_SWITCH.ID = DEST_FABRIC_MEMBER.VIRTUAL_SWITCH_ID and
    DEST_VIRTUAL_SWITCH.DOMAIN_ID = ISL.DEST_DOMAIN_ID and
    DEST_SWITCH_PORT.VIRTUAL_SWITCH_ID = DEST_VIRTUAL_SWITCH.ID and
    DEST_SWITCH_PORT.USER_PORT_NUMBER = ISL.DEST_PORT_NUMBER and
    FABRIC.ID = ISL.FABRIC_ID and
    SOURCE_VIRTUAL_SWITCH.MANAGED_ELEMENT_ID = SOURCE_DEVICE.MANAGED_ELEMENT_ID and
    DEST_VIRTUAL_SWITCH.MANAGED_ELEMENT_ID = DEST_DEVICE.MANAGED_ELEMENT_ID and
    SOURCE_CLUSTER_MEMBER.MEMBER_ME_ID = SOURCE_DEVICE.MANAGED_ELEMENT_ID and
    DEST_CLUSTER_MEMBER.MEMBER_ME_ID = DEST_DEVICE.MANAGED_ELEMENT_ID and
    VCS_DEVICE.MANAGED_ELEMENT_ID = SOURCE_CLUSTER_MEMBER.CLUSTER_ME_ID;
```

ISL_TRUNK_GROUP_MEMBER_INFO

```
CREATE VIEW isl_trunk_group_member_info AS
```

```
    select
     ISL_TRUNK_GROUP.ID,
     ISL_TRUNK_GROUP.VIRTUAL_SWITCH_ID,
     ISL_TRUNK_GROUP.MASTER_USER_PORT,
     ISL_TRUNK_MEMBER.MISSING,
     ISL_TRUNK_MEMBER.TRUSTED,
     ISL_TRUNK_MEMBER.MISSING_TIME,
     ISL_TRUNK_MEMBER.PORT_NUMBER,
     SWITCH_PORT.WWN,
     SWITCH_PORT.TYPE,
     SWITCH_PORT.STATUS,
     SWITCH_PORT.SPEED,
     SWITCH_PORT.ID as SWITCH_PORT_ID,
     SWITCH_PORT.SPEED_TYPE
from
     ISL_TRUNK_GROUP,
     ISL_TRUNK_MEMBER,
     SWITCH_PORT
where
     ISL_TRUNK_GROUP.id = ISL_TRUNK_MEMBER.GROUP_ID
     and ISL_TRUNK_GROUP.VIRTUAL_SWITCH_ID = SWITCH_PORT.VIRTUAL_SWITCH_ID
     and ISL_TRUNK_MEMBER.PORT_NUMBER= SWITCH_PORT.USER_PORT_NUMBER;
```

ISL_TRUNK_INFO

```
create or replace view ISL_TRUNK_INFO as
select
   ISL_TRUNK_GROUP.ID,
   ISL_TRUNK_GROUP.TRUSTED,
   ISL_TRUNK_GROUP.MISSING,
   ISL_TRUNK_GROUP.MISSING_TIME,
   ISL_TRUNK_GROUP.MEMBER_TRACKING_STATUS,
   ISL_INFO.COST,
   ISL_INFO.TYPE,
   ISL_INFO.SOURCE_PORT_NUMBER,
   ISL_INFO.SOURCE_SWITCH_ID,
   ISL_INFO.MISSING_REASON,
   SOURCE_CORE_SWITCH.IP_ADDRESS as SOURCE_SWITCH_IP_ADDRESS,
   SOURCE_VIRTUAL_SWITCH.WWN as SOURCE_SWITCH_WWN,
   SOURCE_VIRTUAL_SWITCH.MANAGEMENT_STATE as SOURCE_SWITCH_MANAGEMENT_STATE,
   SOURCE_VIRTUAL_SWITCH.MONITORED as SOURCE_SWITCH_MONITORED,
   ISL_INFO.SOURCE_DOMAIN_ID as MASTER_PORT,
   ISL_INFO.SOURCE_SWITCH_NAME,
   ISL_INFO.SOURCE_SWITCH_PORT_ID,
   ISL_INFO.DEST_PORT_NUMBER,
   ISL_INFO.DEST_SWITCH_ID,
   DEST_CORE_SWITCH.IP_ADDRESS as DEST_SWITCH_IP_ADDRESS,
   DEST_VIRTUAL_SWITCH.WWN as DEST_SWITCH_WWN,
   DEST_VIRTUAL_SWITCH.MANAGEMENT_STATE as DEST_SWITCH_MANAGEMENT_STATE,
   DEST_VIRTUAL_SWITCH.MONITORED as DEST_SWITCH_MONITORED,
   ISL_INFO.SOURCE_SWITCH_PORT_WWN,
   ISL_INFO.DEST_DOMAIN_ID as REMOTE_MASTER_PORT,
   ISL_INFO.DEST_SWITCH_NAME,
   ISL_INFO.DEST_SWITCH_PORT_ID

from
   ISL_TRUNK_GROUP,
   ISL_INFO,
   CORE_SWITCH SOURCE_CORE_SWITCH,
   CORE_SWITCH DEST_CORE_SWITCH,
   VIRTUAL_SWITCH SOURCE_VIRTUAL_SWITCH,
   VIRTUAL_SWITCH DEST_VIRTUAL_SWITCH
```

```
where
   ISL_INFO.SOURCE_SWITCH_ID = ISL_TRUNK_GROUP.VIRTUAL_SWITCH_ID
   and ISL_INFO.SOURCE_PORT_NUMBER = ISL_TRUNK_GROUP.MASTER_USER_PORT
   and ISL_INFO.SOURCE_SWITCH_ID = SOURCE_VIRTUAL_SWITCH.ID
   and SOURCE_VIRTUAL_SWITCH.CORE_SWITCH_ID = SOURCE_CORE_SWITCH.ID
   and ISL_INFO.DEST_SWITCH_ID = DEST_VIRTUAL_SWITCH.ID
   and DEST_VIRTUAL_SWITCH.CORE_SWITCH_ID = DEST_CORE_SWITCH.ID;
```

L2_NEIGHBOR_INFO

```
create or replace view L2_NEIGHBOR_INFO as
select
   L2_NEIGHBOR.INTERFACE_ID,
   L2_NEIGHBOR.RMT_IP_ADDRESS,
   L2_NEIGHBOR.RMT_IF_NAME,
   LLDP_DATA.DEVICE_ID as RMT_DEVICE_ID,
   LLDP_DATA.INTERFACE_ID as RMT_INTERFACE_ID,
   PHY_INTF.PHYSICAL_ADDRESS as RMT_INTERFACE_MAC,
   RMT_DEVICE.IS_ROUTER
from
   device RMT_DEVICE,
   LLDP_DATA,
   L2_NEIGHBOR,
   physical_interface PHY_INTF
where
   LLDP_DATA.CHASSIS_ID = L2_NEIGHBOR.LLDP_REM_CHASSIS_ID
   and LLDP_DATA.CHASSIS_ID_SUBTYPE = L2_NEIGHBOR.LLDP_REM_CHASSIS_ID_SUBTYPE
   and LLDP_DATA.PORT_ID = L2_NEIGHBOR.LLDP_REM_PORT_ID
   and LLDP_DATA.PORT_ID_SUBTYPE = L2_NEIGHBOR.LLDP_REM_PORT_ID_SUBTYPE
   and LLDP_DATA.DEVICE_ID = RMT_DEVICE.device_id
   and PHY_INTF.interface_id = LLDP_DATA.INTERFACE_ID;
```

MAPS_EVENT_DETAILS_INFO

```
create or replace view MAPS_EVENT_DETAILS_INFO as
select
    MAPS_EVENT.ID,
    MAPS_EVENT.HOST_TIME,
    MAPS_EVENT.CATEGORY,
    MAPS_EVENT.VIOLATION_TYPE,
    MAPS_EVENT.MANAGED_ELEMENT_ID,
    MAPS_EVENT.ORIGIN_FABRIC_ID,
    MAPS_EVENT.SWITCH_PORT_ID,
    MAPS_EVENT.INTERFACE_ID,
    MAPS_EVENT.FCIP_CIRCUIT_ID,
    MAPS_EVENT.FRU_NAME,
    MAPS_EVENT.VM_ID,
    MAPS_EVENT.FLOW_DEFINITION_ID,
    MAPS_EVENT. SUB_FLOW_KEY,
    MAPS_EVENT.FCIP_TUNNEL_ID,
    MAPS_EVENT.PORT_TYPE,
MAPS_EVENT.COLLECTION_NAME,
    MAPS_EVENT_DETAILS.SWITCH_TIME,
    MAPS_EVENT_DETAILS.RULE_NAME,
    MAPS_EVENT_DETAILS.RULE_CONDITION,
    MAPS_EVENT_DETAILS.TIME_BASE,
    MAPS_EVENT_DETAILS.ACTIONS,
    MAPS_EVENT_DETAILS.CURRENT_VALUE,
    MAPS_EVENT_DETAILS.SWITCH_ENABLED_ACTIONS,
MAPS_EVENT_DETAILS.SEVERITY,
```

```
    VIRTUAL_SWITCH.NAME as SWITCH_NAME,
    SWITCH_PORT.NAME as SWITCH_PORT_NAME,
    INTERFACE.NAME AS INTERFACE_NAME,
    SWITCH_PORT.WWN as SWITCH_PORT_WWN,
    SWITCH_PORT.SLOT_NUMBER as SWITCH_PORT_SLOT,
    SWITCH_PORT.PORT_NUMBER as SWITCH_PORT_NUMBER,
    SWITCH_PORT.PORT_ID as SWITCH_PORT_PORT_ID,
    FCIP_TUNNEL_CIRCUIT.CIRCUIT_NUMBER,
    FCIP_TUNNEL_CIRCUIT.SLOT_NUMBER as FCIP_SLOT_NUMBER,
    FCIP_TUNNEL_CIRCUIT.VE_PORT_NUMBER as FCIP_PORT_NUMBER,
    NP_FLOW_DEFINITION.NAME as FLOW_NAME,
    MAPS_EVENT_CAUSE_ACTION.ACTION
from
    MAPS_EVENT_DETAILS
    inner join
        MAPS_EVENT on
        MAPS_EVENT.ID = MAPS_EVENT_DETAILS.MAPS_EVENT_ID
    left outer join MAPS_EVENT_CAUSE_ACTION
       on MAPS_EVENT.VIOLATION_TYPE = MAPS_EVENT_CAUSE_ACTION.VIOLATION_TYPE
    left outer join VIRTUAL_SWITCH
       on MAPS_EVENT.MANAGED_ELEMENT_ID = VIRTUAL_SWITCH.MANAGED_ELEMENT_ID
    left outer join SWITCH_PORT
       on MAPS_EVENT.SWITCH_PORT_ID = SWITCH_PORT.ID
    LEFT JOIN INTERFACE ON MAPS_EVENT.INTERFACE_ID = INTERFACE.INTERFACE_ID
    left outer join FCIP_TUNNEL_CIRCUIT
       on MAPS_EVENT.FCIP_CIRCUIT_ID = FCIP_TUNNEL_CIRCUIT.ID
    left outer join NP_FLOW_DEFINITION
       on MAPS_EVENT.FLOW_DEFINITION_ID = NP_FLOW_DEFINITION.ID
        left outer join FCIP_TUNNEL
       on MAPS_EVENT.FCIP_TUNNEL_ID = FCIP_TUNNEL.ID;
```

N2F_PORT_MAP_REPORT_INFO

```
create or replace view N2F_PORT_MAP_REPORT_INFO as
select
AG_N_PORT.REMOTE_PORT_WWN as EDGE_SWITCH_PORT_WWN,
AG_N_PORT.WWN as AG_N_PORT_WWN,
AG_N_PORT.VIRTUAL_SWITCH_ID as AG_VIRTUAL_SWITCH_ID,
EDGE_SWITCH_PORT.VIRTUAL_SWITCH_ID as EDGE_SWITCH_PORT_VSID,
EDGE_SWITCH_PORT.SLOT_NUMBER as EDGE_SWITCH_PORTSLOT_NUMBER,
EDGE_SWITCH_PORT.PORT_NUMBER as EDGE_SWITCH_PORT_NUMBER,
EDGE_SWITCH_PORT.USER_PORT_NUMBER as EDGE_SWITCH_PORT_USER_PORT_NUMBER,
EDGE_SWITCH_PORT.NAME as EDGE_SWITCH_PORT_NAME,
EDGE_SWITCH_PORT.SPEED as EDGE_SWITCH_PORT_SPEED,
EDGE_SWITCH_PORT.STATUS as EDGE_SWITCH_PORT_STATUS,
EDGE_SWITCH_PORT.STATE as EDGE_SWITCH_PORT_STATE,
EDGE_SWITCH_PORT.TYPE as EDGE_SWITCH_PORT_TYPE,
EDGE_SWITCH_PORT.SPEEDS_SUPPORTED as EDGE_SWITCH_PORT_SPEEDS_SUPPORTED,
EDGE_SWITCH_PORT.PHYSICAL_PORT as EDGE_SWITCH_PORT_PHYSICAL_OR_LOGICAL_PORT,
EDGE_SWITCH_PORT.PORT_INDEX as EDGE_SWITCH_PORT_ZONING_PORT_INDEX,
EDGE_SWITCH_PORT.PORT_ID as EDGE_SWITCH_PORT_ID,
EDGE_SWITCH_PORT.AREA_ID as EDGE_SWITCH_PORT_AREA_ID,
EDGE_SWITCH_PORT.MAC_ADDRESS as EDGE_SWITCH_PORT_MAC_ADDRESS,
EDGE_SWITCH_PORT.PORT_MOD as EDGE_SWITCH_PORT_MOD,
EDGE_SWITCH_PORT.FULL_TYPE as EDGE_SWITCH_PORT_FULL_TYPE,
EDGE_SWITCH_PORT.HEALTH as EDGE_SWITCH_PORT_HEALTH,
EDGE_SWITCH_PORT.STATUS_MESSAGE as EDGE_SWITCH_PORT_STATUS_MESSAGE,
EDGE_SWITCH_PORT.MAX_PORT_SPEED as EDGE_SWITCH_PORT_MAX_PORT_SPEED,
EDGE_SWITCH_PORT.LICENSED as EDGE_SWITCH_PORT_LICENSED,
EDGE_SWITCH_PORT.REMOTE_NODE_WWN as EDGE_SWITCH_PORT_REMOTE_NODE_WWN,
EDGE_SWITCH_PORT.REMOTE_PORT_WWN as EDGE_SWITCH_PORT_REMOTE_PORT_WWN,
EDGE_SWITCH_PORT.TRUNKED as EDGE_SWITCH_PORT_TRUNKED,
```

```
EDGE_SWITCH_PORT.TRUNK_MASTER as EDGE_SWITCH_PORT_TRUNK_MASTER,
EDGE_SWITCH_PORT.FICON_SUPPORTED as EDGE_SWITCH_PORT_FICON_SUPPORTED,
EDGE_SWITCH_PORT.BLOCKED as EDGE_SWITCH_PORT_BLOCKED,
EDGE_SWITCH_PORT.NPIV as EDGE_SWITCH_PORT_NPIV,
EDGE_SWITCH_PORT.NPIV_CAPABLE as EDGE_SWITCH_PORT_NPIV_CAPABLE,
EDGE_SWITCH_PORT.NPIV_ENABLED as EDGE_SWITCH_PORT_NPIV_ENABLED,
EDGE_SWITCH_PORT.QOS_CAPABLE as EDGE_SWITCH_PORT_QOS_CAPABLE,
EDGE_SWITCH_PORT.QOS_ENABLED as EDGE_SWITCH_PORT_QOS_ENABLED,
EDGE_SWITCH_PORT.TUNNEL_CONFIGURED as EDGE_SWITCH_PORT_TUNNEL_CONFIGURED,
EDGE_SWITCH_PORT.FCR_FABRIC_ID as EDGE_SWITCH_PORT_FCR_FABRIC_ID,
EDGE_SWITCH_PORT.FCR_INTEROP_MODE as EDGE_SWITCH_PORT_FCR_INTEROP_MODE,
EDGE_SWITCH_PORT.USER_DEFINED_VALUE1 as EDGE_SWITCH_PORT_USER_DEFINED_VALUE1,
EDGE_SWITCH_PORT.USER_DEFINED_VALUE2 as EDGE_SWITCH_PORT_USER_DEFINED_VALUE2,
EDGE_SWITCH_PORT.USER_DEFINED_VALUE3 as EDGE_SWITCH_PORT_USER_DEFINED_VALUE3,
EDGE_SWITCH_PORT.KIND as EDGE_SWITCH_PORT_KIND,
EDGE_SWITCH_PORT.LAST_UPDATE as EDGE_SWITCH_PORT_LAST_UPDATE,
EDGE_SWITCH.WWN as EDGE_VIRTUAL_SWITCH_WWN,
EDGE_SWITCH.NAME as EDGE_SWITCH_NAME,
EDGE_SWITCH.DOMAIN_ID as EDGE_SWITCH_DOMAIN_ID,
EDGE_SWITCH.SWITCH_MODE as EDGE_SWITCH_MODE,
EDGE_SWITCH.OPERATIONAL_STATUS as EDGE_SWITCH_OPERATIONAL_STATUS,
EDGE_SWITCH.MANAGEMENT_STATE as EDGE_SWITCH_MANAGEMENT_STATE,
EDGE_SWITCH.STATE as EDGE_SWITCH_STATE,
EDGE_SWITCH.STATUS as EDGE_SWITCH_STATUS,
EDGE_SWITCH.STATUS_REASON as EDGE_SWITCH_STATUS_REASON,
EDGE_PHYSICAL_SWITCH.IP_ADDRESS as EDGE_PHYSICAL_SWITCH_IP_ADDRESS,
EDGE_PHYSICAL_SWITCH.ID as EDGE_PHYSICAL_SWITCH_ID,
EDGE_PHYSICAL_SWITCH.WWN as EDGE_PHYSICAL_SWITCH_WWN,
EDGE_PHYSICAL_SWITCH.NAME as EDGE_PHYSICAL_SWITCH_NAME,
EDGE_PHYSICAL_SWITCH.TYPE as EDGE_PHYSICAL_SWITCH_TYPE,
EDGE_PHYSICAL_SWITCH.MODEL as EDGE_PHYSICAL_SWITCH_MODEL,
EDGE_PHYSICAL_SWITCH.VENDOR as EDGE_PHYSICAL_SWITCH_VENDOR,
EDGE_PHYSICAL_SWITCH.REACHABLE as EDGE_PHYSICAL_SWITCH_REACHABLE,
EDGE_PHYSICAL_SWITCH.OPERATIONAL_STATUS as EDGE_PHYSICAL_SWITCH_OPERATIONAL_STATUS,
EDGE_PHYSICAL_SWITCH_DETAILS.MODEL_NUMBER as EDGE_PHYSICAL_SWITCH_MODEL_NUMBER,
EDGE_FABRIC.SEED_SWITCH_WWN as FABRIC_SEED_SWITCH_WWN,
coalesce(EDGE_FABRIC.NAME, EDGE_FABRIC.FABRIC_NAME) as FABRIC_NAME,
EDGE_FABRIC.MANAGED as FABRIC_MANAGED,
EDGE_FABRIC.TYPE as FABRIC_TYPE,
EDGE_FABRIC.SEED_SWITCH_WWN,
EDGE_FABRIC.PRINCIPAL_SWITCH_WWN as PRINCIPAL_WWN,
AG_F_PORT.WWN as AG_F_PORT_WWN,
AG_F_PORT.SLOT_NUMBER as AG_F_PORT_SLOT_NUMBER,
AG_F_PORT.PORT_NUMBER as AG_F_PORT_NUMBER,
AG_F_PORT.USER_PORT_NUMBER as AG_F_PORT_USER_PORT_NUMBER,
AG_F_PORT.NAME as AG_F_PORT_NAME,
AG_F_PORT.SPEED as AG_F_PORT_PORT_SPEED,
AG_F_PORT.STATUS as AG_F_PORT_STATUS,
AG_F_PORT.STATE as AG_F_PORT_STATE,
AG_F_PORT.TYPE as AG_F_PORT_TYPE,
AG_F_PORT.SPEEDS_SUPPORTED as AG_F_PORT_SPEEDS_SUPPORTED,
AG_F_PORT.PHYSICAL_PORT as AG_F_PORT_PHYSICAL_OR_LOGICAL_PORT,
AG_F_PORT.PORT_INDEX as AG_F_PORT_ZONING_PORT_INDEX,
AG_F_PORT.PORT_ID as AG_F_PORT_PORT_ID,
AG_F_PORT.WWN as AG_F_PORT_SWITCH_PORT_WWN,
AG_F_PORT.AREA_ID as AG_F_PORT_AREA_ID,
AG_F_PORT.MAC_ADDRESS as AG_F_PORT_MAC_ADDRESS,
AG_F_PORT.PORT_MOD as AG_F_PORT_MOD,
AG_F_PORT.FULL_TYPE as AG_F_PORT_FULL_TYPE,
AG_F_PORT.HEALTH as AG_F_PORT_HEALTH,
AG_F_PORT.STATUS_MESSAGE as AG_F_SWITCH_PORT_STATUS_MESSAGE,
AG_F_PORT.MAX_PORT_SPEED as AG_F_PORT_MAX_PORT_SPEED,
AG_F_PORT.LICENSED as AG_F_PORT_LICENSED,
```

```
AG_F_PORT.REMOTE_NODE_WWN as AG_F_PORT_NODE_WWN,
AG_F_PORT.REMOTE_PORT_WWN as AG_F_PORT_PORT_WWN,
AG_F_PORT.TRUNKED as AG_F_PORT_TRUNKED,
AG_F_PORT.TRUNK_MASTER as AG_F_PORT_TRUNK_MASTER,
AG_F_PORT.FICON_SUPPORTED as AG_F_PORT_FICON_SUPPORTED,
AG_F_PORT.BLOCKED as AG_F_PORT_BLOCKED,
AG_F_PORT.NPIV as AG_F_PORT_NPIV,
AG_F_PORT.NPIV_CAPABLE as AG_F_PORT_NPIV_CAPABLE,
AG_F_PORT.NPIV_ENABLED as AG_F_PORT_NPIV_ENABLED,
AG_F_PORT.QOS_CAPABLE as AG_F_PORT_QOS_CAPABLE,
AG_F_PORT.QOS_ENABLED as AG_F_PORT_QOS_ENABLED,
AG_F_PORT.TUNNEL_CONFIGURED as AG_F_PORT_TUNNEL_CONFIGURED,
AG_F_PORT.FCR_FABRIC_ID as AG_F_PORT_FCR_FABRIC_ID,
AG_F_PORT.FCR_INTEROP_MODE as AG_F_PORT_FCR_INTEROP_MODE,
AG_F_PORT.USER_DEFINED_VALUE1 as AG_F_PORT_USER_DEFINED_VALUE1,
AG_F_PORT.USER_DEFINED_VALUE2 as AG_F_PORT_USER_DEFINED_VALUE2,
AG_F_PORT.USER_DEFINED_VALUE3 as AG_F_PORT_USER_DEFINED_VALUE3,
AG_F_PORT.KIND as AG_F_PORT_KIND,
AG_F_PORT.LAST_UPDATE as AG_F_PORT_LAST_UPDATE,
AG_SWITCH.WWN as AG_SWITCH_VIRTUAL_SWITCH_WWN,
AG_SWITCH.NAME as  AG_SWITCH_SWITCH_NAME,
AG_SWITCH.DOMAIN_ID as  AG_SWITCH_SWITCH_DOMAIN_ID,
AG_SWITCH.SWITCH_MODE as AG_SWITCH_SWITCH_MODE,
AG_SWITCH.OPERATIONAL_STATUS as AG_SWITCH_SWITCH_OPERATIONAL_STATUS,
AG_SWITCH.MANAGEMENT_STATE as AG_SWITCH_SWITCH_MANAGEMENT_STATE,
AG_SWITCH.STATE as AG_SWITCH_SWITCH_STATE,
AG_SWITCH.STATUS as AG_SWITCH_SWITCH_STATUS,
AG_SWITCH.STATUS_REASON as AG_SWITCH_SWITCH_STATUS_REASON,
AG_PHYSICAL_SWITCH.IP_ADDRESS as  AG_PHYSICAL_SWITCH_IP_ADDRESS,
AG_PHYSICAL_SWITCH.ID as AG_PHYSICAL_SWITCH_ID,
AG_PHYSICAL_SWITCH.WWN as AG_PHYSICAL_SWITCH_WWN,
AG_PHYSICAL_SWITCH.NAME as AG_PHYSICAL_SWITCH_NAME,
AG_PHYSICAL_SWITCH.TYPE as AG_PHYSICAL_SWITCH_TYPE,
AG_PHYSICAL_SWITCH.MODEL as AG_PHYSICAL_SWITCH_MODEL,
AG_PHYSICAL_SWITCH.VENDOR as AG_PHYSICAL_SWITCH_VENDOR,
AG_PHYSICAL_SWITCH.REACHABLE as AG_PHYSICAL_SWITCH_REACHABLE,
AG_PHYSICAL_SWITCH.OPERATIONAL_STATUS as AG_PHYSICAL_SWITCH_OPERATIONAL_STATUS,
AG_PHYSICAL_SWITCH_DETAILS.MODEL_NUMBER as AG_PHYSICAL_SWITCH_MODEL_NUMBER,
END_DEVICE_NODE.FABRIC_ID,
END_DEVICE_NODE.SYMBOLIC_NAME as DEVICE_NODE_SYMBOLIC_NAME,
END_DEVICE_NODE.FDMI_HOST_NAME,
END_DEVICE_NODE.VENDOR,
END_DEVICE_NODE.CAPABILITY_,
END_DEVICE_NODE.TRUSTED as DEVICE_NODE_TRUSTED,
END_DEVICE_NODE.CREATION_TIME as DEVICE_NODE_CREATION_TIME,
END_DEVICE_NODE.MISSING as DEVICE_NODE_MISSING,
END_DEVICE_NODE.MISSING_TIME as DEVICE_NODE_MISSING_TIME,
END_DEVICE_NODE.PROXY_DEVICE,
END_DEVICE_NODE.AG,
END_DEVICE_NODE.PREVIOUS_MISSING_STATE,
END_DEVICE_NODE.SIMULATED,
coalesce(USER_DEFINED_DEVICE_DETAIL.TYPE, END_DEVICE_NODE.TYPE, ''::character varying) as
USER_DEFINED_DEVICE_TYPE,
USER_DEFINED_DEVICE_DETAIL.NAME as USER_DEFINED_NAME,
USER_DEFINED_DEVICE_DETAIL.IP_ADDRESS as USER_DEFINED_IP_ADDRESS,
USER_DEFINED_DEVICE_DETAIL.CONTACT,
USER_DEFINED_DEVICE_DETAIL.LOCATION,
USER_DEFINED_DEVICE_DETAIL.DESCRIPTION,
AG_F_PORT.REMOTE_PORT_WWN as DEVICE_PORT_WWN,
AG_F_PORT.REMOTE_NODE_WWN  as DEVICE_NODE_WWN,
END_DEVICE_PORT.PORT_ID as DEVICE_PORT_FC_ADDRESS,
END_DEVICE_PORT.NUMBER as DEVICE_PORT_NUMBER,
END_DEVICE_PORT.ID as DEVICE_PORT_ID,
```

```
END_DEVICE_PORT.DOMAIN_ID,
END_DEVICE_PORT.TYPE as DEVICE_PORT_TYPE,
END_DEVICE_PORT.SYMBOLIC_NAME as DEVICE_PORT_SYMBOLIC_NAME,
END_DEVICE_PORT.FC4_TYPE,
END_DEVICE_PORT.COS,
END_DEVICE_PORT.IP_PORT,
END_DEVICE_PORT.NPV_PHYSICAL,
END_DEVICE_PORT.HARDWARE_ADDRESS,
END_DEVICE_PORT.TRUSTED as DEVICE_PORT_TRUSTED,
END_DEVICE_PORT.CREATION_TIME as DEVICE_PORT_CREATION_TIME,
END_DEVICE_PORT.MISSING as DEVICE_PORT_MISSING,
END_DEVICE_PORT.MISSING_TIME as DEVICE_PORT_MISSING_TIME,
END_DEVICE_PORT.LOGGED_TO_AG,
END_DEVICE_PORT.AG_NODE_WWN,
END_DEVICE_PORT.AG_N_PORT_WWN As AG_SWITCH_N_PORT_WWN,
END_DEVICE_PORT.MISSING_REASON

from  N2F_PORT_MAP
left join SWITCH_PORT AG_F_PORT on AG_F_PORT.USER_PORT_NUMBER = N2F_PORT_MAP.F_PORT AND
N2F_PORT_MAP.VIRTUAL_SWITCH_ID = AG_F_PORT.VIRTUAL_SWITCH_ID
left join VIRTUAL_SWITCH AG_SWITCH on AG_F_PORT.VIRTUAL_SWITCH_ID = AG_SWITCH.ID AND AG_SWITCH.MONITORED = 1
left join CORE_SWITCH AG_PHYSICAL_SWITCH on AG_SWITCH.CORE_SWITCH_ID = AG_PHYSICAL_SWITCH.ID
left join CORE_SWITCH_DETAILS AG_PHYSICAL_SWITCH_DETAILS on AG_SWITCH.CORE_SWITCH_ID =
AG_PHYSICAL_SWITCH_DETAILS.CORE_SWITCH_ID
left join DEVICE_PORT END_DEVICE_PORT on AG_F_PORT.REMOTE_PORT_WWN = END_DEVICE_PORT.WWN and AG_F_PORT.WWN =
END_DEVICE_PORT.SWITCH_PORT_WWN
left join DEVICE_NODE END_DEVICE_NODE on END_DEVICE_PORT.NODE_ID = END_DEVICE_NODE.ID
left join USER_DEFINED_DEVICE_DETAIL on END_DEVICE_NODE.WWN = USER_DEFINED_DEVICE_DETAIL.WWN

left join SWITCH_PORT AG_N_PORT on AG_N_PORT.USER_PORT_NUMBER = N2F_PORT_MAP.N_PORT AND
N2F_PORT_MAP.VIRTUAL_SWITCH_ID = AG_N_PORT.VIRTUAL_SWITCH_ID
left join SWITCH_PORT EDGE_SWITCH_PORT on AG_N_PORT.REMOTE_PORT_WWN = EDGE_SWITCH_PORT.WWN
left join VIRTUAL_SWITCH EDGE_SWITCH on EDGE_SWITCH_PORT.VIRTUAL_SWITCH_ID = EDGE_SWITCH.ID
left join CORE_SWITCH EDGE_PHYSICAL_SWITCH on EDGE_SWITCH.CORE_SWITCH_ID = EDGE_PHYSICAL_SWITCH.ID
left join CORE_SWITCH_DETAILS EDGE_PHYSICAL_SWITCH_DETAILS on EDGE_SWITCH.CORE_SWITCH_ID =
EDGE_PHYSICAL_SWITCH_DETAILS.CORE_SWITCH_ID
left join FABRIC_MEMBER EDGE_FABRIC_MEMBER ON EDGE_SWITCH.ID = EDGE_FABRIC_MEMBER.VIRTUAL_SWITCH_ID
left join FABRIC EDGE_FABRIC on EDGE_FABRIC_MEMBER.FABRIC_ID = EDGE_FABRIC.ID AND EDGE_FABRIC.MANAGED = 1;
```

MODULE_INFO

```
CREATE VIEW module_info AS
    select distinct
TEMP_MODULE.MODULE_ID,
TEMP_MODULE.NUM_PORTS,
TEMP_MODULE.IS_PRESENT,
case
    when TEMP_MODULE.IS_PRESENT = 1 then 'YES'
    else 'NO'
end as IS_PRESENT_TXT,
TEMP_MODULE.IS_MANAGEMENT_MODULE,
case
    when TEMP_MODULE.IS_MANAGEMENT_MODULE = 1 then 'YES'
    else 'NO'
end as IS_MANAGEMENT_MODULE_TXT,
TEMP_MODULE.NUM_CPUS,
TEMP_MODULE.HW_REVISION,
TEMP_MODULE.SW_REVISION,
TEMP_MODULE.SLOT_NUM,
TEMP_MODULE.DEVICE_ID,
TEMP_MODULE.PHYSICAL_DEVICE_ID,
TEMP_MODULE.UNIT_NUMBER,
```

```
TEMP_MODULE.UNIT_PRESENT,
case
    when TEMP_MODULE.UNIT_PRESENT = 1 then 'YES'
    else 'NO'
end as UNIT_PRESENT_TXT,
TEMP_MODULE.MANAGED_ELEMENT_ID,
TEMP_MODULE.IP_ADDRESS,
TEMP_FOUNDRY_MODULE.SERIAL_NUM,
TEMP_FOUNDRY_MODULE.DRAM_SIZE,
TEMP_FOUNDRY_MODULE.BOOT_FLASH_SIZE,
TEMP_FOUNDRY_MODULE.CODE_FLASH_SIZE,
TEMP_FOUNDRY_MODULE.MODULE_TYPE,
TEMP_MODULE.DESCRIPTION as MODULE_TYPE_TXT,
TEMP_MODULE.MODULE_STATUS,
TEMP_MODULE.REDUNDANT_STATUS
from
    (
    select distinct
    MODULE.MODULE_ID,
    MODULE.NUM_PORTS,
    MODULE.IS_PRESENT,
    MODULE.IS_MANAGEMENT_MODULE,
    MODULE.NUM_CPUS,
    MODULE.HW_REVISION,
    MODULE.SW_REVISION,
    SLOT.SLOT_NUM,
    PHYSICAL_DEVICE.DEVICE_ID,
    PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID,
    PHYSICAL_DEVICE.UNIT_NUMBER,
    PHYSICAL_DEVICE.UNIT_PRESENT,
    DEVICE.MANAGED_ELEMENT_ID,
    DEVICE.IP_ADDRESS,
    MODULE.DESCRIPTION,
    MODULE.MODULE_STATUS,
    MODULE.REDUNDANT_STATUS
    from MODULE, SLOT, MODULE_SLOT_PRESENT, DEVICE, PHYSICAL_DEVICE
    where
    MODULE.MODULE_ID = MODULE_SLOT_PRESENT.MODULE_ID
    and MODULE_SLOT_PRESENT.SLOT_ID = SLOT.SLOT_ID
    and SLOT.PHYSICAL_DEVICE_ID = PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID
    and DEVICE.DEVICE_ID = PHYSICAL_DEVICE.DEVICE_ID
    ) TEMP_MODULE
left join
    (
    select
    FOUNDRY_MODULE.MODULE_ID,
    FOUNDRY_MODULE.SERIAL_NUM,
    FOUNDRY_MODULE.DRAM_SIZE,
    FOUNDRY_MODULE.BOOT_FLASH_SIZE,
    FOUNDRY_MODULE.CODE_FLASH_SIZE,
    FOUNDRY_MODULE.MODULE_TYPE
    from FOUNDRY_MODULE
    ) TEMP_FOUNDRY_MODULE ON TEMP_MODULE.MODULE_ID = TEMP_FOUNDRY_MODULE.MODULE_ID;
```

MON_AOR_INFO

```
create or replace view MON_AOR_INFO as
select
   USER_AOR_MAP.USER_NAME,
   AOR.NAME,
   AOR.DESCRIPTION,
```

```
   AOR_FABRIC_MAP.FABRIC_ID,
   AOR_HOST_MAP.HOST_ID,
   AOR.ID as AOR_ID
from
    USER_AOR_MAP
    left outer join AOR on
    USER_AOR_MAP.AOR_ID=AOR.ID
    left outer join AOR_FABRIC_MAP on
    AOR.ID=AOR_FABRIC_MAP.AOR_ID
    left outer join AOR_HOST_MAP on
    AOR.ID=AOR_HOST_MAP.AOR_ID;
```

MON_DEVICE_CONNECTION_INFO

```
create or replace view MON_DEVICE_CONNECTION_INFO as
select
      distinct on (DEVICE_CONNECTION.ID)
     DEVICE_CONNECTION.ID,
          DEVICE_CONNECTION.FABRIC_ID,
    DEVICE_CONNECTION.DEVICE_PORT_ID,
          DEVICE_CONNECTION.SWITCH_PORT_ID,
    DEVICE_CONNECTION.AG_PORT_ID,
    COALESCE(DEVICE_ENCLOSURE_MEMBER.ENCLOSURE_ID, HBA.HOST_ID, VM_HOST.DEVICE_ENCLOSURE_ID) as
DEVICE_ENCLOSURE_ID,
    DEVICE_PORT.NODE_ID,
    DEVICE_CONNECTION.MISSING,
    SWPORT.VIRTUAL_SWITCH_ID,
    AGPORT.VIRTUAL_SWITCH_ID as AG_SWITCH_ID,
    DEVICE_PORT.WWN as DEVICE_PORT_WWN,
    COALESCE(USERDEFINEDDETAILS.TYPE, DN.TYPE) as DEVICE_TYPE
from DEVICE_CONNECTION
    inner join DEVICE_PORT on DEVICE_CONNECTION.DEVICE_PORT_ID = DEVICE_PORT.ID
    inner join DEVICE_NODE DN on DEVICE_PORT.NODE_ID = DN.ID
    left join SWITCH_PORT SWPORT on DEVICE_CONNECTION.SWITCH_PORT_ID = SWPORT.ID
    left join SWITCH_PORT AGPORT on DEVICE_CONNECTION.AG_PORT_ID = AGPORT.ID
    left join HBA_PORT_DEVICE_PORT_MAP on DEVICE_PORT.ID = HBA_PORT_DEVICE_PORT_MAP.DEVICE_PORT_ID
    left join HBA_PORT on HBA_PORT_DEVICE_PORT_MAP.HBA_PORT_ID = HBA_PORT.DEVICE_PORT_ID
    left join HBA on HBA_PORT.HBA_ID = HBA.ID
    left join VM_FC_HBA_DEVICE_PORT_MAP ON VM_FC_HBA_DEVICE_PORT_MAP.DEVICE_PORT_ID = DEVICE_PORT.ID
    left join VM_FC_HBA ON VM_FC_HBA.ID = VM_FC_HBA_DEVICE_PORT_MAP.VM_FC_HBA_ID
    left join VM_HOST ON VM_HOST.DEVICE_ENCLOSURE_ID = VM_FC_HBA.VM_HOST_ID
    left join DEVICE_ENCLOSURE_MEMBER on DEVICE_PORT.ID = DEVICE_ENCLOSURE_MEMBER.DEVICE_PORT_ID
    left join USER_DEFINED_DEVICE_DETAIL USERDEFINEDDETAILS on DN.WWN = USERDEFINEDDETAILS.WWN;
```

MON_DEVICE_PORT_INFO

```
create or replace view MON_DEVICE_PORT_INFO as
select
        distinct on (DEVICE_PORT.ID)
        DEVICE_PORT.ID,
        DEVICE_PORT.NODE_ID,
        DEVICE_PORT.DOMAIN_ID,
        DEVICE_PORT.WWN,
    DEVICE_PORT.SWITCH_PORT_WWN,
        DEVICE_PORT.NUMBER,
        DEVICE_PORT.PORT_ID,
    DEVICE_PORT.TYPE,
        DEVICE_PORT.SYMBOLIC_NAME,
DEVICE_PORT.MISSING,
    DEVICE_PORT.NPV_PHYSICAL,
    DEVICE_PORT.EDGE_SWITCH_PORT_WWN,
```

```
    DEVICE_PORT.AG_NODE_WWN,
    DEVICE_PORT.AG_N_PORT_WWN,
    FICON_DEVICE_PORT.TYPE_NUMBER,
        FICON_DEVICE_PORT.MODEL_NUMBER,
UDD_PORT.NAME as PORT_NAME,
        UDD_PORT.USER_DEFINED_VALUE1 as PORT_USER_DEFINED_VALUE1,
        UDD_PORT.USER_DEFINED_VALUE2 as PORT_USER_DEFINED_VALUE2,
        UDD_PORT.USER_DEFINED_VALUE3 as PORT_USER_DEFINED_VALUE3,
UDD_NODE.NAME as NODE_NAME,
        UDD_NODE.TYPE as USER_DEFINED_TYPE,
        UDD_NODE.IP_ADDRESS,
        UDD_NODE.CONTACT,
        UDD_NODE.LOCATION,
        UDD_NODE.DESCRIPTION,
        UDD_NODE.USER_DEFINED_VALUE1 as NODE_USER_DEFINED_VALUE1,
        UDD_NODE.USER_DEFINED_VALUE2 as NODE_USER_DEFINED_VALUE2,
        UDD_NODE.USER_DEFINED_VALUE3 as NODE_USER_DEFINED_VALUE3,
DEVICE_NODE.WWN as DEVICE_NODE_WWN,
    DEVICE_NODE.FDMI_HOST_NAME,
    DEVICE_NODE.SYMBOLIC_NAME as DEVICE_SYMBOLIC_NAME,
    DEVICE_NODE.AG as AG_DEVICE,
    DEVICE_NODE.PROXY_DEVICE,
    DEVICE_NODE.TYPE as DEVICE_TYPE,
    coalesce(SWITCH_PORT.NAME, VIRTUAL_FCOE_PORT.NAME) as SWITCH_PORT_NAME,
    coalesce (SWITCH_PORT.TYPE, VIRTUAL_FCOE_PORT.PORT_TYPE) as SWITCH_PORT_TYPE,
    SWITCH_PORT.LOGICAL_PORT_WWN as SWITCH_LOGICAL_PORT_WWN,
    coalesce(VS1.WWN, VS2.WWN) as SWITCH_WWN,
    coalesce(VS1.MANAGEMENT_STATE, VS2.MANAGEMENT_STATE) as MANAGEMENT_STATE,
    coalesce(VS1.MONITORED, VS2.MONITORED) as MONITORED,
    FABRIC.PRINCIPAL_SWITCH_WWN as PRINCIPAL_WWN,
    FABRIC.ID as FABRIC_ID,
        coalesce(VS1.MANAGED_ELEMENT_ID, VS2.MANAGED_ELEMENT_ID) AS SWITCH_ME_ID,
    SWITCH_PORT.ID as SWITCH_PORT_DB_ID,
    SWITCH_PORT.PORT_ID as SWITCH_PORT_ID
from
        DEVICE_PORT
            inner join DEVICE_NODE
            on DEVICE_PORT.NODE_ID = DEVICE_NODE.ID
                        left outer join USER_DEFINED_DEVICE_DETAIL as UDD_PORT
                                    on DEVICE_PORT.WWN = UDD_PORT.WWN
                        left outer join USER_DEFINED_DEVICE_DETAIL as UDD_NODE
                                    on DEVICE_NODE.WWN = UDD_NODE.WWN
                        left outer join FICON_DEVICE_PORT
                                    on DEVICE_PORT.ID = FICON_DEVICE_PORT.DEVICE_PORT_ID
        left outer join SWITCH_PORT
            on DEVICE_PORT.SWITCH_PORT_WWN = SWITCH_PORT.WWN
        left outer join VIRTUAL_FCOE_PORT
            on DEVICE_PORT.SWITCH_PORT_WWN = VIRTUAL_FCOE_PORT.PORT_WWN
        left outer join VIRTUAL_SWITCH VS1
            on SWITCH_PORT.VIRTUAL_SWITCH_ID = VS1.ID
        left outer join VIRTUAL_SWITCH VS2
            on VIRTUAL_FCOE_PORT.VIRTUAL_SWITCH_ID = VS2.ID
        left outer join FABRIC
            on DEVICE_NODE.FABRIC_ID = FABRIC.ID
where SWITCH_PORT.FAKE_PORT  is null or SWITCH_PORT.FAKE_PORT = 0;
```

MON_HBA_PORT_DETAILS_INFO

```
create or replace view MON_HBA_PORT_INFO as
select
    distinct on (HBA_PORT.DEVICE_PORT_ID)
    HBA_PORT.DEVICE_PORT_ID,
```

```
    HBA_PORT.OPERATING_STATE,
    HBA_PORT.MISSING,
    HBA_PORT.OPERATING_SPEED,
    HBA_PORT.PORT_NWWN as NODE_WWN,
    HBA_PORT.PHYSICAL_PORT_WWN,
    HBA_PORT.SWITCH_IP,
    HBA_PORT.PRINCIPAL_SWITCH_WWN,
    HBA_PORT.PORT_NUMBER,
    HBA_PORT.NAME,
    HBA_PORT_DETAIL.V_PORT_COUNT,
    HBA.HOST_ID,
    HBA.NAME as HBA_NAME,
    HBA.MODEL,
    HBA.MODEL_DESCRIPTION,
    HBA.OPERATING_STATUS,
    HBA.SERIAL_NUMBER,
    HBA.TEMPERATURE,
    HBA.MANAGEMENT_STATE,
    HBA.FIRMWARE_VERSION,
    HBA.MAC_ADDRESS
from HBA_PORT
inner join HBA on HBA_PORT.HBA_ID =HBA.ID
left outer join HBA_PORT_DETAIL on HBA_PORT.DEVICE_PORT_ID = HBA_PORT_DETAIL.DEVICE_PORT_ID;
```

MON_HBA_TARGET_INFO

```
create or replace view MON_HBA_TARGET_INFO as
select
    distinct on (HBA_TARGET.DEVICE_PORT_ID)
    HBA_TARGET.DEVICE_PORT_ID,
    HBA_TARGET.HBA_REMOTE_PORT_LUN_ID,
    HBA_TARGET.MISSING as LUN_MISSING,
    HBA_REMOTE_PORT.SYMBOLIC_NAME,
    HBA_REMOTE_PORT.PORT_WWN,
    HBA_REMOTE_PORT.NODE_WWN,
    HBA_REMOTE_PORT.NAME,
    HBA_REMOTE_PORT.FC_ADDRESS,
    HBA_REMOTE_PORT.SPEED,
    HBA_REMOTE_PORT.STATE,
    HBA_REMOTE_PORT.DEVICE_TYPE,
    HBA_REMOTE_PORT_LUN.TARGET_WWN,
    HBA_REMOTE_PORT_LUN.PHYSICAL_LUN,
    HBA_REMOTE_PORT_LUN.LUN_ID
from
    HBA_TARGET, HBA_REMOTE_PORT, HBA_REMOTE_PORT_LUN
where
    HBA_TARGET.HBA_REMOTE_PORT_LUN_ID = HBA_REMOTE_PORT_LUN.ID and
    HBA_REMOTE_PORT.ID = HBA_REMOTE_PORT_LUN.HBA_REMOTE_PORT_ID;
```

MON_MAPS_EVENT_DETAILS_INFO

```
create or replace view MON_MAPS_EVENT_DETAILS_INFO as
select
    distinct on (MAPS_EVENT.ID)
    MAPS_EVENT.ID,
    MAPS_EVENT.HOST_TIME,
    MAPS_EVENT.CATEGORY,
    MAPS_EVENT.VIOLATION_TYPE,
    MAPS_EVENT.MANAGED_ELEMENT_ID,
```

```
    MAPS_EVENT.ORIGIN_FABRIC_ID,
    MAPS_EVENT.SWITCH_PORT_ID,
    MAPS_EVENT.INTERFACE_ID,
    MAPS_EVENT.FCIP_CIRCUIT_ID,
    MAPS_EVENT.FRU_NAME,
    MAPS_EVENT.VM_ID,
    MAPS_EVENT.FLOW_DEFINITION_ID,
    MAPS_EVENT. SUB_FLOW_KEY,
    MAPS_EVENT.FCIP_TUNNEL_ID,
    MAPS_EVENT.PORT_TYPE,
    MAPS_EVENT.COLLECTION_NAME,
    MAPS_EVENT_DETAILS.SWITCH_TIME,
    MAPS_EVENT_DETAILS.RULE_NAME,
    MAPS_EVENT_DETAILS.RULE_CONDITION,
    MAPS_EVENT_DETAILS.TIME_BASE,
    MAPS_EVENT_DETAILS.ACTIONS,
    MAPS_EVENT_DETAILS.CURRENT_VALUE,
    MAPS_EVENT_DETAILS.SWITCH_ENABLED_ACTIONS,
    VIRTUAL_SWITCH.NAME as SWITCH_NAME,
    SWITCH_PORT.NAME as SWITCH_PORT_NAME,
    SWITCH_PORT.WWN as SWITCH_PORT_WWN,
    SWITCH_PORT.SLOT_NUMBER as SWITCH_PORT_SLOT,
    SWITCH_PORT.PORT_NUMBER as SWITCH_PORT_NUMBER,
    SWITCH_PORT.PORT_ID as SWITCH_PORT_PORT_ID,
    FCIP_TUNNEL_CIRCUIT.CIRCUIT_NUMBER as FCIP_TUNNEL_CIRCUIT_NUMBER,
    FCIP_TUNNEL_CIRCUIT.SLOT_NUMBER as FCIP_SLOT_NUMBER,
    FCIP_TUNNEL_CIRCUIT.VE_PORT_NUMBER as FCIP_PORT_NUMBER,
    NP_FLOW_DEFINITION.NAME as FLOW_NAME,
    MAPS_EVENT_CAUSE_ACTION.ACTION
from
    MAPS_EVENT_DETAILS
    inner join
        MAPS_EVENT on
        MAPS_EVENT.ID = MAPS_EVENT_DETAILS.MAPS_EVENT_ID
    left outer join MAPS_EVENT_CAUSE_ACTION
       on MAPS_EVENT.VIOLATION_TYPE = MAPS_EVENT_CAUSE_ACTION.VIOLATION_TYPE
    left outer join VIRTUAL_SWITCH
       on MAPS_EVENT.MANAGED_ELEMENT_ID = VIRTUAL_SWITCH.MANAGED_ELEMENT_ID
    left outer join SWITCH_PORT
       on MAPS_EVENT.SWITCH_PORT_ID = SWITCH_PORT.ID
    left outer join FCIP_TUNNEL_CIRCUIT
       on MAPS_EVENT.FCIP_CIRCUIT_ID = FCIP_TUNNEL_CIRCUIT.ID
    left outer join NP_FLOW_DEFINITION
       on MAPS_EVENT.FLOW_DEFINITION_ID = NP_FLOW_DEFINITION.ID
    left outer join FCIP_TUNNEL
       on MAPS_EVENT.FCIP_TUNNEL_ID = FCIP_TUNNEL.ID;
```

MON_SWITCH_INFO

```
create or replace view MON_SWITCH_INFO as
select
   CORE_SWITCH.ID as PHYSICAL_SWITCH_ID,
   CORE_SWITCH.NAME as PHYSICAL_SWITCH_NAME,
   CORE_SWITCH.IP_ADDRESS,
   CORE_SWITCH.WWN as PHYSICAL_SWITCH_WWN,
   CORE_SWITCH.OPERATIONAL_STATUS as PHYSICAL_OPERATIONAL_STATUS,
   CORE_SWITCH_DETAILS.TYPE,
   CORE_SWITCH.MODEL,
   CORE_SWITCH.MANAGED_ELEMENT_ID as CORE_MANAGED_ELEMENT_ID,
   CORE_SWITCH_DETAILS.FIRMWARE_VERSION,
   VIRTUAL_SWITCH.ID,
   VIRTUAL_SWITCH.NAME,
```

```
   VIRTUAL_SWITCH.WWN,
   VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID,
   VIRTUAL_SWITCH.DOMAIN_ID,
   VIRTUAL_SWITCH.SWITCH_MODE,
   VIRTUAL_SWITCH.OPERATIONAL_STATUS,
VIRTUAL_SWITCH.USER_NAME,
   VIRTUAL_SWITCH.PASSWORD,
   VIRTUAL_SWITCH.MANAGEMENT_STATE,
   VIRTUAL_SWITCH.STATE,
   VIRTUAL_SWITCH.STATUS,

   VIRTUAL_SWITCH.USER_DEFINED_VALUE_1,
   VIRTUAL_SWITCH.USER_DEFINED_VALUE_2,
   VIRTUAL_SWITCH.USER_DEFINED_VALUE_3,
VIRTUAL_SWITCH.DEFAULT_LOGICAL_SWITCH,
   VIRTUAL_SWITCH.DISCOVERED_PORT_COUNT,
   VIRTUAL_SWITCH.MANAGED_ELEMENT_ID,
   VIRTUAL_SWITCH.MONITORED,

FABRIC_MEMBER.FABRIC_ID,
   FABRIC_MEMBER.MISSING,
   FABRIC.NAME as FABRIC_DCM_NAME,
   FABRIC.FABRIC_NAME,
   FABRIC.SEED_SWITCH_WWN,
   FABRIC.TYPE as FABRIC_TYPE,
   FABRIC.MANAGED as FABRIC_MANAGED,
   FABRIC.PRINCIPAL_SWITCH_WWN,
   VIRTUAL_SWITCH.BOUND,
   VIRTUAL_SWITCH.BOUND_BNA_IP_ADDRESS
from
   CORE_SWITCH,
   VIRTUAL_SWITCH,
   FABRIC_MEMBER,
   FABRIC,
   CORE_SWITCH_DETAILS
where
   VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
   and FABRIC_MEMBER.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
   and FABRIC_MEMBER.FABRIC_ID = FABRIC.ID
   and CORE_SWITCH_DETAILS.CORE_SWITCH_ID = CORE_SWITCH.ID;
```

MON_SWITCH_PORT_INFO

```
create or replace view MON_SWITCH_PORT_INFO as
 select SWITCH_PORT.ID,
    SWITCH_PORT.VIRTUAL_SWITCH_ID,
    SWITCH_PORT.WWN,
    SWITCH_PORT.NAME,
    SWITCH_PORT.SLOT_NUMBER,
    SWITCH_PORT.PORT_NUMBER,
    SWITCH_PORT.USER_PORT_NUMBER,
    SWITCH_PORT.PORT_ID,
    SWITCH_PORT.PORT_INDEX,
    SWITCH_PORT.AREA_ID,
    SWITCH_PORT.MAC_ADDRESS,
    SWITCH_PORT.TYPE,
    SWITCH_PORT.STATUS,
    SWITCH_PORT.PHYSICAL_PORT,
    SWITCH_PORT.SPEED,
    SWITCH_PORT.REMOTE_NODE_WWN,
    SWITCH_PORT.REMOTE_PORT_WWN,
```

```
    SWITCH_PORT.NPIV,
    SWITCH_PORT.NPIV_CAPABLE,
    SWITCH_PORT.NPIV_ENABLED,
    SWITCH_PORT.USER_DEFINED_VALUE1,
    SWITCH_PORT.USER_DEFINED_VALUE2,
    SWITCH_PORT.USER_DEFINED_VALUE3,
    SWITCH_PORT.STATE,
    SWITCH_PORT.OCCUPIED
   from SWITCH_PORT;
```

MON_USER_AOR_INFO

```
create or replace view MON_USER_INFO as
select
   USER_.ID,
   USER_.NAME,
   USER_.DESCRIPTION,
   USER_.PASSWORD,
   USER_.EMAIL,
   USER_.NOTIFICATION_ENABLED,
   USER_.FULL_NAME,
   USER_.PHONE_NUMBER,
   USER_.STATUS
from
    USER_;
```

MON_VM_VIRTUAL_MACHINE_INFO

```
create or replace view MON_VM_VIRTUAL_MACHINE_INFO as
select
   distinct on (VM_VIRTUAL_MACHINE.ID)
   VM_VIRTUAL_MACHINE.ID,
   VM_VIRTUAL_MACHINE.HOST_ID,
   VM_VIRTUAL_MACHINE.HYPERVISOR_VM_ID,
   VM_VIRTUAL_MACHINE.NAME as VM_NAME,
   VM_VIRTUAL_MACHINE.DESCRIPTION,
   VM_VIRTUAL_MACHINE.OS,
   VM_VIRTUAL_MACHINE.STATUS as VM_STATUS,
   VM_VIRTUAL_MACHINE.IP_ADDRESS as VM_IP_ADDRESS,
   VM_VIRTUAL_MACHINE.HOSTNAME,
   VM_VIRTUAL_MACHINE.NODE_WWN as VM_NODE_WWN,
   VM_HOST.NODE_WWN as HOST_NODE_WWN,
   VM_HOST.HYPERVISOR_NAME,
   VM_HOST.HYPERVISOR_TYPE,
   VM_HOST.CLUSTER_NAME,
   VM_PATH.VM_PORT_WWN
from VM_VIRTUAL_MACHINE
inner join VM_HOST on VM_VIRTUAL_MACHINE.HOST_ID=VM_HOST.DEVICE_ENCLOSURE_ID
left join VM_PATH on VM_VIRTUAL_MACHINE.ID=VM_PATH.VM_ID;
```

MON_ZONE_DB_INFO

```
create or replace view MON_ZONE_DB_INFO as
select
    ZONE_DB.ID,
    ZONE_DB.FABRIC_ID,
    ZONE_DB.OFFLINE,
    ZONE_DB.NAME,
    ZONE_DB.LAST_MODIFIED,
```

```
    ZONE_DB.ZONE_CONFIG_SIZE,
    ZONE_DB.ZONE_AVAILABLE_SIZE,
    ZONE_DB_CONFIG.ID AS CONFIG_ID,
    ZONE_DB_CONFIG.DEFINED_CONTENT,
    ZONE_DB_CONFIG.ACTIVE_CONTENT,
    ZONE_DB_CONFIG.TI_ZONE_CONTENT
from
    ZONE_DB, ZONE_DB_CONFIG
where
    ZONE_DB.ID = ZONE_DB_CONFIG.ZONE_DB_ID;
```

NPORT_WWN_MAP_INFO

This view provides a consolidation between Nport WWN map and AG''s N and F ports. It considers only those N-Ports that are currently occupied that is having non-empty remote port wwn. This is required because NPort-WWN mapping might exist for NPorts that are not yet online and if a device is connected to AG through some F-Port that is mapped to some other N-Port that is online then AG will use that mapping.

```
create or replace view NPORT_WWN_MAP_INFO as
select
    NPORT_WWN_MAP.VIRTUAL_SWITCH_ID,
    NPORT_WWN_MAP.N_PORT,
    NPORT_WWN_MAP.DEVICE_PORT_WWN,
    AG_N_PORT.REMOTE_PORT_WWN as EDGE_SWITCH_PORT_WWN,
    AG_N_PORT.WWN as AG_N_PORT_WWN,
    AG_F_PORT.USER_PORT_NUMBER as F_PORT,
    AG_F_PORT.WWN as AG_F_PORT_WWN,
    AG_F_PORT.REMOTE_NODE_WWN
from
    NPORT_WWN_MAP,
    SWITCH_PORT AG_N_PORT,
    SWITCH_PORT AG_F_PORT,
    VIRTUAL_SWITCH AG_SWITCH
where
    NPORT_WWN_MAP.VIRTUAL_SWITCH_ID = AG_N_PORT.VIRTUAL_SWITCH_ID
    and NPORT_WWN_MAP.N_PORT = AG_N_PORT.USER_PORT_NUMBER
    and NPORT_WWN_MAP.VIRTUAL_SWITCH_ID =  AG_F_PORT.VIRTUAL_SWITCH_ID
    and NPORT_WWN_MAP.DEVICE_PORT_WWN = AG_F_PORT.REMOTE_PORT_WWN
    AND AG_N_PORT.VIRTUAL_SWITCH_ID = AG_SWITCH.ID
    and AG_SWITCH.MONITORED = 1;
```

NP_FLOW_DEFINITION_INFO

```
create or replace view NP_FLOW_DEFINITION_INFO as
 select NFP.ID as FLOW_ID,
    NFP.NAME as FLOW_NAME,
    NFP.VIRTUAL_SWITCH_ID as VS_ID,
    NFP.SRCDEV,
    NFP.DSTDEV,
    NFP.SRCPORT,
    NFP.DSTPORT,
    NFP.BIDIR,
    NFP.SFID,
    NFP.DFID,
    NFP.SRCDOMAIN,
    NFP.DSTDOMAIN,
    NFP.LUNID,
    NFP.OXID,
    NFP.QOS,
    NFP.OPTION,
```

```
 NFP.SCSICMD,
 NFP.TYPE,
 NFP.RCTL,
 NFP.PROTOCOL_TYPE,
 NFP.FRAME_OFFSET,
 NFP.SIZE,
 NFP.PATTERN,
 NFP.LAST_UPDATED_TIME,
 NFP.MONITOR_FEATURE,
 NFP.GENERATOR_FEATURE,
 NFP.MIRROR_FEATURE,
 NFP.IS_PREDEFINED,
 NFP.MIRROR_PORT,
 NFP.REMOTE_MIRROR_PORT,
 NFP.MAPS_PORT_GROUP,
 NFP.SRC_SWITCH_PORT_ID,
 NFP.DST_SWITCH_PORT_ID,
 NFP.SRC_DEVICE_PORT_ID,
 NFP.DST_DEVICE_PORT_ID,
 NFP.SRC_HBA_ID,
 NFP.DST_HBA_ID,
 NFP.SRC_VM_ID,
 NFP.DST_VM_ID,
 NFP.MIRROR_SWITCH_PORT_ID,
 NFP.SRC_DEVICE_NODE_ID,
 NFP.DST_DEVICE_NODE_ID,
 CS.IP_ADDRESS,
 VS.NAME as VS_NAME,
 VS.WWN as VS_WWN,
VS.DOMAIN_ID as VS_DOMAIN_ID,
 F.ID as FABRIC_ID,
 VS.VIRTUAL_FABRIC_ID,
 SRC_SP.WWN as SRC_SP_WWN,
 SRC_SP.NAME as SRC_SP_NAME,
 (SRC_SP.SLOT_NUMBER || '/'::TEXT) || SRC_SP.PORT_NUMBER as SRC_SP_PORT_NUMBER,
 SRC_SP.USER_PORT_NUMBER as SRC_SP_USER_PORT_NUMBER,
 SRC_SP.STATUS as SRC_SP_STATUS,
 SRC_SP.PORT_ID as SRC_SP_FCADDRESS,
 DST_SP.WWN as DST_SP_WWN,
 DST_SP.NAME as DST_SP_NAME,
 (DST_SP.SLOT_NUMBER || '/'::TEXT) || DST_SP.PORT_NUMBER as DST_SP_PORT_NUMBER,
 DST_SP.USER_PORT_NUMBER as DST_SP_USER_PORT_NUMBER,
 DST_SP.STATUS as DST_SP_STATUS,
 DST_SP.PORT_ID as DST_SP_FCADDRESS,
 MIR_SP.WWN as MIR_SP_WWN,
 MIR_SP.NAME as MIR_SP_NAME,
 (MIR_SP.SLOT_NUMBER || '/'::TEXT) || MIR_SP.PORT_NUMBER as MIR_SP_PORT_NUMBER,
 MIR_SP.USER_PORT_NUMBER as MIR_SP_USER_PORT_NUMBER,
 MIR_SP.PORT_ID as MIR_SP_FCADDRESS,
 SRC_DPI.WWN as SRC_DP_WWN,
 SRC_DPI.NPV_PHYSICAL as SRC_NPV_PHYSICAL,
 SRC_DPI.NUMBER as SRC_DP_NUMBER,
 SRC_DPI.PORT_ID as SRC_FC_ADDRESS,
 SRC_DNI.TYPE as SRC_NODE_TYPE,
 SRC_DNI.WWN as SRC_DN_WWN,
 DST_DPI.WWN as DST_DP_WWN,
 DST_DPI.NPV_PHYSICAL as DST_NPV_PHYSICAL,
 DST_DPI.NUMBER as DST_DP_NUMBER,
 DST_DPI.PORT_ID as DST_FC_ADDRESS,
 DST_DNI.TYPE as DST_NODE_TYPE,
 DST_DNI.WWN as DST_DN_WWN,
 NFP.VM_ENTITY_ID,
 VM.NAME as VM_NAME,
```

```
 APP_DTL.APPLICATIONS_NAME
from NP_FLOW_DEFINITION NFP
  join VIRTUAL_SWITCH VS on NFP.VIRTUAL_SWITCH_ID = VS.ID
  join FABRIC_MEMBER FM on VS.ID = FM.VIRTUAL_SWITCH_ID
 join FABRIC F on F.ID = FM.FABRIC_ID
  join CORE_SWITCH CS on CS.ID = VS.CORE_SWITCH_ID
  left join SWITCH_PORT SRC_SP on SRC_SP.ID = NFP.SRC_SWITCH_PORT_ID
  left join SWITCH_PORT DST_SP on DST_SP.ID = NFP.DST_SWITCH_PORT_ID
  left join SWITCH_PORT MIR_SP on MIR_SP.ID = NFP.MIRROR_SWITCH_PORT_ID
  left join DEVICE_PORT SRC_DPI on SRC_DPI.ID = NFP.SRC_DEVICE_PORT_ID
  left join DEVICE_NODE SRC_DNI on SRC_DNI.ID = NFP.SRC_DEVICE_NODE_ID
  left join DEVICE_PORT DST_DPI on DST_DPI.ID = NFP.DST_DEVICE_PORT_ID
  left join DEVICE_NODE DST_DNI on DST_DNI.ID = NFP.DST_DEVICE_NODE_ID
  left join VM_VIRTUAL_MACHINE VM on VM.INSTANCE_UUID = NFP.VM_ENTITY_ID
  left join VM_APPLICATION_DETAILS APP_DTL on APP_DTL.VM_INSTANCE_UUID = NFP.VM_ENTITY_ID;
```

NP_SUB_FLOW_INFO

```
create or replace view NP_SUB_FLOW_INFO as
 select NSP.ID as SUB_FLOW_ID,
    NSP.FLOW_DEFINITION_ID,
    NSP.FEATURE,
    NSP.SRCDEV,
    NSP.DSTDEV,
    NSP.SRCPORT,
    NSP.DSTPORT,
    NSP.BIDIR,
    NSP.SFID,
    NSP.DFID,
    NSP.SRCDOMAIN,
    NSP.DSTDOMAIN,
    NSP.LUNID,
    NSP.LAST_UPDATED_TIME,
    NSP.IS_MISSING,
    NSP.OXID,
    NSP.RXID,
    NSP.CS_CTL,
    NSP.SIZE,
    NSP.PATTERN,
    NSP.SUB_FLOW_MD5HASH,
    NSP.MIRROR_PORT,
    NSP.KEY as SUB_FLOW_KEY,
    NSP.SRC_VIRTUAL_SWITCH_ID,
    NSP.SRC_SWITCH_PORT_ID,
    NSP.DST_VIRTUAL_SWITCH_ID,
    NSP.DST_SWITCH_PORT_ID,
    NSP.SRC_DEVICE_PORT_ID,
    NSP.DST_DEVICE_PORT_ID,
    NSP.SUB_FLOW_ORIGIN,
    NSP.SRC_HBA_ID,
    NSP.DST_HBA_ID,
    NSP.SRC_VM_ID,
    NSP.DST_VM_ID,
    NSP.MIRROR_SWITCH_PORT_ID,
    NSP.SRC_DEVICE_NODE_ID,
    NSP.DST_DEVICE_NODE_ID,
    NFD.NAME as FLOW_NAME,
    NFD.OPTION,
    NFD.SCSICMD,
    NFD.SRCDEV as FD_SRCDEV,
    NFD.DSTDEV as FD_DSTDEV,
    NFD.SRCPORT as FD_SRCPORT,
```

```
 NFD.DSTPORT as FD_DSTPORT,
 CS.IP_ADDRESS,
 VS.ID as VS_ID,
 VS.NAME as VS_NAME,
 VS.WWN as VS_WWN,
 VS.DOMAIN_ID as VS_DOMAIN_ID,
 F.ID as FABRIC_ID,
 VS.VIRTUAL_FABRIC_ID,
SRC_SP.WWN as SRC_SP_WWN,
 SRC_SP.NAME as SRC_SP_NAME,
 (SRC_SP.SLOT_NUMBER || '/'::TEXT) || SRC_SP.PORT_NUMBER as SRC_SP_PORT_NUMBER,
 SRC_SP.USER_PORT_NUMBER as SRC_SP_USER_PORT_NUMBER,
 SRC_SP.STATUS as SRC_SP_STATUS,
 SRC_SP.PORT_ID as SRC_SP_FCADDRESS,
 DST_SP.WWN as DST_SP_WWN,
 DST_SP.NAME as DST_SP_NAME,
 (DST_SP.SLOT_NUMBER || '/'::TEXT) || DST_SP.PORT_NUMBER as DST_SP_PORT_NUMBER,
 DST_SP.USER_PORT_NUMBER as DST_SP_USER_PORT_NUMBER,
 DST_SP.STATUS as DST_SP_STATUS,
 DST_SP.PORT_ID as DST_SP_FCADDRESS,
 MIR_SP.ID as MIR_SP_ID,
 MIR_SP.WWN as MIR_SP_WWN,
 MIR_SP.NAME as MIR_SP_NAME,
 (MIR_SP.SLOT_NUMBER || '/'::TEXT) || MIR_SP.PORT_NUMBER as MIR_SP_PORT_NUMBER,
 MIR_SP.USER_PORT_NUMBER as MIR_SP_USER_PORT_NUMBER,
 MIR_SP.PORT_ID as MIR_SP_FCADDRESS,
 SRC_DPI.WWN as SRC_DP_WWN,
 SRC_DPI.NPV_PHYSICAL as SRC_NPV_PHYSICAL,
 SRC_DPI.NUMBER as SRC_DP_NUMBER,
 SRC_DPI.PORT_ID as SRC_FC_ADDRESS,
 SRC_DNI.TYPE as SRC_NODE_TYPE,
 SRC_DNI.WWN as SRC_DN_WWN,
 DST_DPI.WWN as DST_DP_WWN,
 DST_DPI.NPV_PHYSICAL as DST_NPV_PHYSICAL,
 DST_DPI.NUMBER as DST_DP_NUMBER,
 DST_DPI.PORT_ID as DST_FC_ADDRESS,
 DST_DNI.TYPE as DST_NODE_TYPE,
 DST_DNI.WWN as DST_DN_WWN,
 USR_DFD_DEV_SRC_DET.NAME as SRC_DP_NAME,
 USR_DFD_DEV_DST_DET.NAME as DST_DP_NAME,
 NSP.VM_ENTITY_ID,
 VM.NAME as VM_NAME,
 APP_DTL.APPLICATIONS_NAME
from NP_SUB_FLOW NSP
  join NP_FLOW_DEFINITION NFD on NSP.FLOW_DEFINITION_ID = NFD.ID
  join VIRTUAL_SWITCH VS on NFD.VIRTUAL_SWITCH_ID = VS.ID
  join FABRIC_MEMBER FM on VS.ID = FM.VIRTUAL_SWITCH_ID
  join FABRIC F on F.ID = FM.FABRIC_ID
  join CORE_SWITCH CS on CS.ID = VS.CORE_SWITCH_ID
  left join SWITCH_PORT SRC_SP on SRC_SP.ID = NSP.SRC_SWITCH_PORT_ID
  left join SWITCH_PORT DST_SP on DST_SP.ID = NSP.DST_SWITCH_PORT_ID
  left join SWITCH_PORT MIR_SP on MIR_SP.ID = NSP.MIRROR_SWITCH_PORT_ID
  left join DEVICE_PORT SRC_DPI on SRC_DPI.ID = NSP.SRC_DEVICE_PORT_ID
  left join DEVICE_NODE SRC_DNI on SRC_DNI.ID = NSP.SRC_DEVICE_NODE_ID
  left join DEVICE_PORT DST_DPI on DST_DPI.ID = NSP.DST_DEVICE_PORT_ID
  left join DEVICE_NODE DST_DNI on DST_DNI.ID = NSP.DST_DEVICE_NODE_ID
  left join USER_DEFINED_DEVICE_DETAIL USR_DFD_DEV_SRC_DET on USR_DFD_DEV_SRC_DET.WWN = SRC_DPI.WWN
  left join USER_DEFINED_DEVICE_DETAIL USR_DFD_DEV_DST_DET on USR_DFD_DEV_DST_DET.WWN = DST_DPI.WWN
  left join VM_VIRTUAL_MACHINE VM on VM.INSTANCE_UUID = NSP.VM_ENTITY_ID
  left join VM_APPLICATION_DETAILS APP_DTL on APP_DTL.VM_INSTANCE_UUID = NSP.VM_ENTITY_ID;
```

PHANTOM_PORT_INFO

```
create or replace view PHANTOM_PORT_INFO as
select
    PHANTOM_PORT.ID,
    PHANTOM_PORT.WWN,
    PHANTOM_PORT.VIRTUAL_SWITCH_ID,
    PHANTOM_PORT.PORT_NUMBER,
    PHANTOM_PORT.PORT_ID,
    PHANTOM_PORT.SPEED,
    PHANTOM_PORT.MAX_SPEED,
    PHANTOM_PORT.TYPE,
    PHANTOM_PORT.REMOTE_NODE_WWN,
    PHANTOM_PORT.REMOTE_PORT_WWN,
    PHANTOM_PORT.PHANTOM_TYPE,
    PHANTOM_PORT.BB_FABRIC_ID,
    VIRTUAL_SWITCH.WWN as VIRTUAL_SWITCH_WWN,
    VIRTUAL_SWITCH.MANAGEMENT_STATE,
    VIRTUAL_SWITCH.MONITORED
from
    PHANTOM_PORT,
    VIRTUAL_SWITCH
where
    PHANTOM_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID;
```

PRODUCT_INFO

```
CREATE VIEW product_info AS
    select distinct
TEMP_DEVICE.DEVICE_ID,
TEMP_DEVICE.MANAGED_ELEMENT_ID,
TEMP_DEVICE.ALIAS_NAME,
TEMP_DEVICE.HOST_NAME,
TEMP_DEVICE.OPER_STATUS,
case
    when TEMP_DEVICE.OPER_STATUS = 1 then
    (case
        when TEMP_DEVICE.FABRIC_WATCH_STATUS = 2 then 'DEGRADED'
        when TEMP_DEVICE.FABRIC_WATCH_STATUS = 3 then 'DOWN'
        else 'REACHABLE'
     end)
    when TEMP_DEVICE.OPER_STATUS = 2 then 'NOT REACHABLE'
    when TEMP_DEVICE.OPER_STATUS = 3 then 'DEGRADED'
    when TEMP_DEVICE.OPER_STATUS = 4 then 'MARGINAL'
    when TEMP_DEVICE.OPER_STATUS = 5 then 'DOWN'
    else 'UNKNOWN'
end as OPER_STATUS_TXT,
TEMP_DEVICE.FABRIC_WATCH_STATUS,
TEMP_DEVICE.FABRIC_WATCH_STATUS_REASON,
TEMP_DEVICE.ADMIN_STATUS,
case
    when TEMP_DEVICE.ADMIN_STATUS = 1 then 'TROUBLESHOOTING'
    else 'NORMAL'
end as ADMIN_STATUS_TXT,
TEMP_DEVICE.ADMIN_STATUS_LAST_UPDATED,
TEMP_DEVICE.MEMO,
TEMP_DEVICE.MEMO_LAST_UPDATED,
TEMP_DEVICE.SYS_OID,
TEMP_DEVICE.RBRIDGE_ID,
TEMP_DEVICE.IP_ADDRESS,
TEMP_FOUNDRY_DEVICE.PRODUCT_TYPE,
case
```

```
    when TEMP_DEVICE.IS_ROUTER = 1 then 'ROUTER'
    else 'L2 SWITCH'
end as PRODUCT_TYPE_TXT,
case
    when TEMP_DEVICE.IS_FOUNDRY = 1 then 'IOS'
    when TEMP_DEVICE.IS_DCB_SWITCH = 1 then 'FOS'
    when TEMP_DEVICE.IS_VCS_CAPABLE = 1 then 'NOS'
    else 'UNKNOWN'
end as SWITCH_OS,
TEMP_DEVICE.IS_ROUTER,
TEMP_DEVICE.IS_SLB,
TEMP_DEVICE.SERIAL_NUMBER,
TEMP_DEVICE.SYS_NAME,
case
    when TEMP_DEVICE.SUB_CATEGORY > 0 then (select distinct VCSD.SYS_NAME from DEVICE as VCSD where
VCSD.MANAGED_ELEMENT_ID
    in (select distinct VM.CLUSTER_ME_ID from VCS_CLUSTER_MEMBER as VM where TEMP_DEVICE.MANAGED_ELEMENT_ID =
VM.MEMBER_ME_ID))
    else null
end as VCS_NAME,

case
    when TEMP_DEVICE.SUB_CATEGORY > 0 then (select distinct VCSD.IP_ADDRESS from DEVICE as VCSD where
VCSD.MANAGED_ELEMENT_ID
    in (select distinct VM.CLUSTER_ME_ID from VCS_CLUSTER_MEMBER as VM where TEMP_DEVICE.MANAGED_ELEMENT_ID =
VM.MEMBER_ME_ID))
    else null
end as VCS_IP_ADDRESS,
TEMP_DEVICE.SYS_CONTACT,
TEMP_DEVICE.SYS_LOCATION,
TEMP_DEVICE.DESCRIPTION,
TEMP_DEVICE.LAST_SEEN_TIME,
TO_TIMESTAMP(TEMP_DEVICE.LAST_SEEN_TIME,'YYYYMMDDHH24MISS') as LAST_SEEN_TIMESTAMP,
TEMP_DEVICE.VendOR,
TEMP_DEVICE.CATEGORY,
case
    when TEMP_DEVICE.CATEGORY = 1 then 'FIXED CONFIGURATION'
    when TEMP_DEVICE.CATEGORY = 2 then 'CHASSIS'
    when TEMP_DEVICE.CATEGORY = 3 then 'STACK'
    when TEMP_DEVICE.CATEGORY = 4 then 'ACCESS POINT'
    when TEMP_DEVICE.CATEGORY = 5 then 'WIRELESS CONTROLLER'
    else 'UNKNOWN'
end as CATEGORY_TXT,
TEMP_DEVICE.SUB_CATEGORY,
case
    when TEMP_DEVICE.SUB_CATEGORY = 1 then 'DCB 8000'
    when TEMP_DEVICE.SUB_CATEGORY = 2 then 'DCB 8470'
    when TEMP_DEVICE.SUB_CATEGORY = 3 then 'DCB M8428'
    when TEMP_DEVICE.SUB_CATEGORY = 4 then 'DCX'
    when TEMP_DEVICE.SUB_CATEGORY = 5 then 'DCX-4S'
    when TEMP_DEVICE.SUB_CATEGORY = 6 then 'VCS/VDX'
    when TEMP_DEVICE.SUB_CATEGORY = 7 then 'VDX 6720-24'
    when TEMP_DEVICE.SUB_CATEGORY = 8 then 'VDX 6720-60'
    when TEMP_DEVICE.SUB_CATEGORY = 9 then 'VDX 6710'
    when TEMP_DEVICE.SUB_CATEGORY = 10 then 'VDX 6730-24'
    when TEMP_DEVICE.SUB_CATEGORY = 11 then 'VDX 6730-60'
    when TEMP_DEVICE.SUB_CATEGORY = 12 then 'VDX 8770-4'
    when TEMP_DEVICE.SUB_CATEGORY = 13 then 'VDX 8770-8'
    when TEMP_DEVICE.SUB_CATEGORY = 14 then 'VDX 8770-16'
    when TEMP_DEVICE.SUB_CATEGORY = 15 then 'VDX 2730'
    else 'IP DEVICE'
end as SUB_CATEGORY_TXT,
TEMP_DEVICE.FIRST_SEEN_TIME,
```

```
TO_TIMESTAMP(TEMP_DEVICE.FIRST_SEEN_TIME,'YYYYMMDDHH24MISS') as FIRST_SEEN_TIMESTAMP,
TEMP_DEVICE.PORT_COUNT,
TEMP_DEVICE.LICENSE_PORT_COUNT,
case
    when TEMP_DEVICE.SUB_CATEGORY = 0 then (select distinct SWITCH_MODEL.MODEL from SWITCH_MODEL where
TEMP_DEVICE.SYS_OID = SWITCH_MODEL.SYS_OID)
    else TEMP_DEVICE.BRIEF_PRODUCT_FAMILY
end as MODEL,
TEMP_FOUNDRY_DEVICE.IMAGE_VERSION as FIRMWARE,
TEMP_PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID,
TEMP_PHYSICAL_DEVICE.NUM_SLOTS,
TEMP_PHYSICAL_DEVICE.UNIT_NUMBER,
TEMP_DEVICE.USER_DEFINED_VALUE_1,
TEMP_DEVICE.USER_DEFINED_VALUE_2,
TEMP_DEVICE.USER_DEFINED_VALUE_3
from DEVICE as TEMP_DEVICE
left join
(
    select
    FOUNDRY_DEVICE.DEVICE_ID,
    FOUNDRY_DEVICE.PRODUCT_TYPE,
    FOUNDRY_DEVICE.IMAGE_VERSION
    from FOUNDRY_DEVICE
) TEMP_FOUNDRY_DEVICE on TEMP_DEVICE.DEVICE_ID = TEMP_FOUNDRY_DEVICE.DEVICE_ID
left join
(
    select
    PHYSICAL_DEVICE.DEVICE_ID,
    PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID,
    PHYSICAL_DEVICE.NUM_SLOTS,
    PHYSICAL_DEVICE.UNIT_NUMBER
    from PHYSICAL_DEVICE
) TEMP_PHYSICAL_DEVICE on TEMP_DEVICE.DEVICE_ID = TEMP_PHYSICAL_DEVICE.DEVICE_ID;
```

PORT_BOTTLENECK_CONF_INFO

This view provides combine port bottleneck configuration and enough information from switch port for the client to identify the port.

```
create or replace view PORT_BOTTLENECK_CONF_INFO as
select
    PORT_BOTTLENECK_CONFIG.SWITCH_PORT_ID,
    PORT_BOTTLENECK_CONFIG.BOTTLENECK_DETECT_ENABLED,
    PORT_BOTTLENECK_CONFIG.ALERTS_ENABLED,
    PORT_BOTTLENECK_CONFIG.CONGESTION_THRESHOLD,
    PORT_BOTTLENECK_CONFIG.LATENCY_THRESHOLD,
    PORT_BOTTLENECK_CONFIG.WINDOW_,
    PORT_BOTTLENECK_CONFIG.QUIET_TIME,
    PORT_BOTTLENECK_CONFIG.CREATION_TIME,
    PORT_BOTTLENECK_CONFIG.LAST_UPDATE_TIME,
    PORT_BOTTLENECK_CONFIG.LATENCY_SEVERITY,
    PORT_BOTTLENECK_CONFIG.LATENCY_TIME,
    SWITCH_PORT.VIRTUAL_SWITCH_ID,
    SWITCH_PORT.USER_PORT_NUMBER,
    SWITCH_PORT.TYPE,
    SWITCH_PORT.WWN
  from
    PORT_BOTTLENECK_CONFIG
        left outer join SWITCH_PORT
            on PORT_BOTTLENECK_CONFIG.SWITCH_PORT_ID = SWITCH_PORT.ID;

comment on view PORT_BOTTLENECK_CONF_INFO is
Combine port bottleneck configuration and enough info from switch port for the client to identify the port.;
```

PORT_BOTTLENECK_STAT_INFO

This view provides combine port bottleneck status and enough information from the switch port for the client to identify the port.

```
create or replace view PORT_BOTTLENECK_STAT_INFO as
select
    PORT_BOTTLENECK_STATUS.SWITCH_PORT_ID,
    PORT_BOTTLENECK_STATUS.STATUS,
    SWITCH_PORT.VIRTUAL_SWITCH_ID,
    SWITCH_PORT.USER_PORT_NUMBER,
    SWITCH_PORT.TYPE
  from
    PORT_BOTTLENECK_STATUS
        left outer join SWITCH_PORT
            on PORT_BOTTLENECK_STATUS.SWITCH_PORT_ID = SWITCH_PORT.ID;
```

PORT_GROUP_INFO

```
create or replace view PORT_GROUP_INFO as
select
    SWITCH_PORT.ID as PORT_ID,
    SWITCH_PORT.NAME as SWITCH_PORT_NAME,
    SWITCH_PORT.WWN,
    SWITCH_PORT.HEALTH,
    SWITCH_PORT.STATUS,
    SWITCH_PORT.PORT_NUMBER,
    SWITCH_PORT.SLOT_NUMBER,
    SWITCH_PORT.FICON_SUPPORTED,
    SWITCH_PORT.STATE,
    SWITCH_PORT.USER_PORT_NUMBER,
    VIRTUAL_SWITCH.NAME as VIRTUAL_SWITCH_NAME,
    VIRTUAL_SWITCH.ID as SWITCH_ID,
    FABRIC.NAME as FABRIC_NAME,
    FABRIC.MANAGED as FABRIC_MANAGED,
    PORT_GROUP.ID as PORT_GROUP_ID,
    PORT_GROUP_MEMBER.ID as PORT_GROUP_MEMBER_ID
from
    SWITCH_PORT, VIRTUAL_SWITCH, FABRIC, FABRIC_MEMBER, PORT_GROUP_MEMBER, PORT_GROUP
where
    VIRTUAL_SWITCH .ID = SWITCH_PORT.VIRTUAL_SWITCH_ID and
    FABRIC_MEMBER.VIRTUAL_SWITCH_ID = SWITCH_PORT.VIRTUAL_SWITCH_ID and
    FABRIC_MEMBER.FABRIC_ID = FABRIC.ID and
    SWITCH_PORT.ID = PORT_GROUP_MEMBER.SWITCH_PORT_ID and
    PORT_GROUP_MEMBER.PORT_GROUP_ID = PORT_GROUP.ID;
```

ROLE_PRIVILEGE_INFO

```
create or replace view ROLE_PRIVILEGE_INFO as
select
    ROLE.ID,
    ROLE.NAME as ROLE_NAME,
    ROLE.DESCRIPTION as ROLE_DESCRIPTION,
    ROLE.HIDDEN as ROLE_HIDDEN,
    PRIVILEGE.ID as PRIVILEGE_ID,
    PRIVILEGE.NAME as PRIVILEGE_NAME,
    PRIVILEGE.AREA as PRIVILEGE_AREA,
    ROLE_PRIVILEGE_MAP.PERMISSION
from
    ROLE,
    ROLE_PRIVILEGE_MAP,
    PRIVILEGE
```

```
where
    ROLE.ID = ROLE_PRIVILEGE_MAP.ROLE_ID and
    PRIVILEGE.ID = ROLE_PRIVILEGE_MAP.PRIVILEGE_ID;
```

PORT_PROFILE_INFO

```
create or replace view PORT_PROFILE_INFO as
select
    PORT_PROFILE.ID,
    PORT_PROFILE.SWITCH_ME_ID,
    PORT_PROFILE.NAME,
    PORT_PROFILE.STATE,
    PORT_PROFILE.SWITCH_PORT_MODE,
    PORT_PROFILE.ACL_PROFILE,
    PORT_PROFILE.QOS_PROFILE,
    PORT_PROFILE.FCOE_PROFILE,
    PORT_PROFILE.VLAN_PROFILE,
    PORT_PROFILE.VLAN_DETAILS,
    PORT_PROFILE.DEFAULT_PROFILE,
    PORT_PROFILE.ACL_NAME,
    PORT_PROFILE.FCOE_MAP_NAME,
        PORT_PROFILE.ACTIVATED,
    PORT_PROFILE_QOS_MAP.DCB_MODE,
    PORT_PROFILE_QOS_MAP.ETHERNET_MODE,
    PORT_PROFILE_QOS_MAP.PAUSE_TX,
    PORT_PROFILE_QOS_MAP.PAUSE_RX,
    PORT_PROFILE_QOS_MAP.COS_COS,
    PORT_PROFILE_QOS_MAP.TRAFFIC_CLASS,
        PORT_PROFILE_QOS_MAP.COS,
        PORT_PROFILE_QOS_MAP.CEE_MAP,
    PORT_PROFILE_QOS_PFC_MAP.COS0_TX,
    PORT_PROFILE_QOS_PFC_MAP.COS0_RX,
    PORT_PROFILE_QOS_PFC_MAP.COS1_TX,
    PORT_PROFILE_QOS_PFC_MAP.COS1_RX,
    PORT_PROFILE_QOS_PFC_MAP.COS2_TX,
    PORT_PROFILE_QOS_PFC_MAP.COS2_RX,
    PORT_PROFILE_QOS_PFC_MAP.COS3_TX,
    PORT_PROFILE_QOS_PFC_MAP.COS3_RX,
    PORT_PROFILE_QOS_PFC_MAP.COS4_TX,
    PORT_PROFILE_QOS_PFC_MAP.COS4_RX,
    PORT_PROFILE_QOS_PFC_MAP.COS5_TX,
    PORT_PROFILE_QOS_PFC_MAP.COS5_RX,
    PORT_PROFILE_QOS_PFC_MAP.COS6_TX,
    PORT_PROFILE_QOS_PFC_MAP.COS6_RX,
    PORT_PROFILE_QOS_PFC_MAP.COS7_TX,
    PORT_PROFILE_QOS_PFC_MAP.COS7_RX,
        PORT_PROFILE_QOS_MAP.TRUST_COS
from PORT_PROFILE
    left join PORT_PROFILE_QOS_MAP
        on PORT_PROFILE.ID = PORT_PROFILE_QOS_MAP.PROFILE_ID
    left join PORT_PROFILE_QOS_PFC_MAP
        on PORT_PROFILE.ID = PORT_PROFILE_QOS_PFC_MAP.PROFILE_ID;
```

PORT_PROFILE_INTERFACE_INFO

```
create or replace view PORT_PROFILE_INTERFACE_INFO as
select
    PORT_PROFILE.ID,
    PORT_PROFILE.SWITCH_ME_ID,
    PORT_PROFILE.NAME,
    PORT_PROFILE.ACL_PROFILE,
```

```
    PORT_PROFILE.QOS_PROFILE,
    PORT_PROFILE.FCOE_PROFILE,
    PORT_PROFILE.VLAN_PROFILE,
    PORT_PROFILE.VLAN_DETAILS,
    PORT_PROFILE.DEFAULT_PROFILE,
    PORT_PROFILE.ACL_NAME,
    PORT_PROFILE.FCOE_MAP_NAME,
    PORT_PROFILE_INTERFACE_MAP.INTERFACE_ID,
    PORT_PROFILE_INTERFACE_MAP.SWITCH_PORT_ID
from
    PORT_PROFILE,
    PORT_PROFILE_INTERFACE_MAP
where
    PORT_PROFILE.ID= PORT_PROFILE_INTERFACE_MAP.PROFILE_ID;
```

PORT_PROFILE_MAC_INFO

```
create or replace view PORT_PROFILE_MAC_INFO as
select
PORT_PROFILE_MAC_MAP.PROFILE_ID,
PORT_PROFILE_MAC_MAP.MAC,
PORT_PROFILE_MAC_MAP.NAME as MAC_NAME,
VM_VIRTUAL_ETHERNET_ADAPTER.PORT_GROUP_NAME,
VM_VIRTUAL_ETHERNET_ADAPTER.DISPLAY_LABEL,
VM_VIRTUAL_MACHINE.NAME as VM_NAME,
VM_VCENTER_MEMBER.HOST_NAME as HOST_NAME,
VM_VCENTER.NAME as VCENTER_NAME,
INTERFACE.IDENTIFIER
from
PORT_PROFILE_MAC_MAP
left outer join VM_VIRTUAL_ETHERNET_ADAPTER on PORT_PROFILE_MAC_MAP.MAC =
VM_VIRTUAL_ETHERNET_ADAPTER.MAC_ADDRESS
left outer join VM_VIRTUAL_MACHINE on VM_VIRTUAL_ETHERNET_ADAPTER.VIRTUAL_MACHINE_ID = VM_VIRTUAL_MACHINE.ID
left outer join VM_VCENTER_MEMBER on VM_VIRTUAL_MACHINE.HOST_ID = VM_VCENTER_MEMBER.VM_HOST_ID
left outer join VM_VCENTER on VM_VCENTER_MEMBER.VM_VCENTER_ID = VM_VCENTER.ID
left outer join L2_NEIGHBOR on PORT_PROFILE_MAC_MAP.MAC = encode(L2_NEIGHBOR.LLDP_REM_PORT_ID, 'base64')
left outer join INTERFACE on L2_NEIGHBOR.INTERFACE_ID = INTERFACE.INTERFACE_ID;
```

PORT_VLAN_INFO

```
create view PORT_VLAN_INFO as
select
    PV.*,
    DEVICE_ID,
    NAME,
    TABLE_SUBTYPE
from
    VLAN V,
    PORT_VLAN PV
where
    V.VLAN_DB_ID = PV.VLAN_DB_ID;
```

PROTOCOL_VLAN_INFO

```
create or replace view PROTOCOL_VLAN_INFO as
select
        V.*,
        port_vlan_db_id,
        is_dynamic,
        protocol
```

```
    from vlan V, sub_port_vlan SPV, protocol_vlan PV
    where V.vlan_db_id = SPV.vlan_db_id AND SPV.vlan_db_id = PV.vlan_db_id;
```

SFLOW

```
create or replace view SFLOW as
  select DEVICE_ID, IN_SLOT, IN_PORT, OUT_SLOT, OUT_PORT, L4_PROTOCOL, IP_TOS, SRC_SUBNET_BITS,
DEST_SUBNET_BITS, IN_PRIORITY, OUT_PRIORITY, IN_VLAN, OUT_VLAN, L3_PROTOCOL, L4_SRC_PORT, L4_DEST_PORT,
TIME_IN_SECONDS, SRC_MAC, DEST_MAC, L3_SRC_ADDR, L3_DEST_ADDR, TCP_FLAGS, LOCAL_AS, SRC_AS, SRC_PEER_AS,
SFLOW_IP_ROUTE_INFO_ID, IP_FLOW_LABEL, SRC_USER, DEST_USER, FRAMES, BYTES, IN_UNIT, OUT_UNIT
  from SFLOW_HOUR_SUMMARY
  where SLNUM <= (select MAX_SLNUM from SFLOW_HOUR_SUMMARY_SLNUM fetch first 1 rows only)
  union all
  select DEVICE_ID, IN_SLOT, IN_PORT, OUT_SLOT, OUT_PORT, L4_PROTOCOL, IP_TOS, SRC_SUBNET_BITS,
DEST_SUBNET_BITS, IN_PRIORITY, OUT_PRIORITY, IN_VLAN, OUT_VLAN, L3_PROTOCOL, L4_SRC_PORT, L4_DEST_PORT,
TIME_IN_SECONDS, SRC_MAC, DEST_MAC, L3_SRC_ADDR, L3_DEST_ADDR, TCP_FLAGS, LOCAL_AS, SRC_AS, SRC_PEER_AS,
SFLOW_IP_ROUTE_INFO_ID, IP_FLOW_LABEL, SRC_USER, DEST_USER, FRAMES, BYTES, IN_UNIT, OUT_UNIT
  from SFLOW_STAGING
  where SLNUM >= (select MIN_SLNUM from SFLOW_STAGING_SLNUM fetch first 1 rows only);
```

SFLOW_MINUTE_L3_VIEW

```
create or replace view SFLOW_MINUTE_L3_VIEW as
  select DEVICE_ID, TIME_IN_SECONDS, L3_SRC_ADDR, L3_DEST_ADDR, L3_PROTOCOL, L4_PROTOCOL, TCP_FLAGS, IN_VLAN,
OUT_VLAN, FRAMES, BYTES
  from SFLOW_MINUTE_L3
  where SLNUM <= (select MAX_SLNUM from SFLOW_MINUTE_L3_SLNUM fetch first 1 rows only)
  union all
  select DEVICE_ID, TIME_IN_SECONDS, L3_SRC_ADDR, L3_DEST_ADDR, L3_PROTOCOL, L4_PROTOCOL, TCP_FLAGS, IN_VLAN,
OUT_VLAN, FRAMES, BYTES
  from SFLOW_STAGING
  where SLNUM >= (select MIN_SLNUM from SFLOW_STAGING_SLNUM fetch first 1 rows only);
```

SFLOW_MINUTE_MAC_VIEW

```
create or replace view SFLOW_MINUTE_MAC_VIEW as
  select DEVICE_ID, TIME_IN_SECONDS, SRC_MAC, DEST_MAC, IN_VLAN, OUT_VLAN, FRAMES, BYTES
  from SFLOW_MINUTE_MAC
  where SLNUM <= (select MAX_SLNUM from SFLOW_MINUTE_MAC_SLNUM fetch first 1 rows only)
  union all
  select DEVICE_ID, TIME_IN_SECONDS, SRC_MAC, DEST_MAC, IN_VLAN, OUT_VLAN, FRAMES, BYTES
  from SFLOW_STAGING
  where SLNUM >= (select MIN_SLNUM from SFLOW_STAGING_SLNUM fetch first 1 rows only);
```

SCOM_EE_MONITOR_INFO

This view provides combined ee_monitor, ee_monitor_stats, device_port and device_node tables to get the EE Monitor information for SCOM plug-in.

```
create or replace view SCOM_EE_MONITOR_INFO as
select distinct
    EE_MONITOR.NAME,
    EE_MONITOR.SWITCH_PORT_ID,
    EE_MONITOR.SOURCE_PORT_ID,
    EE_MONITOR.DEST_PORT_ID,
    EE_MONITOR_STATS.TX,
    EE_MONITOR_STATS.RX,
```

```
    EE_MONITOR_STATS.CRCERRORS,
    EE_MONITOR_STATS.CREATION_TIME,
    SOURCE_PORT.PORT_ID as SID,
    DEST_PORT.PORT_ID as DID,
    SOURCE_NODE.WWN as SOURCE_DEVICE_WWN,
    SOURCE_PORT.WWN as SOURCE_PORT_WWN,
    DEST_NODE.WWN as DEST_DEVICE_WWN,
    DEST_PORT.WWN as DEST_PORT_WWN,
    SOURCE_NODE.FABRIC_ID as SOURCE_FABRIC_ID,
    DEST_NODE.FABRIC_ID as DEST_FABRIC_ID,
    SOURCE_PORT.DOMAIN_ID as SOURCE_SWITCH_DOMAIN_ID,
    DEST_PORT.DOMAIN_ID as DEST_SWITCH_DOMAIN_ID
from
    DEVICE_PORT as SOURCE_PORT,
    DEVICE_PORT as DEST_PORT,
    DEVICE_NODE as DEST_NODE,
    DEVICE_NODE as SOURCE_NODE,
    EE_MONITOR,
    EE_MONITOR_STATS
where
    SOURCE_PORT.ID = EE_MONITOR.SOURCE_PORT_ID
    and EE_MONITOR.ID = EE_MONITOR_STATS.EE_MONITOR_ID
    and SOURCE_PORT.NODE_ID = SOURCE_NODE.ID
    and DEST_PORT.ID = EE_MONITOR.DEST_PORT_ID
    and DEST_PORT.NODE_ID = DEST_NODE.ID
    and EE_MONITOR_STATS.CREATION_TIME in (
        select MAX(CREATION_TIME)
        from EE_MONITOR_STATS
        group by EE_MONITOR_ID);
```

SENSOR_INFO

```
create or replace view SENSOR_INFO as
select
    SENSOR.ID,
    SENSOR.CORE_SWITCH_ID,
    SENSOR.SENSOR_ID,
    SENSOR.CURRENT_READING,
    SENSOR.TYPE,
    SENSOR.SUB_TYPE,
    SENSOR.DESCRIPTION,
    SENSOR.STATUS,
    SENSOR.OPERATIONAL_STATUS,
    SENSOR.PART_NUMBER,
    SENSOR.SERIAL_NUMBER,
    SENSOR.VERSION,
    SENSOR.CREATION_TIME,
    SENSOR.LAST_UPDATE_TIME,
    SENSOR.FRU_TYPE,
    SENSOR.UNIT_NUMBER,
    SENSOR.STATE,
    CORE_SWITCH.WWN as PHYSICAL_SWITCH_WWN,
    VIRTUAL_SWITCH.SWITCH_MODE as VIRTUAL_SWITCH_MODE,
    VIRTUAL_SWITCH.MANAGEMENT_STATE,
    VIRTUAL_SWITCH.MONITORED
from
    SENSOR,
    CORE_SWITCH,
    VIRTUAL_SWITCH
where
    SENSOR.CORE_SWITCH_ID = CORE_SWITCH.ID and
    SENSOR.CORE_SWITCH_ID = VIRTUAL_SWITCH.CORE_SWITCH_ID;
```

SMART_CARD_USAGE_INFO

```
create or replace view SMART_CARD_USAGE_INFO as
select
    SC.ID SMART_CARD_ID,
    SC.CARD_TYPE,
    SC.CARD_INFO,
    SC.CARDCN_ID,
    SC.FIRST_NAME,
    SC.LAST_NAME,
    SC.NOTES,
    SC.CREATION_TIME,
    -1 ENGINE_ID,
    EG.ID ENCRYPTION_GROUP_ID,
    EG.NAME GROUP_NAME,
    -1 CARD_POSITION,
    -1 CRYPTO_SWITCH_ID,
    -1 SLOT_NUMBER
from
    SMART_CARD SC,
    ENCRYPTION_GROUP EG,
    QUORUM_CARD_GROUP_MAPPING QCGM
where
    QCGM.SMART_CARD_ID = SC.ID
    and EG.ID = QCGM.ENCRYPTION_GROUP_ID
    and SC.CARD_TYPE = 0
union
select
    SC.ID SMART_CARD_ID,
    SC.CARD_TYPE,
    SC.CARD_INFO,
    SC.CARDCN_ID,
    SC.FIRST_NAME,
    SC.LAST_NAME,
    SC.NOTES,
    SC.CREATION_TIME,
    -1 ENGINE_ID,
    EG.ID ENCRYPTION_GROUP_ID,
    EG.NAME GROUP_NAME,
    RCGM.POSITION_ CARD_POSITION,
    -1 CRYPTO_SWITCH_ID,
    -1 SLOT_NUMBER
from
    SMART_CARD SC,
    ENCRYPTION_GROUP EG,
    RECOVERY_CARD_GROUP_MAPPING RCGM
where
    SC.ID = RCGM.SMART_CARD_ID
    and EG.ID = RCGM.ENCRYPTION_GROUP_ID
    and SC.CARD_TYPE = 1
union
select
    SC.ID SMART_CARD_ID,
    SC.CARD_TYPE,
    SC.CARD_INFO,
    SC.CARDCN_ID,
    SC.FIRST_NAME,
    SC.LAST_NAME,
    SC.NOTES,
    SC.CREATION_TIME,
    EE.ID ENGINE_ID,
    -1 ENCRYPTION_GROUP_ID,
    '' GROUP_NAME,
```

```
    -1 CARD_POSITION,
    EE.SWITCH_ID CRYPTO_SWITCH_ID,
    EE.SLOT_NUMBER SLOT_NUMBER
from
    SMART_CARD SC,
    ENCRYPTION_ENGINE EE,
    SYSTEM_CARD_ENGINE_MAPPING SCEM
where
    SC.ID = SCEM.SMART_CARD_ID
    and EE.ID = SCEM.ENCRYPTION_ENGINE_ID
    and SC.CARD_TYPE = 2;
```

SWITCH_CONFIG_INFO

```
create or replace view SWITCH_CONFIG_INFO as
select
    SWITCH_CONFIG.ID,
    SWITCH_CONFIG.NAME,
    SWITCH_CONFIG.SWITCH_ID,
    SWITCH_CONFIG.CORE_SWITCH_ID,
    SWITCH_CONFIG.BACKUP_DATE_TIME,
    SWITCH_CONFIG.CONFIG_DATA,
    SWITCH_CONFIG.CEE_CONFIG_DATA,
    SWITCH_CONFIG.KEEP_COPY,
    SWITCH_CONFIG.CREATED_BY,
    SWITCH_CONFIG.COMMENTS,
    SWITCH_CONFIG.CONFIG_TYPE,
    SWITCH_CONFIG. IS_BASELINE,
    SWITCH_CONFIG. BACKUP_TYPE,
    SWITCH_CONFIG. DRIFT_STATUS,
    SWITCH_CONFIG_DETAIL.IP_ADDRESS,
    SWITCH_CONFIG_DETAIL.WWN,
    SWITCH_CONFIG_DETAIL.PHYSICAL_SWITCH_WWN,
    SWITCH_CONFIG_DETAIL.MODEL_NUMBER as SWITCH_MODEL_NUMBER
from
    SWITCH_CONFIG,
    SWITCH_CONFIG_DETAIL
where
    SWITCH_CONFIG.ID= SWITCH_CONFIG_DETAIL.SWITCH_CONFIG_ID;
```

SWITCH_PORT_DETAILS_INFO

```
create or replace view SWITCH_PORT_DETAILS_REPORT_INFO as
with SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW as (
select SWITCH_PORT.ID as REMOTE_SWITCH_PORT_ID,
SWITCH_PORT.VIRTUAL_SWITCH_ID as REMOTE_VIRTUAL_SWITCH_ID,
SWITCH_PORT.SLOT_NUMBER as REMOTE_PORT_SLOT_NUMBER,
SWITCH_PORT.PORT_NUMBER as REMOTE_PORT_NUMBER,
SWITCH_PORT.USER_PORT_NUMBER as REMOTE_USER_PORT_NUMBER,
SWITCH_PORT.NAME as REMOTE_PORT_NAME,
SWITCH_PORT.SPEED as REMOTE_PORT_SPEED,
SWITCH_PORT.STATUS as REMOTE_PORT_STATUS,
SWITCH_PORT.STATE as REMOTE_PORT_STATE,
SWITCH_PORT.TYPE as REMOTE_PORT_TYPE,
SWITCH_PORT.SPEEDS_SUPPORTED as REMOTE_PORT_SPEEDS_SUPPORTED,
SWITCH_PORT.PHYSICAL_PORT as REMOTE_PHYSICAL_OR_LOGICAL_PORT,
SWITCH_PORT.PORT_INDEX as REMOTE_ZONING_PORT_INDEX,
SWITCH_PORT.PORT_ID as REMOTE_PORT_ID,
SWITCH_PORT.WWN as REMOTE_SWITCH_PORT_WWN,
SWITCH_PORT.AREA_ID as REMOTE_PORT_AREA_ID,
SWITCH_PORT.MAC_ADDRESS as REMOTE_PORT_MAC_ADDRESS,
```

```
SWITCH_PORT.PORT_MOD as REMOTE_PORT_MOD,
SWITCH_PORT.FULL_TYPE as REMOTE_PORT_FULL_TYPE,
SWITCH_PORT.HEALTH as REMOTE_PORT_HEALTH,
SWITCH_PORT.STATUS_MESSAGE as REMOTE_SWITCH_PORT_STATUS_MESSAGE,
SWITCH_PORT.MAX_PORT_SPEED as REMOTE_PORT_MAX_PORT_SPEED,
SWITCH_PORT.LICENSED as REMOTE_PORT_LICENSED,
SWITCH_PORT.REMOTE_NODE_WWN as REMOTE_REMOTE_NODE_WWN,
SWITCH_PORT.REMOTE_PORT_WWN as REMOTE_REMOTE_PORT_WWN,
SWITCH_PORT.TRUNKED as REMOTE_PORT_TRUNKED,
SWITCH_PORT.TRUNK_MASTER as REMOTE_PORT_TRUNK_MASTER,
SWITCH_PORT.FICON_SUPPORTED as REMOTE_PORT_FICON_SUPPORTED,
SWITCH_PORT.BLOCKED as REMOTE_PORT_BLOCKED,
SWITCH_PORT.NPIV as REMOTE_PORT_NPIV,
SWITCH_PORT.NPIV_CAPABLE as REMOTE_PORT_NPIV_CAPABLE,
SWITCH_PORT.NPIV_ENABLED as REMOTE_PORT_NPIV_ENABLED,
SWITCH_PORT.QOS_CAPABLE as REMOTE_PORT_QOS_CAPABLE,
SWITCH_PORT.QOS_ENABLED as REMOTE_PORT_QOS_ENABLED,
SWITCH_PORT.TUNNEL_CONFIGURED as REMOTE_PORT_TUNNEL_CONFIGURED,
SWITCH_PORT.FCR_FABRIC_ID as REMOTE_FCR_FABRIC_ID,
SWITCH_PORT.FCR_INTEROP_MODE as REMOTE_PORT_FCR_INTEROP_MODE,
SWITCH_PORT.USER_DEFINED_VALUE1 as REMOTE_PORT_USER_DEFINED_VALUE1,
SWITCH_PORT.USER_DEFINED_VALUE2 as REMOTE_PORT_USER_DEFINED_VALUE2,
SWITCH_PORT.USER_DEFINED_VALUE3 as REMOTE_PORT_USER_DEFINED_VALUE3,
SWITCH_PORT.KIND as REMOTE_PORT_KIND,
SWITCH_PORT.LAST_UPDATE as REMOTE_PORT_LAST_UPDATE,
VIRTUAL_SWITCH.NAME as REMOTE_SWITCH_NAME,
VIRTUAL_SWITCH.SWITCH_MODE as REMOTE_SWITCH_MODE,
VIRTUAL_SWITCH.DOMAIN_ID as REMOTE_SWITCH_DOMAIN_ID,
VIRTUAL_SWITCH.WWN as REMOTE_VIRTUAL_SWITCH_WWN,
VIRTUAL_SWITCH.OPERATIONAL_STATUS as REMOTE_SWITCH_OPERATIONAL_STATUS,
VIRTUAL_SWITCH.MANAGEMENT_STATE as REMOTE_SWITCH_MANAGEMENT_STATE,
VIRTUAL_SWITCH.STATE as REMOTE_SWITCH_STATE,
VIRTUAL_SWITCH.STATUS as REMOTE_SWITCH_STATUS,
VIRTUAL_SWITCH.STATUS_REASON as REMOTE_SWITCH_STATUS_REASON,
CORE_SWITCH.ID as REMOTE_CORE_SWITCH_ID,
CORE_SWITCH.IP_ADDRESS as REMOTE_CORE_SWITCH_IP_ADDRESS,
CORE_SWITCH.WWN as REMOTE_CORE_SWITCH_WWN,
CORE_SWITCH.NAME as REMOTE_CORE_SWITCH_NAME,
CORE_SWITCH.TYPE as REMOTE_CORE_SWITCH_TYPE,
CORE_SWITCH.MODEL as REMOTE_CORE_SWITCH_MODEL,
CORE_SWITCH.VENDOR as REMOTE_CORE_SWITCH_VENDOR,
CORE_SWITCH.REACHABLE as REMOTE_CORE_SWITCH_REACHABLE,
CORE_SWITCH.OPERATIONAL_STATUS as REMOTE_CORE_SWITCH_OPERATIONAL_STATUS,
CORE_SWITCH_DETAILS.MODEL_NUMBER as REMOTE_CORE_SWITCH_MODEL_NUMBER,
SWITCH_MODEL.DESCRIPTION as REMOTE_SWITCH_MODEL_DESCRIPTION
from SWITCH_PORT
join VIRTUAL_SWITCH on SWITCH_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
join CORE_SWITCH on VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
left join CORE_SWITCH_DETAILS on VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH_DETAILS.CORE_SWITCH_ID
left join SWITCH_MODEL on CORE_SWITCH_DETAILS.TYPE = SWITCH_MODEL.SWBD_TYPE and
coalesce(CORE_SWITCH_DETAILS.SUB_TYPE, '0'::character varying)::integer = SWITCH_MODEL.SUBTYPE
where SWITCH_PORT.LICENSED = 1 and SWITCH_PORT.PHYSICAL_PORT = 1 and SWITCH_PORT.KIND::text<> 'ICL'::text and
((SWITCH_PORT.TYPE::text= ANY (array['E-PORT'::character varying, 'G-PORT'::character varying,
'U-PORT'::character varying, 'F-PORT'::character varying, 'L-PORT'::character varying, 'EX-PORT'::character
varying, 'LB-PORT'::character varying, 'FL-PORT'::character varying, 'SIM-PORT'::character varying,
'N-PORT'::character varying]::TEXT[])) or SWITCH_PORT.TYPE::text~~ 'LB-PORT%'::TEXT) and
VIRTUAL_SWITCH.MONITORED = 1 and (VIRTUAL_SWITCH.SWITCH_MODE = ANY (array[0, 2])) and ((CORE_SWITCH.TYPE <>
all (array[62, 63])) or (SWITCH_PORT.SLOT_NUMBER <> all (array[5, 8])))
),DEVPORT_DEVNODE_VIEW as (
SELECT DEVICE_PORT.TYPE as DEVICE_PORT_TYPE,
DEVICE_PORT.WWN as DEVICE_PORT_WWN, DEVICE_PORT.NPV_PHYSICAL,
DEVICE_PORT.EDGE_SWITCH_PORT_WWN as DEVICE_PORT_EDGE_SWITCH_PORT_WWN,
DEVICE_PORT.LOGGED_TO_AG as DEVICE_PORT_LOGGED_TO_AG,
```

```
DEVICE_PORT.AG_NODE_WWN as DEVICE_PORT_AG_NODE_WWN,
DEVICE_PORT.AG_N_PORT_WWN as DEVICE_PORT_AG_N_PORT_WWN,
DEVICE_PORT.SYMBOLIC_NAME as DEVICE_PORT_NAME,
DEVICE_PORT.ID as DEVICE_PORT_ID,
DEVICE_PORT.PORT_ID as DEVICE_PORT_PORT_ID,
DEVICE_PORT.SWITCH_PORT_WWN as DEVICE_PORT_SWITCH_PORT_WWN,
DEVICE_PORT.FC4_TYPE as DEVICE_PORT_FC4_TYPE,
DEVICE_PORT.NODE_ID as DEVICE_NODE_ID,
DEVICE_PORT.COS as DEVICE_PORT_COS,
DEVICE_NODE.WWN as DEVICE_NODE_WWN,
DEVICE_NODE.DEVICE_TYPE as DEVICE_NODE_DEVICE_TYPE,
DEVICE_NODE.SYMBOLIC_NAME as DEVICE_NODE_NAME,
DEVICE_NODE.FDMI_HOST_NAME as DEVICE_FDMI_HOST_NAME,
DEVICE_NODE.CAPABILITY_ as DEVICE_NODE_CAPABILITY,
DEVICE_NODE.TYPE as DEVICE_NODE_TYPE,
DEVICE_NODE.VENDOR as DEVICE_NODE_VENDOR,
DEVICE_NODE.PROXY_DEVICE as DEVICE_NODE_PROXY_DEVICE,
DEVICE_NODE.AG as DEVICE_NODE_IS_AG,
DEVICE_NODE.SIMULATED as DEVICE_NODE_SIMULATED,
DEVICE_NODE.FABRIC_ID as DEVICE_NODE_FABRIC_ID
from DEVICE_PORT
join DEVICE_NODE on DEVICE_PORT.NODE_ID = DEVICE_NODE.ID
)
select FABRIC_MEMBER.FABRIC_ID,
FABRIC.SEED_SWITCH_WWN as FABRIC_SEED_SWITCH_WWN,
FABRIC.NAME as FABRIC_NAME,
FABRIC.MANAGEMENT_STATE as FABRIC_MANAGEMENT_STATE,
FABRIC.PRINCIPAL_SWITCH_WWN as FABRIC_PRINCIPAL_SWITCH_WWN,
FABRIC.FABRIC_NAME as FABRIC_SWITCH_PERSIST_FABRIC_NAME,
FABRIC.STATUS as FABRIC_STATUS,
FABRIC.BOTTLENECK_STATUS as FABRIC_BOTTLENECK_STATUS,
CORE_SWITCH.ID as CORE_SWITCH_ID,
CORE_SWITCH.IP_ADDRESS as CORE_SWITCH_IP_ADDRESS,
CORE_SWITCH.WWN as CORE_SWITCH_WWN, CORE_SWITCH.NAME as CORE_SWITCH_NAME,
CORE_SWITCH.TYPE as CORE_SWITCH_TYPE,
CORE_SWITCH.MODEL as CORE_SWITCH_MODEL,
CORE_SWITCH.VENDOR as CORE_SWITCH_VENDOR,
CORE_SWITCH.REACHABLE as CORE_SWITCH_REACHABLE,
CORE_SWITCH.OPERATIONAL_STATUS as CORE_SWITCH_OPERATIONAL_STATUS,
VIRTUAL_SWITCH.NAME as SWITCH_NAME, VIRTUAL_SWITCH.SWITCH_MODE,
VIRTUAL_SWITCH.DOMAIN_ID, VIRTUAL_SWITCH.WWN as VIRTUAL_SWITCH_WWN,
VIRTUAL_SWITCH.OPERATIONAL_STATUS as SWITCH_OPERATIONAL_STATUS,
VIRTUAL_SWITCH.MANAGEMENT_STATE as SWITCH_MANAGEMENT_STATE,
VIRTUAL_SWITCH.STATE as SWITCH_STATE,
VIRTUAL_SWITCH.STATUS as SWITCH_STATUS,
VIRTUAL_SWITCH.STATUS_REASON as SWITCH_STATUS_REASON,
SWITCH_PORT.ID as SWITCH_PORT_ID, SWITCH_PORT.VIRTUAL_SWITCH_ID,
SWITCH_PORT.SLOT_NUMBER, SWITCH_PORT.PORT_NUMBER,
SWITCH_PORT.USER_PORT_NUMBER, SWITCH_PORT.NAME as PORT_NAME,
SWITCH_PORT.SPEED as PORT_SPEED, SWITCH_PORT.STATUS as PORT_STATUS,
SWITCH_PORT.STATE as PORT_STATE, SWITCH_PORT.TYPE as PORT_TYPE,
SWITCH_PORT.SPEEDS_SUPPORTED,
SWITCH_PORT.PHYSICAL_PORT as PHYSICAL_OR_LOGICAL_PORT,
SWITCH_PORT.PORT_INDEX as ZONING_PORT_INDEX, SWITCH_PORT.OCCUPIED,
SWITCH_PORT.REMOTE_NODE_WWN as SWITCH_PORT_REMOTE_NODE_WWN,
SWITCH_PORT.PORT_ID, SWITCH_PORT.WWN as SWITCH_PORT_WWN,
SWITCH_PORT.AREA_ID, SWITCH_PORT.MAC_ADDRESS, SWITCH_PORT.PORT_MOD,
SWITCH_PORT.FULL_TYPE, SWITCH_PORT.HEALTH,
SWITCH_PORT.STATUS_MESSAGE as SWITCH_PORT_STATUS_MESSAGE,
SWITCH_PORT.MAX_PORT_SPEED, SWITCH_PORT.REMOTE_PORT_WWN,
SWITCH_PORT.LICENSED, SWITCH_PORT.TRUNKED, SWITCH_PORT.TRUNK_MASTER,
SWITCH_PORT.FICON_SUPPORTED, SWITCH_PORT.BLOCKED, SWITCH_PORT.NPIV,
SWITCH_PORT.NPIV_CAPABLE, SWITCH_PORT.NPIV_ENABLED, SWITCH_PORT.QOS_CAPABLE,
```

```
SWITCH_PORT.QOS_ENABLED, SWITCH_PORT.TUNNEL_CONFIGURED,
SWITCH_PORT.FCR_FABRIC_ID, SWITCH_PORT.FCR_INTEROP_MODE,
SWITCH_PORT.USER_DEFINED_VALUE1, SWITCH_PORT.USER_DEFINED_VALUE2,
SWITCH_PORT.USER_DEFINED_VALUE3, SWITCH_PORT.KIND, SWITCH_PORT.LAST_UPDATE,
DEVPORT_DEVNODE_VIEW.DEVICE_PORT_TYPE, DEVPORT_DEVNODE_VIEW.DEVICE_PORT_WWN,
DEVPORT_DEVNODE_VIEW.NPV_PHYSICAL,
DEVPORT_DEVNODE_VIEW.DEVICE_PORT_EDGE_SWITCH_PORT_WWN,
DEVPORT_DEVNODE_VIEW.DEVICE_PORT_LOGGED_TO_AG,
DEVPORT_DEVNODE_VIEW.DEVICE_PORT_AG_NODE_WWN,
DEVPORT_DEVNODE_VIEW.DEVICE_PORT_AG_N_PORT_WWN,
DEVPORT_DEVNODE_VIEW.DEVICE_PORT_NAME, DEVPORT_DEVNODE_VIEW.DEVICE_PORT_ID,
DEVPORT_DEVNODE_VIEW.DEVICE_PORT_PORT_ID,
DEVPORT_DEVNODE_VIEW.DEVICE_PORT_SWITCH_PORT_WWN,
DEVPORT_DEVNODE_VIEW.DEVICE_PORT_FC4_TYPE,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_ID, DEVPORT_DEVNODE_VIEW.DEVICE_PORT_COS,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_WWN,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_DEVICE_TYPE,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_NAME,
DEVPORT_DEVNODE_VIEW.DEVICE_FDMI_HOST_NAME,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_CAPABILITY,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_TYPE,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_VENDOR,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_PROXY_DEVICE,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_IS_AG,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_SIMULATED,
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_FABRIC_ID,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_PORT_ID,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_VIRTUAL_SWITCH_ID,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_SLOT_NUMBER,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_NUMBER,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_USER_PORT_NUMBER,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_NAME,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_SPEED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_STATUS,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_STATE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_TYPE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_SPEEDS_SUPPORTED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PHYSICAL_OR_LOGICAL_PORT,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_ZONING_PORT_INDEX,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_ID,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_PORT_WWN,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_AREA_ID,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_MAC_ADDRESS,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_MOD,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_FULL_TYPE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_HEALTH,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_PORT_STATUS_MESSAGE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_MAX_PORT_SPEED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_LICENSED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_REMOTE_NODE_WWN,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_REMOTE_PORT_WWN,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_TRUNKED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_TRUNK_MASTER,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_FICON_SUPPORTED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_BLOCKED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_NPIV,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_NPIV_CAPABLE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_NPIV_ENABLED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_QOS_CAPABLE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_QOS_ENABLED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_TUNNEL_CONFIGURED,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_FCR_FABRIC_ID,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_FCR_INTEROP_MODE,
```

```
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_USER_DEFINED_VALUE1,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_USER_DEFINED_VALUE2,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_USER_DEFINED_VALUE3,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_KIND,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_PORT_LAST_UPDATE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_NAME,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_MODE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_DOMAIN_ID,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_VIRTUAL_SWITCH_WWN,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_OPERATIONAL_STATUS,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_MANAGEMENT_STATE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_STATE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_STATUS,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_STATUS_REASON,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_ID,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_IP_ADDRESS,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_WWN,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_NAME,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_TYPE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_MODEL,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_VENDOR,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_REACHABLE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_OPERATIONAL_STATUS,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_MODEL_NUMBER,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_MODEL_DESCRIPTION,
USER_DEFINED_DEVICE_DETAIL.NAME as USER_DEFINED_DEVICE_NAME,
USER_DEFINED_DEVICE_DETAIL.IP_ADDRESS as USER_DEFINED_DEVICE_IP_ADDRESS,
USER_DEFINED_DEVICE_DETAIL.CONTACT as USER_DEFINED_DEVICE_CONTACT,
USER_DEFINED_DEVICE_DETAIL.LOCATION as USER_DEFINED_DEVICE_LOCATION,
USER_DEFINED_DEVICE_DETAIL.DESCRIPTION as USER_DEFINED_DEVICE_DESCRIPTION,
PORT_BOTTLENECK_STATUS.STATUS as BOTTLENECK_STATUS,
coalesce(USER_DEFINED_DEVICE_DETAIL.TYPE, DEVPORT_DEVNODE_VIEW.DEVICE_NODE_TYPE, ''::character varying) as
USER_DEFINED_DEVICE_TYPE,
coalesce(DEVPORT_DEVNODE_VIEW.DEVICE_NODE_NAME, SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_NAME,
''::character varying) as DEVICE_NAME,
coalesce(DEVPORT_DEVNODE_VIEW.DEVICE_NODE_VENDOR, SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_VENDOR,
''::character varying) as DEVICE_VENDOR,
coalesce(DEVPORT_DEVNODE_VIEW.DEVICE_NODE_TYPE,
SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_MODEL_DESCRIPTION, ''::CHARACTER VARYING) as DEVICE_TYPE,
coalesce(SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_CORE_SWITCH_MODEL_NUMBER, ''::character varying) as
DEVICE_MODEL,
coalesce(DEVPORT_DEVNODE_VIEW.DEVICE_NODE_WWN, SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_VIRTUAL_SWITCH_WWN,
''::bpchar) as CONNECTED_DEVICE_OR_SWITCH_WWN
from SWITCH_PORT
join VIRTUAL_SWITCH on SWITCH_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
join CORE_SWITCH on VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
join FABRIC_MEMBER on VIRTUAL_SWITCH.ID = FABRIC_MEMBER.VIRTUAL_SWITCH_ID
join FABRIC on FABRIC_MEMBER.FABRIC_ID = FABRIC.ID
left join USER_DEFINED_DEVICE_DETAIL on SWITCH_PORT.REMOTE_NODE_WWN::bpchar = USER_DEFINED_DEVICE_DETAIL.WWN
left join DEVPORT_DEVNODE_VIEW on SWITCH_PORT.WWN = DEVPORT_DEVNODE_VIEW.DEVICE_PORT_SWITCH_PORT_WWN and
SWITCH_PORT.REMOTE_PORT_WWN::bpchar = DEVPORT_DEVNODE_VIEW.DEVICE_PORT_WWN and FABRIC.ID =
DEVPORT_DEVNODE_VIEW.DEVICE_NODE_FABRIC_ID
left join SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW on SWITCH_PORT.REMOTE_PORT_WWN is not null and
SWITCH_PORT.REMOTE_PORT_WWN::TEXT<> ''::TEXT and SWITCH_PORT.WWN is not null and SWITCH_PORT.WWN <>
''::bpchar and SWITCH_PORT.REMOTE_PORT_WWN::bpchar = SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_SWITCH_PORT_WWN
and SWITCH_PORT.WWN = SWPRT_VSW_CSW_CSWDET_SWMDL_VIEW.REMOTE_REMOTE_PORT_WWN::bpchar
left join PORT_BOTTLENECK_STATUS on PORT_BOTTLENECK_STATUS.SWITCH_PORT_ID = SWITCH_PORT.ID
where SWITCH_PORT.LICENSED = 1 and SWITCH_PORT.PHYSICAL_PORT = 1 and SWITCH_PORT.KIND::text<> 'ICL'::text and
((SWITCH_PORT.TYPE::text= any (array['E-Port'::character varying, 'G-Port'::character varying,
'U-Port'::character varying, 'F-Port'::character varying, 'L-Port'::character varying, 'EX-Port'::character
varying, 'LB-Port'::character varying, 'FL-Port'::character varying, 'SIM-Port'::character varying,
```

```
'N-Port'::character varying]::text[])) or SWITCH_PORT.TYPE::text~~ 'LB-Port%'::text) and
VIRTUAL_SWITCH.MONITORED = 1 and (VIRTUAL_SWITCH.SWITCH_MODE = any (array[0, 2])) and ((CORE_SWITCH.TYPE <>
all (array[62, 63])) or (SWITCH_PORT.SLOT_NUMBER <> all (array[5, 8]))) and (FABRIC.MANAGED is null or
FABRIC.MANAGED = 1) and (FABRIC.TYPE is null or (FABRIC.TYPE <> all (array[65, 66, 4])));
```

SWITCH_PORT_DETAILS_REPORT_INFO

```
CREATE OR REPLACE VIEW switch_port_details_report_info AS
 WITH swprt_vsw_csw_cswdet_swmdl_view AS (
         SELECT switch_port.id AS remote_switch_port_id,
            switch_port.virtual_switch_id AS remote_virtual_switch_id,
            switch_port.slot_number AS remote_port_slot_number,
            switch_port.port_number AS remote_port_number,
            switch_port.user_port_number AS remote_user_port_number,
            switch_port.name AS remote_port_name,
            switch_port.speed AS remote_port_speed,
            switch_port.status AS remote_port_status,
            switch_port.state AS remote_port_state,
            switch_port.type AS remote_port_type,
            switch_port.speeds_supported AS remote_port_speeds_supported,
            switch_port.physical_port AS remote_physical_or_logical_port,
            switch_port.port_index AS remote_zoning_port_index,
            switch_port.port_id AS remote_port_id,
            switch_port.wwn AS remote_switch_port_wwn,
            switch_port.area_id AS remote_port_area_id,
            switch_port.mac_address AS remote_port_mac_address,
            switch_port.port_mod AS remote_port_mod,
            switch_port.full_type AS remote_port_full_type,
            switch_port.health AS remote_port_health,
            switch_port.status_message AS remote_switch_port_status_message,
            switch_port.max_port_speed AS remote_port_max_port_speed,
            switch_port.licensed AS remote_port_licensed,
            switch_port.remote_node_wwn AS remote_remote_node_wwn,
            switch_port.remote_port_wwn AS remote_remote_port_wwn,
            switch_port.trunked AS remote_port_trunked,
            switch_port.trunk_master AS remote_port_trunk_master,
            switch_port.ficon_supported AS remote_port_ficon_supported,
            switch_port.blocked AS remote_port_blocked,
            switch_port.npiv AS remote_port_npiv,
            switch_port.npiv_capable AS remote_port_npiv_capable,
            switch_port.npiv_enabled AS remote_port_npiv_enabled,
            switch_port.qos_capable AS remote_port_qos_capable,
            switch_port.qos_enabled AS remote_port_qos_enabled,
            switch_port.tunnel_configured AS remote_port_tunnel_configured,
            switch_port.fcr_fabric_id AS remote_fcr_fabric_id,
            switch_port.fcr_interop_mode AS remote_port_fcr_interop_mode,
            switch_port.user_defined_value1 AS remote_port_user_defined_value1,
            switch_port.user_defined_value2 AS remote_port_user_defined_value2,
            switch_port.user_defined_value3 AS remote_port_user_defined_value3,
            switch_port.kind AS remote_port_kind,
            switch_port.last_update AS remote_port_last_update,
            virtual_switch.name AS remote_switch_name,
            virtual_switch.switch_mode AS remote_switch_mode,
            virtual_switch.domain_id AS remote_switch_domain_id,
            virtual_switch.wwn AS remote_virtual_switch_wwn,
            virtual_switch.operational_status AS remote_switch_operational_status,
            virtual_switch.management_state AS remote_switch_management_state,
            virtual_switch.state AS remote_switch_state,
            virtual_switch.status AS remote_switch_status,
            virtual_switch.status_reason AS remote_switch_status_reason,
            virtual_switch.managed_element_id AS remote_virtual_switch_managed_element_id,
```

```
            core_switch.id AS remote_core_switch_id,
            core_switch.ip_address AS remote_core_switch_ip_address,
            core_switch.wwn AS remote_core_switch_wwn,
            core_switch.name AS remote_core_switch_name,
            core_switch.type AS remote_core_switch_type,
            core_switch.model AS remote_core_switch_model,
            core_switch.vendor AS remote_core_switch_vendor,
            core_switch.reachable AS remote_core_switch_reachable,
            core_switch.operational_status AS remote_core_switch_operational_status,
            core_switch.managed_element_id AS remote_core_switch_managed_element_id,
            core_switch_details.model_number AS remote_core_switch_model_number,
            switch_model.description AS remote_switch_model_description
           FROM switch_port
      JOIN virtual_switch ON switch_port.virtual_switch_id = virtual_switch.id
   JOIN core_switch ON virtual_switch.core_switch_id = core_switch.id
   LEFT JOIN core_switch_details ON virtual_switch.core_switch_id = core_switch_details.core_switch_id
   LEFT JOIN switch_model ON core_switch_details.type = switch_model.swbd_type AND
COALESCE(core_switch_details.sub_type, '0'::character varying)::integer = switch_model.subtype
  WHERE switch_port.licensed = 1 AND switch_port.physical_port = 1 AND switch_port.kind::text <> 'ICL'::text
AND ((switch_port.type::text = ANY (ARRAY['E-Port'::character varying::text, 'G-Port'::character
varying::text, 'U-Port'::character varying::text, 'F-Port'::character varying::text, 'L-Port'::character
varying::text, 'EX-Port'::character varying::text, 'LB-Port'::character varying::text, 'FL-Port'::character
varying::text, 'SIM-Port'::character varying::text, 'N-Port'::character varying::text])) OR
switch_port.type::text ~~ 'LB-Port%'::text) AND virtual_switch.monitored = 1 AND (virtual_switch.switch_mode
= ANY (ARRAY[0, 2])) AND ((core_switch.type <> ALL (ARRAY[62, 63])) OR (switch_port.slot_number <> ALL
(ARRAY[5, 8])))
        ), devport_devnode_view AS (
         SELECT device_port.type AS device_port_type,
            device_port.wwn AS device_port_wwn, device_port.npv_physical,
            device_port.edge_switch_port_wwn AS device_port_edge_switch_port_wwn,
            device_port.logged_to_ag AS device_port_logged_to_ag,
            device_port.ag_node_wwn AS device_port_ag_node_wwn,
            device_port.ag_n_port_wwn AS device_port_ag_n_port_wwn,
            device_port.symbolic_name AS device_port_name,
            device_port.id AS device_port_id,
            device_port.port_id AS device_port_port_id,
            device_port.switch_port_wwn AS device_port_switch_port_wwn,
            device_port.fc4_type AS device_port_fc4_type,
            device_port.node_id AS device_node_id,
            device_port.cos AS device_port_cos,
            device_node.wwn AS device_node_wwn,
            device_node.device_type AS device_node_device_type,
            device_node.symbolic_name AS device_node_name,
            device_node.fdmi_host_name AS device_fdmi_host_name,
            device_node.capability_ AS device_node_capability,
            device_node.type AS device_node_type,
            device_node.vendor AS device_node_vendor,
            device_node.proxy_device AS device_node_proxy_device,
            device_node.ag AS device_node_is_ag,
            device_node.simulated AS device_node_simulated,
            device_node.fabric_id AS device_node_fabric_id
           FROM device_port
      JOIN device_node ON device_port.node_id = device_node.id
        )
SELECT fabric_member.fabric_id,
    fabric.seed_switch_wwn AS fabric_seed_switch_wwn,
    fabric.name AS fabric_name,
    fabric.management_state AS fabric_management_state,
    fabric.principal_switch_wwn AS fabric_principal_switch_wwn,
    fabric.fabric_name AS fabric_switch_persist_fabric_name,
    fabric.status AS fabric_status,
    fabric.bottleneck_status AS fabric_bottleneck_status,
    core_switch.id AS core_switch_id,
```

```
core_switch.ip_address AS core_switch_ip_address,
core_switch.wwn AS core_switch_wwn, core_switch.name AS core_switch_name,
core_switch.type AS core_switch_type,
core_switch.model AS core_switch_model,
core_switch.vendor AS core_switch_vendor,
core_switch.reachable AS core_switch_reachable,
core_switch.operational_status AS core_switch_operational_status,
core_switch.managed_element_id AS core_switch_managed_element_id,
virtual_switch.name AS switch_name, virtual_switch.switch_mode,
virtual_switch.domain_id, virtual_switch.wwn AS virtual_switch_wwn,
virtual_switch.operational_status AS switch_operational_status,
virtual_switch.management_state AS switch_management_state,
virtual_switch.state AS switch_state,
virtual_switch.status AS switch_status,
virtual_switch.status_reason AS switch_status_reason,
virtual_switch.managed_element_id AS virtual_switch_managed_element_id,
switch_port.id AS switch_port_id, switch_port.virtual_switch_id,
switch_port.slot_number, switch_port.port_number,
switch_port.user_port_number, switch_port.name AS port_name,
switch_port.speed AS port_speed, switch_port.status AS port_status,
switch_port.state AS port_state, switch_port.type AS port_type,
switch_port.speeds_supported,
switch_port.physical_port AS physical_or_logical_port,
switch_port.port_index AS zoning_port_index, switch_port.occupied,
switch_port.remote_node_wwn AS switch_port_remote_node_wwn,
switch_port.port_id, switch_port.wwn AS switch_port_wwn,
switch_port.area_id, switch_port.mac_address, switch_port.port_mod,
switch_port.full_type, switch_port.health,
switch_port.status_message AS switch_port_status_message,
switch_port.max_port_speed, switch_port.remote_port_wwn,
switch_port.licensed, switch_port.trunked, switch_port.trunk_master,
switch_port.ficon_supported, switch_port.blocked, switch_port.npiv,
switch_port.npiv_capable, switch_port.npiv_enabled, switch_port.qos_capable,
switch_port.qos_enabled, switch_port.tunnel_configured,
switch_port.fcr_fabric_id, switch_port.fcr_interop_mode,
switch_port.user_defined_value1, switch_port.user_defined_value2,
switch_port.user_defined_value3, switch_port.kind, switch_port.last_update,
devport_devnode_view.device_port_type, devport_devnode_view.device_port_wwn,
devport_devnode_view.npv_physical,
devport_devnode_view.device_port_edge_switch_port_wwn,
devport_devnode_view.device_port_logged_to_ag,
devport_devnode_view.device_port_ag_node_wwn,
devport_devnode_view.device_port_ag_n_port_wwn,
devport_devnode_view.device_port_name, devport_devnode_view.device_port_id,
devport_devnode_view.device_port_port_id,
devport_devnode_view.device_port_switch_port_wwn,
devport_devnode_view.device_port_fc4_type,
devport_devnode_view.device_node_id, devport_devnode_view.device_port_cos,
devport_devnode_view.device_node_wwn,
devport_devnode_view.device_node_device_type,
devport_devnode_view.device_node_name,
devport_devnode_view.device_fdmi_host_name,
devport_devnode_view.device_node_capability,
devport_devnode_view.device_node_type,
devport_devnode_view.device_node_vendor,
devport_devnode_view.device_node_proxy_device,
devport_devnode_view.device_node_is_ag,
devport_devnode_view.device_node_simulated,
devport_devnode_view.device_node_fabric_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_port_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_virtual_switch_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_slot_number,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_number,
```

```
swprt_vsw_csw_cswdet_swmdl_view.remote_user_port_number,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_name,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_speed,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_status,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_state,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_type,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_speeds_supported,
swprt_vsw_csw_cswdet_swmdl_view.remote_physical_or_logical_port,
swprt_vsw_csw_cswdet_swmdl_view.remote_zoning_port_index,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_port_wwn,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_area_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_mac_address,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_mod,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_full_type,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_health,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_port_status_message,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_max_port_speed,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_licensed,
swprt_vsw_csw_cswdet_swmdl_view.remote_remote_node_wwn,
swprt_vsw_csw_cswdet_swmdl_view.remote_remote_port_wwn,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_trunked,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_trunk_master,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_ficon_supported,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_blocked,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_npiv,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_npiv_capable,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_npiv_enabled,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_qos_capable,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_qos_enabled,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_tunnel_configured,
swprt_vsw_csw_cswdet_swmdl_view.remote_fcr_fabric_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_fcr_interop_mode,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_user_defined_value1,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_user_defined_value2,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_user_defined_value3,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_kind,
swprt_vsw_csw_cswdet_swmdl_view.remote_port_last_update,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_name,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_mode,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_domain_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_virtual_switch_wwn,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_operational_status,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_management_state,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_state,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_status,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_status_reason,
swprt_vsw_csw_cswdet_swmdl_view.remote_virtual_switch_managed_element_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_ip_address,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_wwn,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_name,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_type,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_model,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_vendor,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_reachable,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_operational_status,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_managed_element_id,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_model_number,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_model_description,
user_defined_device_detail.name AS user_defined_device_name,
user_defined_device_detail.ip_address AS user_defined_device_ip_address,
user_defined_device_detail.contact AS user_defined_device_contact,
```

```
    user_defined_device_detail.location AS user_defined_device_location,
    user_defined_device_detail.description AS user_defined_device_description,
    port_bottleneck_status.status AS bottleneck_status,
    COALESCE(user_defined_device_detail.type, devport_devnode_view.device_node_type, ''::character varying)
AS user_defined_device_type,
    COALESCE(devport_devnode_view.device_node_name, swprt_vsw_csw_cswdet_swmdl_view.remote_switch_name,
''::character varying) AS device_name,
    COALESCE(devport_devnode_view.device_node_vendor,
swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_vendor, ''::character varying) AS device_vendor,
    COALESCE(devport_devnode_view.device_node_type,
swprt_vsw_csw_cswdet_swmdl_view.remote_switch_model_description, ''::character varying) AS device_type,
    COALESCE(swprt_vsw_csw_cswdet_swmdl_view.remote_core_switch_model_number, ''::character varying) AS
device_model,
   COALESCE(devport_devnode_view.device_node_wwn, swprt_vsw_csw_cswdet_swmdl_view.remote_virtual_switch_wwn,
''::bpchar) AS connected_device_or_switch_wwn
   FROM switch_port
   JOIN virtual_switch ON switch_port.virtual_switch_id = virtual_switch.id
   JOIN core_switch ON virtual_switch.core_switch_id = core_switch.id
   JOIN fabric_member ON virtual_switch.id = fabric_member.virtual_switch_id
   JOIN fabric ON fabric_member.fabric_id = fabric.id
   LEFT JOIN user_defined_device_detail ON switch_port.remote_node_wwn::bpchar =
user_defined_device_detail.wwn
   LEFT JOIN devport_devnode_view ON switch_port.wwn = devport_devnode_view.device_port_switch_port_wwn AND
switch_port.remote_port_wwn::bpchar = devport_devnode_view.device_port_wwn AND fabric.id =
devport_devnode_view.device_node_fabric_id
   LEFT JOIN swprt_vsw_csw_cswdet_swmdl_view ON switch_port.remote_port_wwn IS NOT NULL AND
switch_port.remote_port_wwn::text <> ''::text AND switch_port.wwn IS NOT NULL AND switch_port.wwn <>
''::bpchar AND switch_port.remote_port_wwn::bpchar = swprt_vsw_csw_cswdet_swmdl_view.remote_switch_port_wwn
AND switch_port.wwn = swprt_vsw_csw_cswdet_swmdl_view.remote_remote_port_wwn::bpchar
   LEFT JOIN port_bottleneck_status ON port_bottleneck_status.switch_port_id = switch_port.id
  WHERE switch_port.licensed = 1 AND switch_port.physical_port = 1 AND switch_port.kind::text <> 'ICL'::text
AND ((switch_port.type::text = ANY (ARRAY['E-Port'::character varying::text, 'G-Port'::character
varying::text, 'U-Port'::character varying::text, 'F-Port'::character varying::text, 'L-Port'::character
varying::text, 'EX-Port'::character varying::text, 'LB-Port'::character varying::text, 'FL-Port'::character
varying::text, 'SIM-Port'::character varying::text, 'N-Port'::character varying::text])) OR
switch_port.type::text ~~ 'LB-Port%'::text) AND virtual_switch.monitored = 1 AND (virtual_switch.switch_mode
= ANY (ARRAY[0, 2])) AND ((core_switch.type <> ALL (ARRAY[62, 63])) OR (switch_port.slot_number <> ALL
(ARRAY[5, 8]))) AND (fabric.managed IS NULL OR fabric.managed = 1) AND (fabric.type IS NULL OR (fabric.type <>
ALL (ARRAY[65, 66, 4])));
```

SWITCH_DETAILS_INFO

```
create or replace view SWITCH_DETAILS_INFO as
select
   CORE_SWITCH.ID as PHYSICAL_SWITCH_ID,
   CORE_SWITCH.NAME as PHYSICAL_SWITCH_NAME,
   CORE_SWITCH.IP_ADDRESS,
   CORE_SWITCH.WWN as PHYSICAL_SWITCH_WWN,
   CORE_SWITCH.OPERATIONAL_STATUS as PHYSICAL_OPERATIONAL_STATUS,
   CORE_SWITCH.TYPE,
   CORE_SWITCH.MAX_VIRTUAL_SWITCHES,
   CORE_SWITCH.FIRMWARE_VERSION,
   CORE_SWITCH.VENDOR,
   CORE_SWITCH.REACHABLE,
   CORE_SWITCH.UNREACHABLE_TIME,
   CORE_SWITCH.MODEL,
   CORE_SWITCH.SYSLOG_REGISTERED,
   CORE_SWITCH.SNMP_REGISTERED,
   CORE_SWITCH.USER_IP_ADDRESS,
   CORE_SWITCH.MANAGING_SERVER_IP_ADDRESS,
   CORE_SWITCH.CREATION_TIME as CS_CREATION_TIME,
   CORE_SWITCH.LAST_UPDATE_TIME as CS_LAST_UPDATE_TIME,
```

```
CORE_SWITCH.NUM_VIRTUAL_SWITCHES,
CORE_SWITCH.VF_ENABLED,
CORE_SWITCH.VF_SUPPORTED,
CORE_SWITCH.CALL_HOME_ENABLED,
CORE_SWITCH.MANAGED_ELEMENT_ID as CORE_MANAGED_ELEMENT_ID,
CORE_SWITCH.NAT_PRIVATE_IP_ADDRESS,
CORE_SWITCH.ALTERNATE_IP_ADDRESS,
CORE_SWITCH.MAC_ADDRESS,
VIRTUAL_SWITCH.ID,
VIRTUAL_SWITCH.NAME,
VIRTUAL_SWITCH.OPERATIONAL_STATUS,
VIRTUAL_SWITCH.SWITCH_MODE,
VIRTUAL_SWITCH.AD_CAPABLE,
VIRTUAL_SWITCH.WWN,
VIRTUAL_SWITCH.ROLE,
VIRTUAL_SWITCH.FCS_ROLE,
VIRTUAL_SWITCH.DOMAIN_ID,
VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID,
VIRTUAL_SWITCH.BASE_SWITCH,
VIRTUAL_SWITCH.MAX_ZONE_CONFIG_SIZE,
VIRTUAL_SWITCH.CREATION_TIME,
VIRTUAL_SWITCH.LAST_UPDATE_TIME,
VIRTUAL_SWITCH.USER_NAME,
VIRTUAL_SWITCH.PASSWORD,
VIRTUAL_SWITCH.MANAGEMENT_STATE,
VIRTUAL_SWITCH.STATE,
VIRTUAL_SWITCH.STATUS,
VIRTUAL_SWITCH.STATUS_REASON,
VIRTUAL_SWITCH.FABRIC_IDID_MODE,
VIRTUAL_SWITCH.LOGICAL_ID,
VIRTUAL_SWITCH.USER_DEFINED_VALUE_1,
VIRTUAL_SWITCH.USER_DEFINED_VALUE_2,
VIRTUAL_SWITCH.USER_DEFINED_VALUE_3,
VIRTUAL_SWITCH.FMS_MODE,
VIRTUAL_SWITCH.DYNAMIC_LOAD_SHARING,
VIRTUAL_SWITCH.PORT_BASED_ROUTING,
VIRTUAL_SWITCH.IN_ORDER_DELIVERY,
VIRTUAL_SWITCH.INSISTENT_DID_MODE,
VIRTUAL_SWITCH.FCR_CAPABLE,
VIRTUAL_SWITCH.LAST_PORT_MEMBERSHIP_CHANGE,
VIRTUAL_SWITCH.FCIP_CIRCUIT_CAPABLE,
VIRTUAL_SWITCH.MAX_FCIP_TUNNELS,
VIRTUAL_SWITCH.MAX_FCIP_CIRCUITS,
VIRTUAL_SWITCH.FCIP_LICENSED,
VIRTUAL_SWITCH.MANAGED_ELEMENT_ID,
VIRTUAL_SWITCH.RNID_SEQUENCE_NUMBER as VS_RNID_SEQUENCE_NUMBER,
VIRTUAL_SWITCH.CLUSTER_MODE,
VIRTUAL_SWITCH.VCS_ID,
VIRTUAL_SWITCH.CLUSTER_TYPE,
VIRTUAL_SWITCH.RNID_TAG,
VIRTUAL_SWITCH.SWITCH_ID,
VIRTUAL_SWITCH.MONITORED,
VIRTUAL_SWITCH.MAPS_ENABLED_ACTIONS,
VIRTUAL_SWITCH.FEATURES_ENABLED,
VIRTUAL_SWITCH.ROUTING_POLICY,
VIRTUAL_SWITCH.FABRIC_STATUS,
VIRTUAL_SWITCH.PROTOCOL,
FABRIC_MEMBER.FABRIC_ID,
FABRIC_MEMBER.TRUSTED,
FABRIC_MEMBER.MISSING,
FABRIC_MEMBER.MISSING_TIME,
CORE_SWITCH_DETAILS.ETHERNET_MASK,
CORE_SWITCH_DETAILS.FC_MASK,
```

```
   CORE_SWITCH_DETAILS.FC_IP,
   CORE_SWITCH_DETAILS.FC_CERTIFICATE,
   CORE_SWITCH_DETAILS.SW_LICENSE_ID,
   CORE_SWITCH_DETAILS.SUPPLIER_SERIAL_NUMBER,
   CORE_SWITCH_DETAILS.PART_NUMBER,
   CORE_SWITCH_DETAILS.CHECK_BEACON,
   CORE_SWITCH_DETAILS.TIMEZONE,
   CORE_SWITCH_DETAILS.MAX_PORT,
   CORE_SWITCH_DETAILS.CHASSIS_SERVICE_TAG,
   CORE_SWITCH_DETAILS.BAY_ID,
   CORE_SWITCH_DETAILS.TYPE_NUMBER,
   CORE_SWITCH_DETAILS.MODEL_NUMBER,
   CORE_SWITCH_DETAILS.MANUFACTURER,
   CORE_SWITCH_DETAILS.PLANT_OF_MANUFACTURER,
   CORE_SWITCH_DETAILS.SWITCH_SERIAL_NUMBER,
   CORE_SWITCH_DETAILS.ACT_CP_PRI_FW_VERSION,
   CORE_SWITCH_DETAILS.ACT_CP_SEC_FW_VERSION,
   CORE_SWITCH_DETAILS.STBY_CP_PRI_FW_VERSION,
   CORE_SWITCH_DETAILS.STBY_CP_SEC_FW_VERSION,
   CORE_SWITCH_DETAILS.TYPE as DETAILS_TYPE,
   CORE_SWITCH_DETAILS.EGM_CAPABLE,
   CORE_SWITCH_DETAILS.SUB_TYPE,
   CORE_SWITCH_DETAILS.PARTITION,
   CORE_SWITCH_DETAILS.MAX_NUM_OF_BLADES,
   CORE_SWITCH_DETAILS.SNMP_INFORMS_ENABLED,
   CORE_SWITCH_DETAILS.VENDOR_VERSION,
   CORE_SWITCH_DETAILS.VENDOR_PART_NUMBER,
   CORE_SWITCH_DETAILS.CONTACT,
   CORE_SWITCH_DETAILS.LOCATION,
   CORE_SWITCH_DETAILS.DESCRIPTION,
   CORE_SWITCH_DETAILS.RNID_SEQUENCE_NUMBER,
   CORE_SWITCH_DETAILS.FIRMWARE_VERSION as CSD_FIRMWARE_VERSION,
   CORE_SWITCH_DETAILS.CHASSIS_PACKAGE_TYPE,
   CORE_SWITCH_DETAILS.IP_ADDRESS_PREFIX,
   CORE_SWITCH_DETAILS.DOMAIN_NAME,
   CORE_SWITCH_DETAILS.FRAME_LOG_SIZE,
   CORE_SWITCH_DETAILS.FRAME_LOG_ENABLED,
   CORE_SWITCH_DETAILS.MAPS_ENABLED
from
   CORE_SWITCH,
   VIRTUAL_SWITCH,
   FABRIC_MEMBER,
   CORE_SWITCH_DETAILS
where
   VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
   and FABRIC_MEMBER.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
   and CORE_SWITCH_DETAILS.CORE_SWITCH_ID = CORE_SWITCH.ID;
```

SWITCH_DISCOVERED_MAC_INFO

```
create or replace view SWITCH_DISCOVERED_MAC_INFO as
select
    L2_NEIGHBOR.LLDP_REM_CHASSIS_ID_VALUE,
    L2_NEIGHBOR.INTERFACE_ID,
    INTERFACE.NAME as INTERFACE_NAME,
    DEVICE.SYS_NAME as DEVICE_NAME,
    DEVICE.IP_ADDRESS, DEVICE.DEVICE_ID
from
    L2_NEIGHBOR,
    INTERFACE,
    DEVICE
where
```

```
    L2_NEIGHBOR.LLDP_REM_CHASSIS_ID_SUBTYPE = 4
    and L2_NEIGHBOR.INTERFACE_ID = INTERFACE.INTERFACE_ID
    and INTERFACE.DEVICE_ID = DEVICE.DEVICE_ID;
```

SWITCH_PORT_INFO

```
CREATE OR REPLACE VIEW switch_port_info AS
 SELECT switch_port.id,
    switch_port.virtual_switch_id,
    switch_port.wwn,
    switch_port.name,
    switch_port.slot_number,
    switch_port.port_number,
    switch_port.user_port_number,
    switch_port.port_id,
    switch_port.port_index,
    switch_port.area_id,
    switch_port.mac_address,
    switch_port.port_mod,
    switch_port.type,
    switch_port.full_type,
    switch_port.ext_type,
    switch_port.status,
    switch_port.health,
    switch_port.status_message,
    switch_port.physical_port,
    switch_port.locked_port_type,
    switch_port.category,
    switch_port.protocol,
    switch_port.speed,
    switch_port.speeds_supported,
    switch_port.max_port_speed,
    switch_port.desired_credits,
    switch_port.buffer_allocated,
    switch_port.estimated_distance,
    switch_port.actual_distance,
    switch_port.long_distance_setting,
    switch_port.degraded_port,
    switch_port.remote_node_wwn,
    switch_port.remote_port_wwn,
    switch_port.licensed,
    switch_port.swapped,
    switch_port.trunked,
    switch_port.trunk_master,
    switch_port.persistent_disable,
    switch_port.ficon_supported,
    switch_port.blocked,
    switch_port.prohibit_port_numbers,
    switch_port.prohibit_port_count,
    switch_port.npiv,
    switch_port.npiv_capable,
    switch_port.npiv_enabled,
    switch_port.fc_fast_write_enabled,
    switch_port.isl_rrdy_enabled,
    switch_port.rate_limit_capable,
    switch_port.rate_limited,
    switch_port.qos_capable,
    switch_port.qos_enabled,
    switch_port.tunnel_configured,
    switch_port.fcip_tunnel_up,
    switch_port.fcr_fabric_id,
    switch_port.fcr_interop_mode,
```

```
    switch_port.calculated_status,
    switch_port.user_defined_value1,
    switch_port.user_defined_value2,
    switch_port.user_defined_value3,
    switch_port.kind,
    switch_port.state,
    switch_port.previous_status,
    switch_port.last_update,
    switch_port.occupied,
    switch_port.port_bit_mask,
    switch_port.logical_port_number,
    switch_port.default_area_id,
    switch_port.logical_port_wwn,
    switch_port.latency_detect_supported,
    switch_port.eport_disabled,
    switch_port.speed_negotiated,
    switch_port.identifier,
    switch_port.port_capabilities,
    switch_port.fake_port,
    switch_port.xisl_port_list,
    switch_port.port_commission_state,
    switch_port.features_supported,
    switch_port.features_enabled,
    switch_port.features_active,
    switch_port.disabled_reason_code,
    switch_port.disabled_reason,
    switch_port.fenced,
    switch_port.master_port_number,
    switch_port.speed_type,
    switch_port.qsfp_unit_number,
    virtual_switch.wwn AS virtual_switch_wwn,
    virtual_switch.role AS switch_role,
    virtual_switch.virtual_fabric_id,
    virtual_switch.domain_id,
    virtual_switch.interop_mode,
    virtual_switch.management_state,
    virtual_switch.managed_element_id,
    virtual_switch.monitored,
    virtual_switch.routing_policy,
    core_switch.type AS switch_type,
    core_switch.firmware_version,
    core_switch.ip_address,
    core_switch.wwn AS physical_switch_wwn,
    core_switch.model AS switch_model,
    core_switch_details.model_number AS switch_model_number
   FROM switch_port,
    virtual_switch,
    core_switch
     LEFT JOIN core_switch_details ON core_switch_details.core_switch_id = core_switch.id
  WHERE switch_port.virtual_switch_id = virtual_switch.id AND virtual_switch.core_switch_id = core_switch.id;
```

SWITCH_SNMP_INFO

```
create or replace view SWITCH_SNMP_INFO as
select
    CORE_SWITCH.ID as PHYSICAL_SWITCH_ID,
    CORE_SWITCH.NAME as PHYSICAL_SWITCH_NAME,
    CORE_SWITCH.IP_ADDRESS,
    CORE_SWITCH.WWN as PHYSICAL_SWITCH_WWN,
    CORE_SWITCH.OPERATIONAL_STATUS as PHYSICAL_OPERATIONAL_STATUS,
    CORE_SWITCH.TYPE,
    CORE_SWITCH.MAX_VIRTUAL_SWITCHES,
```

```
    CORE_SWITCH.FIRMWARE_VERSION,
    CORE_SWITCH.VENDOR,
    CORE_SWITCH.REACHABLE,
    CORE_SWITCH.UNREACHABLE_TIME,
    CORE_SWITCH.MODEL,
    CORE_SWITCH_DETAILS.CONTACT,
    CORE_SWITCH_DETAILS.LOCATION,
    CORE_SWITCH_DETAILS.DESCRIPTION,
    VIRTUAL_SWITCH.ID,
    VIRTUAL_SWITCH.NAME,
    VIRTUAL_SWITCH.OPERATIONAL_STATUS,
    VIRTUAL_SWITCH.SWITCH_MODE,
    VIRTUAL_SWITCH.AD_CAPABLE,
    VIRTUAL_SWITCH.FCIP_CAPABLE,
    VIRTUAL_SWITCH.WWN,
    VIRTUAL_SWITCH.ROLE,
    VIRTUAL_SWITCH.FCS_ROLE,
    VIRTUAL_SWITCH.DOMAIN_ID,
    VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID,
    VIRTUAL_SWITCH.BASE_SWITCH,
    VIRTUAL_SWITCH.MAX_ZONE_CONFIG_SIZE,
    VIRTUAL_SWITCH.CREATION_TIME,
    VIRTUAL_SWITCH.LAST_UPDATE_TIME,
    VIRTUAL_SWITCH.USER_NAME,
    VIRTUAL_SWITCH.PASSWORD,
    VIRTUAL_SWITCH.MANAGEMENT_STATE,
    VIRTUAL_SWITCH.STATE,
    VIRTUAL_SWITCH.STATUS,
    VIRTUAL_SWITCH.STATUS_REASON,
    VIRTUAL_SWITCH.MONITORED,
    VIRTUAL_SWITCH.USER_DEFINED_VALUE_1,
    VIRTUAL_SWITCH.USER_DEFINED_VALUE_2,
    VIRTUAL_SWITCH.USER_DEFINED_VALUE_3,
    FABRIC_MEMBER.FABRIC_ID,
    FABRIC_MEMBER.TRUSTED,
    FABRIC_MEMBER.MISSING,
    FABRIC_MEMBER.MISSING_TIME,
    coalesce(SNMP_CREDENTIALS.PORT_NUMBER, (select SNMP_PROFILE.PORT_NUMBER from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_PORT_NUMBER,
    coalesce(SNMP_CREDENTIALS.RETRY_COUNT, (select SNMP_PROFILE.RETRY_COUNT from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_RETRY_COUNT,
    coalesce(SNMP_CREDENTIALS.TIMEOUT, (select SNMP_PROFILE.TIMEOUT from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_TIMEOUT,
    coalesce(SNMP_CREDENTIALS.VERSION, (select SNMP_PROFILE.VERSION from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_VERSION,
    coalesce(SNMP_CREDENTIALS.READ_COMMUNITY_STRING, (select SNMP_PROFILE.READ_COMMUNITY_STRING from
SNMP_PROFILE where SNMP_PROFILE.NAME='default')) as SNMP_READ_COMMUNITY_STRING,
    coalesce(SNMP_CREDENTIALS.WRITE_COMMUNITY_STRING, (select SNMP_PROFILE.WRITE_COMMUNITY_STRING from
SNMP_PROFILE where SNMP_PROFILE.NAME='default')) as SNMP_WRITE_COMMUNITY_STRING,
    coalesce(SNMP_CREDENTIALS.USER_NAME, (select SNMP_PROFILE.USER_NAME from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_USER_NAME,
    coalesce(SNMP_CREDENTIALS.CONTEXT_NAME, (select SNMP_PROFILE.CONTEXT_NAME from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_CONTEXT_NAME,
    coalesce(SNMP_CREDENTIALS.AUTH_PROTOCOL, (select SNMP_PROFILE.AUTH_PROTOCOL from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_AUTH_PROTOCOL,
    coalesce(SNMP_CREDENTIALS.AUTH_PASSWORD, (select SNMP_PROFILE.AUTH_PASSWORD from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_AUTH_PASSWORD,
    coalesce(SNMP_CREDENTIALS.PRIV_PROTOCOL, (select SNMP_PROFILE.PRIV_PROTOCOL from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_PRIV_PROTOCOL,
    coalesce(SNMP_CREDENTIALS.PRIV_PASSWORD, (select SNMP_PROFILE.PRIV_PASSWORD from SNMP_PROFILE where
SNMP_PROFILE.NAME='default')) as SNMP_PRIV_PASSWORD,
    coalesce(SNMP_CREDENTIALS.SNMP_INFORMS_ENABLED, (select SNMP_PROFILE.SNMP_INFORMS_ENABLED from
SNMP_PROFILE where SNMP_PROFILE.NAME='default')) as SNMP_INFORMS_ENABLED
```

```
from
    VIRTUAL_SWITCH
        left outer join CORE_SWITCH
            on VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
        left outer join CORE_SWITCH_DETAILS
            on CORE_SWITCH.ID = CORE_SWITCH_DETAILS.CORE_SWITCH_ID
        left outer join FABRIC_MEMBER
            on FABRIC_MEMBER.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
        left outer join SNMP_CREDENTIALS
            on VIRTUAL_SWITCH.ID = SNMP_CREDENTIALS.VIRTUAL_SWITCH_ID;
```

TIME_SERIES_DATA_INFO

```
CREATE VIEW time_series_data_info AS
    (  ( ( ( ( (
select * from TIME_SERIES_DATA_1
union all
select TIME_SERIES_DATA_1_30MIN.TIME_IN_SECONDS,
       TIME_SERIES_DATA_1_30MIN.TARGET_TYPE,
       TIME_SERIES_DATA_1_30MIN.MEASURE_ID,
       TIME_SERIES_DATA_1_30MIN.TARGET_ID,
       TIME_SERIES_DATA_1_30MIN.COLLECTOR_ID,
       TIME_SERIES_DATA_1_30MIN.MEASURE_INDEX,
       TIME_SERIES_DATA_1_30MIN.ME_ID,
       TIME_SERIES_DATA_1_30MIN.VALUE,
       TIME_SERIES_DATA_1_30MIN.SUM_VALUE
from TIME_SERIES_DATA_1_30MIN)
union all
select TIME_SERIES_DATA_1_2HOUR.TIME_IN_SECONDS,
       TIME_SERIES_DATA_1_2HOUR.TARGET_TYPE,
       TIME_SERIES_DATA_1_2HOUR.MEASURE_ID,
       TIME_SERIES_DATA_1_2HOUR.TARGET_ID,
       TIME_SERIES_DATA_1_2HOUR.COLLECTOR_ID,
       TIME_SERIES_DATA_1_2HOUR.MEASURE_INDEX,
       TIME_SERIES_DATA_1_2HOUR.ME_ID,
       TIME_SERIES_DATA_1_2HOUR.VALUE,
       TIME_SERIES_DATA_1_2HOUR.SUM_VALUE
from TIME_SERIES_DATA_1_2HOUR)
union all
select TIME_SERIES_DATA_1_1DAY.TIME_IN_SECONDS,
       TIME_SERIES_DATA_1_1DAY.TARGET_TYPE,
       TIME_SERIES_DATA_1_1DAY.MEASURE_ID,
       TIME_SERIES_DATA_1_1DAY.TARGET_ID,
       TIME_SERIES_DATA_1_1DAY.COLLECTOR_ID,
       TIME_SERIES_DATA_1_1DAY.MEASURE_INDEX,
       TIME_SERIES_DATA_1_1DAY.ME_ID,
       TIME_SERIES_DATA_1_1DAY.VALUE,
       TIME_SERIES_DATA_1_1DAY.SUM_VALUE
from TIME_SERIES_DATA_1_1DAY)
union all
select * from TIME_SERIES_DATA_2)
union all
select TIME_SERIES_DATA_2_30MIN.TIME_IN_SECONDS,
       TIME_SERIES_DATA_2_30MIN.TARGET_TYPE,
       TIME_SERIES_DATA_2_30MIN.MEASURE_ID,
       TIME_SERIES_DATA_2_30MIN.TARGET_ID,
       TIME_SERIES_DATA_2_30MIN.COLLECTOR_ID,
       TIME_SERIES_DATA_2_30MIN.MEASURE_INDEX,
       TIME_SERIES_DATA_2_30MIN.ME_ID,
       TIME_SERIES_DATA_2_30MIN.VALUE,
       TIME_SERIES_DATA_2_30MIN.SUM_VALUE
from TIME_SERIES_DATA_2_30MIN)
```

```
union all
select TIME_SERIES_DATA_2_2HOUR.TIME_IN_SECONDS,
       TIME_SERIES_DATA_2_2HOUR.TARGET_TYPE,
       TIME_SERIES_DATA_2_2HOUR.MEASURE_ID,
       TIME_SERIES_DATA_2_2HOUR.TARGET_ID,
       TIME_SERIES_DATA_2_2HOUR.COLLECTOR_ID,
       TIME_SERIES_DATA_2_2HOUR.MEASURE_INDEX,
       TIME_SERIES_DATA_2_2HOUR.ME_ID,
       TIME_SERIES_DATA_2_2HOUR.VALUE,
       TIME_SERIES_DATA_2_2HOUR.SUM_VALUE
from TIME_SERIES_DATA_2_2HOUR)
union all
select TIME_SERIES_DATA_2_1DAY.TIME_IN_SECONDS,
       TIME_SERIES_DATA_2_1DAY.TARGET_TYPE,
       TIME_SERIES_DATA_2_1DAY.MEASURE_ID,
       TIME_SERIES_DATA_2_1DAY.TARGET_ID,
       TIME_SERIES_DATA_2_1DAY.COLLECTOR_ID,
       TIME_SERIES_DATA_2_1DAY.MEASURE_INDEX,
       TIME_SERIES_DATA_2_1DAY.ME_ID,
       TIME_SERIES_DATA_2_1DAY.VALUE,
       TIME_SERIES_DATA_2_1DAY.SUM_VALUE
from TIME_SERIES_DATA_2_1DAY
```

TIME_SERIES_DATA_VIEW

```
create or replace view TIME_SERIES_DATA_VIEW as
(         SELECT de.device_id, cast (de.ip_address as varchar(255)) AS device_ip,
                    tsd.target_type, de.device_id AS target_id,
                    de.sys_name AS target_name,
                    measure.measure_type AS collectible_type,
                    tsd.measure_id AS collectible_id, tsd.collector_id,
                    pdc.name AS collector_name,
                    (measure.name::text || '.'::text) || tsd.measure_index::text AS collectible_name,
                    measure.detail AS collectible_detail, tsd.value,
                    tsd.time_in_seconds, tsd.measure_index
                   FROM time_series_data_info tsd
               JOIN device de ON tsd.target_id = de.device_id
          JOIN pm_data_collector pdc ON pdc.id = tsd.collector_id
     JOIN measure ON measure.id = tsd.measure_id
    WHERE tsd.target_type = 0 OR tsd.target_type = 18
UNION ALL
                 SELECT de.device_id, cast (de.ip_address as varchar(255)) AS device_ip,
                    tsd.target_type, ifs.interface_id AS target_id,
                    ifs.if_name AS target_name,
                    measure.measure_type AS collectible_type,
                    tsd.measure_id AS collectible_id, tsd.collector_id,
                    pm_data_collector.name AS collector_name,
                    (measure.name::text || '.'::text) || tsd.measure_index::text AS collectible_name,
                    measure.detail AS collectible_detail, tsd.value,
                    tsd.time_in_seconds, tsd.measure_index
                   FROM time_series_data_info tsd
               JOIN interface ifs ON (tsd.target_type = 1 OR tsd.target_type = 2 OR tsd.target_type =15) AND
tsd.target_id = ifs.interface_id
          JOIN device de ON ifs.device_id = de.device_id
     JOIN pm_data_collector ON pm_data_collector.id = tsd.collector_id
    JOIN measure ON measure.id = tsd.measure_id)
UNION ALL
          SELECT de.device_id, cast (de.ip_address as varchar(255)) AS device_ip, tsd.target_type,
             sp.id AS target_id, sp.name AS target_name,
             measure.measure_type AS collectible_type,
             tsd.measure_id AS collectible_id, tsd.collector_id,
```

```
            pm_data_collector.name AS collector_name,
            (measure.name::text || '.'::text) || tsd.measure_index::text AS collectible_name,
            measure.detail AS collectible_detail, tsd.value,
            tsd.time_in_seconds, tsd.measure_index
           FROM time_series_data_info tsd
      JOIN switch_port sp ON tsd.target_type = 4 AND tsd.target_id = sp.id
   JOIN virtual_switch vs ON sp.virtual_switch_id = vs.id
   JOIN device de ON vs.managed_element_id = de.managed_element_id
   JOIN pm_data_collector ON pm_data_collector.id = tsd.collector_id
   JOIN measure ON measure.id = tsd.measure_id
UNION ALL
         SELECT 0 as device_id, cast (vs.ip_address as varchar(255)) AS device_ip, tsd.target_type,
            sp.id AS target_id, sp.name AS target_name,
            measure.measure_type AS collectible_type,
            tsd.measure_id AS collectible_id, tsd.collector_id,
            pm_data_collector.name AS collector_name,
            (measure.name::text || '.'::text) || tsd.measure_index::text AS collectible_name,
            measure.detail AS collectible_detail, tsd.value,
            tsd.time_in_seconds, tsd.measure_index
           FROM time_series_data_info tsd
      JOIN switch_port sp ON (tsd.target_type = 4 OR tsd.target_type = 5 OR tsd.target_type = 6) AND
tsd.target_id = sp.id
   JOIN switch_info vs ON sp.virtual_switch_id = vs.id
   JOIN pm_data_collector ON pm_data_collector.id = tsd.collector_id
   JOIN measure ON measure.id = tsd.measure_id
UNION ALL
SELECT 0 as device_id, cast (de.ip_address as varchar(255)) AS device_ip,
                    tsd.target_type, de.id AS target_id,
                    cast (de.physical_switch_name as text) AS target_name,
                    measure.measure_type AS collectible_type,
                    tsd.measure_id AS collectible_id, tsd.collector_id,
                    pdc.name AS collector_name,
                    (measure.name::text || '.'::text) || tsd.measure_index::text AS collectible_name,
                    measure.detail AS collectible_detail, tsd.value,
                    tsd.time_in_seconds, tsd.measure_index
                   FROM time_series_data_info tsd
              JOIN switch_info de ON tsd.target_id = de.id
         JOIN pm_data_collector pdc ON pdc.id = tsd.collector_id
    JOIN measure ON measure.id = tsd.measure_id
   WHERE tsd.target_type = 3;
```

TRILL_INFO

```
create or replace view TRILL_INFO as
select distinct
    TRILL.ID,
    VCS_DEVICE.DEVICE_ID as VCS_DEVICE_ID,
    TRILL.CLUSTER_ME_ID,
    TRILL.COST,
    TRILL.TYPE as LINK_TYPE,
    TRILL.MISSING,
    TRILL.TRUNKED,
    TRILL.SOURCE_DOMAIN_ID,
    TRILL.SOURCE_PORT_NUMBER,
    TRILL.SOURCE_PORT_NAME as SOURCE_SWITCH_PORT_NAME,
    TRILL.SOURCE_ME_ID,
    SOURCE_DEVICE.DEVICE_ID AS SOURCE_DEVICE_ID,
    TRILL.DEST_DOMAIN_ID,
    TRILL.DEST_PORT_NUMBER,
    TRILL.DEST_PORT_NAME as DEST_SWITCH_PORT_NAME,
    TRILL.DEST_ME_ID,
    DEST_DEVICE.DEVICE_ID AS DEST_DEVICE_ID
```

```
from
    TRILL,
    device VCS_DEVICE,
    device SOURCE_DEVICE,
    VIRTUAL_SWITCH SOURCE_VIRTUAL_SWITCH,
    device DEST_DEVICE,
    VIRTUAL_SWITCH DEST_VIRTUAL_SWITCH
where
    SOURCE_DEVICE.MANAGED_ELEMENT_ID = TRILL.SOURCE_ME_ID and
    DEST_DEVICE.MANAGED_ELEMENT_ID = TRILL.DEST_ME_ID and
    VCS_DEVICE.MANAGED_ELEMENT_ID = TRILL.CLUSTER_ME_ID and
    SOURCE_VIRTUAL_SWITCH.MANAGED_ELEMENT_ID = TRILL.SOURCE_ME_ID and
    DEST_VIRTUAL_SWITCH.MANAGED_ELEMENT_ID = TRILL.DEST_ME_ID;
```

TRILL_TRUNK_INFO

```
create or replace view TRILL_TRUNK_INFO as
select
    TRILL_TRUNK_GROUP.ID,
    TRILL_TRUNK_GROUP.ME_ID,
    TRILL_TRUNK_GROUP.MASTER_PORT_NUMBER,
    TRILL_TRUNK_MEMBER.PORT_NUMBER as MEMBER_PORT_NUMBER,
    MEMBER_DEVICE.DEVICE_ID,
    INTERFACE.INTERFACE_ID,
    VCS_CLUSTER_MEMBER.CLUSTER_ME_ID,
    CLUSTER_DEVICE.DEVICE_ID as CLUSTER_DEVICE_ID
from
    TRILL_TRUNK_GROUP
    inner join
      TRILL_TRUNK_MEMBER on
      TRILL_TRUNK_MEMBER.GROUP_ID = TRILL_TRUNK_GROUP.ID
    inner join
      DEVICE as MEMBER_DEVICE on
      MEMBER_DEVICE.MANAGED_ELEMENT_ID = TRILL_TRUNK_GROUP.ME_ID
    left outer join
      INTERFACE on
      INTERFACE.DEVICE_ID = MEMBER_DEVICE.DEVICE_ID and
      INTERFACE.IDENTIFIER =  TRILL_TRUNK_MEMBER.PORT_NUMBER
    left outer join
      VCS_CLUSTER_MEMBER on
      VCS_CLUSTER_MEMBER.MEMBER_ME_ID = TRILL_TRUNK_GROUP.ME_ID
    left outer join
      DEVICE as CLUSTER_DEVICE on
      CLUSTER_DEVICE.MANAGED_ELEMENT_ID = VCS_CLUSTER_MEMBER.CLUSTER_ME_ID;
```

USER_ROLE_RESOURCE_INFO

```
create or replace view USER_ROLE_RESOURCE_INFO as
select
  RESOURCE_GROUP.ID RESOURCE_GROUP_ID,
  RESOURCE_GROUP.NAME RESOURCE_GROUP_NAME,
  ROLE.ID ROLE_ID,
  ROLE.NAME ROLE_NAME,
  USER_.NAME USER_NAME
from
  USER_,
  RESOURCE_GROUP,
  ROLE,
  USER_RESOURCE_MAP,
  USER_ROLE_MAP
```

```
where
  USER_ROLE_MAP.USER_NAME = USER_.NAME
  and USER_ROLE_MAP.ROLE_ID = ROLE.ID
  and USER_RESOURCE_MAP.RESOURCE_GROUP_ID = RESOURCE_GROUP.ID
  and USER_RESOURCE_MAP.USER_NAME = USER_.NAME;
```

VIRTUAL_FCOE_PORT_INFO

```
create or replace view VIRTUAL_FCOE_PORT_INFO as
select
    VIRTUAL_FCOE_PORT.ID,
    VIRTUAL_FCOE_PORT.VIRTUAL_SWITCH_ID,
    VIRTUAL_FCOE_PORT.PORT_WWN,
    VIRTUAL_FCOE_PORT.PORT_SPEED,
    VIRTUAL_FCOE_PORT.PORT_TYPE,
    VIRTUAL_FCOE_PORT.ENABLED,
    VIRTUAL_FCOE_PORT.STATUS,
    VIRTUAL_FCOE_PORT.TRUNK_INDEX,
    VIRTUAL_FCOE_PORT.PORT_NUMBER,
    VIRTUAL_FCOE_PORT.NAME,
    VIRTUAL_FCOE_PORT.SLOT_NUMBER,
    VIRTUAL_FCOE_PORT.VLAN_ID,
    VIRTUAL_FCOE_PORT.DEVICE_COUNT,
    VIRTUAL_FCOE_PORT.PEER_MAC,
    VIRTUAL_SWITCH.WWN as VIRTUAL_SWITCH_WWN,
    VIRTUAL_SWITCH.ROLE as SWITCH_ROLE,
    VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID as VIRTUAL_FABRIC_ID,
    VIRTUAL_SWITCH.DOMAIN_ID as DOMAIN_ID,
    VIRTUAL_SWITCH.INTEROP_MODE as INTEROP_MODE,
    VIRTUAL_SWITCH.MANAGEMENT_STATE,
    VIRTUAL_SWITCH.MONITORED,
    CORE_SWITCH.TYPE as SWITCH_TYPE,
    CORE_SWITCH.FIRMWARE_VERSION as FIRMWARE_VERSION,
    CORE_SWITCH.IP_ADDRESS as IP_ADDRESS,
    CORE_SWITCH.WWN as PHYSICAL_SWITCH_WWN,
    CORE_SWITCH.MODEL as SWITCH_MODEL,
    CORE_SWITCH_DETAILS.MODEL_NUMBER as SWITCH_MODEL_NUMBER
from
    VIRTUAL_FCOE_PORT, CORE_SWITCH, VIRTUAL_SWITCH, CORE_SWITCH_DETAILS
where
    VIRTUAL_FCOE_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID and
    VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID and
    CORE_SWITCH_DETAILS.CORE_SWITCH_ID = CORE_SWITCH.ID;
```

VIRTUAL_PORT_WWN_DETAILS_INFO

```
create or replace view VIRTUAL_PORT_WWN_DETAILS_INFO as
select distinct
    VIRTUAL_PORT_WWN_DETAILS.SWITCH_ID,
    VIRTUAL_PORT_WWN_DETAILS.SWITCH_PORT_NUMBER,
    VIRTUAL_PORT_WWN_DETAILS.SLOT_NUMBER,
    coalesce(CS1.IP_ADDRESS, CS2.IP_ADDRESS, UDDD.IP_ADDRESS) as IP_ADDRESS,
    coalesce(VS1.NAME, VS2.NAME, UDDD.NAME) as SWITCH_NAME,
    coalesce(VS1.WWN, VS2.WWN) as SWITCH_WWN,
    VIRTUAL_PORT_WWN_DETAILS.AG_NODE_WWN,
    VIRTUAL_PORT_WWN_DETAILS.AG_PORT_NUMBER,
    VIRTUAL_PORT_WWN_DETAILS.STATUS,
    VIRTUAL_PORT_WWN_DETAILS.TYPE,
    VIRTUAL_PORT_WWN_DETAILS.USER_VPWWN,
    VIRTUAL_PORT_WWN_DETAILS.AUTO_VPWWN,
    VIRTUAL_PORT_WWN_DETAILS.DEVICE_PORT_WWN,
```

```
    coalesce(SP1.ID, SP2.ID) as SWITCH_PORT_ID,
    coalesce(SP1.WWN, SP2.WWN) as PORT_WWN,
    coalesce(SP1.TYPE, SP2.TYPE) AS PORT_TYPE,
    coalesce(SP1.NAME, SP2.NAME) as PORT_NAME
from
    VIRTUAL_PORT_WWN_DETAILS
        left outer join VIRTUAL_SWITCH VS1
            on (VIRTUAL_PORT_WWN_DETAILS.AG_PORT_NUMBER = -1
                and VIRTUAL_PORT_WWN_DETAILS.SWITCH_ID = VS1.ID)
        left outer join VIRTUAL_SWITCH VS2
            on VIRTUAL_PORT_WWN_DETAILS.AG_NODE_WWN = VS2.WWN
        left outer join CORE_SWITCH CS1
            on VS1.CORE_SWITCH_ID = CS1.ID
        left outer join CORE_SWITCH CS2
            on VS2.CORE_SWITCH_ID = CS2.ID
        left outer join SWITCH_PORT SP1
            on (SP1.VIRTUAL_SWITCH_ID=VS1.ID
                and VIRTUAL_PORT_WWN_DETAILS.SLOT_NUMBER = SP1.SLOT_NUMBER
                and VIRTUAL_PORT_WWN_DETAILS.SWITCH_PORT_NUMBER = SP1.PORT_NUMBER
                and SP1.TYPE NOT IN ('GigE-Port','TE-Port'))
        left outer join SWITCH_PORT SP2
            on (SP2.VIRTUAL_SWITCH_ID=VS2.ID
                and VIRTUAL_PORT_WWN_DETAILS.AG_PORT_NUMBER = SP2.PORT_NUMBER
                and SP2.TYPE NOT IN ('GigE-Port','TE-Port'))
        left outer join USER_DEFINED_DEVICE_DETAIL UDDD
            on VIRTUAL_PORT_WWN_DETAILS.AG_NODE_WWN = UDDD.WWN;
```

VM_ADDRESS_INFO

```
create or replace view VM_ADDRESS_INFO AS
 select
    DECODE(VM_VIRTUAL_ETHERNET_ADAPTER.MAC_ADDRESS::TEXT, 'HEX'::TEXT) AS MAC_ADDRESS,
     VM_VIRTUAL_MACHINE.NAME AS VM_NAME,
     DECODE(VM_VIRTUAL_MACHINE.IP_ADDRESS::TEXT, 'HEX'::TEXT) AS VM_ADDRESS,
     VM_VCENTER_MEMBER.HOST_NAME AS VM_HOST_NAME,
     DECODE(VM_VIRTUAL_MACHINE.IP_ADDRESS::TEXT, 'HEX'::TEXT) AS VM_HOST_ADDRESS,
     VM_VIRTUAL_ETHERNET_ADAPTER.VIRTUAL_MACHINE_ID AS VM_ID,
     VM_VIRTUAL_MACHINE.HOST_ID AS VM_HOST_ID

 FROM
    VM_VIRTUAL_MACHINE,
    VM_VIRTUAL_ETHERNET_ADAPTER,
    VM_VCENTER_MEMBER

  WHERE
    VM_VIRTUAL_MACHINE.ID = VM_VIRTUAL_ETHERNET_ADAPTER.VIRTUAL_MACHINE_ID
    AND VM_VIRTUAL_MACHINE.HOST_ID = VM_VCENTER_MEMBER.VM_HOST_ID;
```

VLAN_INT_CLASSIFIER_INFO

```
CREATE VIEW vlan_int_classifier_info AS
    select VLAN_INTERFACE_RELATION.VLAN_INTERFACE_RELATION_ID,
       VLAN_INTERFACE_RELATION.VLAN_DB_ID,
       VLAN_INTERFACE_RELATION.INTERFACE_ID,
       VLAN_INT_C_TAG_RELATION.C_TAG_ID,
       MAC_GROUP.NAME,
       MAC_GROUP.MAC_GROUP_ID,
       MAC_GROUP.TYPE,
       MAC_GROUP_MEMBER.MAC_ADDRESS,
       MAC_GROUP_MEMBER.MASK,
       MAC_GROUP.ID AS MAC_GROUP_DB_ID,
```

```
       DEVICE.DEVICE_ID
from INTERFACE ,DEVICE ,PORT_VLAN
        left outer join VLAN_INTERFACE_RELATION  on VLAN_INTERFACE_RELATION.VLAN_DB_ID = PORT_VLAN.VLAN_DB_ID
        left outer join VLAN_INT_MAC_GROUP_RELATION TEMP_MAC_RELATION on
VLAN_INTERFACE_RELATION.VLAN_INTERFACE_RELATION_ID =        TEMP_MAC_RELATION.VLAN_INTERFACE_RELATION_ID
        left outer join VLAN_INT_C_TAG_RELATION on VLAN_INT_C_TAG_RELATION.VLAN_INTERFACE_RELATION_ID =
VLAN_INTERFACE_RELATION.VLAN_INTERFACE_RELATION_ID
        left outer join MAC_GROUP on TEMP_MAC_RELATION.MAC_GROUP_DB_ID =  MAC_GROUP.ID
        left outer join MAC_GROUP_MEMBER on MAC_GROUP_MEMBER.MAC_GROUP_DB_ID =  MAC_GROUP.ID
       left outer join DEVICE_MAC_GROUP_MAPPING  on DEVICE_MAC_GROUP_MAPPING.MAC_GROUP_DB_ID  =  MAC_GROUP.ID
where VLAN_INTERFACE_RELATION.INTERFACE_ID = INTERFACE.INTERFACE_ID and INTERFACE.DEVICE_ID =
DEVICE.DEVICE_ID;
```

VM_CONNECTIVITY_INFO

This view combines fabric and VM information to derive end to end connectivity information for the VM.

```
create or replace view VM_CONNECTIVITY_INFO as
select
  VM_VCENTER.HOST AS VCENTER_HOST,
  DEVICE_PORT.SWITCH_PORT_WWN,
  DEVICE_PORT.DOMAIN_ID,
  device_port.id as device_port_id,
  DEVICE_PORT.NUMBER,
  CORE_SWITCH.IP_ADDRESS,
  CORE_SWITCH.NAME AS CORE_NAME,
  VM_VCENTER.ID AS VCENTER_ID,
  DEVICE_ENCLOSURE.ID AS HOST_DB_ID,
  DEVICE_ENCLOSURE.IP_ADDRESS AS HYPERVISOR_HOST,
  VM_VIRTUAL_MACHINE.ID as VM_ID,
  VM_VIRTUAL_MACHINE.IP_ADDRESS AS VM_IP_ADDRESS,
  VM_VIRTUAL_MACHINE.HOSTNAME AS VM_HOST_NAME,
  VM_VIRTUAL_MACHINE.UUID AS VM_UUID,
  VM_VIRTUAL_MACHINE.NAME AS VM_NAME,
  VM_PATH.NAME AS PATH_NAME,
  VM_PATH.HBA_PORT AS ADAPTER_PORT_WWN,
  VM_PATH.TARGET_PORT AS TARGET_PORT_WWN,
  VM_STORAGE.NAME AS LUN_CAN_NAME,
  VM_PATH.FS_TYPE,
  VM_VIRTUAL_MACHINE.HYPERVISOR_VM_ID,
  DEVICE_ENCLOSURE.MANAGED_ELEMENT_ID AS HOST_ME_ID,
  DEVICE_ENCLOSURE.IP_ADDRESS AS HOST_IP_ADDRESS,
  DEVICE_ENCLOSURE.HOST_NAME AS HYPERVISOR_HOST_NAME,
  FABRIC.NAME AS FABRIC_NAME,
  VIRTUAL_SWITCH.NAME AS VIRTUAL_NAME,
  SWITCH_PORT.STATUS AS SWITCH_PORT_STATUS,
  SWITCH_PORT.ID as SWITCH_PORT_ID,
  SWITCH_PORT.PORT_ID,
  SWITCH_PORT.PORT_NUMBER,
  SWITCH_PORT.SLOT_NUMBER,
  USER_DEFINED_DEVICE_DETAIL.NAME AS ADAPTER_PORT_NAME,
  VM_PATH.FABRIC_ID,
  VM_PATH.VM_PORT_WWN,
  VM_STORAGE.MODEL,
  VM_STORAGE.VENDOR
from
  DEVICE_PORT
    left join USER_DEFINED_DEVICE_DETAIL
      on DEVICE_PORT.WWN = USER_DEFINED_DEVICE_DETAIL.WWN,
  CORE_SWITCH,
  SWITCH_PORT,
  VIRTUAL_SWITCH,
```

```
    DEVICE_NODE,
    FABRIC,
    VM_STORAGE,
    VM_PATH,
    DEVICE_ENCLOSURE,
    VM_VIRTUAL_MACHINE,
    VM_VCENTER,
    VM_DATA_CENTER,
    VM_HOST
  where
    VM_PATH.HBA_PORT = DEVICE_PORT.WWN
    and VM_PATH.VM_ID = VM_VIRTUAL_MACHINE.ID
    and VM_PATH.STORAGE_ID = VM_STORAGE.ID
    and VM_STORAGE.HOST_ID = DEVICE_ENCLOSURE.ID
    and DEVICE_ENCLOSURE.ID = VM_HOST.DEVICE_ENCLOSURE_ID
    and DEVICE_ENCLOSURE.ID = VM_VIRTUAL_MACHINE.HOST_ID
    and VM_HOST.VM_DATACENTER_ID = VM_DATA_CENTER.ID
    and VM_DATA_CENTER.VCENTER_ID = VM_VCENTER.ID
    and DEVICE_PORT.SWITCH_PORT_WWN = SWITCH_PORT.WWN
    and SWITCH_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
    and VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
    and DEVICE_PORT.NODE_ID = DEVICE_NODE.ID
    and DEVICE_NODE.FABRIC_ID = FABRIC.ID

  union all

  select
    VM_VCENTER.HOST AS VCENTER_HOST,
    DEVICE_PORT.SWITCH_PORT_WWN,
    DEVICE_PORT.DOMAIN_ID,
    device_port.id as device_port_id,
    DEVICE_PORT.NUMBER,
    CORE_SWITCH.IP_ADDRESS,
    CORE_SWITCH.NAME as CORE_NAME,
    VM_VCENTER.ID as VCENTER_ID,
    DEVICE_ENCLOSURE.ID AS HOST_DB_ID,
    DEVICE_ENCLOSURE.IP_ADDRESS as HYPERVISOR_HOST,
    VM_VIRTUAL_MACHINE.ID as VM_ID,
    VM_VIRTUAL_MACHINE.IP_ADDRESS AS VM_IP_ADDRESS,
    VM_VIRTUAL_MACHINE.HOSTNAME AS VM_HOST_NAME,
    VM_VIRTUAL_MACHINE.UUID as VM_UUID,
    VM_VIRTUAL_MACHINE.NAME as VM_NAME,
    VM_PATH.NAME as PATH_NAME,
    VM_PATH.HBA_PORT as ADAPTER_PORT_WWN,
    VM_PATH.TARGET_PORT as TARGET_PORT_WWN,
    VM_STORAGE.NAME as LUN_CAN_NAME,
    VM_PATH.FS_TYPE,
    VM_VIRTUAL_MACHINE.HYPERVISOR_VM_ID,
    DEVICE_ENCLOSURE.MANAGED_ELEMENT_ID AS HOST_ME_ID,
    DEVICE_ENCLOSURE.IP_ADDRESS AS HOST_IP_ADDRESS,
    DEVICE_ENCLOSURE.HOST_NAME AS HYPERVISOR_HOST_NAME,
    FABRIC.NAME as FABRIC_NAME,
    VIRTUAL_SWITCH.NAME as VIRTUAL_NAME,
    SWITCH_PORT.STATUS as SWITCH_PORT_STATUS,
    SWITCH_PORT.ID as SWITCH_PORT_ID,
    SWITCH_PORT.PORT_ID,
    SWITCH_PORT.PORT_NUMBER,
    SWITCH_PORT.SLOT_NUMBER,
    USER_DEFINED_DEVICE_DETAIL.NAME as ADAPTER_PORT_NAME,
    VM_PATH.FABRIC_ID,
    VM_PATH.VM_PORT_WWN,
    VM_STORAGE.MODEL,
    VM_STORAGE.VENDOR
```

```
from
  DEVICE_PORT
    LEFT JOIN USER_DEFINED_DEVICE_DETAIL
      on DEVICE_PORT.WWN = USER_DEFINED_DEVICE_DETAIL.WWN,
  CORE_SWITCH,
  SWITCH_PORT,
  VIRTUAL_SWITCH,
  DEVICE_NODE,
  FABRIC,
  DEVICE_PORT_MAC_ADDRESS_MAP,
  GIGE_PORT,
  VM_STORAGE,
  VM_PATH,
  DEVICE_ENCLOSURE,
  VM_VIRTUAL_MACHINE,
  VM_VCENTER,
  VM_DATA_CENTER,
  VM_HOST
where
  VM_PATH.HBA_PORT = DEVICE_PORT.WWN
  and VM_PATH.VM_ID = VM_VIRTUAL_MACHINE.ID
  and VM_PATH.STORAGE_ID = VM_STORAGE.ID
  and VM_STORAGE.HOST_ID = DEVICE_ENCLOSURE.ID
  and DEVICE_ENCLOSURE.ID = VM_HOST.DEVICE_ENCLOSURE_ID
  and DEVICE_ENCLOSURE.ID = VM_VIRTUAL_MACHINE.HOST_ID
  and VM_HOST.VM_DATACENTER_ID = VM_DATA_CENTER.ID
  and VM_DATA_CENTER.VCENTER_ID = VM_VCENTER.ID
  and DEVICE_PORT.ID = DEVICE_PORT_MAC_ADDRESS_MAP.DEVICE_PORT_ID
  and DEVICE_PORT_MAC_ADDRESS_MAP.MAC_ADDRESS::TEXT = GIGE_PORT.REMOTE_MAC_ADDRESS::TEXT
  and GIGE_PORT.SWITCH_PORT_ID = SWITCH_PORT.ID
  and SWITCH_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
  and VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
  and DEVICE_PORT.NODE_ID = DEVICE_NODE.ID
  and DEVICE_NODE.FABRIC_ID = FABRIC.ID;

comment on view VM_CONNECTIVITY_INFO is
'Combine fabric and VM info to derive end to end connectivity information for the VM';
```

VM_NETWORK_CONNECTIVITY_INFO

```
CREATE VIEW vm_network_connectivity_info AS
    select VM_VIRTUAL_ETHERNET_ADAPTER.ID as VNIC_ID,
VM_VIRTUAL_ETHERNET_ADAPTER.MAC_ADDRESS,VM_VIRTUAL_ETHERNET_ADAPTER.IP_ADDRESS as VM_IP_ADDRESS,
VM_VIRTUAL_ETHERNET_ADAPTER.DISPLAY_LABEL,
VM_VIRTUAL_ETHERNET_ADAPTER.PORT_GROUP_NAME, VM_VIRTUAL_MACHINE.ID as VM_ID, VM_VIRTUAL_MACHINE.NAME as
VIRTUAL_MACHINE_NAME, VM_VIRTUAL_MACHINE.HOST_ID as HOST_ID, VM_VCENTER_MEMBER.HOST_NAME,
VM_HOST.CLUSTER_NAME, VM_DATA_CENTER.ID as DATA_CENTER_ID,
VM_VIRTUAL_ETHERNET_ADAPTER.VM_STD_VSWITCH_PORT_GROUP_ID as STD_PORT_GROUP_ID,
VM_VIRTUAL_ETHERNET_ADAPTER.VM_DV_PORT_ID as DV_PORT_ID,
VM_STD_VSWITCH_PORT_GROUP.NAME as UPLINK_PORT_GROUP_NAME, VM_STANDARD_VIRTUAL_SWITCH.NAME as VM_SWITCH_NAME,
VM_PHYSICAL_NIC.MAC_ADDRESS as PNIC_MAC,
VM_HOST_END_DEV_CONNECTIVITY.INTERFACE_ID, INTERFACE.NAME as INTERFACE_NAME, INTERFACE.DEVICE_ID as
SWITCH_ID, DEVICE.IP_ADDRESS as SWITCH_IP, DEVICE.SYS_NAME as SWITCH_NAME,
DEVICE.OPER_STATUS as SWITCH_STATUS, CLUSTER_DEVICE.VCS_LICENSED, PORT_PROFILE.NAME as PORT_PROFILE_NAME,
PROFILE_DOMAINS.DOMAIN_NAMES as PORT_PROFILE_DOMAIN_NAMES, PROFILE_VLAN_MAP.VLAN as PORT_PROFILE_VLAN,
VM_NETWORK_SETTINGS.VLAN_IDS as PORT_GROUP_VLAN

from VM_VIRTUAL_MACHINE, VM_HOST, VM_DATA_CENTER, VM_VCENTER_MEMBER, VM_VIRTUAL_ETHERNET_ADAPTER
join VM_STD_VSWITCH_PORT_GROUP on VM_VIRTUAL_ETHERNET_ADAPTER.VM_STD_VSWITCH_PORT_GROUP_ID =
VM_STD_VSWITCH_PORT_GROUP.ID
left join VM_STANDARD_VIRTUAL_SWITCH on VM_STANDARD_VIRTUAL_SWITCH.ID =
VM_STD_VSWITCH_PORT_GROUP.VM_STANDARD_VIRTUAL__SWITCH_ID
```

```
left join VM_PHYSICAL_NIC on VM_PHYSICAL_NIC.VM_STANDARD_VIRTUAL_SWITCH_ID = VM_STANDARD_VIRTUAL_SWITCH.ID
join VM_HOST_END_DEV_CONNECTIVITY on VM_HOST_END_DEV_CONNECTIVITY.VM_PHYSICAL_NIC_ID = VM_PHYSICAL_NIC.ID
left join INTERFACE on INTERFACE.INTERFACE_ID = VM_HOST_END_DEV_CONNECTIVITY.INTERFACE_ID
left join DEVICE on DEVICE.DEVICE_ID = INTERFACE.DEVICE_ID
left join VCS_CLUSTER_MEMBER on VCS_CLUSTER_MEMBER.MEMBER_ME_ID = DEVICE.MANAGED_ELEMENT_ID
left join DEVICE as CLUSTER_DEVICE on CLUSTER_DEVICE.MANAGED_ELEMENT_ID = VCS_CLUSTER_MEMBER.CLUSTER_ME_ID
left join PORT_PROFILE_INTERFACE_MAP on PORT_PROFILE_INTERFACE_MAP.INTERFACE_ID = INTERFACE.INTERFACE_ID
left join PORT_PROFILE on PORT_PROFILE.ID = PORT_PROFILE_INTERFACE_MAP.PROFILE_ID
left join (select PORT_PROFILE_DOMAIN_MAP.PROFILE_ID, array_to_string(array_agg(PORT_PROFILE_DOMAIN.NAME),
',') DOMAIN_NAMES from PORT_PROFILE_DOMAIN_MAP join PORT_PROFILE_DOMAIN on
PORT_PROFILE_DOMAIN_MAP.PROFILE_DOMAIN_ID = PORT_PROFILE_DOMAIN.ID group by
PORT_PROFILE_DOMAIN_MAP.PROFILE_ID) PROFILE_DOMAINS on PROFILE_DOMAINS.PROFILE_ID = PORT_PROFILE.ID
left join (select PROFILE_ID, array_agg(VLANID)::varchar as VLAN from PORT_PROFILE_VLAN_MAP group by
PROFILE_ID) PROFILE_VLAN_MAP on PROFILE_VLAN_MAP.PROFILE_ID = PORT_PROFILE.ID
left join VM_NETWORK_SETTINGS on VM_NETWORK_SETTINGS.VM_STD_VSWITCH_PORT_GROUP_ID =
VM_VIRTUAL_ETHERNET_ADAPTER.VM_STD_VSWITCH_PORT_GROUP_ID
where VM_VIRTUAL_ETHERNET_ADAPTER.VIRTUAL_MACHINE_ID = VM_VIRTUAL_MACHINE.ID and VM_VIRTUAL_MACHINE.HOST_ID =
VM_VCENTER_MEMBER.VM_HOST_ID and
VM_VIRTUAL_MACHINE.HOST_ID = VM_HOST.DEVICE_ENCLOSURE_ID and VM_HOST.VM_DATACENTER_ID = VM_DATA_CENTER.ID

union

select VNIC_DV_PORT.VNIC_ID as VNIC_ID,
VNIC_DV_PORT.MAC_ADDRESS,VNIC_DV_PORT.VM_IP_ADDRESS,VNIC_DV_PORT.DISPLAY_LABEL,
VNIC_DV_PORT.PORT_GROUP_NAME, VNIC_DV_PORT.VM_ID, VNIC_DV_PORT.VIRTUAL_MACHINE_NAME, VNIC_DV_PORT.HOST_ID,
VNIC_DV_PORT.HOST_NAME, VNIC_DV_PORT.HOST_NAME, VNIC_DV_PORT.DATA_CENTER_ID,
VNIC_DV_PORT.VM_STD_VSWITCH_PORT_GROUP_ID as STD_PORT_GROUP_ID, VNIC_DV_PORT.DV_PORT_ID,
PNIC_DV_PORT.PORT_GROUP_NAME as UPLINK_PORT_GROUP_NAME,
VNIC_DV_PORT.SWITCH_NAME as VM_SWITCH_NAME, PNIC_DV_PORT.PNIC_MAC, VM_HOST_END_DEV_CONNECTIVITY.INTERFACE_ID,
INTERFACE.NAME as INTERFACE_NAME,
INTERFACE.DEVICE_ID as SWITCH_ID, DEVICE.IP_ADDRESS as SWITCH_IP, DEVICE.SYS_NAME as SWITCH_NAME,
DEVICE.OPER_STATUS as SWITCH_STATUS, CLUSTER_DEVICE.VCS_LICENSED, PORT_PROFILE.NAME as PORT_PROFILE_NAME,
PROFILE_DOMAINS.DOMAIN_NAMES as PORT_PROFILE_DOMAIN_NAMES, PROFILE_VLAN_MAP.VLAN as PORT_PROFILE_VLAN,
VM_NETWORK_SETTINGS.VLAN_IDS as PORT_GROUP_VLAN
from

(select VNIC.ID as VNIC_ID, VNIC.MAC_ADDRESS,VNIC.IP_ADDRESS as VM_IP_ADDRESS, VNIC.DISPLAY_LABEL,
VNIC.PORT_GROUP_NAME, VNIC.VM_STD_VSWITCH_PORT_GROUP_ID, VM_VIRTUAL_MACHINE.ID as VM_ID,
VM_VIRTUAL_MACHINE.NAME as VIRTUAL_MACHINE_NAME, VM_VIRTUAL_MACHINE.HOST_ID as HOST_ID,
VM_VCENTER_MEMBER.HOST_NAME, VM_HOST.CLUSTER_NAME, VM_DATA_CENTER.ID as DATA_CENTER_ID,
 DVPORT.ID as DV_PORT_ID, DVPORT.VM_DV_SWITCH_ID, VM_DV_SWITCH.NAME as SWITCH_NAME
from VM_VIRTUAL_MACHINE, VM_VCENTER_MEMBER, VM_HOST, VM_DATA_CENTER, VM_VIRTUAL_ETHERNET_ADAPTER VNIC,
VM_DV_PORT DVPORT, VM_DV_SWITCH
where VNIC.VM_DV_PORT_ID = DVPORT.ID and VNIC.VIRTUAL_MACHINE_ID = VM_VIRTUAL_MACHINE.ID and
VM_VIRTUAL_MACHINE.HOST_ID = VM_VCENTER_MEMBER.VM_HOST_ID and DVPORT.VM_DV_SWITCH_ID = VM_DV_SWITCH.ID and
VM_VIRTUAL_MACHINE.HOST_ID = VM_HOST.DEVICE_ENCLOSURE_ID AND
VM_HOST.VM_DATACENTER_ID = VM_DATA_CENTER.ID) as VNIC_DV_PORT
left join VM_DV_PORT on VM_DV_PORT.ID = VNIC_DV_PORT.DV_PORT_ID
left join VM_NETWORK_SETTINGS on VM_NETWORK_SETTINGS.VM_DV_PORT_GROUP_ID = VM_DV_PORT.VM_DV_PORT_GROUP_ID,

(select DVPORT.VM_DV_SWITCH_ID, DVPORTGROUP.ID as DV_PORT_GROUP_ID, DVPORTGROUP.NAME as PORT_GROUP_NAME,
PNIC.ID as PNIC_ID, PNIC.MAC_ADDRESS as PNIC_MAC
from VM_PHYSICAL_NIC PNIC, VM_DV_PORT DVPORT, VM_DV_PORT_GROUP DVPORTGROUP
where PNIC.VM_DV_PORT_ID = DVPORT.ID and DVPORT.VM_DV_PORT_GROUP_ID = DVPORTGROUP.ID) as PNIC_DV_PORT
join VM_HOST_END_DEV_CONNECTIVITY on VM_HOST_END_DEV_CONNECTIVITY.VM_PHYSICAL_NIC_ID = PNIC_DV_PORT.PNIC_ID
left join INTERFACE on INTERFACE.INTERFACE_ID = VM_HOST_END_DEV_CONNECTIVITY.INTERFACE_ID
left join DEVICE on DEVICE.DEVICE_ID = INTERFACE.DEVICE_ID
left join VCS_CLUSTER_MEMBER on VCS_CLUSTER_MEMBER.MEMBER_ME_ID = DEVICE.MANAGED_ELEMENT_ID
left join DEVICE as CLUSTER_DEVICE on CLUSTER_DEVICE.MANAGED_ELEMENT_ID = VCS_CLUSTER_MEMBER.CLUSTER_ME_ID
left join PORT_PROFILE_INTERFACE_MAP on PORT_PROFILE_INTERFACE_MAP.INTERFACE_ID = INTERFACE.INTERFACE_ID
left join PORT_PROFILE on PORT_PROFILE.ID = PORT_PROFILE_INTERFACE_MAP.PROFILE_ID
```

```
left join (select PORT_PROFILE_DOMAIN_MAP.PROFILE_ID, array_to_string(array_agg(PORT_PROFILE_DOMAIN.NAME),
',') DOMAIN_NAMES from PORT_PROFILE_DOMAIN_MAP join PORT_PROFILE_DOMAIN on
PORT_PROFILE_DOMAIN_MAP.PROFILE_DOMAIN_ID = PORT_PROFILE_DOMAIN.ID group by
PORT_PROFILE_DOMAIN_MAP.PROFILE_ID) PROFILE_DOMAINS on PROFILE_DOMAINS.PROFILE_ID = PORT_PROFILE.ID
left join (select PROFILE_ID, array_agg(VLANID)::varchar as VLAN from PORT_PROFILE_VLAN_MAP group by
PROFILE_ID) PROFILE_VLAN_MAP on PROFILE_VLAN_MAP.PROFILE_ID = PORT_PROFILE.ID

where VNIC_DV_PORT.VM_DV_SWITCH_ID = PNIC_DV_PORT.VM_DV_SWITCH_ID;
```

VM_DATASTORE_DETAILS_INFO

```
create or replace view VM_DATASTORE_DETAILS_INFO as
select vm_virtual_machine_datastore_map.virtual_machine_id,
vm_virtual_machine_datastore_map.vm_datastore_details_id,
vm_datastore_details.datacenter_id, vm_virtual_machine_datastore_map.provisioned_storage,
vm_virtual_machine_datastore_map.not_shared_storage, vm_virtual_machine_datastore_map.used_storage,
vm_datastore_details.name, vm_datastore_details.accessible, vm_datastore_details.status,
vm_datastore_details.file_system_type, vm_datastore_details.total_capacity, vm_datastore_details.free_space,
vm_datastore_details.last_update_time, vm_datastore_details.rdm_supported,
vm_datastore_details.perfile_thin_provisioning_supported, vm_datastore_details.storage_iorm_supported,
vm_datastore_details.directory_hierarchy_supported, vm_datastore_details.location
from vm_virtual_machine_datastore_map, vm_datastore_details
where vm_virtual_machine_datastore_map.vm_datastore_details_id = vm_datastore_details.id;
```

VM_EE_MONITOR_INFO

This view provides combined ee_monitor, ee_monitor_stats, device_port and device_node tables to get the EE Monitor information for vmplug-in.

```
create or replace view VM_EE_MONITOR_INFO as
select distinct
    EE_MONITOR.NAME,
    EE_MONITOR.SWITCH_PORT_ID,
    EE_MONITOR.SOURCE_PORT_ID,
    EE_MONITOR.DEST_PORT_ID,
    EE_MONITOR_STATS.TX,
    EE_MONITOR_STATS.RX,
    EE_MONITOR_STATS.CRCERRORS,
    EE_MONITOR_STATS.CREATION_TIME,
    SOURCE_PORT.PORT_ID as SID,
    DEST_PORT.PORT_ID as DID,
    SOURCE_NODE.WWN as SOURCE_DEVICE_WWN,
    SOURCE_PORT.WWN as SOURCE_PORT_WWN,
    DEST_NODE.WWN as DEST_DEVICE_WWN,
    DEST_PORT.WWN as DEST_PORT_WWN,
    SOURCE_NODE.FABRIC_ID as SOURCE_FABRIC_ID,
    DEST_NODE.FABRIC_ID as DEST_FABRIC_ID,
    SOURCE_PORT.DOMAIN_ID as SOURCE_SWITCH_DOMAIN_ID,
    DEST_PORT.DOMAIN_ID as DEST_SWITCH_DOMAIN_ID,
    VM_VIRTUAL_MACHINE.HYPERVISOR_VM_ID,
    VM_VIRTUAL_MACHINE.NAME as VM_NAME
from
    VM_PATH,
    VM_VIRTUAL_MACHINE,
    DEVICE_PORT as SOURCE_PORT,
    DEVICE_PORT as DEST_PORT,
    DEVICE_NODE as DEST_NODE,
    DEVICE_NODE as SOURCE_NODE,
    EE_MONITOR,
    EE_MONITOR_STATS
```

```
where
    VM_PATH.HBA_PORT::BPCHAR = SOURCE_PORT.WWN
    and VM_PATH.VM_ID = VM_VIRTUAL_MACHINE.ID
    and SOURCE_PORT.ID = EE_MONITOR.SOURCE_PORT_ID
    and EE_MONITOR.ID = EE_MONITOR_STATS.EE_MONITOR_ID
    and SOURCE_PORT.NODE_ID = SOURCE_NODE.ID
    and DEST_PORT.ID = EE_MONITOR.DEST_PORT_ID
    and DEST_PORT.NODE_ID = DEST_NODE.ID
    and EE_MONITOR_STATS.CREATION_TIME in (select MAX(CREATION_TIME) from EE_MONITOR_STATS group by
EE_MONITOR_ID);

comment on view VM_EE_MONITOR_INFO is
'Combined ee_monitor, ee_monitor_stats, device_port and device_node tables to get the EE Monitor info for
vmplug-in';
```

VM_HOST_INFO

```
CREATE VIEW vm_host_info AS
    select vm_data_center.vcenter_
   VM_DATA_CENTER.VCENTER_ID as VCENTER_ID,
   VM_HOST.DEVICE_ENCLOSURE_ID as HOST_ID,
   VM_HOST.VM_DATACENTER_ID as DATACENTER_ID,
   VM_HOST.NODE_WWN         as HOST_NODE_WWN,
   VM_HOST.HYPERVISOR_NAME,
   VM_HOST.HYPERVISOR_TYPE,
   VM_HOST.CPU_COUNT,
   VM_HOST.CPU_TYPE,
   VM_HOST.CPU_RESOURCES    as HOST_CPU_RESOURCES,
   VM_HOST.MEM_RESOURCES    as HOST_MEM_RESOURCES,
   VM_HOST.LICENSE_SERVER,
   VM_HOST.BOOT_TIME        as HOST_BOOT_TIME,
   VM_HOST.CLUSTER_NAME
   VM_VIRTUAL_MACHINE.ID                  as VM_ID,
   VM_VIRTUAL_MACHINE.HYPERVISOR_VM_ID,
   VM_VIRTUAL_MACHINE.NAME                as VM_NAME,
   VM_VIRTUAL_MACHINE.DESCRIPTION         as VM_DESCRIPTION,
   VM_VIRTUAL_MACHINE.OS                  as VM_OS,
   VM_VIRTUAL_MACHINE.STATUS              as VM_STATUS,
   VM_VIRTUAL_MACHINE.VCPU_COUNT,
   VM_VIRTUAL_MACHINE.CPU_RESOURCES       as VM_CPU_RESOURCES,
   VM_VIRTUAL_MACHINE.MEM_RESOURCES       as VM_MEM_RESOURCES,
   VM_VIRTUAL_MACHINE.IP_ADDRESS          as VM_IP_ADDRESS,
   VM_VIRTUAL_MACHINE.HOSTNAME            as VM_HOSTNAME,
   VM_VIRTUAL_MACHINE.BOOT_TIME           as VM_BOOT_TIME,
   VM_VIRTUAL_MACHINE.DATASTORE_NAME,
   VM_VIRTUAL_MACHINE.DATASTORE_LOCATION,
   VM_VIRTUAL_MACHINE.NODE_WWN            as VM_NODE_WWN
vm_virtual_machine.instance_uuid,
    vm_application_details.applications_name
from
   VM_DATA_CENTER,
   VM_HOST
      left join VM_VIRTUAL_MACHINE
         on VM_HOST.DEVICE_ENCLOSURE_ID = VM_VIRTUAL_MACHINE.HOST_ID
LEFT JOIN vm_application_details ON vm_application_details.vm_instance_uuid=vm_virtual_machine.instance_uuid
  WHERE vm_data_center.id = vm_host.vm_datacenter_id;
```

VM_LUN_INFO

```
create or replace view VM_LUN_INFO as
select
```

```
   VM_STORAGE.HOST_ID,
   VM_STORAGE.ID                as LUN_ID,
   VM_STORAGE.NAME              as LUN_NAME,
   VM_STORAGE.TARGET_NODE,
   VM_STORAGE.VENDOR,
   VM_STORAGE.MODEL,
   VM_STORAGE.SERIAL_NUMBER,
   VM_STORAGE.TYPE,
   VM_STORAGE.CAPACITY,
   VM_STORAGE.STATUS            as LUN_STATUS,
   VM_STORAGE.PATH_POLICY,
   VM_STORAGE.ISCSI_TARGET_ADDRESS,
   VM_STORAGE.ISCSI_TARGET_PORT,
   VM_STORAGE.NAS_REMOTE_HOST,
   VM_STORAGE.NAS_REMOTE_PATH,
   VM_PATH.FS_TYPE,
   VM_PATH.ID               as PATH_ID,
   VM_PATH.VM_ID            as PATH_VM_ID,
   VM_PATH.NAME             as PATH_NAME,
   VM_PATH.FABRIC_ID,
   VM_PATH.HBA_PORT,
   VM_PATH.VM_PORT_WWN,
   VM_PATH.TARGET_PORT,
   VM_PATH.HBA_NODE,
   VM_PATH.VM_NODE_WWN,
   VM_PATH.TARGET_NODE      as PATH_TARGET_NODE,
   VM_PATH.HBA_NAME,
   VM_PATH.USAGE            as PATH_USAGE,
   VM_PATH.ENABLED          as PATH_ENABLED,
   VM_PATH.ACTIVE           as PATH_ACTIVE,
   VM_PATH.PREFERRED        as PATH_PREFERRED
from
   VM_STORAGE join VM_PATH on VM_STORAGE.ID = VM_PATH.STORAGE_ID;
```

VM_STATISTICS_INFO

This view gets the FC port statistics for the VM Connectivity data.

```
create or replace view VM_STATISTICS_INFO as
select distinct
    DEVICE_PORT.SWITCH_PORT_WWN,
    DEVICE_PORT.DOMAIN_ID,
    DEVICE_ENCLOSURE.IP_ADDRESS as HYPERVISOR_HOST,
    VM_PATH.HBA_PORT as ADAPTER_PORT_WWN,
    VM_VIRTUAL_MACHINE.HYPERVISOR_VM_ID,
    VM_VIRTUAL_MACHINE.NAME as VM_NAME,
    CORE_SWITCH.IP_ADDRESS,
    CORE_SWITCH.NAME as CORE_NAME,
    FC_PORT_STATS.TX,
    FC_PORT_STATS.RX,
    FC_PORT_STATS.TX_UTILIZATION,
    FC_PORT_STATS.RX_UTILIZATION,
    FC_PORT_STATS.SYNCLOSSES,
    FC_PORT_STATS.SIGNALLOSSES,
    FC_PORT_STATS.SEQUENCEERRORS,
    FC_PORT_STATS.INVALIDTRANSMISSIONS,
    FC_PORT_STATS.CRCERRORS,
    FC_PORT_STATS.CREATION_TIME,
    VIRTUAL_SWITCH.NAME as VIRTUAL_NAME,
    SWITCH_PORT.STATUS as SWITCH_PORT_STATUS,
    SWITCH_PORT.PORT_ID,
    SWITCH_PORT.PORT_NUMBER
```

```
from
    VM_STORAGE,
    VM_HOST,
    DEVICE_ENCLOSURE,
    VM_VIRTUAL_MACHINE,
    VM_PATH,
    DEVICE_PORT,
    SWITCH_PORT,
    CORE_SWITCH,
    FC_PORT_STATS,
    VIRTUAL_SWITCH
where
    VM_PATH.HBA_PORT::BPCHAR = DEVICE_PORT.WWN
    and VM_PATH.VM_ID = VM_VIRTUAL_MACHINE.ID
    and VM_PATH.STORAGE_ID = VM_STORAGE.ID
    and VM_STORAGE.HOST_ID = DEVICE_ENCLOSURE.ID
    and DEVICE_ENCLOSURE.ID = VM_HOST.DEVICE_ENCLOSURE_ID
    and DEVICE_ENCLOSURE.ID = VM_VIRTUAL_MACHINE.HOST_ID
    and DEVICE_PORT.SWITCH_PORT_WWN = SWITCH_PORT.WWN
    and SWITCH_PORT.ID = FC_PORT_STATS.PORT_ID
    and SWITCH_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
    and VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
    and FC_PORT_STATS.CREATION_TIME in (select MAX(CREATION_TIME) from FC_PORT_STATS group by PORT_ID)

union

select
    DEVICE_PORT.SWITCH_PORT_WWN,
    DEVICE_PORT.DOMAIN_ID,
    DEVICE_ENCLOSURE.IP_ADDRESS as HYPERVISOR_HOST,
    VM_PATH.HBA_PORT as ADAPTER_PORT_WWN,
    VM_VIRTUAL_MACHINE.HYPERVISOR_VM_ID,
    VM_VIRTUAL_MACHINE.NAME as VM_NAME,
    CORE_SWITCH.IP_ADDRESS,
    CORE_SWITCH.NAME as CORE_NAME,
    SWITCH_TE_PORT_STATS.TRANSMIT_OK,
    SWITCH_TE_PORT_STATS.RECEIVE_OK,
    SWITCH_TE_PORT_STATS.TRANSMIT_OK_PERCENT_UTIL,
    SWITCH_TE_PORT_STATS.RECEIVE_OK_PERCENT_UTIL,
    (-1) AS SYNCLOSSES,
    (-1) AS SIGNALLOSSES,
    (-1) AS SEQUENCEERRORS,
    (-1) AS INVALIDTRANSMISSIONS,
    (-1) AS CRCERRORS,
    SWITCH_TE_PORT_STATS.CREATION_TIME,
    VIRTUAL_SWITCH.NAME as VIRTUAL_NAME,
    SWITCH_PORT.STATUS as SWITCH_PORT_STATUS,
    SWITCH_PORT.PORT_ID,
    SWITCH_PORT.PORT_NUMBER
from
    VM_STORAGE,
    VM_HOST,
    DEVICE_ENCLOSURE,
    VM_VIRTUAL_MACHINE,
    VM_PATH,
    DEVICE_PORT,
    SWITCH_PORT,
    CORE_SWITCH,
    SWITCH_TE_PORT_STATS,
    VIRTUAL_SWITCH,
    DEVICE_PORT_MAC_ADDRESS_MAP,
    DEVICE_PORT_GIGE_PORT_LINK,
    GIGE_PORT
```

```
where
    VM_PATH.HBA_PORT::BPCHAR = DEVICE_PORT.WWN
    and VM_PATH.VM_ID = VM_VIRTUAL_MACHINE.ID
    and VM_PATH.STORAGE_ID = VM_STORAGE.ID
    and VM_STORAGE.HOST_ID = DEVICE_ENCLOSURE.ID
    and DEVICE_ENCLOSURE.ID = VM_HOST.DEVICE_ENCLOSURE_ID
    and DEVICE_ENCLOSURE.ID = VM_VIRTUAL_MACHINE.HOST_ID
    and DEVICE_PORT.ID = DEVICE_PORT_MAC_ADDRESS_MAP.DEVICE_PORT_ID
    and DEVICE_PORT_MAC_ADDRESS_MAP.MAC_ADDRESS = GIGE_PORT.REMOTE_MAC_ADDRESS
    and GIGE_PORT.SWITCH_PORT_ID = SWITCH_TE_PORT_STATS.PORT_ID
    and GIGE_PORT.SWITCH_PORT_ID = SWITCH_PORT.ID
    and SWITCH_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
    and VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID
    and SWITCH_TE_PORT_STATS.CREATION_TIME in (select max(CREATION_TIME) from SWITCH_TE_PORT_STATS group by
PORT_ID);
```

VR_CONN_MODULE_INFO

```
create or replace view VR_CONN_MODULE_INFO as
select distinct
    VR_CONN_MODULE.ID,
    VR_CONN_MODULE.VR_CONN_DOMAIN_ID,
    VR_CONN_MODULE.VCEM_ASSIGNED_ID,
    VR_CONN_MODULE.WWN,
    VR_CONN_MODULE.PRODUCT_NAME,
    VR_CONN_MODULE.SERIAL_NUMBER,
    VR_CONN_MODULE.STATUS,
    VR_CONN_MODULE.IO_BAY,
    VR_CONN_MODULE.VENDOR,
    VR_CONN_MODULE.CREATION_TIME,
    VR_CONN_MODULE.LAST_UPDATE_TIME,
    VR_CONN_DOMAIN.NAME as DOMAIN_NAME,
    VR_CONN_DOMAIN.GUID as DOMAIN_GUID,
    VR_CONN_DOMAIN.FIRMWARE_VERSION,
    VR_CONN_DOMAIN_GROUP.NAME as DOMAIN_GROUP_NAME,
    VCEM_PROFILE.ID as VCEM_PROFILE_ID,
    VCEM_PROFILE.DISCOVERY_STATUS,
    VCEM_PROFILE.LAST_FAILURE_TIMESTAMP as VCEM_LAST_FAILED_TIME,
    VCEM_PROFILE.LAST_SUCCESSFUL_TIMESTAMP as VCEM_LAST_SUCCESSFUL_TIME,
    VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
    VIRTUAL_SWITCH.MANAGED_ELEMENT_ID as VIRTUAL_SWITCH_ME_ID,
    VIRTUAL_SWITCH.NAME,
    CORE_SWITCH.IP_ADDRESS,
    FABRIC_MEMBER.FABRIC_ID,
    FABRIC.MANAGED as FABRIC_MANAGED
from
    VR_CONN_MODULE
    inner join
        VR_CONN_DOMAIN on
        VR_CONN_DOMAIN.ID = VR_CONN_MODULE.VR_CONN_DOMAIN_ID
    inner join
        VCEM_PROFILE on
        VCEM_PROFILE.ID = VR_CONN_DOMAIN.VCEM_PROFILE_ID
    left outer join
        VR_CONN_DOMAIN_GROUP on
        VR_CONN_DOMAIN_GROUP.ID = VR_CONN_DOMAIN.VR_CONN_DOMAIN_GROUP_ID
    left outer join
        VIRTUAL_SWITCH on
        VIRTUAL_SWITCH.WWN = VR_CONN_MODULE.WWN
    left outer join
        CORE_SWITCH on
        CORE_SWITCH.ID = VIRTUAL_SWITCH.CORE_SWITCH_ID
```

```
        inner join
            FABRIC_MEMBER on
            FABRIC_MEMBER.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
        inner join
            FABRIC on
            FABRIC_MEMBER.FABRIC_ID = FABRIC.ID
    union
    select distinct
        VR_CONN_MODULE.ID,
        VR_CONN_MODULE.VR_CONN_DOMAIN_ID,
        VR_CONN_MODULE.VCEM_ASSIGNED_ID,
        VR_CONN_MODULE.WWN,
        VR_CONN_MODULE.PRODUCT_NAME,
        VR_CONN_MODULE.SERIAL_NUMBER,
        VR_CONN_MODULE.STATUS,
        VR_CONN_MODULE.IO_BAY,
        VR_CONN_MODULE.VENDOR,
        VR_CONN_MODULE.CREATION_TIME,
        VR_CONN_MODULE.LAST_UPDATE_TIME,
        VR_CONN_DOMAIN.NAME as DOMAIN_NAME,
        VR_CONN_DOMAIN.GUID as DOMAIN_GUID,
        VR_CONN_DOMAIN.FIRMWARE_VERSION,
        VR_CONN_DOMAIN_GROUP.NAME as DOMAIN_GROUP_NAME,
        VCEM_PROFILE.ID as VCEM_PROFILE_ID,
        VCEM_PROFILE.DISCOVERY_STATUS,
        VCEM_PROFILE.LAST_FAILURE_TIMESTAMP as VCEM_LAST_FAILED_TIME,
        VCEM_PROFILE.LAST_SUCCESSFUL_TIMESTAMP as VCEM_LAST_SUCCESSFUL_TIME,
        VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
        VIRTUAL_SWITCH.MANAGED_ELEMENT_ID as VIRTUAL_SWITCH_ME_ID,
        VIRTUAL_SWITCH.NAME,
        CORE_SWITCH.IP_ADDRESS,
        DEVICE_NODE.FABRIC_ID,
        FABRIC.MANAGED as FABRIC_MANAGED
    from
        VR_CONN_MODULE
        inner join
            VR_CONN_DOMAIN on
            VR_CONN_DOMAIN.ID = VR_CONN_MODULE.VR_CONN_DOMAIN_ID
        inner join
            VCEM_PROFILE on
            VCEM_PROFILE.ID = VR_CONN_DOMAIN.VCEM_PROFILE_ID
        left outer join
            VR_CONN_DOMAIN_GROUP on
            VR_CONN_DOMAIN_GROUP.ID = VR_CONN_DOMAIN.VR_CONN_DOMAIN_GROUP_ID
        left outer join
            VIRTUAL_SWITCH on
            VIRTUAL_SWITCH.WWN = VR_CONN_MODULE.WWN
        left outer join
            CORE_SWITCH on
            CORE_SWITCH.ID = VIRTUAL_SWITCH.CORE_SWITCH_ID
        left outer join
            DEVICE_NODE on
            DEVICE_NODE.WWN = VR_CONN_MODULE.WWN
        left outer join
            FABRIC on
            DEVICE_NODE.FABRIC_ID = FABRIC.ID;
```

VR_CONN_MODULE_PORT_INFO

```
    create or replace view VR_CONN_MODULE_PORT_INFO as
    select
        VR_CONN_MODULE_PORT.ID,
```

```
    VR_CONN_MODULE_PORT.VR_CONN_MODULE_ID,
    VR_CONN_MODULE_PORT.WWN,
    VR_CONN_MODULE_PORT.POSITION_,
    VR_CONN_MODULE_PORT.FABRIC_NAME,
    VR_CONN_MODULE_PORT.SPEED,
    VR_CONN_MODULE_PORT.STATUS,
    VR_CONN_MODULE_PORT.LAST_STATUS,
    VR_CONN_MODULE_PORT.REMOTE_NODE_WWN,
    VR_CONN_MODULE_PORT.CREATION_TIME,
    VR_CONN_MODULE_PORT.LAST_UPDATE_TIME,
    VR_CONN_MODULE.IO_BAY,
    VR_CONN_DOMAIN.ID as VR_CONN_DOMAIN_ID,
    VCEM_PROFILE.ID as VCEM_PROFILE_ID,
    SWITCH_PORT.ID as SWITCH_PORT_ID,
    VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID
from
    VR_CONN_MODULE_PORT
    inner join
        VR_CONN_MODULE on
        VR_CONN_MODULE.ID = VR_CONN_MODULE_PORT.VR_CONN_MODULE_ID
    inner join
        VR_CONN_DOMAIN on
        VR_CONN_DOMAIN.ID = VR_CONN_MODULE.VR_CONN_DOMAIN_ID
    inner join
        VCEM_PROFILE on
        VCEM_PROFILE.ID = VR_CONN_DOMAIN.VCEM_PROFILE_ID
    left outer join
        SWITCH_PORT on
        SWITCH_PORT.WWN = VR_CONN_MODULE_PORT.WWN
    left outer join
        VIRTUAL_SWITCH on
        VIRTUAL_SWITCH.WWN = VR_CONN_MODULE.WWN;
```

VR_CONN_NPIV_INFO

```
create or replace view VR_CONN_NPIV_INFO as
select
    VR_CONN_WWN.ID,
    VR_CONN_WWN.VR_CONN_FC_CONNECTION_ID,
    VR_CONN_WWN.PORT_ADDRESS as PORT_WWN,
    VR_CONN_WWN.NODE_ADDRESS as NODE_WWN,
    VR_CONN_SERVER_PROFILE.NAME as SERVER_PROFILE_NAME,
    VR_CONN_SERVER_PROFILE.BAY_NAME,
    coalesce(VR_CONN_SERVER_PROFILE.BAY_NUMBER, VR_CONN_FC_CONNECTION.CONNECTION_BAY) as BAY_NUMBER,
    VR_CONN_SERVER_PROFILE.VIRTUAL_SERIAL_NUMBER,
    VCEM_PROFILE.ID as VCEM_PROFILE_ID,
    VR_CONN_DOMAIN.ID as VIRTUAL_CONNECT_DOMAIN_ID,
    VR_CONN_MODULE.ID as VIRTUAL_CONNECT_MODULE_ID,
    VR_CONN_MODULE_PORT.ID as VIRTUAL_CONNECT_MODULE_PORT_ID,
    VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
    coalesce(SWITCH_PORT.WWN, VR_CONN_MODULE_PORT.WWN) as UPLINK_PORT_WWN,
    coalesce(SWITCH_PORT.PORT_NUMBER, VR_CONN_MODULE_PORT.POSITION_) as UPLINK_PORT_NUMBER,
    DEVICE_PORT.ID as DEVICE_PORT_ID,
    DEVICE_PORT.NUMBER as DEVICE_PORT_NUMBER,
    DEVICE_PORT.TYPE as DEVICE_PORT_TYPE,
    DEVICE_NODE.ID as DEVICE_NODE_ID,
    DEVICE_NODE.FABRIC_ID,
    USER_DEFINED_DEVICE_DETAIL.NAME
from
    VR_CONN_WWN
    inner join
        VR_CONN_FC_CONNECTION on
```

```
        VR_CONN_FC_CONNECTION.ID = VR_CONN_WWN.VR_CONN_FC_CONNECTION_ID
    inner join
        VR_CONN_SERVER_PROFILE on
        VR_CONN_SERVER_PROFILE.ID = VR_CONN_FC_CONNECTION.VR_CONN_SERVER_PROFILE_ID
    inner join
        VR_CONN_DOMAIN on
        VR_CONN_DOMAIN.GUID = VR_CONN_SERVER_PROFILE.BAY_ENCLOSURE_UUID
    inner join
        VCEM_PROFILE on
        VCEM_PROFILE.ID = VR_CONN_SERVER_PROFILE.VCEM_PROFILE_ID
    inner join
        VR_CONN_MODULE on
        VR_CONN_MODULE.VR_CONN_DOMAIN_ID = VR_CONN_DOMAIN.ID and
        VR_CONN_MODULE.IO_BAY = VR_CONN_FC_CONNECTION.CONNECTION_BAY
    inner join
        VR_CONN_MODULE_PORT on
        VR_CONN_MODULE_PORT.VR_CONN_MODULE_ID = VR_CONN_MODULE.ID and
        VR_CONN_MODULE_PORT.POSITION_ = VR_CONN_FC_CONNECTION.PORT_NUMBER
    left outer join
        VIRTUAL_SWITCH on
        VIRTUAL_SWITCH.WWN = VR_CONN_MODULE.WWN
    left outer join
        SWITCH_PORT on
        SWITCH_PORT.WWN = VR_CONN_MODULE_PORT.WWN
    left outer join
        DEVICE_PORT on
        DEVICE_PORT.WWN = VR_CONN_WWN.PORT_ADDRESS
    left outer join
        DEVICE_NODE on
        DEVICE_NODE.WWN = VR_CONN_WWN.NODE_ADDRESS
    left outer join
        USER_DEFINED_DEVICE_DETAIL on
        USER_DEFINED_DEVICE_DETAIL.WWN = VR_CONN_WWN.PORT_ADDRESS;
```

VMM_DISCOVERED_MAC_INFO

```
create or replace view VMM_DISCOVERED_MAC_INFO AS
select
      VM_VIRTUAL_ETHERNET_ADAPTER.MAC_ADDRESS,
      VM_VIRTUAL_ETHERNET_ADAPTER.DISPLAY_LABEL,
      VM_VIRTUAL_ETHERNET_ADAPTER.PORT_GROUP_NAME,
      VM_VIRTUAL_MACHINE.NAME AS VIRTUAL_MACHINE_NAME,
      VM_VCENTER_MEMBER.HOST_NAME,
      VM_VCENTER.NAME AS VCENTER_NAME
from
      VM_VIRTUAL_ETHERNET_ADAPTER,
      VM_VIRTUAL_MACHINE,
      VM_VCENTER_MEMBER,
      VM_VCENTER
where
     VM_VIRTUAL_ETHERNET_ADAPTER.VIRTUAL_MACHINE_ID = VM_VIRTUAL_MACHINE.ID
     AND VM_VIRTUAL_MACHINE.HOST_ID = VM_VCENTER_MEMBER.VM_HOST_ID
     AND VM_VCENTER_MEMBER.VM_VCENTER_ID = VM_VCENTER.ID

union all
select
     VM_HOST_VIRTUAL_NIC.MAC AS MAC_ADDRESS,
     VM_HOST_VIRTUAL_NIC.DEVICE_NAME AS DISPLAY_LABEL,
     VM_HOST_VIRTUAL_NIC.PORT_GROUP_KEY AS PORT_GROUP_NAME,
     NULL::UNKNOWN AS VIRTUAL_MACHINE_NAME,
     VM_VCENTER_MEMBER.HOST_NAME,
     VM_VCENTER.NAME AS VCENTER_NAME
```

```
from
     VM_HOST_VIRTUAL_NIC,
     VM_VCENTER_MEMBER,
     VM_VCENTER
where
     VM_HOST_VIRTUAL_NIC.VM_HOST_ID = VM_VCENTER_MEMBER.VM_HOST_ID AND VM_VCENTER_MEMBER.VM_VCENTER_ID =
VM_VCENTER.ID

union all
select
      VM_PHYSICAL_NIC.MAC_ADDRESS,
      VM_PHYSICAL_NIC.DEVICE_NAME,
      NULL::UNKNOWN AS PORT_GROUP_NAME,
      NULL::UNKNOWN AS VIRTUAL_MACHINE_NAME,
      VM_VCENTER_MEMBER.HOST_NAME,
      VM_VCENTER.NAME AS VCENTER_NAME
from
      VM_PHYSICAL_NIC,
      VM_VCENTER_MEMBER,
      VM_VCENTER
where
      VM_PHYSICAL_NIC.VM_HOST_ID = VM_VCENTER_MEMBER.VM_HOST_ID AND VM_VCENTER_MEMBER.VM_VCENTER_ID =
VM_VCENTER.ID;
```

VM_VIRTUAL_ETHERNET_ADAPTER_INFO

```
create or replace view VM_VIRTUAL_ETHERNET_ADAPTER_INFO as
select
VM_VIRTUAL_ETHERNET_ADAPTER.ID,
VM_VIRTUAL_ETHERNET_ADAPTER.MAC_ADDRESS,
VM_VIRTUAL_ETHERNET_ADAPTER.DISPLAY_LABEL,
VM_VIRTUAL_ETHERNET_ADAPTER.PORT_GROUP_NAME,
VM_VIRTUAL_MACHINE.NAME as VIRTUAL_MACHINE_NAME,
VM_VCENTER_MEMBER.HOST_NAME,
VM_VCENTER.NAME as VCENTER_NAME

From
VM_VIRTUAL_ETHERNET_ADAPTER,
VM_VIRTUAL_MACHINE,
VM_VCENTER_MEMBER,
VM_VCENTER

Where
VM_VIRTUAL_ETHERNET_ADAPTER.VIRTUAL_MACHINE_ID = VM_VIRTUAL_MACHINE.ID
And  VM_VIRTUAL_MACHINE.HOST_ID = VM_VCENTER_MEMBER.VM_HOST_ID
And  VM_VCENTER_MEMBER.VM_VCENTER_ID = VM_VCENTER.ID;
```

ZONE_DB_INFO

```
create or replace view ZONE_DB_INFO as
select
    ZONE_DB.ID,
    ZONE_DB.FABRIC_ID,
    ZONE_DB.OFFLINE,
    ZONE_DB.NAME,
    ZONE_DB.CREATED,
    ZONE_DB.CREATED_BY,
    ZONE_DB.LAST_MODIFIED,
    ZONE_DB.LAST_MODIFIED_BY,
    ZONE_DB.LAST_APPLIED,
    ZONE_DB.LAST_APPLIED_BY,
```

```
    ZONE_DB.DEFAULT_ZONE_STATUS,
    ZONE_DB.MCDATA_DEFAULT_ZONE,
    ZONE_DB.MCDATA_SAFE_ZONE,
    ZONE_DB.ZONE_TXN_SUPPORTED,
    ZONE_DB.ZONE_CONFIG_SIZE,
    ZONE_DB.ZONE_AVAILABLE_SIZE,
    ZONE_DB_CONFIG.ID AS CONFIG_ID,
    ZONE_DB_CONFIG.DEFINED_CONTENT,
    ZONE_DB_CONFIG.ACTIVE_CONTENT,
    ZONE_DB_CONFIG.TI_ZONE_CONTENT
from
    ZONE_DB, ZONE_DB_CONFIG
where
    ZONE_DB.ID = ZONE_DB_CONFIG.ZONE_DB_ID;
```

ZONE_DB_REPORT_INFO

```
CREATE OR REPLACE VIEW zone_db_report_info AS
 SELECT zone_db.id, zone_db.fabric_id, zone_db.offline, zone_db.name,
    zone_db.created, zone_db.created_by, zone_db.last_modified,
    zone_db.last_modified_by, zone_db.last_applied, zone_db.last_applied_by,
    zone_db.default_zone_status, zone_db.mcdata_default_zone,
    zone_db.mcdata_safe_zone, zone_db.zone_txn_supported,
    zone_db.zone_config_size, zone_db.txn_status, zone_db.zone_available_size,
    zone_db_config.id AS config_id, zone_db_config.defined_content,
    zone_db_config.active_content, zone_db_config.ti_zone_content,
    virtual_switch.name AS seed_vs_name, virtual_switch.wwn AS vs_wwn,
    virtual_switch.id AS vs_id, core_switch.name AS cs_name,
    core_switch.ip_address AS seed_cs_switch_ip, core_switch.id AS cs_id,
    core_switch.wwn AS cs_wwn, fabric.name AS fabric_name,
    fabric.status AS fabric_status
   FROM zone_db
   JOIN zone_db_config ON zone_db.id = zone_db_config.zone_db_id
   JOIN fabric ON zone_db.fabric_id = fabric.id
   JOIN virtual_switch ON fabric.seed_switch_wwn = virtual_switch.wwn
   JOIN core_switch ON virtual_switch.core_switch_id = core_switch.id
  WHERE fabric.type !=  4 AND fabric.managed = 1;
```

AP_USAGE

```
CREATE VIEW ap_usage AS
    SELECT ap_station.device_id, ap_station.time_stamp, count(*) AS num_clients FROM ap_station WHERE
(ap_station.radio > 0) GROUP BY ap_station.device_id, ap_station.time_stamp;
```

EVENTS

```
CREATE VIEW events AS
    SELECT emain.trap_log_id, emain.trap_sender, emain."timestamp", emain.severity, emsgs.messages,
emain.is_ack, emain.log_type, emain.slot, emain.port, emain.device_id, emain.event_action_id,
emain.device_group_id, emain.port_group_id, emain.trap_device_ip, emain.log_sub_type, emain.unit FROM
(events_main emain LEFT JOIN events_messages emsgs ON ((emain.messages_id = emsgs.messages_id)));
```

SFLOW_MINUTE_BGP_VIEW

```
create or replace view SFLOW_MINUTE_BGP_VIEW as
  select DEVICE_ID, TIME_IN_SECONDS, SRC_AS, SFLOW_IP_ROUTE_INFO_ID, IN_VLAN, OUT_VLAN, FRAMES, BYTES
  from SFLOW_MINUTE_BGP
  where SLNUM <= (select MAX_SLNUM from SFLOW_MINUTE_BGP_SLNUM fetch first 1 rows only)
  union all
```

```
  select DEVICE_ID, TIME_IN_SECONDS, SRC_AS, SFLOW_IP_ROUTE_INFO_ID, IN_VLAN, OUT_VLAN, FRAMES, BYTES
  from SFLOW_STAGING
  where SLNUM >= (select MIN_SLNUM from SFLOW_STAGING_SLNUM fetch first 1 rows only)
  and SRC_AS != 0 OR SFLOW_IP_ROUTE_INFO_ID != 0;
```

SFLOW_MINUTE_VLAN_VIEW

```
create or replace view SFLOW_MINUTE_VLAN_VIEW as
  select DEVICE_ID, TIME_IN_SECONDS, IN_VLAN, OUT_VLAN, FRAMES, BYTES
  from SFLOW_MINUTE_VLAN
  where SLNUM <= (select MAX_SLNUM from SFLOW_MINUTE_VLAN_SLNUM fetch first 1 rows only)
  union all
  select DEVICE_ID, TIME_IN_SECONDS, IN_VLAN, OUT_VLAN, FRAMES, BYTES
  from SFLOW_STAGING
  where SLNUM >= (select MIN_SLNUM from SFLOW_STAGING_SLNUM fetch first 1 rows only);
```

PHYSICAL_DEVICE_INFO

```
create or replace view PHYSICAL_DEVICE_INFO as
select
    PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID as PD_PHYSICAL_DEVICE_ID,
    PHYSICAL_DEVICE.DEVICE_ID,
    PHYSICAL_DEVICE.DESCRIPTION,
    PHYSICAL_DEVICE.NUM_SLOTS,
    PHYSICAL_DEVICE.TABLE_SUBTYPE,
    PHYSICAL_DEVICE.UNIT_NUMBER,
    PHYSICAL_DEVICE.UNIT_NEIGHBOR1,
    PHYSICAL_DEVICE.UNIT_NEIGHBOR2,
    PHYSICAL_DEVICE.UNIT_PRESENT,
    FOUNDRY_PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID as FPD_PHYSICAL_DEVICE_ID,
    FOUNDRY_PHYSICAL_DEVICE.SERIAL_NUMBER,
    FOUNDRY_PHYSICAL_DEVICE.PRODUCT_TYPE,
    DEVICE.IP_ADDRESS
from
    PHYSICAL_DEVICE,
    FOUNDRY_PHYSICAL_DEVICE,
    DEVICE
where
    DEVICE.DEVICE_ID = PHYSICAL_DEVICE.DEVICE_ID
    and PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID = FOUNDRY_PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID;
```

SLOT_INFO

```
create or replace view SLOT_INFO as
select
    SLOT.*,
    PHYSICAL_DEVICE.UNIT_NUMBER,
    DEVICE.IP_ADDRESS
from
    PHYSICAL_DEVICE,
    SLOT,
    DEVICE
where
    DEVICE.DEVICE_ID = PHYSICAL_DEVICE.DEVICE_ID
    and SLOT.PHYSICAL_DEVICE_ID = PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID;
```

MANAGED_ELEMENT_INFO

Common managed element data used by custom DTO methods to identify the managed element type, and provide a link to the details table for the managed element. Some common managed element fields are included in this view so Fault Management can use this view to identify the managed element ID for an event source.

```
create or replace view MANAGED_ELEMENT_INFO as
select
    MANAGED_ELEMENT.ID as MANAGED_ELEMENT_ID,
    DEVICE.DEVICE_ID as IP_DEVICE_ID,
    coalesce(CS_ME.ID, CS_VS.ID) as CORE_SWITCH_ID,
    VIRTUAL_SWITCH.ID as VIRTUAL_SWITCH_ID,
    DEVICE_ENCLOSURE.ID as DEVICE_ENCLOSURE_ID,
    DEVICE.IP_ADDRESS as LAN_IP_ADDRESS,
    coalesce (CS_VS.IP_ADDRESS, CS_ME.IP_ADDRESS, DEVICE_ENCLOSURE.IP_ADDRESS) as SAN_IP_ADDRESS,
    VIRTUAL_SWITCH.VIRTUAL_FABRIC_ID,
    coalesce (VIRTUAL_SWITCH.WWN, CS_ME.WWN, DEVICE.NODE_WWN) as NODE_WWN
from
    MANAGED_ELEMENT
        left outer join VIRTUAL_SWITCH on MANAGED_ELEMENT.ID = VIRTUAL_SWITCH.MANAGED_ELEMENT_ID
        left outer join CORE_SWITCH CS_ME on (MANAGED_ELEMENT.ID = CS_ME.MANAGED_ELEMENT_ID)
        left outer join CORE_SWITCH CS_VS on (CS_VS.ID = VIRTUAL_SWITCH.CORE_SWITCH_ID)
        left outer join DEVICE on MANAGED_ELEMENT.ID = DEVICE.MANAGED_ELEMENT_ID
        left outer join DEVICE_ENCLOSURE on MANAGED_ELEMENT.ID = DEVICE_ENCLOSURE.MANAGED_ELEMENT_ID;
```

SNMP_DATA_INFO

```
create or replace view SNMP_DATA_INFO as
select * from SNMP_DATA
union all
select * from SNMP_DATA_30MIN
union all
select * from SNMP_DATA_2HOUR
union all
select * from SNMP_DATA_1DAY;
```

SNMP_EXPR_DATA_INFO

```
create or replace view SNMP_EXPR_DATA_INFO as
select * from SNMP_EXPR_DATA
union all
select * from SNMP_EXPR_DATA_30MIN
union all
select * from SNMP_EXPR_DATA_2HOUR
union all
select * from SNMP_EXPR_DATA_1DAY;
```

SNMP_DATA_VIEW

```
create or replace view snmp_data_view as
        (       (       (       (         SELECT de.device_id, de.ip_address AS device_ip, se.target_type,
de.device_id AS target_id, de.sys_name AS target_name, 1 AS collectible_type, se.expression_id AS
collectible_id, se.collector_id, ( SELECT perf_collector.name AS collector_name
                                                    FROM perf_collector
                                                   WHERE perf_collector.collector_id = se.collector_id) AS
collector_name, ( SELECT snmp_expression.name AS collectible_name
                                                    FROM snmp_expression
                                                  WHERE snmp_expression.expression_id = se.expression_id) AS
collectible_name, ( SELECT snmp_expression.equation AS collectible_detail
```

```
                                                    FROM snmp_expression
                                                   WHERE snmp_expression.expression_id = se.expression_id) AS
collectible_detail, se.value, se.time_in_seconds, '' AS mib_index
                                            FROM snmp_expr_data_info se
                                      JOIN device de ON se.target_id = de.device_id
                                     WHERE se.target_type = 0
                                UNION ALL
                                    SELECT de.device_id, de.ip_address AS device_ip, sd.target_type,
de.device_id AS target_id, de.sys_name AS target_name, 0 AS collectible_type, sd.mib_object_id AS
collectible_id, sd.collector_id, ( SELECT perf_collector.name AS collector_name
                                                    FROM perf_collector
                                                   WHERE perf_collector.collector_id = sd.collector_id) AS
collector_name, ( SELECT (mib_object.name::text || '.'::text) || sd.mib_index::text AS collectible_name
                                                    FROM mib_object
                                                   WHERE mib_object.mib_object_id = sd.mib_object_id) AS
collectible_name, ( SELECT mib_object.oid AS collectible_detail
                                                    FROM mib_object
                                                   WHERE mib_object.mib_object_id = sd.mib_object_id) AS
collectible_detail, sd.value, sd.time_in_seconds, sd.mib_index
                                            FROM snmp_data_info sd
                                      JOIN device de ON sd.target_id = de.device_id
                                     WHERE sd.target_type = 0::numeric)
                        UNION ALL
                                 SELECT de.device_id, de.ip_address AS device_ip, sd.target_type,
ifs.interface_id AS target_id, ifs.if_name AS target_name, 0 AS collectible_type, sd.mib_object_id AS
collectible_id, sd.collector_id, ( SELECT perf_collector.name AS collector_name
                                            FROM perf_collector
                                      WHERE perf_collector.collector_id = sd.collector_id) AS collector_name,
( SELECT (mib_object.name::text || '.'::text) || sd.mib_index::text AS collectible_name
                                            FROM mib_object
                                      WHERE mib_object.mib_object_id = sd.mib_object_id) AS collectible_name,
( SELECT mib_object.oid AS collectible_detail
                                            FROM mib_object
                                           WHERE mib_object.mib_object_id = sd.mib_object_id) AS
collectible_detail, sd.value, sd.time_in_seconds, sd.mib_index
                                     FROM snmp_data_info sd
                                JOIN interface ifs ON sd.target_type = 1::numeric AND sd.target_id =
ifs.interface_id
                          JOIN device de ON ifs.device_id = de.device_id)
                UNION ALL
                         SELECT de.device_id, de.ip_address AS device_ip, se.target_type, ifs.interface_id AS
target_id, ifs.if_name AS target_name, 1 AS collectible_type, se.expression_id AS collectible_id,
se.collector_id, ( SELECT perf_collector.name AS collector_name
                                      FROM perf_collector
                                     WHERE perf_collector.collector_id = se.collector_id) AS collector_name, (
SELECT snmp_expression.name AS collectible_name
                                      FROM snmp_expression
                                    WHERE snmp_expression.expression_id = se.expression_id) AS collectible_name,
( SELECT snmp_expression.equation AS collectible_detail
                                      FROM snmp_expression
                                   WHERE snmp_expression.expression_id = se.expression_id) AS collectible_detail,
se.value, se.time_in_seconds, '' AS mib_index
                             FROM snmp_expr_data_info se
                        JOIN interface ifs ON se.target_type = 1 AND se.target_id = ifs.interface_id
                   JOIN device de ON ifs.device_id = de.device_id)
        UNION ALL
                  SELECT de.device_id, de.ip_address AS device_ip, sd.target_type, sp.id AS target_id, sp.name
AS target_name, 0 AS collectible_type, sd.mib_object_id AS collectible_id, sd.collector_id, ( SELECT
perf_collector.name AS collector_name
                             FROM perf_collector
                            WHERE perf_collector.collector_id = sd.collector_id) AS collector_name, ( SELECT
(mib_object.name::text || '.'::text) || sd.mib_index::text AS collectible_name
                             FROM mib_object
```

```
                          WHERE mib_object.mib_object_id = sd.mib_object_id) AS collectible_name, ( SELECT
mib_object.oid AS collectible_detail
                           FROM mib_object
                         WHERE mib_object.mib_object_id = sd.mib_object_id) AS collectible_detail, sd.value,
sd.time_in_seconds, sd.mib_index
                  FROM snmp_data_info sd
              JOIN switch_port sp ON sd.target_type = 4::numeric AND sd.target_id = sp.id
         JOIN device de ON (( SELECT sw.managed_element_id
                 FROM virtual_switch sw
                WHERE sw.id = sp.virtual_switch_id)) = de.managed_element_id)
UNION ALL
         SELECT de.device_id, de.ip_address AS device_ip, se.target_type, sp.id AS target_id, sp.name AS
target_name, 1 AS collectible_type, se.expression_id AS collectible_id, se.collector_id, ( SELECT
perf_collector.name AS collector_name
                  FROM perf_collector
                 WHERE perf_collector.collector_id = se.collector_id) AS collector_name, ( SELECT
snmp_expression.name AS collectible_name
                  FROM snmp_expression
                 WHERE snmp_expression.expression_id = se.expression_id) AS collectible_name, ( SELECT
snmp_expression.equation AS collectible_detail
                  FROM snmp_expression
                 WHERE snmp_expression.expression_id = se.expression_id) AS collectible_detail, se.value,
se.time_in_seconds, '' AS mib_index
           FROM snmp_expr_data_info se
      JOIN switch_port sp ON se.target_type = 4 AND se.target_id = sp.id
   JOIN device de ON (( SELECT sw.managed_element_id
          FROM virtual_switch sw
         WHERE sw.id = sp.virtual_switch_id)) = de.managed_element_id;
```

SUBFLOW_DETAILS_INFO

```
CREATE OR REPLACE VIEW subflow_details_info AS
 SELECT nsf.id AS sub_flow_id, nfd.virtual_switch_id AS amp_vs_id,
    nfd.name AS flow_name, nfd.id AS flow_definition_id,
    scs.id AS src_core_switch_id, scs.ip_address AS src_ip_address,
    svs.id AS src_switch_id, svs.name AS src_switch_name,
    svs.wwn AS src_switch_wwn, svs.domain_id AS src_domain_id,
    ssp.id AS src_switch_port_id, ssp.wwn AS src_switch_port_wwn,
    ssp.name AS src_switch_port_name, ssp.slot_number AS src_switch_slot_number,
    ssp.port_number AS src_switch_port_number,
    ssp.user_port_number AS src_switch_user_port_number,
    ssp.port_id AS src_switch_port_fc_address, sdn.id AS src_device_node_id,
    sdn.wwn AS src_device_node_wwn, sdn.type AS src_device_node_type,
    sdp.id AS src_device_port_id, sdp.wwn AS src_device_port_wwn,
    sdp.port_id AS src_device_fcaddress, sdp.number AS src_device_port_number,
    nsf.lunid, dcs.id AS dst_core_switch_id, dcs.ip_address AS dst_ip_address,
    dvs.id AS dst_switch_id, dvs.name AS dst_switch_name,
    dvs.wwn AS dst_switch_wwn, dvs.domain_id AS dst_domain_id,
    dsp.id AS dst_switch_port_id, dsp.wwn AS dst_switch_port_wwn,
    dsp.name AS dst_switch_port_name, dsp.slot_number AS dst_switch_slot_number,
    dsp.port_number AS dst_switch_port_number,
    dsp.user_port_number AS dst_switch_user_port_number,
    dsp.port_id AS dst_switch_port_fc_address, ddn.id AS dst_device_node_id,
    ddn.wwn AS dst_device_node_wwn, ddn.type AS dst_device_node_type,
    ddp.id AS dst_device_port_id, ddp.wwn AS dst_device_port_wwn,
    ddp.port_id AS dst_device_fcaddress, ddp.number AS dst_device_port_number,
    nsf.sub_flow_origin, nsf.key AS sub_flow_key, nsf.is_missing,
    ssp.port_index AS src_switch_port_index,
    dsp.port_index AS dst_switch_port_index
   FROM np_sub_flow nsf
   JOIN np_flow_definition nfd ON nsf.flow_definition_id = nfd.id
   JOIN virtual_switch svs ON nsf.src_virtual_switch_id = svs.id
```

```
  JOIN core_switch scs ON svs.core_switch_id = scs.id
  JOIN switch_port ssp ON nsf.src_switch_port_id = ssp.id
  JOIN device_port sdp ON nsf.src_device_port_id = sdp.id
  JOIN device_node sdn ON nsf.src_device_node_id = sdn.id
  JOIN virtual_switch dvs ON nsf.dst_virtual_switch_id = dvs.id
  JOIN core_switch dcs ON dvs.core_switch_id = dcs.id
  JOIN switch_port dsp ON nsf.dst_switch_port_id = dsp.id
  JOIN device_port ddp ON nsf.dst_device_port_id = ddp.id
  JOIN device_node ddn ON nsf.dst_device_node_id = ddn.id
 WHERE nsf.feature = 5;
```

VM_VNETWORK_INFO

This view provideds combine VM and device information to derive VM to the ingress switch port information.

```
create or replace view VM_VNETWORK_INFO as
select
    VM_HOST.HYPERVISOR_NAME as VHOST_NAME,
    VM_VIRTUAL_MACHINE.NAME as VM_NAME,
    VM_VIRTUAL_MACHINE.HYPERVISOR_VM_ID as VM_ID,
    VM_VIRTUAL_ETHERNET_ADAPTER.MAC_ADDRESS as VNIC_MAC,
    VM_DV_PORT_GROUP.NAME as PGRP_NAME,
    VM_DV_SWITCH.NAME as VSWITCH_NAME,
    VNIC_DV_PORT.NAME as DVPORT_NAME,
    VM_PHYSICAL_NIC.DEVICE_NAME as PNIC_NAME,
    VM_PHYSICAL_NIC.MAC_ADDRESS as PNIC_MAC,
    DEVICE.SYS_NAME as SWITCH_NAME,
    DEVICE.IP_ADDRESS as SWITCH_IP,
    PHYSICAL_PORT.PORT_NUM as SWITCH_PORT,
    INTERFACE.PORT_STATUS as SWITCH_PORT_STATUS
from
    VM_HOST
        left join VM_VIRTUAL_MACHINE on VM_HOST.DEVICE_ENCLOSURE_ID = VM_VIRTUAL_MACHINE.HOST_ID,
    VM_VIRTUAL_ETHERNET_ADAPTER,
    VM_DV_PORT VNIC_DV_PORT,
    VM_DV_PORT PNIC_DV_PORT,
    VM_DV_PORT_GROUP,
    VM_DV_SWITCH,
    VM_PHYSICAL_NIC,
    VM_HOST_END_DEV_CONNECTIVITY,
    INTERFACE,
    DEVICE,
    PHYSICAL_INTERFACE,
    PHYSICAL_PORT
where
    VM_VIRTUAL_MACHINE.ID = VM_VIRTUAL_ETHERNET_ADAPTER.VIRTUAL_MACHINE_ID and
    VM_VIRTUAL_ETHERNET_ADAPTER.VM_DV_PORT_ID is not null and
    VM_VIRTUAL_ETHERNET_ADAPTER.VM_DV_PORT_ID = VNIC_DV_PORT.ID and
    VNIC_DV_PORT.VM_DV_PORT_GROUP_ID = VM_DV_PORT_GROUP.ID and
    VNIC_DV_PORT.VM_DV_SWITCH_ID = VM_DV_SWITCH.ID and
    PNIC_DV_PORT.VM_DV_SWITCH_ID = VM_DV_SWITCH.ID and
    PNIC_DV_PORT.ID = VM_PHYSICAL_NIC.VM_DV_PORT_ID and
    VM_PHYSICAL_NIC.ID = VM_HOST_END_DEV_CONNECTIVITY.VM_PHYSICAL_NIC_ID and
    VM_HOST_END_DEV_CONNECTIVITY.INTERFACE_ID = INTERFACE.INTERFACE_ID and
    INTERFACE.DEVICE_ID = DEVICE.DEVICE_ID and
    VM_HOST_END_DEV_CONNECTIVITY.INTERFACE_ID = PHYSICAL_INTERFACE.INTERFACE_ID and
    PHYSICAL_INTERFACE.PHYSICAL_PORT_ID = PHYSICAL_PORT.PHYSICAL_PORT_ID

union all

select
    VM_HOST.HYPERVISOR_NAME as VHOST_NAME,
```

```
    VM_VIRTUAL_MACHINE.NAME as VM_NAME,
    VM_VIRTUAL_MACHINE.HYPERVISOR_VM_ID as VM_ID,
    VM_VIRTUAL_ETHERNET_ADAPTER.MAC_ADDRESS as VNIC_MAC,
    VM_STD_VSWITCH_PORT_GROUP.NAME as PGRP_NAME,
    VM_STANDARD_VIRTUAL_SWITCH.NAME as VSWITCH_NAME,
    null as DVPORT_NAME,
    VM_PHYSICAL_NIC.DEVICE_NAME as PNIC_NAME,
    VM_PHYSICAL_NIC.MAC_ADDRESS as PNIC_MAC,
    DEVICE.SYS_NAME as SWITCH_NAME,
    DEVICE.IP_ADDRESS as SWITCH_IP,
    PHYSICAL_PORT.PORT_NUM as SWITCH_PORT,
    INTERFACE.PORT_STATUS as SWITCH_PORT_STATUS
from
    VM_HOST
        left join VM_VIRTUAL_MACHINE on VM_HOST.DEVICE_ENCLOSURE_ID = VM_VIRTUAL_MACHINE.HOST_ID,
    VM_VIRTUAL_ETHERNET_ADAPTER,
    VM_STD_VSWITCH_PORT_GROUP,
    VM_STANDARD_VIRTUAL_SWITCH,
    VM_PHYSICAL_NIC,
    VM_HOST_END_DEV_CONNECTIVITY,
    INTERFACE,
    DEVICE,
    PHYSICAL_INTERFACE,
    PHYSICAL_PORT
where
    VM_VIRTUAL_MACHINE.ID = VM_VIRTUAL_ETHERNET_ADAPTER.VIRTUAL_MACHINE_ID and
    VM_VIRTUAL_ETHERNET_ADAPTER.VM_STD_VSWITCH_PORT_GROUP_ID is not null and
    VM_VIRTUAL_ETHERNET_ADAPTER.VM_STD_VSWITCH_PORT_GROUP_ID = VM_STD_VSWITCH_PORT_GROUP.ID and
    VM_STD_VSWITCH_PORT_GROUP.VM_STANDARD_VIRTUAL__SWITCH_ID = VM_STANDARD_VIRTUAL_SWITCH.ID and
    VM_STANDARD_VIRTUAL_SWITCH.ID = VM_PHYSICAL_NIC.VM_STANDARD_VIRTUAL_SWITCH_ID and
    VM_PHYSICAL_NIC.ID = VM_HOST_END_DEV_CONNECTIVITY.VM_PHYSICAL_NIC_ID and
    VM_HOST_END_DEV_CONNECTIVITY.INTERFACE_ID = INTERFACE.INTERFACE_ID and
    INTERFACE.DEVICE_ID = DEVICE.DEVICE_ID and
    VM_HOST_END_DEV_CONNECTIVITY.INTERFACE_ID = PHYSICAL_INTERFACE.INTERFACE_ID and
    PHYSICAL_INTERFACE.PHYSICAL_PORT_ID = PHYSICAL_PORT.PHYSICAL_PORT_ID;
```

VCS_CLUSTER_MEMBER_INFO

```
CREATE VIEW vcs_cluster_member_info AS
    select
    VCS_DEVICE.DEVICE_ID as VCS_DEVICE_ID,
    VCS_DEVICE.MANAGED_ELEMENT_ID as VCS_ME_ID,
    MEMBER_DEVICE.DEVICE_ID as MEMBER_DEVICE_ID,
    MEMBER_DEVICE.MANAGED_ELEMENT_ID as MEMBER_ME_ID,
    VCS_MEMBER.CREATION_TIME,
    VCS_MEMBER.TRUSTED,
    VCS_MEMBER.MISSING,
    VCS_MEMBER.MISSING_TIME,
    VCS_MEMBER.STATE,
    VCS_MEMBER.FABRIC_STATUS
from
    device VCS_DEVICE,
    device MEMBER_DEVICE,
    VCS_CLUSTER_MEMBER VCS_MEMBER
where
    VCS_MEMBER.CLUSTER_ME_ID = VCS_DEVICE.MANAGED_ELEMENT_ID AND
    VCS_MEMBER.MEMBER_ME_ID = MEMBER_DEVICE.MANAGED_ELEMENT_ID;
```

RESET_VCS_LICENSED

```
CREATE OR REPLACE FUNCTION reset_vcs_licensed(no_of_licenses integer)
```

```
  RETURNS void AS
$BODY$
begin
    UPDATE fabric set vcs_licensed = 0;
    UPDATE device set vcs_licensed = 0;
    UPDATE fabric set vcs_licensed = 1 WHERE fabric.id in (SELECT id FROM fabric ORDER BY creation_time LIMIT
no_of_licenses);
    UPDATE device set vcs_licensed = 1 WHERE device.managed_element_id in (SELECT vcs_cluster_me_id FROM
fabric_vcs_cluster_map WHERE fabric_id in (SELECT id FROM fabric WHERE vcs_licensed=1));
end;
$BODY$
  LANGUAGE plpgsql VOLATILE
  COST 100;
ALTER FUNCTION reset_vcs_licensed(integer)
  OWNER TO dcmadmin;
```

TRILL_TRUNK_INFO

```
create or replace view TRILL_TRUNK_INFO as
select
    TRILL_TRUNK_GROUP.ID,
    TRILL_TRUNK_GROUP.ME_ID,
    TRILL_TRUNK_GROUP.MASTER_PORT_NUMBER,
    TRILL_TRUNK_MEMBER.PORT_NUMBER as MEMBER_PORT_NUMBER,
    MEMBER_DEVICE.DEVICE_ID,
    MASTER_INTERFACE.INTERFACE_ID as MASTER_INTERFACE_ID,
    MASTER_INTERFACE.IF_NAME as MASTER_IF_NAME,
    MEMBER_INTERFACE.INTERFACE_ID as MEMBER_INTERFACE_ID,
    MEMBER_INTERFACE.IF_NAME as MEMBER_IF_NAME,
    VCS_CLUSTER_MEMBER.CLUSTER_ME_ID,
    CLUSTER_DEVICE.DEVICE_ID as CLUSTER_DEVICE_ID
from
    TRILL_TRUNK_GROUP
        inner join TRILL_TRUNK_MEMBER on
            TRILL_TRUNK_MEMBER.GROUP_ID = TRILL_TRUNK_GROUP.ID
        inner join device MEMBER_DEVICE on
            MEMBER_DEVICE.MANAGED_ELEMENT_ID = TRILL_TRUNK_GROUP.ME_ID
        left outer join INTERFACE MASTER_INTERFACE on
            MASTER_INTERFACE.DEVICE_ID = MEMBER_DEVICE.DEVICE_ID and MASTER_INTERFACE.IDENTIFIER =
TRILL_TRUNK_GROUP.MASTER_PORT_NUMBER
        left outer join INTERFACE MEMBER_INTERFACE on
            MEMBER_INTERFACE.DEVICE_ID = MEMBER_DEVICE.DEVICE_ID and MEMBER_INTERFACE.IDENTIFIER =
TRILL_TRUNK_MEMBER.PORT_NUMBER
        left outer join VCS_CLUSTER_MEMBER on
            VCS_CLUSTER_MEMBER.MEMBER_ME_ID = TRILL_TRUNK_GROUP.ME_ID
        left outer join DEVICE CLUSTER_DEVICE on
            CLUSTER_DEVICE.MANAGED_ELEMENT_ID = VCS_CLUSTER_MEMBER.CLUSTER_ME_ID;
```

WIRELESS_INTERFACE

```
create or replace view wireless_interface as
SELECT
      l2.device_id,
      l2.device_ip_address,
      l2.physical_device_id,
      l2.unit_number,
      l2.slot_id,
      l2.slot_num,
      l2.module_id,
```

```
       l2.physical_port_id,
       l2.port_num,
       l2.interface_id,
       l2.name,
       l2.if_name,
       l2.identifier,
       l2.speed_in_mb,
       l2.physical_address,
       l2.interface_id AS radioif_id,
       wireless.radio_type,
       wireless.is_enabled,
       wireless.is_auto_channel,
       wireless.tx_power,
       wireless.channel_number,
       wireless.max_data_rate,
       wireless.beacon_rate,
       wireless.dtim,
       wireless.rts_threshold,
       wireless.is_turbo_mode,
       wireless.radio_g_mode,
       wireless.max_associated_clients
    FROM
     ((SELECT DISTINCT d.device_id, d.ip_address AS device_ip_address, pd.physical_device_id, pd.unit_number,
s.slot_id, s.slot_num, msp.module_id, pp.physical_port_id, pp.port_num, i.interface_id, i.name, i.if_name,
i.identifier, pi.speed_in_mb, pi.physical_address
     FROM device d, physical_device pd, slot s, module_slot_present msp, physical_port pp, physical_interface
pi, interface i
      WHERE
        (((((((d.device_id = pd.device_id)
          AND (pd.physical_device_id = s.physical_device_id))
          AND (s.slot_id = msp.slot_id))
          AND (msp.module_id = pp.module_id))
          AND (pp.physical_port_id = pi.physical_port_id))
          AND (pi.interface_id = i.interface_id))
          AND ((i.table_subtype)::text = 'RADIO_INTERFACE'::text)))
        l2 LEFT JOIN
          (SELECT radio_interface.interface_id
           AS radioif_id, radio_interface.radio_type, radio_interface.is_enabled,
radio_interface.is_auto_channel, radio_interface.tx_power, radio_interface.channel_number,
radio_interface.max_data_rate, radio_interface.beacon_rate, radio_interface.dtim,
radio_interface.rts_threshold, radio_interface.is_turbo_mode, radio_interface.radio_g_mode,
radio_interface.max_associated_clients
           FROM radio_interface)
        wireless ON ((l2.interface_id = wireless.radioif_id)));
```

WIRED_INTERFACE

```
CREATE VIEW wired_interface AS
    select
L2.DEVICE_ID,
L2.MANAGED_ELEMENT_ID,
L2.DEVICE_IP_ADDRESS,
L2.PHYSICAL_DEVICE_ID,
L2.UNIT_NUMBER,
L2.SLOT_ID,
L2.SLOT_NUM,
L2.MODULE_ID,
L2.PHYSICAL_PORT_ID,
L2.PORT_NUM,
L2.INTERFACE_ID,
L2.NAME,
```

```
L2.IF_NAME,
L2.IDENTIFIER,
L2.TABLE_SUBTYPE,
case
    when L2.TABLE_SUBTYPE like 'GBIT_ETHERNET_INTERFACE' then 'GIGABIT ETHERNET'
    when L2.TABLE_SUBTYPE like 'POS_INTERFACE' then 'POS'
    else L2.TABLE_SUBTYPE
end as TABLE_SUBTYPE_TXT,
L2.TAG_MODE,
case
    when L2.TAG_MODE = 1 then 'TAGGED'
    when L2.TAG_MODE = 2 then 'UNTAGGED'
    when L2.TAG_MODE = 3 then 'DUAL'
    else null
end as TAG_MODE_TXT,
L2.USER_DEFINED_VALUE_1,
L2.USER_DEFINED_VALUE_2,
L2.USER_DEFINED_VALUE_3,
L2.SPEED_IN_MB,
L2.PHYSICAL_ADDRESS,
L2.DUPLEX_MODE,
case
    when L2.DUPLEX_MODE = 1 then 'HALF-DUPLEX'
    when L2.DUPLEX_MODE = 2 then 'FULL-DUPLEX'
    when L2.DUPLEX_MODE = 3 then 'AUTO-SENSE'
    else null
end as DUPLEX_MODE_TXT,
L3.IP_ID,
L3.IP_INTERFACE_ID,
L3.IP_ADDRESS,
L3.SUBNET_MASK
from ( select distinct D.DEVICE_ID, D.MANAGED_ELEMENT_ID, D.IP_ADDRESS as DEVICE_IP_ADDRESS,
            PD.PHYSICAL_DEVICE_ID, PD.UNIT_NUMBER, S.SLOT_ID, S.SLOT_NUM,
            MSP.MODULE_ID, PP.PHYSICAL_PORT_ID, PP.PORT_NUM, I.INTERFACE_ID,
            I.NAME, I.IF_NAME, I.IDENTIFIER, I.TABLE_SUBTYPE, I.TAG_MODE,
           I.USER_DEFINED_VALUE_1, I.USER_DEFINED_VALUE_2, I.USER_DEFINED_VALUE_3,
            PI.SPEED_IN_MB, PI.PHYSICAL_ADDRESS, PI.DUPLEX_MODE
           from DEVICE D, PHYSICAL_DEVICE PD, SLOT S, MODULE_SLOT_PRESENT MSP,
            PHYSICAL_PORT PP, PHYSICAL_INTERFACE PI, INTERFACE I
          where D.DEVICE_ID = PD.DEVICE_ID and PD.PHYSICAL_DEVICE_ID = S.PHYSICAL_DEVICE_ID and S.SLOT_ID =
MSP.SLOT_ID and MSP.MODULE_ID = PP.MODULE_ID and PP.PHYSICAL_PORT_ID = PI.PHYSICAL_PORT_ID and
PI.INTERFACE_ID = I.INTERFACE_ID and I.TABLE_SUBTYPE::TEXT <> 'RADIO_INTERFACE'::TEXT) L2
   left join ( select INM_IP_INTERFACE.INTERFACE_ID as IP_ID,
            INM_IP_INTERFACE.IP_INTERFACE_ID, INM_IP_INTERFACE.IP_ADDRESS,
            INM_IP_INTERFACE.SUBNET_MASK
           from INM_IP_INTERFACE) L3 on L2.INTERFACE_ID = L3.IP_ID;
```

CEE_PORT_INFO

```
create or replace view CEE_PORT_INFO as
select
    GIGE_PORT.ID,
    GIGE_PORT.SWITCH_PORT_ID,
    GIGE_PORT.PORT_NUMBER,
    CEE_PORT.ID AS CEE_PORT_ID,
    CEE_PORT.VIRTUAL_SWITCH_ID,
    CEE_PORT.IF_INDEX,
    CEE_PORT.IF_NAME,
    CEE_PORT.IF_MODE,
    CEE_PORT.L2_MODE,
    CEE_PORT.VLAN_ID,
    CEE_PORT.LAG_ID,
```

```
    CEE_PORT.IP_ADDRESS,
    CEE_PORT.MAC_ADDRESS,
    CEE_PORT.PORT_SPEED,
    CEE_PORT.ENABLED,
    CEE_PORT.OCCUPIED,
    CEE_PORT.LAST_UPDATE,
    CEE_PORT.NET_MASK,
    CEE_PORT.PROTOCOL_DOWN_REASON,
    CEE_PORT.MAC_ACL_POLICY,
    CEE_PORT.QOS_TYPE,
    CEE_PORT.QOS_NAME,
    CEE_PORT.DOT1X_ENABLED,
    CEE_PORT.PORT_ROLE,
    CEE_PORT.AMPP_PROFILE_MODE,
    CEE_PORT.EDGE_TYPE,
    CEE_PORT.CONNECTED_STORAGE_TYPE,
    CORE_SWITCH.IP_ADDRESS as PHYSICAL_SWITCH_IP,
    CORE_SWITCH.WWN as PHYSICAL_SWITCH_WWN,
    GIGE_PORT.OPERATIONAL_STATUS,
    GIGE_PORT.MAX_SPEED,
    GIGE_PORT.PORT_TYPE,
    GIGE_PORT.REMOTE_MAC_ADDRESS,
    GIGE_PORT.SLOT_NUMBER,
    VIRTUAL_SWITCH.WWN,
    VIRTUAL_SWITCH.MANAGEMENT_STATE,
    VIRTUAL_SWITCH.MANAGED_ELEMENT_ID,
    VIRTUAL_SWITCH.MONITORED,
    SWITCH_PORT.USER_PORT_NUMBER,
    SWITCH_PORT.STATE,
    SWITCH_PORT.STATUS,
    SWITCH_PORT.NAME,
    SWITCH_PORT.LICENSED,
    SWITCH_PORT.TRUNKED,
    SWITCH_PORT.TRUNK_MASTER,
    SWITCH_PORT.SPEED_TYPE
from
    CEE_PORT, GIGE_PORT, SWITCH_PORT, VIRTUAL_SWITCH, CORE_SWITCH
where
    CEE_PORT.GIGE_PORT_ID = GIGE_PORT.ID
    and GIGE_PORT.SWITCH_PORT_ID = SWITCH_PORT.ID
    and SWITCH_PORT.VIRTUAL_SWITCH_ID = VIRTUAL_SWITCH.ID
    and VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID;
```

REST_ETHERNET_PORT

```
create or replace view REST_ETHERNET_PORT AS

select distinct
i.INTERFACE_ID,
i.DEVICE_ID,
i.NAME,
i.TABLE_SUBTYPE,
i.TAG_MODE,
i.UNTAGGED_VLAN_ID,
i.IF_NAME,
i.PORT_STATUS,
i.PORT_STATE,
i.IF_INDEX,
i.IS_MANAGEMENT_INTERFACE,

pi.PHYSICAL_ADDRESS,
pi.SPEED_IN_MB,
```

```
pi.DUPLEX_MODE,
pi.IS_STACKING_INTERFACE,
pi.IS_PORT_PRESENT,
pi.UNIT_NUMBER,
pi.SLOT_NUMBER,
pi.PORT_NUMBER,

cp.IF_MODE,
cp.L2_MODE,
cp.VLAN_ID,
cp.LAG_ID,
cp.IP_ADDRESS,
cp.ENABLED,
cp.MAC_ACL_POLICY

FROM INTERFACE i
LEFT OUTER JOIN PHYSICAL_INTERFACE pi ON i.INTERFACE_ID = pi.INTERFACE_ID
LEFT OUTER JOIN CEE_PORT cp ON pi.PHYSICAL_ADDRESS = cp.MAC_ADDRESS
```

SPX_PORT_DETAILS_INFO

```
create or replace view SPX_PORT_DETAILS_INFO as
select  PHYSICAL_INTERFACE.INTERFACE_ID,
SPX_PORT_DETAILS.PE_GROUP_NAME, PHYSICAL_INTERFACE.UNIT_NUMBER, INTERFACE.IDENTIFIER,
PHYSICAL_INTERFACE.IS_STACKING_INTERFACE, PHYSICAL_DEVICE.UNIT_ROLE, DEVICE.DEVICE_ID,
PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID, INTERFACE. IF_INDEX, SPX_PORT_DETAILS.CONNECTED_PE_UNIT_NUMBER
from PHYSICAL_DEVICE
right join PHYSICAL_INTERFACE on PHYSICAL_INTERFACE.PHYSICAL_DEVICE_ID  = PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID
right join DEVICE on DEVICE.DEVICE_ID = PHYSICAL_DEVICE.DEVICE_ID
right join INTERFACE on INTERFACE.INTERFACE_ID  = PHYSICAL_INTERFACE.INTERFACE_ID
left join SPX_PORT_DETAILS on SPX_PORT_DETAILS.DEVICE_ID  = PHYSICAL_DEVICE.DEVICE_ID and
SPX_PORT_DETAILS.INTERFACE_ID = PHYSICAL_INTERFACE.INTERFACE_ID
where PHYSICAL_INTERFACE.IS_STACKING_INTERFACE = 3
and PHYSICAL_DEVICE.DEVICE_ID = DEVICE.DEVICE_ID
and PHYSICAL_INTERFACE.PHYSICAL_DEVICE_ID = PHYSICAL_DEVICE.PHYSICAL_DEVICE_ID;
```

Printed in USA