

IBM[®] System Storage[®]



IBM Network Advisor SAN User Manual

Supporting IBM Network Advisor version 12.1

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.

IBM[®] Sytem Storage[®]



IBM Network Advisor SAN User Manual

Supporting IBM Network Advisor version 12.1

Copyright © 2010 - 2013 Brocade Communications Systems, Inc. All Rights Reserved.

The following paragraph does not apply to any country (or region) where such provisions are inconsistent with local law.

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states (or regions) do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

© **Copyright IBM Corporation 2012, 2013.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About This Document

In this chapter	xli
How this document is organized	xli
Supported hardware and software	xlili
What's new in this document	xlvi
Document conventions	xlvi
Text formatting	xlvi
Notes, cautions, and warnings	xlvi
Key terms	xlvi
Additional information	xlvii
Getting technical help	xlvii
How to send your comments	xlviii

Chapter 1

Getting Started

User interface components	1
Management server and client	3
Logging into a server	3
Launching a remote client	4
Clearing previous versions of the remote client	5
Launching the Configuration Wizard	5
Viewing active sessions	9
Disconnecting users	10
Viewing server properties	10
Viewing port status	11
Server and client ports	12
Accessibility features for the Management application	16
Keyboard shortcuts	16
Look and feel customization	17
PostgreSQL database	19
Connecting to the database using pgAdmin III	19
Connecting to the database using the ODBC client (Windows systems)	20
Connecting to the database using the ODBC client (Linux systems)	21
Changing the database user password	23
Supported open source software products	24
SAN feature-to-firmware requirements	27

Chapter 2	Licenses	
	Licenses overview	29
	Managed count	29
	Managed SAN port count calculation	30
	Viewing the license key	30
	License downgrade	31
	Downgrading the edition	31
Chapter 3	Patches	
	Installing a patch	33
	Uninstalling a patch	34
Chapter 4	Discovery	
	SAN discovery overview	37
	FCS policy and seed switches	38
	Backbone Chassis discovery requirements	38
	Discovering fabrics	39
	Editing the password for multiple devices	43
	Configuring SNMP credentials	44
	Reverting to a default SNMP community string	45
	Rediscovering a fabric	46
	Removing a fabric from active discovery	46
	Rediscovering a previously discovered fabric	46
	Deleting a fabric	47
	Viewing the fabric discovery state	47
	Troubleshooting fabric discovery	48
	Managed count exceeded troubleshooting	49
	Virtual Fabric discovery troubleshooting	50
	SAN Fabric monitoring	51
	Stop monitoring of discovered fabrics	52
	Stop monitoring of discovered switches	53
	Resume monitoring of discovered fabrics	54
	Resume monitoring of discovered switches	55
	SAN Seed switch	55
	Seed switch requirements	56
	Seed switch failover	56
	Changing the seed switch	57
	Host discovery	58
	Discovering Hosts by Network address or host name	58
	Importing Hosts from a CSV file	60
	Importing Hosts from a Fabric	61
	Importing Hosts from a VM manager	63
	Editing Host adapter credentials	64
	Removing a host from active discovery	65
	Rediscovering a previously discovered fabric	66
	Deleting a host adapter from discovery	66

Viewing the host discovery state.....	66
Troubleshooting host discovery.....	67
VM Manager discovery.....	68
VM Manager discovery requirements.....	68
Discovering a VM manager.....	68
Editing a VM manager.....	70
Excluding a host from VM manager discovery.....	71
Including a host in VM manager discovery.....	71
Removing a VM manager from active discovery.....	71
Rediscovering a previously discovered VM manager.....	72
Deleting a VM manager from discovery.....	72
Viewing the VM manager discovery state.....	72
Troubleshooting VM manager discovery.....	73

Chapter 5

Application Configuration

Configurable preferences.....	75
Server Data backup.....	77
What is backed up?.....	77
Management server backup.....	77
Configuring backup.....	78
Enabling backup.....	80
Disabling backup.....	81
Viewing the backup status.....	81
Changing the backup interval.....	81
Starting immediate backup.....	82
Reviewing backup events.....	82
Server Data restore.....	83
Restoring data.....	83
Restoring data to a new server.....	84
SAN display settings.....	84
Setting your FICON display.....	84
Resetting your display.....	85
SAN End node display.....	86
Displaying end nodes.....	86
SAN Ethernet loss events.....	87
Enabling SAN Ethernet loss events.....	87
Disabling SAN Ethernet loss events.....	87
Event storage settings.....	88
Configuring event storage.....	88
Storing historical events purged from repository.....	89
Flyover settings.....	89
Configuring flyovers.....	89
Turning flyovers on or off.....	92
Viewing flyovers.....	92
Name settings.....	92
Setting names to be unique.....	93
Setting names to be non-unique.....	93
Fixing duplicate names.....	93

Viewing names	95
Adding a name to an existing device	96
Adding a name to a new device	97
Applying a name to a detached WWN	97
Removing a name from a device	97
Editing names	98
Exporting names	98
Importing Names	99
Searching for a device by name	99
Searching for a device by WWN	100
Miscellaneous security settings	101
Configuring the server name	101
Enforcing MD5 file during import	102
Configuring login security	102
Configuring the login banner display	102
Disabling the login banner	103
Syslog Registration settings	103
Registering a server as a Syslog recipient automatically	103
Configuring the Syslog listing port number	104
SNMP Trap Registration settings	104
Registering a server as a SNMP trap recipient automatically	104
Configuring the SNMP trap listing port number	104
SNMP Trap forwarding credential settings	105
Configuring SNMP v1 and v2c credentials	105
Configuring SNMP v3 credentials	106
Software Configuration	106
Certificates	107
Client export port settings	113
Client/Server IP	114
Memory allocation settings	118
Product communication settings	122
FTP/SCP/SFTP server settings	123
Server port settings	128
Support mode settings	129
FIPS Support	131
Fabric tracking	131
Enabling fabric tracking	132
Disabling fabric tracking	132
Accepting changes for a fabric	133
Accepting changes for all fabrics	134
Accepting changes for a switch, access gateway, or phantom domain	135

Chapter 6

User Account Management

Users overview	137
Configuration requirements	137
Viewing configured users	138
User accounts	140
Creating a new user account	140
Editing a user account	142
Copying a user account	143
Copying and pasting user preferences	143
Assigning roles and areas of responsibility to a user account	144
Removing roles and areas of responsibility from a user account	144
Disabling a user account	145
Enabling a user account	145
Deleting a user account	145
Unlocking a user account	146
Roles	146
Creating a new role	146
Editing a role	147
Copying a role	148
Deleting a role	148
Adding privileges to a role	148
Removing privileges from a role	149
Areas of responsibility	150
Creating an AOR	150
Editing an AOR	151
Copying an AOR	152
Deleting an AOR	152
Assigning products to an AOR	153
Removing products from an AOR	153
Password policies	154
Configuring a password policy	154
Viewing password policy violators	156
Authentication Server Groups on the Management server	156
Assigning roles and AORs to an AD group	157
Removing roles and AORs from an AD group	157
Loading an AD group	158
Deleting an AD group	158
Creating an AD user account	159
Assigning an AD user to an AD group	159
Defining user accounts on the external LDAP server	159
User profiles	161
Viewing your user profile	161
Editing your user profile	162
Changing your password	163
Viewing your password policy	163
Resetting optional messages	164
Configuring e-mail notification	164

Chapter 7

Dashboard Management

Dashboard overview	167
Dashboard toolbar	169
Dashboard messages	170
Dashboards expand navigation bar	170
General dashboard functions	171
Accessing a dashboard	171
Filtering the dashboards list	172
Creating a user-defined dashboard	172
Deleting a user-defined dashboard	173
Setting the dashboard display	173
Customizing the dashboard widgets and monitors	173
Exporting the dashboard display	175
Printing the dashboard display	175
Attaching and detaching the Dashboard tab	175
Setting the network scope	176
Creating a customized network scope	177
Editing a user-defined network scope	178
Deleting a user-defined network scope	178
Setting the data display time frame	179
Default dashboards	179
Product Status and Traffic dashboard	179
SAN Ports Health dashboard	180
Status widgets	180
Bottlenecked Ports widget	181
Events widget	182
Host Adapter Inventory widget	184
SAN Inventory widget	185
SAN Status widget	187
Viewing additional SAN product data	188
Status widget	189
VM Alarms widget	189
Monitoring and Alerting Policy Suite widgets	190
Out of Range Violations widget	191
Port Health Violations widget	193
Performance monitors	194
Displaying monitors on the Performance Dashboard	196
Top Port Alignment Errors monitor	196
Top Port C3 Discards monitor	197
Top Port C3 Discards RX TO monitor	199
Top Port CRC Errors monitor	200
Top Port Encode Error Out monitor	202
Top Port Link Failures monitor	203
Top Port Link Resets monitor	204
Top Port Overflow Errors monitor	206
Top Port Receive EOF monitor	207
Top Port Runtime Errors monitor	208
Top Port Sync Losses monitor	209
Top Port Too Long Errors monitor	210
Top Port Traffic monitor	211

Top Port Underflow Errors monitor	212
Top Port Utilization Percentage monitor	213
Bottom Port Utilization Percentage monitor	214
Top Product CPU Utilization monitor	215
Top Product Memory Utilization monitor	216
Top Product Response Time monitor	217
Top Product Temperature monitor	219
Top Products with Unused Ports monitor	220
Editing a preconfigured performance monitor	221
User-defined performance monitors	222
Monitor types	222
Measures	222
Top or bottom product performance monitors	225
Top or bottom port performance monitors	226
Distribution performance monitors	227
Time series performance monitors	229
Configuring a user-defined product performance monitor	229
Adding targets to a user-defined performance monitor	232
Configuring a user-defined port performance monitor	233
Viewing product distribution data details	235
Viewing port distribution data details	236
Traffic flow dashboard monitors	238
Traffic flow monitor types	238
Traffic flow measures	238
Traffic flow performance graph monitor	239
Top or bottom traffic flow performance monitor	240
Time series traffic flow performance monitor	241
Configuring a traffic flows monitor from a performance graph	242
Configuring a user-defined traffic flow performance monitor	242
Adding targets to a traffic flow performance monitor	244

Chapter 8 View Management

SAN tab overview	247
SAN main toolbar	249
View All list	250
Port Display buttons	250
Connectivity Map toolbar	251
Product List	251
Connectivity Map	253
Utilization Legend	254
Master Log	255
Minimap	256
Status bar	257
Icon legend	258
SAN product icons	258
Host product icons	259
SAN group icons	259
Host group icons	260

SAN port icons	260
SAN product status icons	260
Event icons	261
Customizing the main window	262
Zooming in and out of the Connectivity Map	262
Showing levels of detail on the Connectivity Map	263
Exporting the topology	263
Customizing application tables	264
Product List customization	267
Adding a property label	267
Editing a property label	268
Deleting a property label	268
Search	268
Searching for a device	269
Restricting a search by node	270
Searching for an exact match	270
Clearing search results	271
SAN view management overview	271
Creating a customized view	271
Editing a customized view	273
Deleting a customized view	274
Copying a view	275
SAN topology layout	276
Customizing the layout of devices on the topology	277
Customizing the layout of connections on the topology	278
Changing a group background color	278
Reverting to the default background color	279
Changing the product label	280
Changing the port label	280
Changing the port display	280
Grouping on the topology	281
Collapsing groups	281
Expanding groups	281
Viewing connections	281
Configuring custom connections	282
Deleting a custom connection configuration	282

Chapter 9

Call Home

Call Home overview	284
System requirements	285
Viewing Call Home configurations	285
Showing a Call Home center	288
Hiding a Call Home center	288
Editing a Call Home center	289
Editing the IBM Call Home center	289
Editing an e-mail Call Home center	290

Editing the EMC Call Home center	294
Editing the HP LAN Call Home center	295
Enabling a Call Home center	296
Enabling supportSave	296
Testing the Call Home center connection	297
Disabling a Call Home center	297
Viewing Call Home status	298
Assigning a device to the Call Home center	299
Removing a device from a Call Home center	299
Removing all devices and filters from a Call Home center	299
Defining an event filter	300
Call Home for virtual switches	300
Assigning an event filter to a Call Home center	301
Assigning an event filter to a device	301
Overwriting an assigned event filter	302
Removing all event filter from a Call Home center	302
Removing an event filter from a device	303
Removing an event filter from the Call Home Event Filters list	303
Searching for an assigned event filter	303

Chapter 10

Third-party tools

About third-party tools	305
Starting third-party tools from the application	306
Launching a Telnet session	306
Launching an Telnet session from the SAN tab	306
Launching an Element Manager	307
Launching Web Tools	307
Launching FCR configuration	308
Launching Name Server	309
Launching HCM Agent	310
Launching Fabric Watch	310
Single sign on support for IBM	311
Launch in context support for IBM	312
Available LIC points	313
Adding a tool	314
Entering the server IP address of a tool	315
Adding an option to the Tools menu	316
Changing an option on the Tools menu	317

Removing an option from the Tools menu	317
Adding an option to a device's shortcut menu	318
Changing an option on a device's shortcut menu	319
Removing an option from a device's shortcut menu	320
Microsoft System Center Operations Manager (SCOM) plug-in	320
Registering a SCOM server	321
Editing a SCOM server	322
Removing a SCOM server	322

Chapter 11

Server Management Console

Server Management Console overview	323
Launching the SMC on Windows	323
Launching the SMC on Linux	324
Services tab	324
Monitoring and managing Management application services	324
Refreshing the server status	325
Stopping all services	325
Stopping the CIMOM services	325
Starting all services	326
Restarting all services	326
Changing the database password	326
Ports tab	327
Viewing server port numbers	327
AAA Settings tab	328
Configuring Radius server authentication	328
Configuring LDAP server authentication	331
Configuring TACACS+ server authentication	334
Configuring Common Access Card authentication	337
Configuring switch authentication	339
Configuring Windows authentication	340
Configuring local database authentication	340
Displaying the client authentication audit trail	341
Restore tab	342
Restoring the database	342
Technical Support Information tab	343
Capturing technical support information	343
HCM Upgrade tab	344
Upgrading HCM on the Management server	344
SMI Agent Configuration Tool	345
Launching the SMIA configuration tool on Windows	345
Launching the SMIA configuration tool on Unix	346
Launching a remote SMIA configuration tool	347
Service Location Protocol (SLP) support	347
Home tab	351
Authentication tab	352
CIMOM tab	354

Certificate Management tab	357
Summary tab	359

Chapter 12 SAN Device Configuration

Configuration repository management	363
Saving switch configurations on demand	364
Restoring a switch configuration for a selected device.	365
Scheduling switch configuration backup	366
Restoring a configuration from the repository	368
Viewing configuration file content.	370
Searching the configuration file content	371
Deleting a configuration	372
Exporting a configuration	372
Importing a configuration	372
Keeping a copy past the defined age limit.	373
Replicating configurations.	373
Replicating security configurations.	377
Enhanced group management.	380
Firmware management.	380
Downloading firmware	381
Displaying the firmware repository.	383
Importing a firmware file	384
Deleting a firmware file	385
Frame viewer	386
Viewing discarded frames from a port	388
Clearing the discarded frame log	389
Refreshing the discarded frame log	389
Ports.	389
Viewing port connectivity	389
Refreshing the port connectivity view.	393
Enabling a port.	393
Filtering port connectivity	393
Viewing port details	395
Viewing ports	395
Port types	396
Showing connected ports	396
Viewing port connection properties	396
Determining inactive iSCSI devices	400
Determining port status	400
Viewing port optics.	400
Port commissioning overview.	403
Viewing existing CIMOM servers.	403
Registering a CIMOM server	405
Editing CIMOM server credentials	405
Importing CIMOM servers and credentials	406
Exporting CIMOM servers and credentials.	406
Changing CIMOM server credentials	407
Testing CIMOM server credentials	408
Deleting CIMOM server credentials	408

Decommissioning an F-Port	409
Decommissioning an E-Port	410
Decommissioning all ports on a switch	410
Decommissioning all ports on a blade	411
Recommissioning all ports on a switch	412
Recommissioning all ports on a blade	412
Port commissioning deployment report	413
Administrative Domain-enabled fabric support	413
AD-enabled fabric discovery	413
Management application behavior for AD-enabled fabrics	414
Management application support for AD-enabled fabrics	414
Port Auto Disable	416
Viewing Port Auto Disable status	417
Configuring Port Auto Disable event triggers	418
Enabling Port Auto Disable on individual ports	419
Enabling Port Auto Disable on all ports on a device	419
Disabling Port Auto Disable on individual ports	420
Disabling Port Auto Disable on all ports on a device	420
Stopping Port Auto Disable on a device	421
Resuming Port Auto Disable on a device	421
Unblocking ports	422

Chapter 13

Host Port Mapping

Host port mapping overview	423
Creating a new Host	424
Renaming an HBA Host	425
Deleting an HBA Host	425
Viewing Host properties	425
Associating an HBA with a Host	426
Importing HBA-to-Host mapping	426
Removing an HBA from a Host	428
Exporting Host port mapping	428

Chapter 14

Storage Port Mapping

In this chapter	431
Storage port mapping overview	431
Creating a storage array	432
Adding storage ports to a storage array	432
Unassigning a storage port from a storage array	433
Reassigning mapped storage ports	433
Editing storage array properties	434
Deleting a storage array	434
Viewing storage port properties	434

Viewing storage array properties	435
Importing storage port mapping	435
Exporting storage port mapping.....	437

Chapter 15

Host Management

Host management	439
Brocade adapters	440
Host Bus Adapters	440
Converged Network Adapters	441
Fabric Adapters	441
AnyIO™ technology	442
HCM software	442
HCM features	443
Host adapter discovery.....	444
VM Manager	444
Adding a VM Manager	444
Editing a VM Manager	445
Deleting a VM Manager.....	445
HCM and Management application support on ESXi systems. . .	445
ESXi CIM listener ports	445
Connectivity map.....	447
View management	447
Host port mapping	447
Adapter software.....	448
Driver repository.....	449
Boot image repository	451
Bulk port configuration.....	454
Configuring host adapter ports.....	454
Adapter port WWN virtualization	458
Configuring FAWWNs on switch ports.....	458
FAWWNs on attached AG ports.....	461
Role-based access control	463
Host adapter management privileges	463
Host adapter administrator privileges	463
Host performance management	464
Host security authentication	465
Configuring security authentication using the Management application	465
supportSave on adapters	467
Host fault management	467
Adapter events	467
Filtering event notifications.....	468
Syslog forwarding.....	468

Backup support	469
Configuring backup to a hard drive	469
Enabling backup	470
Disabling backup	470

Chapter 16

Fibre Channel over Ethernet

In this chapter	471
FCoE overview	471
DCBX protocol	472
Enhanced Ethernet features	472
Enhanced Transmission Selection	472
Priority-based flow control	472
Ethernet jumbo frames	473
FCoE protocols supported	473
Ethernet link layer protocols supported	473
FCoE protocols	473
FCoE licensing	474
Saving running configurations	474
Copying switch configurations to selected switches	474
DCB configuration management	475
Switch policies	476
DCB map and Traffic Class map	476
LLDP profiles	476
802.1x policy	476
DCB configuration	477
Minimum DCB configuration for FCoE traffic	477
Adding a LAG	482
Editing a DCB switch	484
Editing a DCB port	486
Editing a LAG	487
Enabling a DCB port or LAG	489
Deleting a LAG	489
QoS configuration	490
Priority-based flow control	490
Creating a DCB map	490
Editing a DCB map	492
Deleting a DCB map	493
Assigning a DCB map to a port or link aggregation group	494
Creating a Traffic Class map	494
Editing a Traffic Class map	495
Deleting a Traffic Class map	495
Assigning a Traffic Class map to a port or link aggregation group	496
FCoE provisioning	496
Changing the VLAN ID on the default FCoE map	497
Enabling or disabling the FCoE map on the port	497

VLAN classifier configuration	498
Adding a VLAN classifier rule	499
Editing a VLAN classifier rule	500
Deleting a VLAN classifier rule	501
Creating a VLAN classifier group	501
Deleting a VLAN classifier group	502
LLDP-DCBX configuration	502
Configuring LLDP for FCoE	502
Adding an LLDP profile	503
Editing an LLDP profile	504
Deleting an LLDP profile	504
Assigning an LLDP profile to a port or ports in a LAG	505
802.1x authentication	506
Enabling 802.1x authentication	506
Disabling 802.1x authentication	506
Setting 802.1x parameters for a port	507
Switch, port, and LAG deployment	508
Deploying DCB product, port, and LAG configurations	508
Source to target switch Fabric OS version compatibility for deployment	512
DCB performance	513
Real-time performance graph	513
Historical performance report	514
FCoE login groups	514
Adding an FCoE login group	515
Editing an FCoE login group	517
Deleting one or more FCoE login groups	518
Disabling the FCoE login management feature on a switch	518
Enabling the FCoE login management feature on a switch	518
Virtual FCoE port configuration	519
Viewing virtual FCoE ports	519
Clearing a stale entry	520

Chapter 17

Security Management

Layer 2 access control list management	523
Fabric OS Layer 2 ACL configuration	524
Creating a Layer 2 ACL from a saved configuration	530
Deleting a Layer 2 ACL configuration from the application	531
Deleting a Layer 2 ACL configuration from the switch	531
Security configuration deployment	532
Deploying a security configuration on demand	533
Saving a security configuration deployment	534
Scheduling a security configuration deployment	535

Chapter 18	FC-FC Routing Service Management	
	Devices that support Fibre Channel routing	539
	Fibre Channel routing overview	540
	Guidelines for setting up Fibre Channel routing	541
	Connecting edge fabrics to a backbone fabric	542
	Configuring routing domain IDs	544
Chapter 19	Virtual Fabrics	
	Virtual Fabrics overview	545
	Terminology for Virtual Fabrics	546
	Virtual Fabrics requirements	547
	FICON best practices for Virtual Fabrics	548
	Configuring Virtual Fabrics	550
	Enabling Virtual Fabrics	551
	Disabling Virtual Fabrics	551
	Creating a logical switch or base switch	552
	Finding the physical chassis for a logical switch	554
	Finding the logical switch from a physical chassis	554
	Assigning ports to a logical switch	555
	Removing ports from a logical switch	556
	Deleting a logical switch	557
	Configuring fabric-wide parameters for a logical fabric	557
	Applying logical fabric settings to all associated logical switches	558
	Moving a logical switch to a different fabric	559
	Changing a logical switch to a base switch	560
Chapter 20	SAN Encryption Configuration	
	Encryption Center features	564
	Encryption user privileges	565
	Smart card usage	566
	Using authentication cards with a card reader	566
	Registering authentication cards from a card reader	567
	Registering authentication cards from the database	569
	Deregistering an authentication card	570
	Setting a quorum for authentication cards	570
	Using system cards	571
	Enabling or disabling the system card requirement	572
	Registering system cards from a card reader	572
	Deregistering system cards	573
	Using smart cards	573
	Tracking smart cards	573
	Editing smart cards	575
	Network connections	576

Blade processor links	577
Configuring blade processor links	577
Encryption node initialization and certificate generation.....	578
Setting encryption node initialization	578
Key Management Interoperability Protocol	578
Configuration parameters	579
Key vault type and vendor	580
Steps for connecting to a DPM appliance	581
Exporting the KAC certificate signing request (CSR)	582
Submitting the CSR to a certificate authority	583
KAC certificate registration expiry	583
Importing the signed KAC certificate	584
Uploading the CA certificate onto the DPM appliance (and first-time configurations)	584
Uploading the KAC certificate onto the DPM appliance (manual identity enrollment)	585
DPM key vault high availability deployment	586
Loading the CA certificate onto the encryption group leader	586
Steps for connecting to an LKM/SSKM appliance	587
Launching the NetApp DataFort Management Console	588
Establishing the trusted link	588
Obtaining and importing the LKM/SSKM certificate	589
Exporting and registering the switch KAC certificates on LKM/SSKM	589
LKM/SSKM key vault high availability deployment	590
Data Encryption Keys	591
Steps for connecting to an ESKM/SKM appliance	592
Configuring a Brocade group on ESKM/SKM	593
Registering the ESKM/SKM Brocade group user name and password	594
Setting up the local Certificate Authority (CA) on ESKM/SKM	595
Downloading the local CA certificate from ESKM/SKM	596
Creating and installing the ESKM/SKM server certificate	596
Enabling SSL on the Key Management System (KMS) Server	598
Creating an ESKM/SKM High Availability cluster	598
Copying the local CA certificate for a clustered ESKM/SKM appliance	599
Adding ESKM/SKM appliances to the cluster	599
Signing the encryption node KAC certificates	600
Importing a signed KAC certificate into a switch	601
ESKM/SKM key vault high availability deployment	601
Data Encryption Keys	602
ESKM/SKM key vault deregistration	603
Steps for connecting to a TEKA appliance	603
Setting up TEKA network connections	604
Creating a client on TEKA	605
Establishing TEKA key vault credentials on the switch	606

Signing the encryption node KAC CSR on the TEKA appliance	607
Importing a signed KAC certificate into a switch	607
Steps for connecting to a TKLM appliance	608
Exporting the Fabric OS node self-signed KAC certificates	609
Converting the KAC certificate format	609
Establishing a default key store and device group on TKLM	609
Adding a device to the device group	609
Creating a self-signed certificate for TKLM	610
Importing the Fabric OS encryption node KAC certificates to TKLM	610
Exporting the TKLM self-signed server certificate	611
Importing the TKLM certificate into the group leader	611
Steps for connecting to a KMIP-compliant SafeNet KeySecure	612
Setting FIPS compliance	613
Creating a local CA	614
Creating a server certificate	615
Creating a cluster	620
Configuring a Brocade group on the KeySecure	621
Registering the KeySecure Brocade group user name and password	622
Signing the encryption node KAC CSR on KMIP	623
Importing a signed KAC certificate into a switch	625
Backing up the certificates	626
Configuring the KMIP server	628
Adding a node to the cluster	629
Steps for connecting to a KMIP-compliant keyAuthority	631
Encryption preparation	632
Creating a new encryption group	633
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)	638
Configuring key vault settings for NetApp Link Key Manager (LKM/SSKM)	643
Configuring key vault settings for HP Enterprise Secure Key Manager (ESKM/SKM)	649
Configuring key vault settings for Thales e_Security keyAuthority (TEKA)	653
Configuring key vault settings for IBM Tivoli Key Lifetime Manager (TKLM)	658
Configuring key vault settings for Key Management Interoperability Protocol	663
Understanding configuration status results	669
Adding a switch to an encryption group	670
Replacing an encryption engine in an encryption group	676
High availability clusters	677
HA cluster configuration rules	677
Creating HA clusters	677

Removing engines from an HA cluster	679
Swapping engines in an HA cluster	679
Failback option	680
Invoking failback	680
Configuring encryption storage targets	680
Adding an encryption target	681
Configuring hosts for encryption targets	689
Adding target disk LUNs for encryption	691
Configuring storage arrays	696
Remote replication LUNs	696
SRDF pairs	697
Metadata requirements and remote replication	697
Adding target tape LUNs for encryption	698
Moving targets	701
Configuring encrypted tape storage in a multi-path environment	702
Tape LUN write early and read ahead	703
Enabling and disabling tape LUN write early and read ahead	703
Tape LUN statistics	704
Viewing and clearing tape container statistics	705
Viewing and clearing tape LUN statistics for specific tape LUNs	706
Viewing and clearing statistics for tape LUNs in a container	707
Encryption engine rebalancing	709
Rebalancing an encryption engine	710
Master keys	710
Active master key	711
Alternate master key	711
Master key actions	711
Saving the master key to a file	712
Saving a master key to a key vault	713
Saving a master key to a smart card set	714
Restoring a master key from a file	715
Restoring a master key from a key vault	716
Restoring a master key from a smart card set	717
Creating a new master key	718
Security settings	719
Zeroizing an encryption engine	719
Setting zeroization	720
Using the Encryption Targets dialog box	720
Redirection zones	721
Disk device decommissioning	722
Decommissioning disk LUNs	723
Displaying and deleting decommissioned key IDs	723

Displaying Universal IDs	725
Rekeying all disk LUNs manually	725
Setting disk LUN Re-key All	726
Viewing disk LUN rekeying details	727
Viewing the progress of manual rekey operations.	728
Thin provisioned LUNs	730
Thin Provisioning support	730
Viewing time left for auto rekey	731
Viewing and editing switch encryption properties	732
Exporting the public key certificate signing request from properties	736
Importing a signed public key certificate from properties ...	736
Enabling and disabling the encryption engine state from Properties	737
Viewing and editing encryption group properties	737
General tab	738
Members tab	742
Security tab	744
HA Clusters tab	746
Link Keys tab	748
Tape Pools tab	749
Engine Operations tab	752
Encryption-related acronyms in log messages	753

Chapter 21

Zoning

Zoning overview	755
Types of zones	756
Online zoning	757
Offline zoning	757
Zoning naming conventions	758
Zoning and FICON	758
Zone database size	758
Zoning configuration	759
Configuring zoning	759
Creating a zone	760
Viewing zone properties	760
Adding members to a zone	761
Creating a member in a zone	762
Removing a member from a zone	763
Renaming a zone	763
Deleting a zone	764
Duplicating a zone	764
Customizing the zone member display	765
Enabling or disabling the default zone for fabrics	765
Creating a zone alias	766
Editing a zone alias	767
Removing an object from a zone alias	768
Exporting zone aliases	768

Renaming a zone alias	768
Deleting a zone alias	769
Duplicating a zone alias	769
Creating a zone configuration	769
Viewing zone configuration properties	770
Adding zones to a zone configuration	770
Removing a zone from a zone configuration	771
Activating a zone configuration	771
Deactivating a zone configuration	773
Renaming a zone configuration	773
Deleting a zone configuration	774
Duplicating a zone configuration	774
Creating an offline zone database	775
Deleting an offline zone database	776
Refreshing a zone database	776
Merging fabrics	777
Merging two zone databases	777
Creating a common active zone configuration in two fabrics	779
Saving a zone database to a switch	780
Exporting an offline zone database	780
Importing an offline zone database	780
Rolling back changes to the offline zone database	781
LSAN zones	781
Supported configurations for LSAN zoning	781
Configuring LSAN zoning	782
Creating an LSAN zone	783
Adding members to the LSAN zone	784
Creating a new member in an LSAN zone	785
Activating LSAN zones	785
LSAN tagging	786
Traffic Isolation zones	786
Failover options	787
Enhanced TI zones	787
Configuring Traffic Isolation zoning	788
Creating a Traffic Isolation zone	788
Adding members to a Traffic Isolation zone	789
Enabling a Traffic Isolation zone	790
Disabling a Traffic Isolation zone	790
Enabling failover on a Traffic Isolation zone	791
Disabling failover on a Traffic Isolation zone	791
Boot LUN zones	792
Creating a Boot LUN zone	792
Modifying a Boot LUN zone	793
Deleting a Boot LUN zone	793
Zoning administration	794
Comparing zone databases	794
Managing zone configuration comparison alerts	796
Setting change limits on zoning activation	796
Clearing the fabric zone database	797

Removing all user names from a zone database	797
Finding a member in one or more zones	798
Finding a zone member in the potential member list	798
Finding zones in a zone configuration	799
Finding a zone configuration member in the zones list	799
Listing zone members	799
Listing un-zoned members	800
Removing an offline device	800
Replacing zone members	801
Replacing an offline device by WWN	801
Replacing an offline device by name	802

Chapter 22

Fibre Channel over IP

FCIP services licensing	806
FCIP Concepts	806
IP network considerations	806
FCIP platforms and supported features	807
FCIP trunking	808
Design for redundancy and fault tolerance	808
FCIP tunnel restrictions for FCP and FICON emulation features	809
FCIP Trunk configuration considerations	809
FCIP circuit failover capabilities	810
Bandwidth calculation during failover	811
Circuit Failover Grouping	811
Adaptive Rate Limiting	814
FSPF link cost calculation when ARL is used	814
QoS SID/DID priorities over an FCIP trunk	814
Configuring QoS Priorities	815
IPsec and IKE implementation over FCIP	816
IPsec for the 4 Gbps platforms	817
IPSec for the 8 Gbps platforms	818
QOS, DSCP, and VLANs	818
DSCP quality of service	819
VLANs and layer two quality of service	819
When both DSCP and L2CoS are used	819
Open systems tape pipelining	820
FCIP Fastwrite and Tape Acceleration	820
FICON emulation features	821
IBM z/OS Global Mirror (z Gm) emulation	821
Tape write pipelining	822
Tape read pipelining	822
Teradata pipelining	822
Connecting cascaded FICON fabrics over FCIP	823
Planning the configuration	824
Configuring IP links and merging the fabrics	825
Configuring DWDM links to use R_RDYs	827

Extending RDR applications over FCIP	827
FCIP configuration guidelines.....	829
Virtual Port Types.....	829
Configuring an FCIP tunnel.....	830
Logical switch function on FCIP Tunnels dialog box	832
Adding an FCIP circuit.....	833
Logical switch function in FCIP Add Circuit dialog box	836
Circuit configuration failure	837
Configuring FCIP tunnel advanced settings	837
Enabling and disabling compression	837
Enabling Open Systems Tape Pipelining (OSTP)	838
Enabling Tperf test mode	839
Configuring QoS percentages	839
Configuring IPsec and IKE policies.....	839
Configuring FICON emulation	841
Viewing FCIP connection properties	843
Viewing General FCIP properties	844
Viewing FCIP port properties	845
Editing FCIP circuits	847
Disabling FCIP tunnels	848
Enabling FCIP tunnels.....	848
Deleting FCIP tunnels	849
Disabling FCIP circuits	849
Enabling FCIP circuits	849
Deleting FCIP Circuits	849
Displaying FCIP performance graphs.....	850
Displaying performance graphs for FC ports	850
Displaying FCIP performance graphs for Ethernet ports....	850
Displaying tunnel properties from the FCIP tunnels dialog box...	851
Displaying FCIP circuit properties from the FCIP tunnels dialog box.....	852
Displaying switch properties from the FCIP Tunnels dialog box..	853
Displaying fabric properties from the FCIP Tunnels dialog box ..	854
Troubleshooting FCIP Ethernet connections	854

Chapter 23

Fabric Binding

Fabric Binding overview	855
Viewing fabric binding membership	855
Enabling fabric binding	857
Disabling fabric binding.....	858
Adding switches to the fabric binding membership list	859

Adding detached devices to the fabric binding membership list	859
Removing switches from fabric binding membership	860
High integrity fabrics overview	860
Activating high integrity fabrics	861
Deactivating high integrity fabrics	862

Chapter 24

Port Fencing

In this chapter	863
About port fencing.	863
Viewing port fencing configurations	864
Thresholds	866
C3 Discard Frames threshold	867
Invalid CRCs threshold.	868
Invalid words threshold	868
Link Reset threshold	868
Protocol error threshold	868
State Change threshold.	868
Adding thresholds	869
Adding a C3 Discard Frames threshold	869
Adding an Invalid CRCs threshold.	870
Adding an Invalid Words threshold	872
Adding a Link Reset threshold	873
Adding a Protocol Error threshold.	874
Adding a State Change threshold	875
Assigning thresholds	877
Unblocking a port.	877
Avoiding port fencing inheritance	878
Editing thresholds	878
Editing a C3 Discard Frames threshold	878
Editing an Invalid CRCs threshold.	879
Editing an Invalid Words threshold.	880
Editing a Link Reset threshold	880
Editing a Protocol Error threshold.	881
Editing a State Change threshold	882
Finding assigned thresholds.	882
Viewing thresholds.	883
Viewing all thresholds on a specific Fabric OS device.	883
Removing thresholds	884
Removing thresholds from individual objects	884
Removing thresholds from the thresholds table	884

Chapter 25

FICON Environments

FICON configurations	887
Configuring a switch for FICON operation	888
Planning the configuration	888
Configuring the switch	890
Configuring FICON display	894
Configuring an Allow/Prohibit Matrix	895
Configuring an Allow/Prohibit Matrix manually	896
Saving or copying Allow/Prohibit Matrix configurations to another device	898
Copying an Allow/Prohibit Matrix configuration	898
Saving an Allow/Prohibit Matrix configuration to another device	899
Activating an Allow/Prohibit Matrix configuration	900
Deleting an Allow/Prohibit Matrix configuration	901
Changing the Allow/Prohibit Matrix display	901
Changing window arrangement	901
Clearing port names	901
Cascaded FICON fabric	902
Configuring a cascaded FICON fabric	903
Cascaded FICON fabric merge	905
Merging two cascaded FICON fabrics	907
Resolving merge conflicts	909
Port groups	910
Creating a port group	910
Viewing port groups	912
Editing a port group	912
Deleting a port group	913
Swapping blades	913

Chapter 25

Deployment Manager

Introduction to the Deployment Manager	915
Editing a deployment configuration	916
Duplicating a deployment configuration	916
Deleting a deployment configuration	917
Deploying a configuration	917
Viewing deployment logs	917
Generating a deployment report	918
Generating a deployment configuration snapshot report	918
Searching the configuration snapshots	918

Chapter 26

Fibre Channel Troubleshooting

In this chapter	921
FC troubleshooting	921
Tracing FC routes	922
Troubleshooting device connectivity	923
Confirming Fabric Device Sharing	925
Troubleshooting port diagnostics	925
Configuring link traffic test parameters	929
FCIP troubleshooting	930
Configuring IP ping	930
Tracing IP routes	932
Viewing FCIP tunnel performance	933

Chapter 27

Performance Data

SAN performance overview	935
SAN performance measures	936
SAN performance management requirements	938
SAN real-time performance data	942
Generating a real-time performance graph	942
Filtering real-time performance data	944
Exporting real-time performance data	945
Clearing port counters	945
SAN historical performance data	946
Enabling SAN-wide historical performance collection	946
Enabling historical performance collection for selected fabrics	946
Disabling historical performance collection	948
Generating and saving a historical performance graph	948
Exporting historical performance data	952
Deleting a favorite graph configuration	953
Performance database views	953
How to extract performance statistics data from the database	954
Performance statistics counters	954
SAN end-to-end monitoring	957
Configuring an end-to-end monitor pair	958
Displaying end-to-end monitor pairs in a real-time graph	959
Displaying end-to-end monitor pairs in a historical graph	960
Refreshing end-to-end monitor pairs	960
Deleting an end-to-end monitor pair	961
SAN Top Talker monitoring	961
Configuring a fabric mode Top Talker monitor	962
Configuring an F_Port mode Top Talker monitor	964
Deleting a Top Talker monitor	965
Pausing a Top Talker monitor	966
Restarting a Top Talker monitor	966

Bottleneck detection	967
Supported configurations for bottleneck detection	967
How bottlenecks are reported	968
Limitations of bottleneck detection	968
Enabling bottleneck alerts and configuring alert parameters	968
Inheriting alert parameters from a switch	971
Copying alert parameters from one switch or port to another	972
Displaying bottleneck statistics	972
Displaying devices that could be affected by an F_Port or FL_Port bottleneck	973
Disabling bottleneck detection	973
Thresholds and event notification	974
Creating and editing a threshold policy	974
Duplicating a threshold policy	977
Assigning a threshold policy	978
Deleting a threshold policy	978
SAN connection utilization	979
Enabling connection utilization	980
Disabling connection utilization	981
Changing connection utilization percentages	981
Viewing Historical Graphs/Tables	986
Mouse functions for graphs	989

Chapter 28

Flow Vision

Overview	991
Why Flow Vision exists	992
Fabric Vision components	992
Flow Vision licensing	992
Flow Vision support	993
Flows	993
Flow provisioning and monitoring	994
Provisioning flows	996
Flow definition examples	1003
Monitoring Flows	1005
Using Flow Vision dialog box options	1006
Flow Vision dialog box overview	1008
Using the Performance Graph	1013
Dashboard flow performance monitor	1015
Flow Vision features	1016
Flow Mirror	1016
Flow Monitor	1017
Flow Generator	1018
Flow parameter support	1022
Context-based flow definitions	1023
Flow parameter and configuration rules and limitations	1024
General flow parameter rules	1024

Supported basic flow parameter combinations	1025
Flow Generator supported flow identification parameter combinations	1025
Flow Mirror supported flow identification parameter combinations	1026
Flow Monitor supported flow parameter combinations ...	1026
Accessing Flow Vision from other management application features	1027
Frame Viewer	1027
MAPS	1028
Bottleneck Detection	1028
Trace route and ping	1029
Port connectivity	1030
Top Talkers	1031
End-to-End Monitors	1031

Chapter 29

Frame Monitor

Frame Monitor	1033
Frame types	1033
Frame Monitoring requirements	1035
Creating a custom frame monitor	1035
Editing a frame monitor	1037
Assigning a frame monitor to a port	1037
Finding frame monitor assignments	1038
Removing a frame monitor from a port	1038
Removing a frame monitor from a switch	1039

Chapter 30

Policy Monitor

Policy monitor overview	1041
Fabric policy monitors	1042
Switch and router policy monitors	1043
Host policy monitors	1045
Management policy monitor	1047
Preconfigured policy monitors	1047
Viewing policy monitor status	1048
Viewing existing policy monitors	1048
Adding a policy monitor	1049
Policy monitor scheduling	1055
Editing a policy monitor	1056
Deleting a policy monitor	1057
Running a policy monitor	1057
Viewing a policy monitor report	1058
Exporting a policy monitor report	1061

Viewing historical reports for all policy monitors	1061
Viewing historical reports for a policy monitor	1062

Chapter 31

Fault Management

Fault management overview	1063
Restrictions.	1064
Event notification	1064
Configuring e-mail notification	1064
Defining filters.	1066
Setting up basic event filtering	1066
Setting up advanced event filtering	1067
Viewing events	1069
SNMP traps	1069
Adding a trap recipient to one or more switches.	1070
Removing a trap recipient from one or more switches	1071
SNMP trap forwarding	1071
Event reception	1075
Adding an SNMP v3 credential	1077
Adding an SNMP v1 or v2c community string	1078
Importing a new MIB into the Management application.	1079
Trap customization.	1080
SNMP informs	1082
Enabling or disabling SNMP informs	1083
Syslogs	1083
Adding a syslog recipient.	1083
Removing a syslog recipient	1084
Syslog forwarding.	1085
Adding a syslog filter	1086
Snort message forwarding	1088
Event action definitions	1088
Creating an event action definition.	1088
Creating a new event action definition by copying an existing definition	1100
Modifying an event action definition	1100
Deleting an event action definition	1101
Configuring event actions for Snort messages	1101
Pseudo events.	1103
Displaying pseudo event definitions	1103
Creating pseudo event definitions	1103
Setting pseudo event policies.	1104
Filtering pseudo event traps	1105
Creating a pseudo event definition by copying an existing definition	1107
Editing a pseudo event definition.	1107
Deleting a pseudo event definition.	1107
Adding a pseudo event on the escalation policy	1108
Creating an event action with a pseudo event on the escalation policy	1109

Adding a pseudo event on the resolving policy	1110
Creating an event action with a pseudo event on the resolving policy	1111
Adding a pseudo event on the flapping policy	1112
Creating an event action with a pseudo event on the flapping policy.	1112
Event custom reports	1114
Defining report settings.	1115
Defining the report identity	1116
Filtering a report definition	1118
Filtering report events by date and time	1120
Creating a new report definition by copying an existing definition	1122
Editing a report definition	1122
Deleting a report definition	1123
Event custom report schedules	1123
Adding or editing an event report schedule	1124
Event logs	1126
Viewing event logs	1126
Copying part of a log entry	1127
Copying an entire log entry	1127
Exporting the entire log	1128
E-mailing all event details from the Master Log	1128
E-mailing selected event details from the Master Log	1128
Displaying event properties from the Master Log	1129
Copying part of the Master Log.	1130
Copying the entire Master Log	1130
Exporting the Master Log	1131
Filtering events in the Master Log.	1131

Chapter 32

Monitoring and Alerting Policy Suite

Monitoring and Alerting Policy Suite overview	1135
MAPS role-based access control.	1136
Enabling MAPS on a device.	1137
MAPS interoperability with other features.	1138
Fabric Watch.	1138
MAPS category, object, and measure hierarchy	1143
MAPS categories, measures, and actions	1144
MAPS monitoring categories	1146
Switch Status monitoring category.	1147
Fabric monitoring category	1148
FRU monitoring category	1149
Security monitoring category.	1149
Resource monitoring category	1150
FCIP monitoring category	1150
Traffic/Flows monitoring category	1151

MAPS policies	1152
User-defined policies	1152
MAPS rules	1154
MAPS conditions	1154
MAPS actions	1155
Fence	1155
SNMP traps	1156
Enabling or disabling policy actions for all policies	1157
Configuring e-mail notification	1158
Viewing MAPS policy data	1159
Configuring a MAPS policy	1161
Editing a MAPS policy	1164
Cloning a MAPS policy	1164
Importing Flow definitions	1165
Removing imported Flows	1166
Activating a MAPS policy	1166
Replicating a policy to other devices	1167
Exporting a MAPS policy	1167
Importing a MAPS policy	1168
Deleting a MAPS policy	1168
Viewing MAPS policy rules	1168
MAPS groups	1171
User-defined groups	1172
Editing a group	1174
Deleting a group	1175
Viewing all groups on a fabric or device	1175
Creating multiple groups	1177
Editing multiple groups	1177
Deleting a group	1178
MAPS violations	1179
MAPS events	1181
Viewing MAPS events	1181
MAPS integration with other features	1184

Chapter 33

Technical Support

In this chapter	1185
Server and client support save	1185
Capturing Server and Client support save data	1185
Capturing Server support save data	1186
Capturing Client support save data	1187
Client support save using a command line interface	1188
Device technical support	1189
Scheduling technical support information collection	1189
Starting immediate technical support information collection	1191
Viewing the technical support repository	1192
Saving technical support information to another location	1193
E-mailing technical support information	1194

	Copying technical support information to an external FTP server	1194
	Deleting technical support files from the repository	1195
	Upload failure data capture	1195
	Enabling upload failure data capture	1195
	Disabling upload failure data capture	1197
	Purging upload failure data capture files	1197
	Configuring the upload failure data capture FTP server	1197
	Saving the upload failure data capture repository	1198
Chapter 34	Reports	
	In this chapter	1199
	Reports overview	1199
	Browser requirements	1199
	SAN report types	1200
	Generating SAN reports	1200
	Viewing SAN reports	1201
	Exporting SAN reports	1202
	Printing SAN reports	1202
	Deleting SAN reports	1203
	Generating SAN performance reports	1203
	Generating SAN zoning reports	1205
	Exporting reports to e-mail recipients	1206
Appendix A	Application menus	
	Dashboard main menus	1207
	SAN main menus	1208
	SAN shortcut menus	1218
Appendix B	Call Home Event Tables	
Appendix C	Event Categories	
	Link incident events	1237
	Product status events	1237
	Product audit events	1238
	Security events	1239
	Security events for FC devices	1239
	Security events for IP devices	1239
	User action events	1240
	Management server events	1240

	Product events.	1241
	IP Performance monitoring events.	1241
Appendix D	User Privileges	
	About user privileges	1243
	About Roles and Access Levels	1260
Appendix E	Device Properties	
	SAN device properties.	1264
	Viewing Fabric properties	1264
	Viewing SAN device properties	1265
	Viewing Storage properties	1268
	Viewing iSCSI Properties dialog box	1270
	Viewing port properties	1271
	Viewing VC module properties	1276
	Host properties	1278
	Viewing adapter port properties	1278
	Properties customization	1280
	Adding a property field	1281
	Editing a property field	1281
	Deleting a property field	1282
	Editing a property field directly	1282
Appendix F	Regular Expressions	
Chapter G	Troubleshooting	
	In this chapter	1289
	Application Configuration Wizard troubleshooting	1290
	Browser troubleshooting.	1290
	Client browser troubleshooting	1291
	Fabric tracking troubleshooting.	1291
	FICON troubleshooting	1292
	Firmware download troubleshooting	1292
	Launch Client troubleshooting.	1294
	Names troubleshooting	1296
	Patch troubleshooting.	1296
	Performance troubleshooting.	1297
	Port Fencing troubleshooting	1301
	Professional edition login troubleshooting	1301
	Server troubleshooting	1301
	Server Management Console troubleshooting	1302

Supportsave troubleshooting	1303
Technical support data collection troubleshooting	1304
View All list troubleshooting	1304
Zoning troubleshooting	1305

Appendix H

Database Fields

Database tables and fields	1307
Views	1490
ADAPTER_PORT_CONFIG_INFO	1490
AG_CONNECTION_INFO	1490
BOOT_IMAGE_FILE_DETAILS_INFO	1491
CNA_ETH_PORT_CONFIG_INFO	1491
CNA_PORT_DETAILS_INFO	1492
CNA_PORT_INFO	1492
CORE_SWITCH_DETAILS_INFO	1493
CRYPTO_HOST_LUN_INFO	1494
CRYPTO_TARGET_ENGINE_INFO	1495
DASHBOARD_PREFERENCES_INFO	1495
DEPLOYMENT_INFO	1496
DEPLOYMENT_LOG	1496
DEVICE_CONNECTION_INFO	1497
EE_MONITOR_STATS_5MIN_INFO	1498
EE_MONITOR_STATS_30MIN_INFO	1498
EE_MONITOR_STATS_2HOUR_INFO	1498
EE_MONITOR_STATS_1DAY_INFO	1499
TE_PORT_STATS_5MIN_INFO	1499
TE_PORT_STATS_30MIN_INFO	1499
TE_PORT_STATS_2HOUR_INFO	1500
TE_PORT_STATS_1DAY_INFO	1500
SWITCH_INFO	1501
DEVICE_INFO	1503
N2F_PORT_MAP_INFO	1504
DEVICE_NODE_INFO	1504
DEVICE_PORT_INFO	1505
DEV_PORT_GIGE_PORT_LINK_INFO	1506
DEV_PORT_MAC_ADDR_MAP_INFO	1507
ISL_CONNECTION_INFO	1507
ISL_INFO	1508
ETHERNET_ISL_INFO	1509
EVENT_DETAILS_INFO	1509
EVENT_INFO	1510
FABRIC_INFO	1511
FCIP_TUNNEL_CIRCUIT_INFO	1512
FCIP_TUNNEL_INFO	1513
FCOE_DEVICE_INFO	1514
FRU_INFO	1515
GIGE_PORT_ECLOUD_LINK_INFO	1515
GIGE_PORT_INFO	1516
HBA_PORT_DETAILS_INFO	1516
HBA_TARGET_INFO	1518

HEALTH_STATUS_INFO	1519
HOST_DISCOVERY_REQUEST_INFO	1520
IFL_INFO	1521
ISL_INFO	1521
ISL_TRILL_INFO	1522
ISL_TRUNK_GROUP_MEMBER_INFO	1524
ISL_TRUNK_INFO	1524
L2_NEIGHBOR_INFO	1525
MAPS_EVENT_DETAILS_INFO	1525
MODULE_INFO	1526
NPORT_WWN_MAP_INFO	1527
PHANTOM_PORT_INFO	1528
PRODUCT_INFO	1528
PORT_BOTTLENECK_CONF_INFO	1531
PORT_BOTTLENECK_STAT_INFO	1531
PORT_GROUP_INFO	1531
ROLE_PRIVILEGE_INFO	1532
SCOM_EE_MONITOR_INFO	1532
SENSOR_INFO	1533
SMART_CARD_USAGE_INFO	1534
SWITCH_CONFIG_INFO	1535
SWITCH_DETAILS_INFO	1535
SWITCH_PORT_INFO	1537
SWITCH_SNMP_INFO	1539
TIME_SERIES_DATA_INFO	1541
TIME_SERIES_DATA_VIEW	1542
USER_ROLE_RESOURCE_INFO	1543
VIRTUAL_FCOE_PORT_INFO	1544
VIRTUAL_PORT_WWN_DETAILS_INFO	1544
VM_ADDRESS_INFO	1545
VM_CONNECTIVITY_INFO	1546
VM_DATASTORE_DETAILS_INFO	1548
VM_EE_MONITOR_INFO	1548
VM_HOST_INFO	1549
VM_LUN_INFO	1550
VM_STATISTICS_INFO	1551
VR_CONN_MODULE_INFO	1552
VR_CONN_MODULE_PORT_INFO	1554
VR_CONN_NPIV_INFO	1555
VMM_DISCOVERED_MAC_INFO	1556
VM_VIRTUAL_ETHERNET_ADAPTER_INFO	1557
ZONE_DB_INFO	1557
PHYSICAL_DEVICE_INFO	1560
SLOT_INFO	1561
MANAGED_ELEMENT_INFO	1561
SNMP_DATA_INFO	1562
SNMP_EXPR_DATA_INFO	1562
SNMP_DATA_VIEW	1562
VM_VNETWORK_INFO	1564
CEE_PORT_INFO	1566

Index

About This Document

In this chapter

- [How this document is organized](#) xli
- [Supported hardware and software](#)..... xliii
- [What's new in this document](#)..... xlv
- [Document conventions](#) xlvi
- [Additional information](#)..... xlvii
- [Getting technical help](#)..... xlvii

How this document is organized

This document is organized to help you find the information that you want as quickly and easily as possible. This document supports IBM Network Advisor 12.1.0 and later.

The document contains the following components:

- [Chapter 1, "Getting Started,"](#) provides a high-level overview of the user interface.
- [Chapter 2, "Licenses,"](#) provides information about the Management application license and upgrading your license.
- [Chapter 3, "Patches,"](#) provides information about installing patches.
- [Chapter 4, "Discovery,"](#) describes how to discover SANs.
- [Chapter 5, "Application Configuration,"](#) provides Management application configuration instructions.
- [Chapter 6, "User Account Management,"](#) provides information on how to manage users.
- [Chapter 7, "Dashboard Management,"](#) provides details about the Dashboard tab.
- [Chapter 8, "View Management,"](#) provides view and topology configuration instructions.
- [Chapter 9, "Call Home,"](#) provides call home configuration instructions.
- [Chapter 10, "Third-party tools,"](#) provides instructions for adding and launching third-party tools.
- [Chapter 11, "Server Management Console,"](#) provides information on using the Server Management Console to stop and start the Management application services, backup the Management application database, and capture technical support information.
- [Chapter 12, "SAN Device Configuration,"](#) provides device configuration instructions.
- [Chapter 13, "Host Port Mapping,"](#) provides instructions about how to create Hosts and assign the HBAs to them and import an externally created Host port mapping file (.CSV) to the Management application.

- [Chapter 14, “Storage Port Mapping,”](#) provides instructions about how to create and assign properties to a Storage Device.
- [Chapter 15, “Host Management,”](#) provides information on how to configure an HBA.
- [Chapter 16, “Fibre Channel over Ethernet,”](#) provides information on how to configure an FCoE.
- [Chapter 17, “Security Management,”](#) provides security configuration instructions.
- [Chapter 18, “FC-FC Routing Service Management,”](#) provides information on how to manage Fibre Channel Routing.
- [Chapter 19, “Virtual Fabrics,”](#) provides logical switch configuration instructions.
- [Chapter 20, “SAN Encryption Configuration,”](#) provides information on encryption.
- [Chapter 21, “Zoning,”](#) provides zoning configuration instructions.
- [Chapter 22, “Fibre Channel over IP,”](#) provides information on how to configure an FCIP.
- [Chapter 23, “Fabric Binding,”](#) provides fabric binding instructions.
- [Chapter 24, “Port Fencing,”](#) provides information on how to configure port fencing.
- [Chapter 25, “FICON Environments,”](#) provides information on how to manage FICON.
- [Chapter 25, “Deployment Manager,”](#) provides information about how to view, deploy, and manage deployment configurations.
- [Chapter 26, “Fibre Channel Troubleshooting,”](#) provides troubleshooting details.
- [Chapter 27, “Performance Data,”](#) provides information on how to manage performance.
- [Chapter 28, “Flow Vision,”](#) is a network diagnostic tool that provides a unified platform to manage traffic-related applications on Fabric OS devices.
- [Chapter 29, “Frame Monitor,”](#) provides information on how to monitor frames.
- [Chapter 30, “Policy Monitor,”](#) provides information on how to configure best practice guidelines.
- [Chapter 31, “Fault Management,”](#) provides event management instructions.
- [Chapter 32, “Monitoring and Alerting Policy Suite,”](#) provides 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 alerts you to problems long before they become costly failures.
- [Chapter 33, “Technical Support,”](#) provides server, client, and device support save instructions.
- [Chapter 34, “Reports,”](#) provides generating report instructions.
- [Appendix A, “Application menus,”](#) provides information about the main and shortcut menus.
- [Appendix B, “Call Home Event Tables,”](#) provides supplemental information about call home event tables.
- [Appendix C, “Event Categories,”](#) provides events that display in the application.
- [Appendix D, “User Privileges,”](#) provides supplemental information about user privileges and access levels.
- [Appendix E, “Device Properties,”](#) provides reference information related to fabric, product, and port properties.
- [Appendix F, “Regular Expressions,”](#) provides a summary of Unicode regular expression constructs that you can use in the Management application.
- [Appendix G, “Troubleshooting,”](#) provides general troubleshooting details.
- [Appendix H, “Database Fields,”](#) provides reference information related to databases.

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 12.1.0, 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 12.1.0:

- Fabric OS 5.0 or later in a pure Fabric OS fabric
- Fabric OS 6.0 or later in a Mixed Fabric

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.

The hardware platforms in the following table are supported by this release of IBM Network Advisor 12.1.X.

TABLE 1 Supported Hardware

IBM Name	Terminology used in documentation	Firmware level required
SAN16B-2	16-port, 4 Gbps FC Switch	Fabric OS 5.0.0 to Fabric OS 6.2.0
SAN24B-4	24-port, 8 Gbps FC Switch	Fabric OS v6.1.0 or later
SAN32B-2	32-port, 4 Gbps FC Switch	Fabric OS 4.4.0 or later
SAN64B-2	64-port, 4 Gbps FC Switch	Fabric OS v5.2.0 or later
SAN32B-3	32-port, 4 Gbps FC Interop Switch	Fabric OS v5.2.1 or later
SAN40B-4	40-port, 8 Gbps FC Switch	Fabric OS v6.1.0 or later
SAN80B-4	80-port, 8 Gbps FC Switch	Fabric OS v6.1.0 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
SAN18B-R	4 Gbps Router, Extension Switch	Fabric OS v5.1.0 or later
SAN04B-R	4 Gbps Extension Switch	Fabric OS v5.1.0 or later
FR4-18i Extension Blade	4 Gbps Router, Extension Blade	Fabric OS v5.1.0 or later
FR8-24 Extension Blade	8 Gbps Router, Extension Blade	Fabric OS 6.4.0 or later

TABLE 1 Supported Hardware

IBM Name	Terminology used in documentation	Firmware level required
SAN06B-R	8 Gbps Extension Switch	Fabric OS v6.3.0 or later
IBM Converged Switch B32	8 Gbps 8-FC-port, 10 GbE 24-CEE port Switch	Fabric OS v6.1.2_CEE
SAN256B	Director Chassis	Fabric OS v5.0.0 to Fabric OS 7.0.0
SAN256B with FC4-16, FC4-32, and FC4-48 Blades	Director Chassis with 4 Gbps 16-FC port, 4 Gbps 32-FC port, and 4 Gbps 48-FC port Blades	Fabric OS v5.2.0 or later (FC4-48)
SAN256B with FR4-18i Blade	Director Chassis with 4 Gbps router, Extension Blade	Fabric OS v5.1.0 or later (FR4-18i)
SAN256B with FC4-16IP Blade	Director Chassis with 4 Gbps 8-FC port and 8 GbE iSCSI Blade	Fabric OS v5.2.0 or later (FC4-16IP)
SAN256B with FC10-6 Blade	Director Chassis with 10 Gbps 6-port ISL Blade	Fabric OS v5.3.0 or later (FC10-6)
SAN768B ^{1, 2}	8-slot Backbone Chassis	Fabric OS v6.0.0 or later
SAN768B ^{1, 2} with FC8-16, FC8-32, and FC8-48 Blades	8-slot Backbone Chassis with 8 Gbps 16-FC port, 8 Gbps 32-FC port, and 8 Gbps 48-FC port Blades	Fabric OS v6.0.0 or later
SAN768B ^{1, 2} with FC8-64 Blade	8-slot Backbone Chassis with 8 Gbps 64-port Blade	Fabric OS v6.4.0 or later
SAN768B ^{1, 2} with FR4-18i Blade	8-slot Backbone Chassis with 4 Gbps Router, Extension Blade	Fabric OS v6.0.0 or later
SAN768B ^{1, 2} with FC10-6 Blade	8-slot Backbone Chassis with FC 10 - 6 ISL Blade	Fabric OS v6.2.0
SAN768B ^{1, 2} with FX8-24 Extension Blade	8-slot Backbone Chassis with 8 Gbps Extension Blade	Fabric OS v6.3.1_CEE
SAN768B ^{1, 2} with FCoE10-24 Blade	8-slot Backbone Chassis with 8 Gbps 24-port FCoE Blade	Fabric OS v6.3.1_CEE
SAN384B ¹	4-slot Backbone Chassis	Fabric OS v6.0.0 or later
SAN384B ¹ with FC8-16, FC8-32, and FC8-48 Blades	4-slot Backbone Chassis with 8 Gbps 16-FC port, 8 Gbps 32-FC port, and 8 Gbps 48-FC port Blades	Fabric OS v6.2.0 or later
SAN384B ¹ with FC8-64 Blade	4-slot Backbone Chassis with 8 Gbps 64-port Blade	Fabric OS v6.4.0 or later
SAN384B ¹ with FR4-18i Blade	4-slot Backbone Chassis with 4 Gbps Router, Extension Blade	Fabric OS v6.2.0 or later
SAN384B ¹ with FC10-6 Blade	4-slot Backbone Chassis with FC 10 - 6 ISL Blade	Fabric OS v6.2.0 or later
SAN384B ¹ with FX8-24 Extension Blade	4-slot Backbone Chassis with 8 Gbps 12-FC port, 10 GbE ports, 2-10 GbE ports Extension Blade	Fabric OS v6.3.1_CEE
SAN384B ¹ with FCoE 10-24 Blades	4-slot Backbone Chassis with 8 Gbps 24-port FCoE Blade	Fabric OS v6.3.0 or later
SAN384B-2 ¹	16 Gbps 4-slot Backbone Chassis	Fabric OS v7.0.0 or later
SAN768B-2 ¹	16 Gbps 8-slot Backbone Chassis	Fabric OS v7.0.0 or later
SAN32B-E4 Encryption Switch	8 Gbps Encryption Switch	Fabric OS v6.1.1_enc or later
FS8-18 Encryption Blade	Encryption Blade	Fabric OS v6.1.1_enc or later
FC8-16 Blade	FC 8 GB 16-port Blade	Fabric OS v6.2.0 or later

TABLE 1 Supported Hardware

IBM Name	Terminology used in documentation	Firmware level required
FC8-32 Blade	FC 8 GB 32-port Blade	Fabric OS v6.2.0 or later
FC8-32E Blade ³	FC 8 GB 32-port Blade	Fabric OS v7.0.1 or later
FC8-48 Blade	FC 8 GB 48-port Blade	Fabric OS v6.2.0 or later
FC8-48E Blade	FC 8 GB 48-port Blade	Fabric OS v7.0.1 or later
FC8-64 Blade	FC 8 GB 64-port Blade	Fabric OS v6.4.0 or later
FC10-6 Blade	FC 10 - 6 ISL Blade	Fabric OS v6.2.0 or later
FC16-32 Blade	16 Gbps 32-port blade	Fabric OS v7.0.0 or later
FC16-48 Blade	16 Gbps 48-port blade	Fabric OS v7.0.0 or later
FCoE10-24 Blade	10 Gig FCoE Port Router Blade	Fabric OS v6.3.0 or later
FX8-24 Extension Blade ^{1, 2}	8 Gbps Extension Blade	Fabric OS v6.3.1_CEE

1. Professional can discover, but not manage this device. 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.
2. Professional Plus Trial and Licensed version can discover, but not manage this device. 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.
3. Only supported on the SAN384B-2 and SAN768B-2 chassis.

What's new in this document

The following changes have been made since this document was last released:

- Information that was added:
 - Flow Vision
 - Monitoring and Alerting Policy Suite
 - Dashboard
 - Product Status and Traffic dashboard
 - Out of Range Violations widget
 - Port Health Violations widget
 - SAN Port Health dashboard
 - Expand navigation bar
 - Network Scope
 - Time Scope
- Information that was changed:
 - Dashboard
 - Port status widgets and performance monitors enhancements
 - Firmware Management
 - Fault Management enhancements
 - Database tables
- Information that was deleted:
 - None.

For further information about new features and documentation updates for this release, refer to the release notes.

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
<i>italic text</i>	Provides emphasis Identifies variables Identifies paths and Internet addresses Identifies document titles
<code>code text</code>	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 more information about IBM SAN products, see the following Web site:

www.ibm.com/servers/storage/support/san

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. IBM Network Advisor Version Number
2. 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

3. 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:

- SAN16B-2—On the nonport side of the chassis
- SAN24B-4, SAN24B-5, SAN32B-2, SAN64B-2, SAN40B-4, SAN80B-4, SAN96B-5, SAN18B-R, SAN04B-R, 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
- SAN32B-3—On the switch ID pull-out tab located on the bottom of the port side of the switch

- SAN48B-5 – On the pull-out tab on the front of the switch
- SAN256B—Inside the chassis next to the power supply bays
- SAN768B—On the bottom right on the port side of the chassis
- SAN384B—On the bottom right on the port side of the chassis, directly above the cable management comb

4. World Wide Name (WWN)

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

If you cannot use the **wwn** command because the switch is inoperable, you can get the WWN from the same place as the serial number, except for the SAN768B. For the SAN768B, 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

In this chapter

- [User interface components](#) 1
- [Management server and client](#) 3
- [Accessibility features for the Management application](#) 16
- [PostgreSQL database](#) 19
- [Supported open source software products](#) 24
- [SAN feature-to-firmware requirements](#) 27

User interface components

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 graphic illustrates the various areas, and descriptions of them are listed below.

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 173).

1 User interface components




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 [Appendix A, “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 249, or [“Dashboard toolbar”](#) on page 169.
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 167.
 - **SAN tab** — Displays the Master Log, Minimap, Connectivity Map (topology), and Product List. For more information, refer to the [“SAN tab overview”](#) on page 247.
4. **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 into a server

You must log into a server to monitor your network.

NOTE

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

To log into a server, complete the following steps.

1. 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. Log into 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.

3. 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.

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 163.

1 Management server and client

5. Click **Login**.
6. 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

To launch a remote client, 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:Port_Number*.

The Management application web start screen displays.

2. Click the Management application web start link.

The **Log In** dialog box displays.

3. Log into 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.

4. Remove a server from the **Network Address** list by selected the IP address and clicking **Delete**.
5. 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 163.

6. Click **Login**.

7. 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 4.

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 12.X.X > Server Management Console**).

1. Choose one of the following options:
 - On Windows systems, select **Start > Programs > Management_Application_Name 12.X.X > Management_Application_Name Configuration**.
 - On UNIX systems, execute `sh Install_Home/bin/configwizard` on the terminal.
2. Click **Next** on the **Welcome** screen.

1 Management server and client

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 124.
 - 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 125.
 - 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 126.
 - 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 2221 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

If the “hostname” contains invalid characters, the host name does not display in the list. Valid characters include alphanumeric and dash (-) characters. The IP address is selected by default. If an IPv6 address is selected, server start up will fail.

If DNS is not configured for your network, do not select the ‘hostname’ option from either the **Server IP Configuration** or **Switch - Server IP Configuration Preferred Address** list. Selecting the ‘hostname’ option prevents clients and devices from communicating with the Server.

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 116.

- c. Click **Next**.
6. Complete the following steps on the **Server Configuration** screen.

NOTE

Do not use port 2638 for any of these port numbers. Port 2638 is used internally by the server.

Network Advisor requires Web Server, Database, Syslog and SNMP port numbers, as well as 15 consecutive port numbers from a Starting port #. On enabling HTTP redirection, port # 80 is used to redirect the HTTP requests to HTTPS.

Web Server Port # (HTTPS)

Redirect HTTP Requests to HTTPS

Database Port #

Starting Port #

Syslog Port #

SNMP Port #

Change this configuration by selecting Server > Options > Server Port from the application.

FIGURE 3 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 128.

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

1 Management server and client

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

NOTE

For Professional software, the server requires 15 consecutive free ports beginning with the starting port number.

NOTE

For Trial and Licensed software, the server requires 18 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. Verify your configuration information on the **Server Configuration Summary** screen and click **Next**.
9. 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.

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

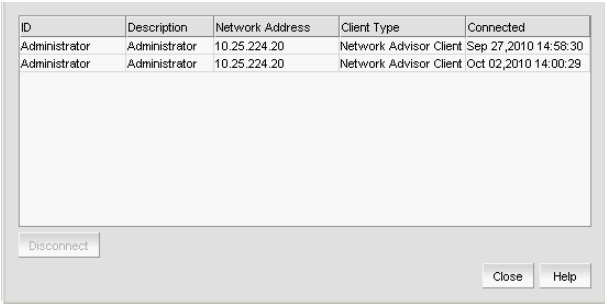
When you launch the Management application or navigate to a new view, the **SAN** tab displays with a gray screen over the Product List and Topology Map while data is loading.

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).



ID	Description	Network Address	Client Type	Connected
Administrator	Administrator	10.25.224.20	Network Advisor Client	Sep 27, 2010 14:58:30
Administrator	Administrator	10.25.224.20	Network Advisor Client	Oct 02, 2010 14:00:29

Disconnect Close Help

FIGURE 4 Active Sessions dialog box

2. Review the active session information.

The following information displays:

- **ID** – Displays the name of the user (for example, Administrator).
- **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.
4. The user you disconnected receives the following message:

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

5. Click **Close**.

When you disconnect a client from 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 5 Server Properties dialog box

2. Review the information.

TABLE 2 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.
Java VM Name	The Java Virtual Machine name.

TABLE 2 Server Properties

Field/Component	Description
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.
Win32 Service	Specifies whether the Win32 service is available on the server. On Unix servers, displays as 'No'.

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.

1 Management server and client

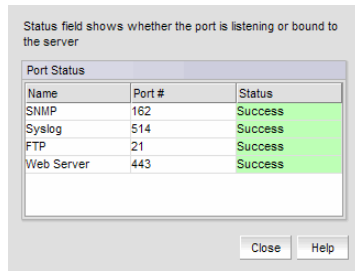


FIGURE 6 Port Status dialog box

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 occupied 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 application). Blank when the port is only used by the Management application 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 3 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.

TABLE 3 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 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	Server–Product Client–Product	Yes 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

1 Management server and client

TABLE 3 Port usage and firewall requirements (Continued)

Port Number	Ports	Transport	Description	Communication Path	Open in Firewall
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 Server-Managed Host	Yes 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 Server-Managed Host	Yes Yes

TABLE 3 Port usage and firewall requirements (Continued)

Port Number	Ports	Transport	Description	Communication Path	Open in Firewall
6343 ²	sFlow	UDP	Receives sFlow data from products if you are monitoring with sFlow	Product-Server	Yes
24600 ^{1,2}	JNP (Java Naming Protocol) port	TCP	Use for service location. Uses SSL for privacy.	Client-Server	Yes
24601 ^{1,2}	EJB (Enterprise Java Bean) connection port	TCP	Client requests to server. Uses SSL for privacy.	Client-Server	Yes
24602 ^{1,2}	HornetQ Netty port	TCP	Use for JMS (Java Message Service), async messages from server to client. Uses SSL for privacy.	Client-Server	Yes
24603 ^{1,2}	JMX RMI port	TCP	Use for JMS control. Uses SSL for privacy.	Client-Server	Yes
24604 ^{1,2}	RMI naming service port	TCP		Client-Server	Yes
24605 ^{1,2}	RMI/JRMP invoker port	TCP		Client-Server	Yes
24606 ^{1,2}	Event Handling CIM Indication listener port	TCP	Used for HBA management	Managed Host - Server	Yes
24607 ^{1,2}	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
24611 ²	JBoss Transaction Services Recovery Manager port	TCP	Not used remotely	Server	Yes
24612 ²	JBoss Transaction Status Manager port	TCP	Not used remotely	Server	Yes
24613 ²	JBoss Pooled invoker port	TCP	Not used remotely	Server	Yes
24614 ²	JBoss Socket invoker port	TCP	Not used remotely	Server	Yes
24615 ²	JBoss RMI dynamic class loading port	TCP	Web service port, not used remotely	Server	Yes
24616 ²	Apache JServ port	TCP	Proxys web server requests, not used remotely	Server	Yes
24617 ²	Remote Management application connector access port	TCP	Not used remotely	Server	Yes
34568	HCM Agent discovery port	TCP	Used for HBA management via JSON	Server - Managed Host	Yes

1 Accessibility features for the Management application

TABLE 3 Port usage and firewall requirements (Continued)

Port Number	Ports	Transport	Description	Communication Path	Open in Firewall
55556 ¹	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 edition.
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 4 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

TABLE 4 Keyboard shortcuts

Menu Item or Function	Keyboard Shortcut
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 5 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.

1 Accessibility features for the Management application

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 affect until after you restart the client.

Changing the font size

The **Options** dialog box enables you to change the font size for all components including the Connectivity map of the Management application interface.

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.

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 affect 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. Selecting **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 username (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 options:
 - (32-bit OS) Select **Start > Settings > Control Panel > Administrative Tools > Data Sources (ODBC)**.
 - (62-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 Network Advisor database.

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

TABLE 6 ODBC RedHat Package Manager (RPM) file requirements

Linux version	RedHat Package Manager file
SUSE	Rpm -l 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 and , 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 psqIODBC** screen click **Next**.
3. Install the file to the usual location for your system's application files (for example, /opt/PostgreSQL/psqIODBC) 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 psqIODBC 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 8.4
Driver = /opt/PostgreSQL/psqlODBC/lib/psqlodbcw.so
Database = dcldb
Servername = 172.26.1.54
UserName = dcadmin
Password = passwOrd
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 `passwOrd` (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.

1 Supported open source software products

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 products

[Table 7](#) lists the open source software third-party software products used in this release.

TABLE 7 Open source software third-party software products

Open Source Software	License Type
7-ZipLZMASDK 4.65	public domain
Abator 1.1	Apache License v2.0
ApacheAnt 1.7.1	Apache License v2.0
ApacheCommonsBeanUtils 1.8.1	Apache License v2.0
ApacheCommonsCodec 1.4	Apache License v2.0
ApacheCommonsCollections 3.2.1	Apache License v2.0
ApacheCommonsCompress 1.0	Apache License v2.0
ApacheCommonsConfiguration 1.6	Apache License v2.0
ApacheCommonsDBCP 1.2.2	Apache License v2.0
ApacheCommonsDigester 2.0	Apache License v2.0
ApacheCommonsDiscovery 0.4	Apache License v2.0
ApacheCommonsFileUpload 1.2.1	Apache License v2.0
ApacheCommonsHTTPClient 3.1	Apache License v2.0
ApacheCommonsIO 1.4	Apache License v2.0
ApacheCommonsJXPath 1.3	Apache License v2.0
ApacheCommonsLang 2.4	Apache License v2.0
ApacheCommonsLogging 0.4	Apache License v2.0
ApacheCommonsMath 2.0	Apache License v2.0

TABLE 7 Open source software third-party software products

Open Source Software	License Type
ApacheCommonsNet 2.0	Apache License v2.0
ApacheCommonsPool 1.5.4	Apache License v2.0
ApacheCommonsValidator 1.3.1	Apache License v2.0
Apache Extras Companion for Apache log4j 1.1	Apache License v2.0
ApacheFTPServer 1.0.3	Apache License v2.0
Apache Log4j 1.2.16	Apache License v2.0
ASM 3.2	Custom License
Axis 1.4	Apache License v2.0
AXL Radius Client API 3.29	AXL Radius Client License
BeanScriptingFramework 2.4.0	Apache License v2.0
BeanShell 2.0b4	Sun Public License / Gnu Lesser Public License
BouncyCastleCryptoProvider 1.45	Bouncy Castle License
CastorBindingFramework 0.9.9.1	Apache License v2.0
Conf M 1.9.3	Java-based software library
DNSJava 2.0.7	BeanShell Software License
dom4j 1.6.1	dom4j License
EnterpriseDFTP 1.5.6	LGPL
GlazedLists 1.8.0	LGPL or MPL
GoogleGuice 1.0	Apache
HPInsightSoftwareVCEMWebClientSDK 6.2	HP SOFTWARE DEVELOPMENT KIT LICENSE AGREEMENT
HornetQ 2.0.0	Apache License v2.0
iBATISDAOFramework 2.2.0	Apache
iBatisforJava 2.3.4	Apache License v2.0
Infinispan 4.0.0 FINAL	LGPL v2.1
InstallAnywhere 2010	Commercial
Ireasoning SNMP API 4.0	IREASONING
iTextJavaPDFLibrary 2.1.7	Affero General Public License
JasperReports 3.6.1	GNU Lesser General Public License version 3
JavaCIFSCientLibrary 1.3.12	LGPL v2.1
JavaServiceWrapper 3.3.9	Custom License
JavaTar2.5andTarTool1.4	public domain
JaxenXPathLibrary 1.1.1	Jaxen License
JbcParser 3.7	Math Parser License
JBossApplicationServer 5.1.0 GA	LGPL
JBossWeb 2.1.9	GNU Lesser General Public License version 3

1 Supported open source software products

TABLE 7 Open source software third-party software products

Open Source Software	License Type
JCalendar 1.3.3	LGPL v2.1
JCommon 1.0.16	LGPL v2.1
JDOM 1.1.1	Apache Style
JFreeChart 1.0.13	LGPL v2.1
JGoodiesForms 1.2.1	BSD
JGoodiesLooks 2.2.2	BSD
JGraph 5.13.0.1	BSD Style
JIDE 2.10.1	JIDE Software License
Jmesa 2.4.5	Apache
JSON-RPCJava 1.0.1	Apache License v2.0
KajabilityTools 0.1	Apache License v2.0
L2Fprod.comCommonComponents 7.3	Apache License v2.0
MaverickJavaSSHAPI 1.4.25	SSH Tools License
MimeTypeDetectionUtility 2.1.2	Apache License v2.0
MyBatisPersistenceFrameworkandSchhemaMigrationsforJava 3.0.2 GA	Apache License v2.0
OpenSAML 2.3.0	Apache License v2.0
OpenSSLforLinux 1.0.0a	OpenSSL License
PostgreSQL 9.2.1	PostgreSQL License
QualityFirstLibrary 0.99.0	Mozilla License V1.1 and qflib License
Quartz Enterprise Job Scheduler 1.66	Apache License v2.0
RockSawRawSocketLibrary 1.0.0	Apache License v2.0
SafeNet Sentinel Caffé 1.6.1	SafeNet License
SafeNet Sentinel RMS SDK 8.2.2	SafeNet License
Sblim-cim-client 1.3.9.3	HCM Sblim CIM Client
SimpleLoggingFacadeForJava 1.5.8	SLF4J License
SunJavaRuntimeEnvironment 1.6.0_31	Commercial
TableLayout 2009-06-10	Custom License
VIJavaAPI 2.1	BSD License
WBEM Solutions J WBEM Server 3.4.4	Commercial
WebNMSSNMPAPI 4.0.6	WebNMS License
XML RPC 1.2-B1	Open Source
YourKitJavaProfiler 9.5.1	YourKit License

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 8 SAN feature to firmware requirements

Feature	Fabric OS
Access Gateway (AG)	AG connected to Fabric OS devices requires firmware 5.2 or later.
Call Home (Trial and Licensed version Only)	Requires Fabric OS 5.2 or later for supportSave. Requires Fabric Watch license for SNMP traps.
Configuration Management	Requires Fabric OS 5.3 or later
Discovery	Requires Fabric OS 5.0 or later for the seed switch in a pure Fabric OS fabric. Requires Fabric OS 6.0 or later for the seed switch in a mixed Fabric OS and M-EOS fabric.
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 4.4 or later for SNMP traps
Fabric Binding (Trial and Licensed version Only)	Requires Fabric OS 5.2 or later in a pure Fabric OS fabric. Requires Fabric OS 6.0 or later in a mixed Fabric OS and M-EOS fabric.
FCIP Management	Requires Fabric OS 5.1 or later to modify. Requires Fabric OS 5.3 or later for FCIP tunnels. Requires FCIP license. Requires Fabric OS 6.0 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.
FICON (Trial and Licensed version Only)	Requires Fabric OS 5.2 or later for cascaded FICON. Requires Fabric OS 6.0 or later for advanced FICON. Requires Fabric OS 6.1.1 or later to configure multiple Allow/Prohibit matrices. Requires FICON CUP license to allow CUP management features.
Firmware Management	Requires Fabric OS 5.0 or later. Requires Fabric OS 6.1.1 or later on 8G devices. Requires Enhanced Group Management license to perform group actions.
High Integrity Fabric	Requires Fabric OS 5.2 or later in a pure Fabric OS fabric. Requires Fabric OS 6.0 or later in a mixed Fabric OS and M-EOS fabric.
Meta SAN	Requires Fabric OS 5.2 or later for FC router and router domain ID configuration. Requires Fabric OS 6.0 or later in a mixed Fabric OS and M-EOS fabric. Requires Integrated Routing license.
Performance	Requires Fabric OS 5.0 or later for FC_ports, -end monitors, and marching ants. Requires Fabric OS 5.3 or later for GE_ports and FCIP tunnels. Requires Fabric OS 6.2 or later for Top Talkers. Requires Advanced Performance Monitoring (APM) license for -end Monitoring and Top Talkers. Requires Enhanced Group Management license for Historical graphs and tables. Requires Fabric Watch license for Performance thresholds.

1 SAN feature-to-firmware requirements

TABLE 8 SAN feature to firmware requirements

Feature	Fabric OS
Port Fencing (Trial and Licensed version Only)	Requires Fabric OS 6.2 or later. Requires Fabric OS 6.3 or later for State Change and C3 Discard Frames violation types.
Security Management	Requires Fabric OS 5.2 and later for SCC Policy. Requires Fabric OS 5.2 and later for DCC Policy. Requires Fabric OS 5.3 and later for IP Filter Policy. Requires Fabric OS 6.0 and later for AD/LDAP Server Configuration. Requires Fabric OS 5.0 and later for RADIUS Server Configuration.
Technical Support Data Collection	Requires Fabric OS 5.2 or later.
Troubleshooting and Diagnostics	Requires Fabric OS 5.2 or later.
Virtual Fabrics (Trial and Licensed version Only)	Requires at least one Virtual Fabrics-enabled physical chassis running Fabric OS 6.2 or later.
Zoning	Requires Fabric OS 5.0 or later for pure Fabric OS fabrics. Requires Fabric OS 6.0 or later for McDATA Fabric Mode. Requires Adaptive Networking license for Quality of Service zones.

Licenses

In this chapter

- Licenses overview 29
- Managed count 29
- Viewing the license key 30
- Viewing the license key 30
- License downgrade 31

Licenses overview

NOTE

If your installation does not require a license key, the **License** dialog box does not display.

License keys consist of an asterisk (*) followed by unique string of alphanumeric characters. 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.

Managed count

The Management application audits and verifies the managed count against the maximum limit for your license under the following conditions:

- Every three hours from server start time. Note that you may be able to manage more products or ports than the maximum licensed limit briefly (maximum of three hours) between these periodic checks.
- When a new client logs in to the server.
- When you access the **License** dialog box (**Help > License**).

NOTE

SAN Professional Plus Trial and Licensed versions can manage up to 2,560 ports.

NOTE

SAN Enterprise Trial and Licensed versions can manage up to 9,000 ports and 36 fabrics.

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.

NOTE

Virtual Fabrics are counted as Fabrics when calculating the managed count limits.

Managed SAN port count calculation

NOTE

If you exceed the maximum port count for your version, software functionality is impacted and you must reduce the port count using the **Discover Fabrics** dialog box or contact your vendor to purchase an additional license for your version.

The managed SAN port count is calculated using the following rules:

- Only switches discovered from the **SAN** tab are counted.
- The switch port must be licensed.
- The ports must belong to a currently monitored fabric.
- ICL ports are not counted.
- The port must be a physical port (for example, VE Ports are not counted the 4 Gbps router extension switch; however, the gigabit ports are counted).
- Access Gateway ports are counted.
- The ports from discovered Virtual Fabrics are counted.
- The ports from managed Fabric OS switches are counted.
- The ports from unmanaged, unreachable, and missing switches are not counted.

Viewing the license key

A license key is required to run the Management application. The license key specifies the expiration date of a trial license, as well as the number of ports allowed.

NOTE

You are not required to enter a license key for a trial license. If you selected 75 Days Trial during installation, you can use the Management application, including all of its features, for a trial period of 75 days. At the termination of the trial period, a “license expired” confirmation message displays. You must enter a license key to continue using the Management application.

To the license key, complete the following steps.

1. Select **Help > License**.

The **License** dialog box displays.

Review the new information in the **License** dialog box fields.

- **Remaining Trial Period** — The number of days remaining in the trial period.
- **License Details** — The managed and licensed count for products and ports. The items to be counted:

- **Managed Count** – The number of managed ports, products, and fabrics.

NOTE

Virtual Fabrics are counted as Fabrics when calculating the managed count limits.

- SAN Enterprise edition supports a maximum of 9,000 ports and 36 fabrics.
 - SAN Professional Plus edition supports a maximum of 2,560 ports and 36 fabrics.
 - Only fabrics and devices discovered from the **SAN** tab are counted.
- **Maximum Limit** – The number of licensed ports or products allowed.

2. Click **OK** to close the **License** dialog box.

License downgrade

You can downgrade from a higher Trial configuration to a licensed version with a lower configuration.

NOTE

You cannot downgrade to Professional Edition.

NOTE

Downgrading to a Trial version is not supported.

NOTE

You cannot downgrade during migration (Configuration Wizard).

Downgrading the edition

You can downgrade from Enterprise to Professional Plus.

Before you downgrade the edition, make sure your application meets the following requirements:

- Make sure that your application configuration is within the limit of the licensed version.
- Make sure that application is not using a Backbone chassis as a seed switch.

NOTE

If you combine more than one downgrade option, you must meet the requirements for all downgrade options.

To downgrade the edition, complete the following steps.

1. Select **Help > License**.

The **License** dialog box displays.

2. Browse to the license key file (.xml) in the **License Key** field and click **Update**.

A message displays that details the support that will no longer be available after the license update.

2 License downgrade

3. Click **Yes** on the message to continue.

The client closes after updating the license successfully. Restart the server through the Server Management Console for the changes to take effect, then log back into the application.

After you downgrade from Enterprise to Professional Plus, the network size changes to small and all network size related parameters (such as, asset collection thread pool size and client and server heap size) are updated.

Patches

In this chapter

- [Installing a patch](#) 33
- [Uninstalling a patch.](#) 34

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:

3 Uninstalling a patch

- 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.

3 Uninstalling a patch

Discovery

In this chapter

- [SAN discovery overview](#) 37
- [Viewing the fabric discovery state](#) 47
- [Troubleshooting fabric discovery](#) 48
- [SAN Fabric monitoring](#) 51
- [SAN Seed switch](#) 55
- [Host discovery](#) 58
- [VM Manager discovery](#) 68

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 123.

NOTE

Professional edition can discover only 1 fabric.

NOTE

Professional Plus edition can discover up to 2,560 ports.

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 9](#) 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 9 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

TABLE 9 Backbone Chassis discovery

Device	Professional	Professional Plus	Enterprise
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
16 Gbps 4-slot Backbone Chassis as seed switch	Yes	Yes	Yes
16 Gbps 4-slot Backbone Chassis as member switch	Yes	Yes	Yes

Discovering fabrics

NOTE

Fabric OS devices must be running Fabric OS 5.0 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.

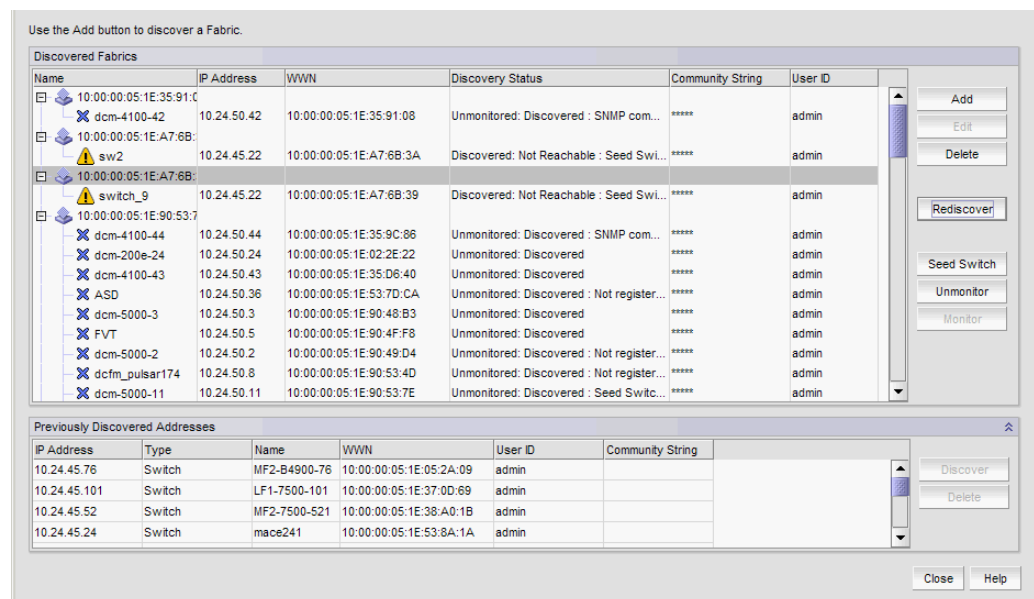


FIGURE 7 Discover Fabrics dialog box

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

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

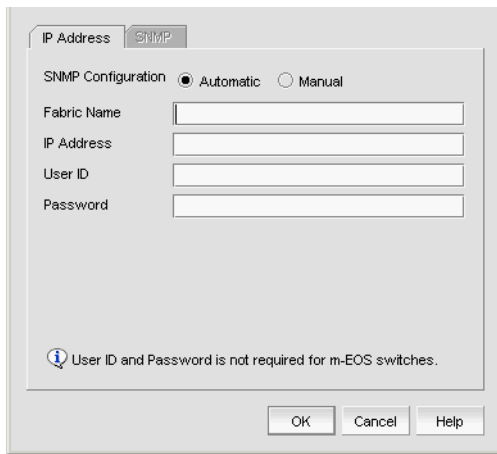


FIGURE 8 Add Fabric Discovery dialog box (IP Address tab)

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 123. 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 56.

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

Professional edition can discover only 1 fabric.

NOTE

Professional Plus edition can discover up to 2,560 ports.

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 547.

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 9 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 10 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.

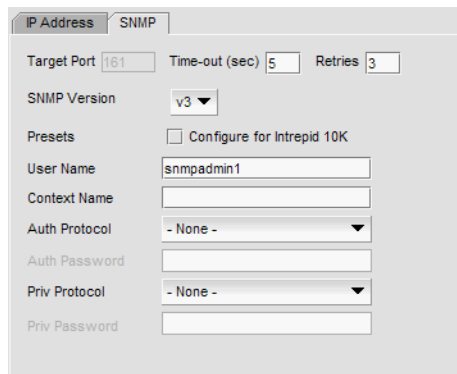


FIGURE 11 Add Fabric Discovery dialog box (SNMP tab)

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.

Go to step 7.
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 10 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 marginal.
	Displays when the fabric, switch, or host is not managed or not monitored.

4 Troubleshooting fabric discovery

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
<p>If you exceed your managed count limit, the Management application displays a “licensed exceeded” message on the topology.</p>	<p>Perform one or more of the following actions to</p> <ul style="list-style-type: none"> • “Changing your network size” • “Remove a device from active discovery” • “Deleting a fabric” <p>Changing your network size</p> <p>If you are at the maximum network size for your license, contact your preferred network provider.</p> <p>To change the size of your network, complete the following steps.</p> <ol style="list-style-type: none"> 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. <p>NOTE: Changes to this option take effect after an application restart.</p> <p>NOTE: You can only restart the server using the Server Management Console (Start > Programs > Management_Application_Name 12.X.X > Server Management Console).</p> <ol style="list-style-type: none"> 6 Click OK on the “changes take effect after application restart” message.
	<p>Remove a device from active discovery</p> <p>To remove a fabric from active discovery, complete the following steps.</p> <ol style="list-style-type: none"> 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.

4 Troubleshooting fabric discovery

Problem	Resolution
	<p>Deleting a fabric</p> <p>Before you can delete a fabric permanently from discovery, you must remove it from active discovery. Refer to “Remove a device from active discovery”.</p> <p>To delete a fabric permanently from discovery, complete the following steps.</p> <ol style="list-style-type: none"> 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
<p>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.</p> <p>At the time of discovery, the user does not have the Chassis Admin role for all other switches in the fabric.</p> <p>After discovery, a device is upgraded to Fabric OS 6.2 or later and is Virtual Fabric-enabled; however, the user does not have Chassis Admin role.</p>	<p>Make sure the user account has Chassis Admin role on the Fabric OS device.</p>
<p>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).</p> <p>At the time of discovery, the user does not have access to all possible logical switches for all other devices in the fabric.</p> <p>After discovery, a device is upgraded to Fabric OS 6.2 or later and is Virtual Fabric-enabled; however, the user does not have access to all possible logical switches.</p>	<p>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.</p>
<p>At the time of discovery, SNMP v3 is not configured.</p> <p>At the time of discovery, SNMP v3 is not configured for all other switches in the fabric.</p> <p>After discovery, a device is upgraded to Fabric OS 6.2 or later and is Virtual Fabric-enabled; however, SNMP v3 is not configured</p>	<p>Configure the SNMP v3 information for the Virtual Fabric-enabled device.</p>
<p>At the time of discovery or fabric refresh, the SNMP v3 user account does not have the Chassis Admin role.</p>	<p>Make sure the SNMP v3 user account has the Chassis Admin role on the Fabric OS device.</p>
<p>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).</p> <p>This access is required to obtain performance statistics from all logical switches.</p>	<p>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.</p>

Problem	Resolution
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 11 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 12 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 120.

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 **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 switch 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 5.X 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.

ATTENTION

If the seed switch firmware is downgraded from Fabric OS 5.2.X to an earlier version, then all RBAC-related data is discarded from the Management application.

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 5.0 or later. For a complete list of all supported Fabric OS hardware, refer to [“Supported hardware and software”](#) on page lii.

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 131.

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 Host 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.

NOTE

SMI and WMI discovery are not supported.

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.

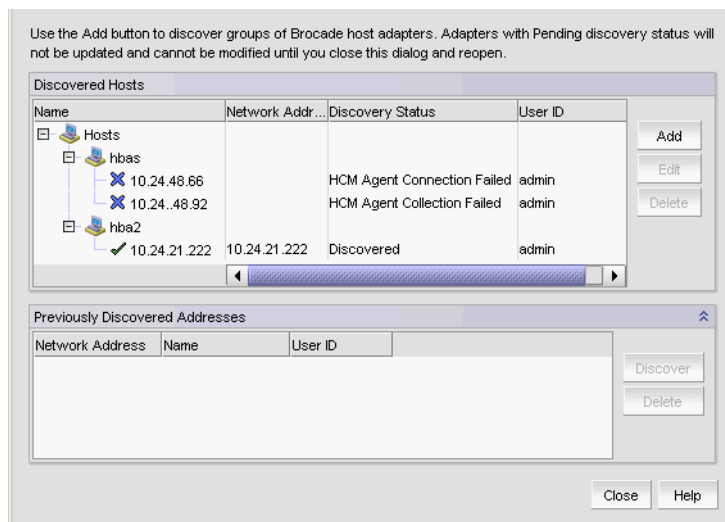


FIGURE 12 Discover Host Adapters dialog box

2. Click **Add**.

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

FIGURE 13 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](#).

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** from the **Protocol** list.
 - To configure discovery without authentication, select the **HTTP** from the **Protocol** list.
9. Enter the port number in the **Port** field.

HCM agent default is 34568. CIM server HTTPS default is 5989. CIM server HTTP default is 5988.

10. Enter your username in the **User ID** field.

HCM agent default is admin. Leave this field blank for the CIM server.
11. Enter your password **Password** field.

HCM agent default is password. Leave this field blank for the CIM server.
12. Repeat [step 5](#) through [step 11](#) for each Host you want to discover.

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 **Discovered Hosts** table with pending status. To update the status from pending you must close and reopen the **Discover Host Adapters** dialog box.

14. 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 14 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 address 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](#) .

If you do not need to configure Host credentials, skip to [step 13](#).

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

- To configure discovery with authentication, select the **HTTPS** from the **Protocol** list.
- To configure discovery without authentication, select the **HTTP** from the **Protocol** list.

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

HCM agent default is 34568. CIM server HTTPS default is 5989. CIM server HTTP default is 5988.

11. Enter your username in the **User ID** field.

HCM agent default is admin. Leave this field blank for the CIM server.

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

HCM agent default is password. Leave this field blank for the CIM server.

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 **Discovered Hosts** table with pending status. To update the status from pending you must close and reopen the **Discover Host Adapters** dialog box.

14. 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.

4 Host discovery

The screenshot shows a dialog box titled "Add Host Adapters". It contains the following fields and controls:

- Discovery Request Name:** A text input field.
- Network Address:** A dropdown menu followed by a text input field and an "Add" button.
- Host List:** A large empty list area, a "Remove" button, and an "Import" button.
- Contact:** Radio buttons for "HCM agent" (selected) and "CIM server (ESXi only)". Below is a "Protocol" dropdown menu set to "HTTPS".
- Port:** A text input field containing the value "34568".
- User ID:** A text input field.
- Password:** A text input field.
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom right.

FIGURE 15 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 list.
6. Click **Add**.

All hosts which are part of a managed fabric 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](#).

If you do not need to configure Host credentials, skip to [step 12](#).

8. Configure discovery authentication by choosing one of the following options:
 - To configure discovery with authentication, select the **HTTPS** from the **Protocol** list.
 - To configure discovery without authentication, select the **HTTP** from the **Protocol** list.
9. Enter the port number in the **Port** field.

HCM agent default is 34568. CIM server HTTPS default is 5989. CIM server HTTP default is 5988.

10. Enter your username in the **User ID** field.

HCM agent default is admin. Leave this field blank for the CIM server.
11. Enter your password **Password** field.

HCM agent default is password. Leave this field blank for the CIM server.

12. 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 **Discovered Hosts** table with pending status. To update the status from pending you must close and reopen the **Discover Host Adapters** dialog box.

13. 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 16 Add Host Adapters dialog box

3. Enter a discovery request name (such as, MyVMMManager) in the **Discovery Request Name** field.
4. Select **Hosts from VM Manager** from the import by list.
5. Select **All VM** or an individual VM from the list.
6. Click **Add**.

All hosts which 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](#).

If you do not need to configure Host credentials, skip to [step 12](#).

8. Configure discovery authentication by choosing one of the following options:
 - To configure discovery with authentication, select the **HTTPS** from the **Protocol** list.
 - To configure discovery without authentication, select the **HTTP** from the **Protocol** list.

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

HCM agent default is 34568. CIM server HTTPS default is 5989. CIM server HTTP default is 5988.

10. Enter your username in the **User ID** field.

HCM agent default is admin. Leave this field blank for the CIM server.

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

HCM agent default is password. Leave this field blank for the CIM server.

12. 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 **Discovered Hosts** table with pending status. To update the status from pending you must close and reopen the **Discover Host Adapters** dialog box.

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

Editing Host adapter credentials

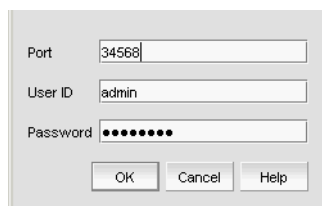
To edit Host credentials, complete the following steps.

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

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

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

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



The image shows a dialog box titled "Edit Host Discovery". It has three text input fields. The first is labeled "Port" and contains the number "34568". The second is labeled "User ID" and contains the text "admin". The third is labeled "Password" and contains a series of dots. Below the input fields are three buttons: "OK", "Cancel", and "Help".

FIGURE 17 Edit Host Discovery 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](#).

If you do not need to configure Host credentials, skip to [step 8](#).

4. Configure discovery authentication by choosing one of the following options:
 - To configure discovery with authentication, select the **HTTPS** from the **Protocol** list.
 - To configure discovery without authentication, select the **HTTP** from the **Protocol** list.
5. Enter the port number in the **Port** field.
HCM agent default is 34568. CIM server HTTPS default is 5989. CIM server HTTP default is 5988.
6. Enter your username in the **User ID** field.
HCM agent default is admin. Leave this field blank for the CIM server.
7. Enter your password **Password** field.
HCM agent default is password. Leave this field blank for the CIM server.
8. 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.
9. 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 previously discovered fabric

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.
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 adapter 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 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.

TABLE 13 Discovery Status Icons

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.

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
- New Discovery Pending
- Created host structure differs from discovered host; Discovery ignored
- Brocade HBA Discovery Failed: HCM Agent connection failed
- HCM Agent collection failed
- CIM Server Authentication failed
- CIM Server connection failed

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 68, [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 68, [step 6](#)), requires a vCenter user with, at minimum, the following read-write privileges on the vCenter server node and all objects in the inventory below the vCenter server:
 - Extension > Register extension
 - Extension > Unregister extension
 - Extension > Update extension

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 68).

To discover a VM manager, complete the following steps.

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

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




FIGURE 18 Discover VM Managers dialog box

2. Click **Add**.

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




FIGURE 19 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.

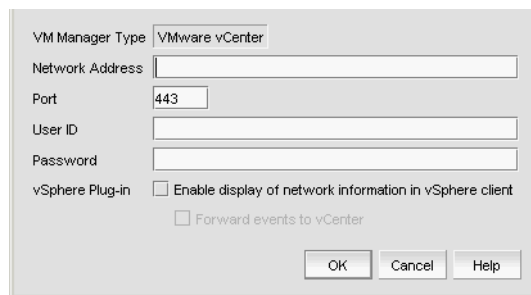


FIGURE 20 Edit VM Manager dialog box

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`.

4 VM Manager discovery

Application Configuration

In this chapter

• Server Data backup	77
• Server Data restore	83
• SAN display settings	84
• SAN End node display	86
• SAN Ethernet loss events	87
• Event storage settings	88
• Flyover settings	89
• Name settings	92
• Miscellaneous security settings	101
• Syslog Registration settings	103
• SNMP Trap Registration settings	104
• SNMP Trap forwarding credential settings	105
• Software Configuration	106
• FIPS Support	131
• Fabric tracking	131

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 88.
- **Flyovers** — Use to customize the properties display in product and connection flyovers. For more information, refer to [“Flyover settings”](#) on page 89.
- **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 17.
- **Performance Graph Style** — Use to configure the color scheme and to display data points for all performance graphics in the management application. For more information, refer to [“Performance Data”](#) on page 935.
- **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.

5 Configurable preferences

- 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 86.
- 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 87.
- SAN Names — Use to set whether unique names are required. For more information, refer to [“Name settings”](#) on page 92.
- Miscellaneous Security — Use to configure server security configurations and the login banner. For more information, refer to [“Miscellaneous security settings”](#) on page 101.
- 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 77 and [“Server Data restore”](#) on page 83.
- Syslog Registration — Use to automatically register the server as the syslog recipient on products. For more information, refer to [“Syslog Registration settings”](#) on page 103.
- 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 104.
- 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 105.
- Certificates — Use to manage keystore and truststore certificates as well as enable or disable certificate validation. For more information, refer to [“Certificates”](#) on page 107.
- 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 113.
- Client/Server IP — Use to configure IP address of the Management application server. For more information, refer to [“Client/Server IP”](#) on page 114.
- Memory Allocation — Use to configure memory allocation for the client and server. For more information, refer to [“Memory allocation settings”](#) on page 118.
- 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 122.
- 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 123.
- Server Port — Use to configure server port settings. For more information, refer to [“Server port settings”](#) on page 128.
- Support Mode — Use to configure support settings to enable enhanced diagnostics. For more information, refer to [“Support mode settings”](#) on page 129.

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

Backing up data takes some time. 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 three options for backing up data to the Management server:

- Configuring backup to a CD drive
- Configuring backup to a hard drive
- Configuring backup to a network drive

The Management server is backed up to D:\Backup (Windows systems) by default. If there is not second hard disk, this is a rewritable (CD-RW) compact disk. Make sure you have a CD-RW disk in the CD recorder drive to ensure that backup can occur. Critical information from the Management application is automatically backed up to the CD-RW when the data directory contents change or when you restart the Management application.

Note that backing up to CD is not the recommended method. The usable capacity of a CD is approximately 700 MB and needs to be replaced when full. Also, CD media has a limited number of re-writes before the medium is exhausted, and write errors occur. It is recommended that you configure the backup system to target a hard drive or a network drive as described in the procedures below.

Back up 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 CD-RW. 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 21) with the currently defined directory displays in the **Backup Output Directory** field.

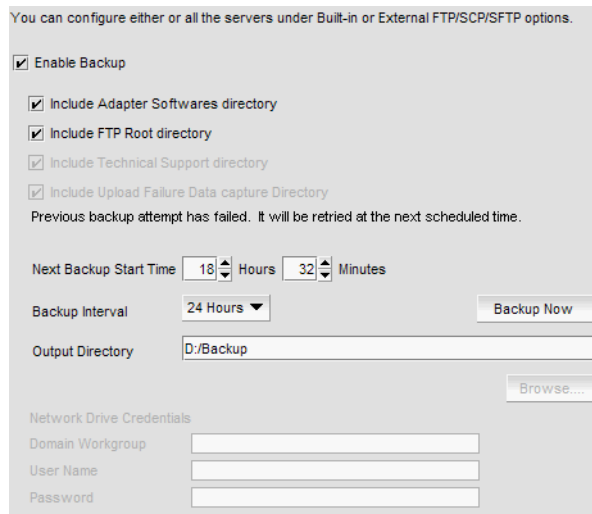


FIGURE 21 Options dialog box (Server Backup pane)

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. Back up data to a hard drive by browsing to the hard drive and directory to which you want to back up your data.

NOTE

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

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

To back up 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 back up 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. Back up data to a CD by completing the following steps.

NOTE

This is not recommended on a permanent basis. CDs have a limited life, and may only last a month. An error message occurs if your Management application can no longer back up to the disc.

- a. Verify that the CD backup directory is correct (default directory is D:\Backup).

It is assumed that drive D is a CD-RW drive.

You can change the directory or use the **Browse** button to select another directory.

- b. Install the formatted disc into the CD drive.

To back up to a writable CD, you must have CD-writing software installed. The disc must be formatted by the CD-writing software so that it behaves like a drive.

10. Click **Apply** or **OK**.

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

For back up to a hard drive or writable CD, if the device does not exist or is not writable, an error message displays that says you have entered an invalid device.

For back up 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 1243.

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 [“Back up directory structure overview”](#) on page 78).

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 5.

Restoring data to a new server

NOTE

The restore data files must use the exact directory structure as the backup directory structure (refer to [“Back up directory structure overview”](#) on page 78).

If your Management application server fails and you must recover information to a new server, restore the data (Refer to [“Restoring data”](#) on page 83 for complete instructions).

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 22](#)).

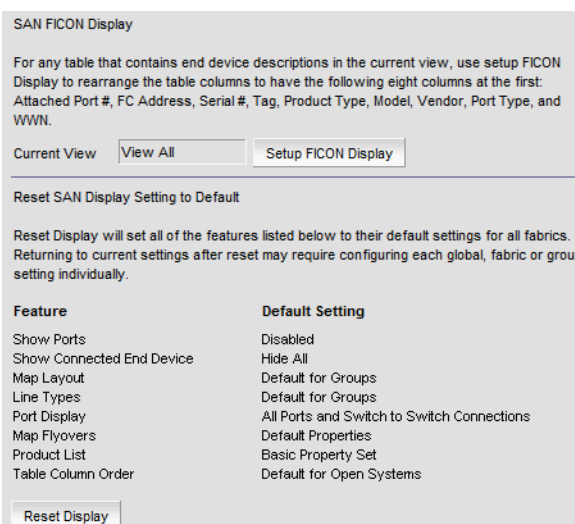


FIGURE 22 Options dialog box (SAN Display pane)

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 15) details the settings that change with reset and the associated default state.

TABLE 15 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: <ul style="list-style-type: none"> 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 Property 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 15)).

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 23).

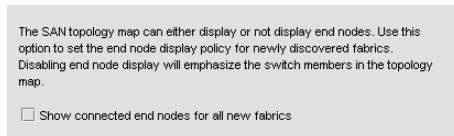


FIGURE 23 Options dialog box (SAN End Node Display pane)

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 24 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 25).

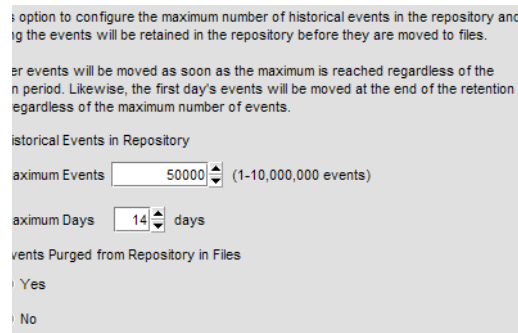


FIGURE 25 Options dialog box (Event Storage pane)

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 26) and completing the following steps.

5 Flyover settings

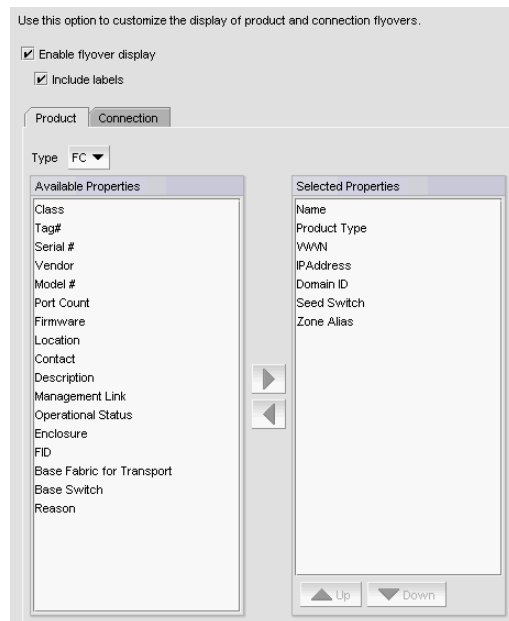


FIGURE 26 Options dialog box (Flyovers pane, Product tab)

- 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 27) and completing the following steps.




FIGURE 27 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

5 Name settings

FCoE

- Name
- Node WWN
- MAC
- Port#
- Port Type
- FCoE Index #

- Click the right arrow to move the selected properties to the **Selected Properties** table.
- 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.

- 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.
- 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 28).

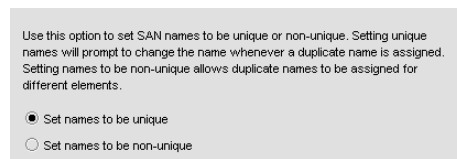


FIGURE 28 Options dialog box (SAN Names pane)

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 93.

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 29).

5 Name settings

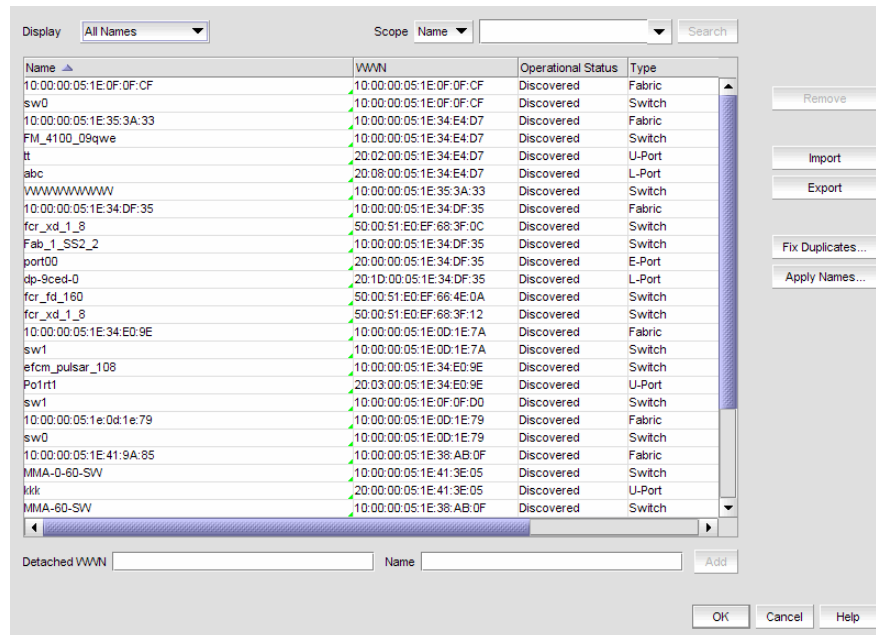


FIGURE 29 Duplicated Names dialog box

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 30).




FIGURE 30 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 99.
- **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 97.
- **Import** button — Click to import name data. For more information, refer to [“Importing Names”](#) on page 99.

- **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 98.
- **Fix Duplicates** button — Click to launch the Fix Duplicates dialog box. For more information, refer to [“Fixing duplicate names”](#) on page 93.
- **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 97.
- **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 97.

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 99 or [“Searching for a device by WWN”](#) on page 100 in the **Configure Names** dialog box or [“Searching for a device”](#) on page 269 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 16 Character limits for names

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)
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 269.

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 269.

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 31).

Use this option to configure various security configurations applicable to the server.

Server Name

CHAP Secret

Retype Secret

Login Security

Display login banner upon client login

Banner Message

This login banner can be configured to adhere to your corporate security policies

Use this option to enforce the MD5 checksum file import while importing the Fabric OS image into the repository.

Enforce Fabric OS MD5 Checksum File Import

FIGURE 31 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 32](#)).




FIGURE 32 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 32).
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 33).




FIGURE 33 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 34).

Use this option to configure the SNMP credentials for the traps forwarded by this server

SNMP v1 / v2c

Community:

Confirm Community:

SNMP v3

User Name:

Context Name:

Auth Protocol:

Auth Password:

Confirm Password:

Priv Protocol:

Priv Password:

Confirm Password:

Engine ID:

FIGURE 34 Options dialog box (Trap Forwarding Credentials pane)

3. Enter the unique community string (case sensitive, 1 to 16 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 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 34).
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 one 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](#))
 - 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 122.

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 35](#)).

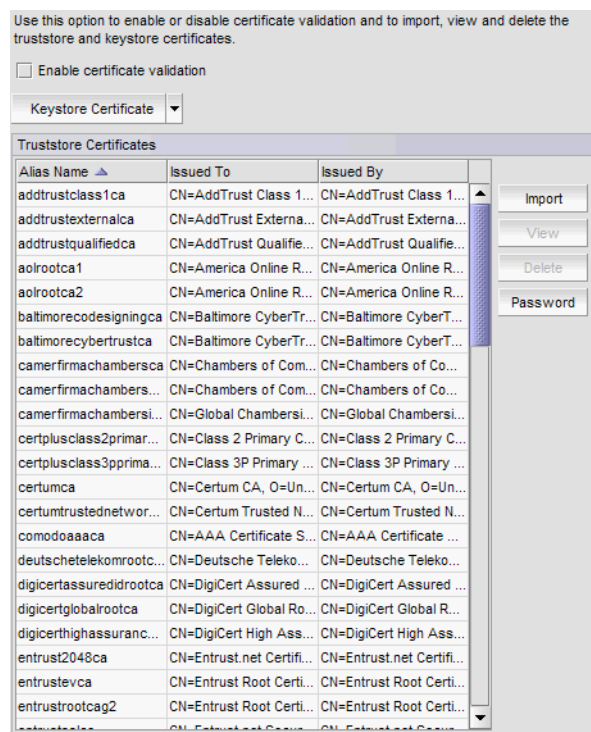


FIGURE 35 Options dialog box (Certificates pane)

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 108.
 - **Export** – Click to export a keystore certificate. For more information, refer to [“Importing a truststore certificate”](#) on page 109.
 - **Replace** – Click to replace the keystore certificate. For more information, refer to [“Deleting a truststore certificate”](#) on page 110.
 - **Change Password** – Click to change the password for the keystore. For more information, refer to [“Changing the keystore password”](#) on page 112.
 - **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 109.
 - **View** button – Click to view the certificate details. For more information, refer to [“Viewing a truststore certificate”](#) on page 108.
 - **Delete** button – Click to delete the certificate. For more information, refer to [“Deleting a truststore certificate”](#) on page 110.
 - **Password** button – Click to change the password for the truststore. For more information, refer to [“Changing the password for the truststore repository”](#) on page 110.
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 36](#)).

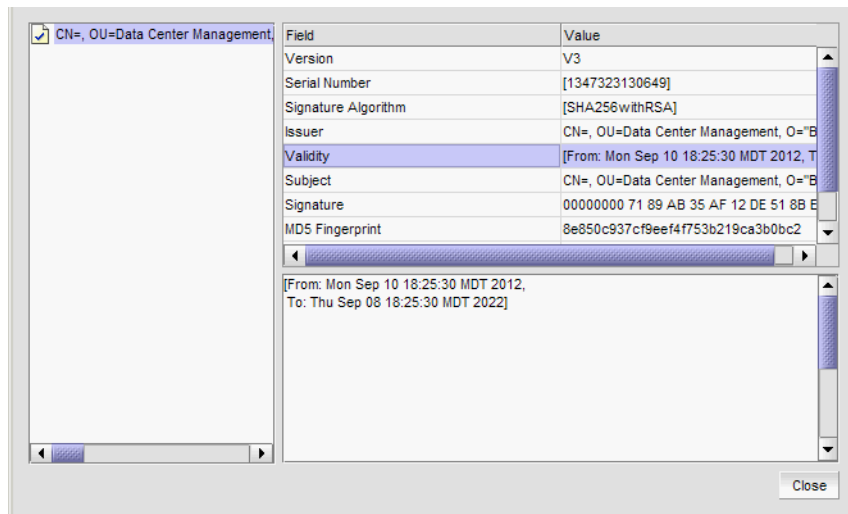


FIGURE 36 Details - Certificate *Name* dialog box

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

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 112.

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 122.

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.

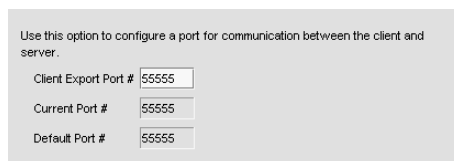
Client export port settings

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 37](#)).



Use this option to configure a port for communication between the client and server.

Client Export Port #	55555
Current Port #	55555
Default Port #	55555

FIGURE 37 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 38](#)).

Use this option to configure the IP Configuration settings.

Server IP Configuration: All

Default: All

Server IP: 10.25.224.133

Server Name: 5A11-16233234

Client - Server IP Configuration

Return Address: 5A11-16233234

Current Return Address: 5A11-16233234

Switch - Server IP Configuration

Preferred Address: 10.25.224.133

If DNS is not configured in your network, do not choose the Return Address as hostname and the Network Advisor Server IP must bind with the host IP Address and not the hostname.

FIGURE 38 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.

4. Select the return IP address in the **Client - Server IP Configuration Return Address** list.
5. Select the preferred IP address in the **Switch - Server IP Configuration Preferred Address** list.

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.

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 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 39).

FIGURE 39 Server IP Configuration screen

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

NOTE

If the “hostname” contains invalid characters, the host name does not display in the list. Valid characters include alphanumeric and dash (-) characters. The IP address is selected by default. If the an IPv6 address is selected, server start up will fail.

If DNS is not configured for your network, do not select the ‘hostname’ option from either the **Server IP Configuration** or **Switch - Server IP Configuration Preferred Address** list. Selecting the ‘hostname’ option prevents clients and devices from communicating with the Server.

- 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 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.

3. (Enterprise only) In the **SAN Network Size is** list, complete the following steps:

For other editions, the SAN Network size is small. You cannot change the SAN 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 — 90 products, 5,000 ports

- Large SAN — 160 products, 9,000 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.

For all network sizes, the default minimum Client Heap Size is 950 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 where 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.

To support more than 8 clients on a 64-bit server, increase the memory allocation for the server to 3076 MB.

Minimum values are as follows:

For a 32-bit Windows/Linux Server:

- Professional: 768 MB
- Professional Plus: 1024 MB
- Enterprise Small : 768 MB
- Enterprise Medium : 1024 MB
- Enterprise Large : 1024 MB

Default values for SAN only Server (**Server Heap Size**):

For a 32-bit Windows/Linux Server:

- Small : 768 MB
- Medium : 1024 MB
- Large : 1024 MB

For all 64-bit servers, the default minimum Server Heap Size for all network sizes is 2048 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.

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 (balance 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.

3. Enter how often you want to check for state changes in the **Check for state change every** field.

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 (Professional): 60 seconds
- Medium: 120 seconds
- Large: 180 seconds

4. Enter how often you want to check for state changes in the **If no state change, Poll switch every** field.

Valid values are from 1 through 3,600 seconds. Default values are as follows:

- Small (Professional): 120 seconds
- Medium: 900 seconds
- Large: 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.

The following table illustrates and describes the icons that indicate the current network size status.

TABLE 17

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: <ul style="list-style-type: none"> • 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 40).

Use this option to configure HTTP or HTTPS connections between the Network Advisor Server and SAN switches.

Connect using HTTP HTTPS (HTTP over SSL) only

Port #

Current Port #

Default Port #

Use this option to configure connections between the Network Advisor Server and IP Products.

Product Communication

SSH only Telnet only SSH then Telnet SSH Port

Configuration File Transfers

SCP only TFTP only SCP then TFTP TFTP then SCP

Web Element Manager

HTTP HTTPS HTTPS then HTTP

Use this option to set the user preferred IP format for the Network Advisor to connect with the products.

User Preferred IP Format (SAN and Network OS products only)

IPv4 IPv6

FIGURE 40 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. Go to [step 5](#).
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** 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 5](#).
The default HTTPS port number is 443.
5. Select **IPv4** or **IPv6** to set the preferred IP format.

6. Click **Apply** or **OK** to save your work.
Changes to this option take effect after an application restart.
7. Click **OK** on the “changes take effect after application restart” message.

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 40).
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.
Changes to this option take effect after an application restart.
5. Click **OK** on the “changes take effect after application restart” message.

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 5.3 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 11.

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 41).

	Value
<input type="checkbox"/> Built-in FTP Server	<input checked="" type="checkbox"/>
User Name	admin
Password	••••••••
Confirm Password	••••••••
<input type="checkbox"/> SCP / SFTP Server	<input checked="" type="checkbox"/>
User Name	admin
Password	••••••••
Confirm Password	••••••••
Preferred Protocol	SCP
Root Directory	C:\Program Files\Network Advisor 12.0.0\data\ftp\root

FIGURE 41 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 passwOrd (where 0 is a zero).

- 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.

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

Configuring an internal SCP or SFTP server

NOTE

SCP is supported on Fabric OS devices running 5.3 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.

- Select **Server > Options**.

The **Options** dialog box displays.

- Select **FTP/SCP/SFTP** in the **Category** list.

The **FTP/SCP/SFTP** pane displays (Figure 41).

	Value
<input type="checkbox"/> Built-in FTP Server	<input checked="" type="checkbox"/>
User Name	admin
Password	••••••••
Confirm Password	••••••••
<input type="checkbox"/> SCP / SFTP Server	<input checked="" type="checkbox"/>
User Name	admin
Password	••••••••
Confirm Password	••••~•••
Preferred Protocol	SCP
Root Directory	C:\Program Files\Network Advisor 12.0.0\data/ftproot

FIGURE 42 Options dialog box (FTP/SCP/SFTP pane)

- 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.

- Select the **SCP/SFTP Server** check box.
- Change your password by entering a new password in the **Password** and **Confirm Password** fields.

The default password is passwOrd (where 0 is a zero).

- 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 5.3 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.

The **FTP/SCP/SFTP** pane displays (Figure 43).

	Value
<input type="checkbox"/> FTP Server	<input type="checkbox"/>
FTP Host IP	
FTP Host User Name	
FTP Directory Path	
Password for FTP	
<input type="checkbox"/> SCP Server	<input type="checkbox"/>
SCP Host IP	
SCP Host User Name	
SCP Directory Path	
Password for SCP	
<input type="checkbox"/> SFTP Server	<input type="checkbox"/>
SFTP Host IP	
SFTP Host User Name	
SFTP Directory Path	
Password for SFTP	
Preferred Protocol(Secured)	SCP

FIGURE 43 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 124.
 - 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 126.
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 44](#)).

Use this option to configure the server port settings. On enabling HTTP redirection, port # 80 is used to redirect HTTP requests to HTTPS.

Server IP	10.25.224.20
Server Name	TechOPS2008
Web Server Port # (HTTPS)	443
Current Port #	443
Default Port #	443
Redirect HTTP Requests to HTTPS	<input checked="" type="checkbox"/>
The server requires 18 consecutive free ports	
Starting Port #	24600

FIGURE 44 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 45).

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 44).

The screenshot shows a dialog box titled "Use this to configure support settings for enhanced diagnostics." It contains two sections: "Log client support data" and "Log server support data".

- Log client support data:** Log Level is set to INFO (dropdown menu).
- Log server support data:** Log Level is set to INFO (dropdown menu), Log Purging Limit is 14 (spin box), Server IP is 10.25.224.133 (text box), and Server Name is 5A11-16233234 (text box).

FIGURE 45 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 122) 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 126) 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



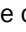



When you enable fabric tracking and a switch is missing from the fabric, a warning level call home event (Switch *Switch_WWN* is missing from the Fabric *Fabric_Name*) is generated in the Master Log and a call home alert is sent to the corresponding call center for this event.

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 300.

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 (Figure 46). This message includes the following information:

The below listed switches, devices and connections with - as status will be removed and + as status will remain in the respective fabrics.
Do you want to continue?
 Do not show me this again

Switches					
Status ▲	Name	Fabric Name	IP Address	WWN	Domain ID

Device Ports							
Status ▲	Reason	Product Type	Port	Fabric Name	Port WWN	Node WWN	Attached Port #
-		Target	22:00:00:04:CF:BD:71:1B	10:00:00:05:1E:90:1B:27	22:00:00:04:CF:BD:71:1B	20:00:00:04:CF:BD:71:1B	20:05:00:05:1E:90:52:FA

Connections											
Status ▲	Reason	Type	Fabric Name	1-WWN	1-Domain ID	1-IP Address	1-Port	2-WWN	2-Domain ID	2-IP Address	2-Port

Yes No

FIGURE 46 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 (Figure 47). This message includes the following information:

The below listed switches, devices and connections with ● as status will be removed and ● as status will remain in the respective fabrics.
Do you want to continue?
 Do not show me this again

Switches					
Status	Name	Fabric Name	IP Address	WWN	Domain ID
●	sw_45_nameadded123568	10:00:00:05:1E:38:A0:1B	10.24.45.13	10:00:00:05:1E:A6:C2:E6	25
●	switch_92	10:00:00:05:1E:38:A0:1B	10.24.45.92	10:00:00:05:1E:40:40:00	33
●	sw01	10:00:00:05:1E:38:A0:1B	10.24.45.95	10:00:00:05:1E:4B:AA:00	2



Device Ports							
Status	Reason	Product Type	Port	Fabric Name	Port WWN	Node WWN	Attached Port #
●		Target	22:00:00:04:CF:BD:71:1B	10:00:00:05:1E:90:1B:27	22:00:00:04:CF:BD:71:1B	20:00:00:04:CF:BD:71:1B	20:05:00:05:1E:90:52:FA
●		Initiator	10:00:00:05:1E:56:5F:B1	10:00:00:05:1E:38:A0:1B	10:00:00:05:1E:56:5F:B1	20:00:00:05:1E:56:5F:B1	50:00:53:31:CA:F0:5A:F6
●		Initiator	10:00:00:05:33:26:88:3E	10:00:00:05:1E:38:A0:1B	10:00:00:05:33:26:88:3E	20:00:00:05:33:26:88:3E	50:00:53:31:CA:F0:5A:F2
●		Target	1B:86:00:11:0D:06:00:00	10:00:00:05:1E:38:A0:1B	1B:86:00:11:0D:06:00:00	1B:86:00:11:0D:06:00:00	gdgdfg
●		Initiator	10:00:00:05:33:26:6C:E5	10:00:00:05:1E:38:A0:1B	10:00:00:05:33:26:6C:E5	20:00:00:05:33:26:6C:E5	20:C2:00:05:1E:4B:AA:00

Connections										
Status	Reason	Type	Fabric Name	1-WWN	1-Domain ID	1-IP Address	1-Port	2-WWN	2-Domain ID	2-IP
●		ISL	10:00:00:05:1E:38:A0:1B	10:00:00:05:1E:40:40:00	33	10.24.45.92	slot11 port13	10:00:00:05:1E:4B:AA:00	2	10.
●		ISL	10:00:00:05:1E:38:A0:1B	10:00:00:05:1E:A6:C2:E6	25	10.24.45.13	Testing1234567	10:00:00:05:1E:4B:AA:00	2	10.
●		ISL	10:00:00:05:1E:38:A0:1B	10:00:00:05:1E:40:40:00	33	10.24.45.92	slot11 port37	10:00:00:05:1E:4B:AA:00	2	10.
●		ISL	10:00:00:05:1E:38:A0:1B	10:00:00:05:1E:40:40:00	33	10.24.45.92	slot9 port15	10:00:00:05:1E:4B:AA:00	2	10.

Yes No

FIGURE 47 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.

5 Fabric tracking

User Account Management

In this chapter

- [Users overview](#) 137
- [User accounts](#) 140
- [Roles](#) 146
- [Areas of responsibility](#) 150
- [Password policies](#) 154
- [Authentication Server Groups on the Management server](#) 156
- [User profiles](#) 161

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.

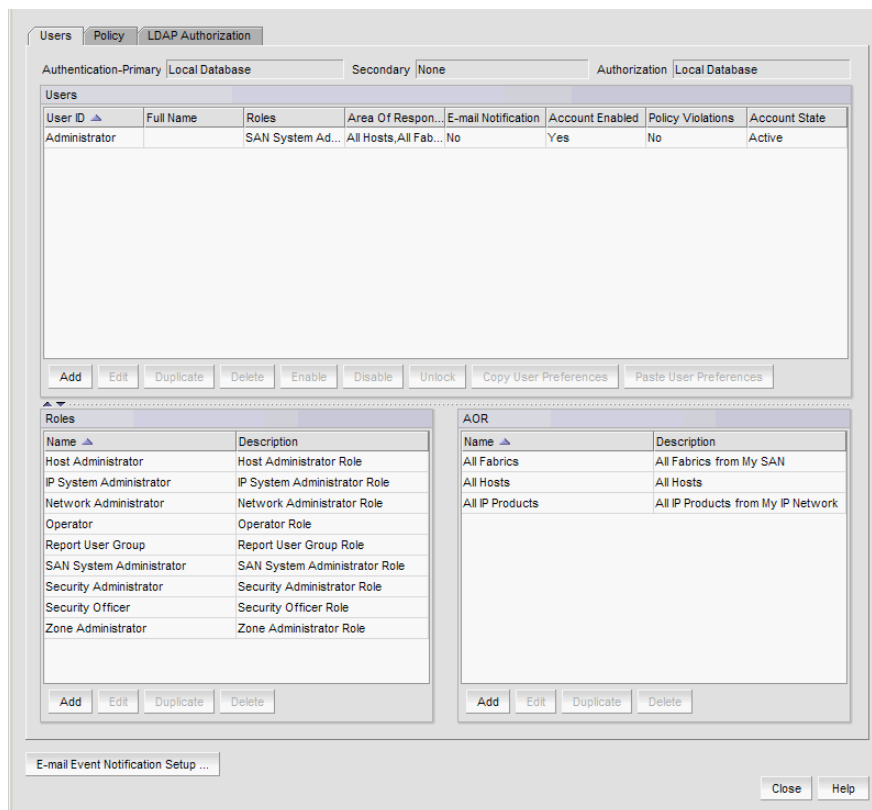


FIGURE 48 Users dialog box - Users tab

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 identify a user.
 - **Full Name** – The user’s full name.
 - **Roles** – List of Roles the user belongs to separated by comma.
 - **Area Of Responsibility** – List of AORs the user belongs to separated by comma.
 - **E-mail Notification** – Whether e-mail notification is enabled for 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
 - **Add** button – Click to launch the **Add Users** dialog box and configure a new user (refer to [“Creating a new user account”](#) on page 140).
 - **Edit** button – Click to launch the **Edit Users** dialog box for the selected user (refer to [“Editing a user account”](#) on page 142).
 - **Duplicate** button – Click to launch the Duplicate Users dialog box for the selected user (refer to [“Copying a user account”](#) on page 143).
 - **Delete** button – Click to delete the selected users (refer to [“Deleting a user account”](#) on page 145).
 - **Enable** button – Select to enable the selected users (refer to [“Enabling a user account”](#) on page 145). Disabled if the selected user is already enabled.
 - **Disable** button – Select to disable the selected users (refer to [“Disabling a user account”](#) on page 145). Disabled if the selected user is already disabled.
 - **Unlock** button – Select to unlock the selected users account (refer to [“Unlocking a user account”](#) on page 146).
 - **Copy User Preferences** button – Select to copy user preference from the selected users account (refer to [“Copying and pasting user preferences”](#) on page 143).
 - **Paste User Preferences** button – Select to paste user preference from the selected users account (refer to [“Copying and pasting user preferences”](#) on page 143).
- **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 146).
 - **Edit** button – Click to edit the selected role (refer to [“Editing a role”](#) on page 147).
 - **Duplicate** button – Click to copy the selected role (refer to [“Copying a role”](#) on page 148).
 - **Delete** button – Click to delete the selected role (refer to [“Deleting a role”](#) on page 148).
 - **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 164).
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 1243.

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 an 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 49 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 163.
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 1066.

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 into 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 140.
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 140.

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

Enables you to 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 “[Assigning roles and areas of responsibility to a user account](#)” on page 144).

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 142.

4. Click **Yes** on the confirmation message.
5. Click **Close** to close the **Users** dialog box.

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 **Users** 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 **Users** 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.

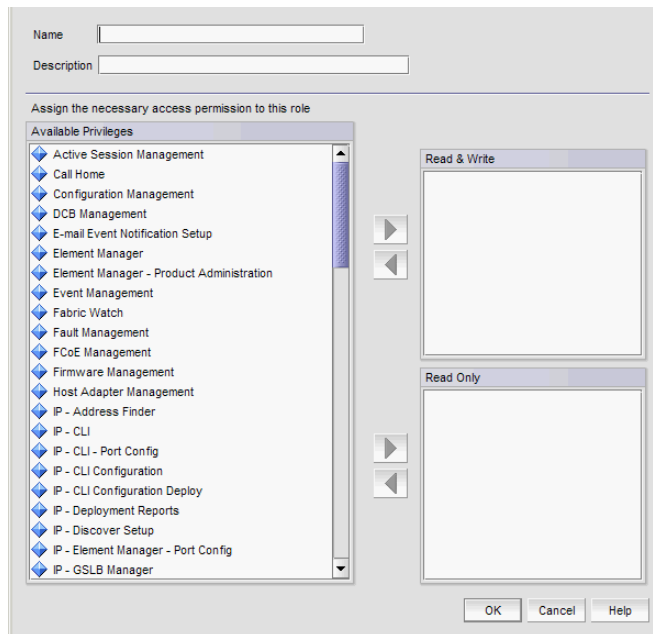


FIGURE 50 Add Role dialog box

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 148 or [“Removing privileges from a role”](#) on page 149.

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 [“Assigning roles and areas of responsibility to a user account”](#) on page 144.

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 146.
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 146.
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 “[Assigning roles and areas of responsibility to a user account](#)” on page 144.
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 FabricSand Hosts into management groups that can be assigned to an Management application user. Users can manage only the FabricSand 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.

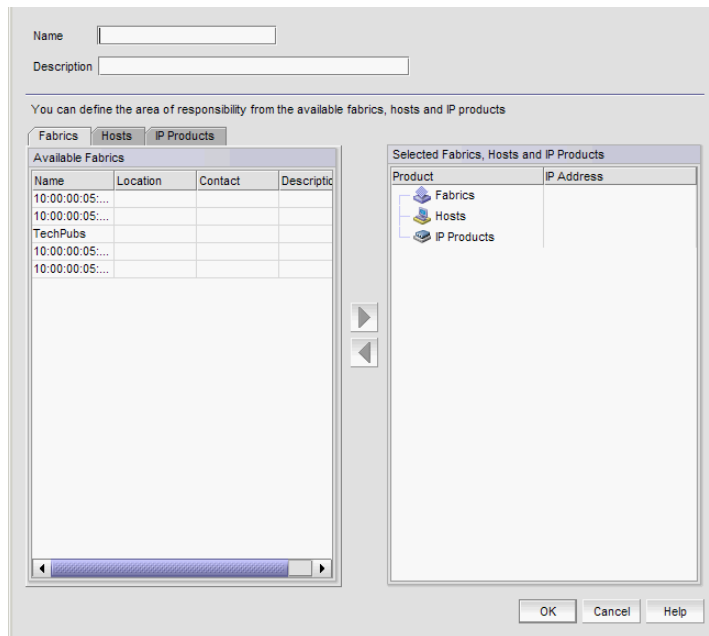


FIGURE 51 Users dialog box - Users tab

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 153 or [“Removing products from an AOR”](#) on page 153.

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 150.

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 150.

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 “[Assigning roles and areas of responsibility to a user account](#)” on page 144.

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 an 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**.

6 Authentication Server Groups on the Management server

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.

Authentication Server Groups on the Management server

NOTE

You must have User Management Read and Write privileges to map roles and AORs to Active Directory (AD) groups.

NOTE

You must configure an 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 331).

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.

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 158.

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.

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.

Default is 389 if security is not enabled. Default is 636 if security is enabled.

6. Select the authentication protocol **MD5** from the **Authentication** 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. 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 server (Local database).

To assign or remove roles and AORs, refer to [“Assigning roles and AORs to an AD group”](#) on page 157 or [“Removing roles and AORs from an AD group”](#) on page 157.

11. 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.

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 XP, select **Start > Programs > Administrative Tools > Active Directory Users and Computers**.
2. Right-click the **Users** folder and select **New > User**.
3. Enter a name in the **First name** field.
4. Enter a name in the **Full name** field.
5. Enter a logon name in the **User logon name** field.
6. Click **Next**.
7. Select the **Password Never Expires** option and click **Next**.
8. Click **Finish**.
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 XP, select **Start > Programs > Administrative Tools > Active Directory Users and Computers**.
2. Right-click the new user in the **Users** pane and select **Add to a Group**.
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**.

6 Authentication Server Groups on the Management server

3. Select **File > Add/Remove Snap-in**.
4. Click **Add**.
5. Select **Active Directory Schema** from the **Available standalone snap-ins list** and click **Add**.
6. Click **Close**.
7. Right-click the **Attributes** folder (Console Root/Active Directory Schema/ Attributes) and select **New > Attribute**.
8. 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**.
9. Right-click the **Attributes** folder (Console Root/Active Directory Schema/ Attributes) and select **New > Attribute**.
10. 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**.
11. Close the Management console.

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**.
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 159.
4. Select **NmAors** in the **Attributes** list and click **Edit**.
5. Enter the areas of responsibility (such as, All Fabrics, All IP Products) 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, IP System 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.

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 162.

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 163.
- **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 164).
 - **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 163).
 - **Optional Messages Reset** button — Click to reset all optional messages to the default behavior. For more information, refer to [“Resetting optional messages”](#) on page 164.
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 1066.

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 1066.
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.

6 User profiles

Dashboard Management

In this chapter

- [Dashboard overview](#) 167
- [Default dashboards](#)..... 179
- [Status widgets](#)..... 180
- [Monitoring and Alerting Policy Suite widgets](#)..... 190
- [Performance monitors](#) 194
- [User-defined performance monitors](#) 222
- [Traffic flow dashboard monitors](#)..... 238

Dashboard overview

NOTE

Only devices in your area of responsibility (AOR) display in the dashboard.

The **Dashboard** tab ([Figure 52](#)) 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 and 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 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.

7 Dashboard overview

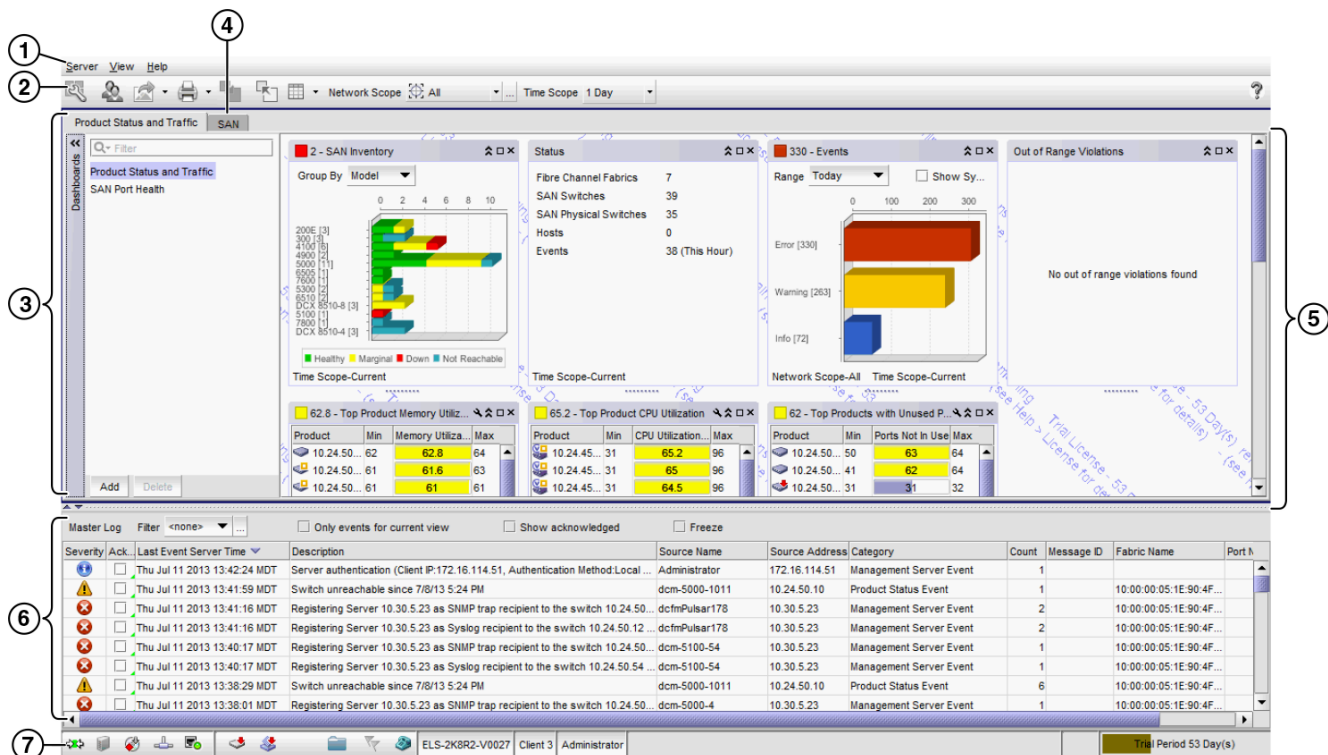


FIGURE 52 Dashboard tab

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 1207.
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 169.
3. **Dashboard tab** — Provides a high-level overview of the network managed by Management application server.
4. **SAN tab** — Displays the Master Log, Minimap, Connectivity Map (topology), and Product List. For more information, refer to the “[SAN tab overview](#)”.
5. **Dashboard expand navigation bar** — 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 170.
6. **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 180 and “[Performance monitors](#)” on page 194.
7. **Master Log** — Displays all events that have occurred on the Management application. For more information, refer to “[Master Log](#)” on page 255.

8. **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 257.

Dashboard toolbar

The toolbar ([Figure 53](#)) is located beneath the menu bar and provides icons and buttons to perform various functions.




FIGURE 53 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 and **Performance Dashboard**. For more information, refer to [“Customizing the dashboard widgets and monitors”](#) on page 173
2. **Users** — Displays the **Users** dialog box. Use to configure users, user groups, and permissions. For more information, refer to [“User accounts”](#) on page 140.
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 175.
4. **Print list** — Prints the dashboard display (all widgets) or a selected widget. For more information, refer to [“Printing the dashboard display”](#) on page 175.
5. **Attach** — Returns the dashboard to the main window. For more information, refer to [“Attaching and detaching the Dashboard tab”](#) on page 175.
6. **Detach** — Detaches the dashboard to a separate window. For more information, refer to [“Attaching and detaching the Dashboard tab”](#) on page 175.
7. **Dashboard display list** — Use to select how to display the status widgets and performance monitors in the dashboard. For more information, refer to [“Setting the dashboard display”](#) on page 173.
8. **Network Scope** — Use to select the network you want to display in the dashboard. For more information, refer to [“Setting the network scope”](#) on page 176.
9. **Network Scope ellipsis button** — Displays the **Edit Scopes** dialog box. Use to configure or delete product and port scopes. For more information, refer to [“Creating a customized network scope”](#) on page 177.
10. **Time Scope** list — Use to select the specific duration for which you want to display data. For more information, refer to [“Setting the data display time frame”](#) on page 179.
11. **Help** — Displays the online help.

Dashboard messages

The dashboard message bar (Figure 54) only displays when the Network Scope or Time Scope has changed. You can also view all dashboard messages and clear them.




FIGURE 54 Dashboard message bar

The toolbar contains the following fields and 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 55) 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 55 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 171.
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 172.
4. **Delete** button — Use to delete the selected user-defined dashboard. For more information, refer to [“Deleting a user-defined dashboard”](#) on page 173.

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 are persisted in that dashboard. For example, if you customize a dashboard to display the **Events** widget and set the **Range to This Hour** in the **Dashboard** tab and set it to **Last 30 Days** in the **Performance Dashboard**, then these preferences persist 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 selected **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.

Accessing a dashboard

From the **Dashboards** expand navigation bar, double-click the dashboard you want to view. Options include:

- **IP Port Health** — Displays preconfigured IP performance monitors. You can display additional status widgets and performance monitors in this dashboard.
- **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.
- *User_defined* dashboard — Displays a user-defined dashboard.

The dashboard you selected displays.

Filtering the dashboards list

You can filter the list of dashboards to only display 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.
4. Select the **Copy active dashboard widgets** to include all widget in the current dashboard to this dashboard.

5. Click **OK**.

The new dashboard displays in the **Dashboards** expand navigation bar and becomes the active dashboard.

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.

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

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 180.
- **Description** – A general description of the status widget.

- Click the **Performance** tab (Figure 56).

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 194.

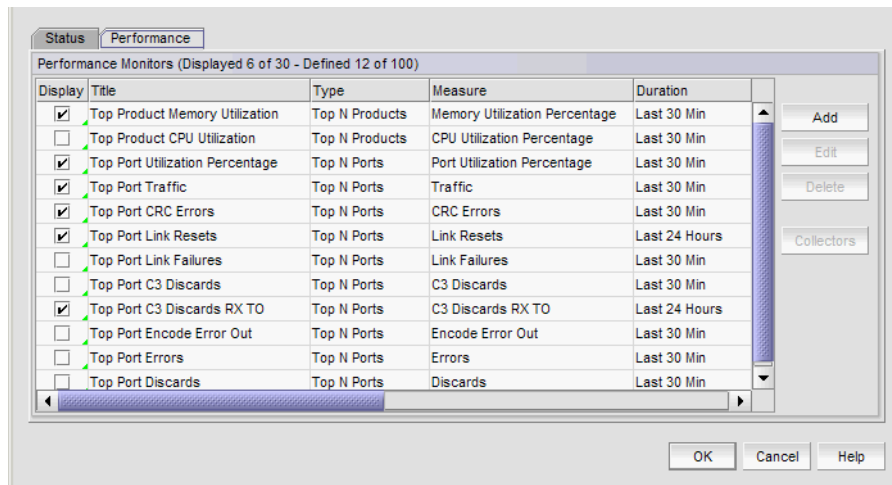


FIGURE 56 Customize Dashboard dialog box, Performance tab

- 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 194
 - **Type** – The type of monitor.
 - **Measure** – The performance measures included in the monitor.
 - **Data Collectors** – The data collectors that provide data for the monitor.
- Click **Add** to add a new performance monitor. For more information, refer to “[Configuring a user-defined product performance monitor](#)” on page 229.
 - Click **Edit** to edit an existing performance monitor. For more information, refer to “[Configuring a user-defined product performance monitor](#)” on page 229 or “[Editing a preconfigured performance monitor](#)” on page 221.
 - Select one or more user-defined monitors and click **Delete** to delete the user-defined performance monitors.
 - 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 a .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.

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).

NOTE

Network scope does not affect the Events widget. The Events widget always includes all objects in your AOR.

From the dashboard, select a network from the **Network Scope** list. Options include:

- All
- Any SAN fabric
- Any Ethernet fabric
- Any system-defined group
- Any user-defined group
- Any user-defined customized network

If you select a fabric scope, violation counts display for all products and ports in the fabric.

If you select a product scope, violation counts display for the selected products and the ports that belong to the selected products.

If you select a port scope, violation counts display for the specified ports and the products to which the ports belong. If any of the selected ports are initiator or target ports, violation counts display for the attached switch port.

Select **All** to include all managed and monitored fabrics or groups in your AOR. The default is **All**. If the fabric or group you select is deleted from discovery, the widget refreshes and returns to the default (**All**).

Creating a customized network scope

You can create a network scope from any objects in your AOR. You can create network scopes based on Fabrics, Products, Product Groups, or Ports.

1. Click the **Network Scope** ellipsis button.

The **Edit Scope** dialog box displays with a list of existing user-defined network scopes in the **Network Scopes** list.

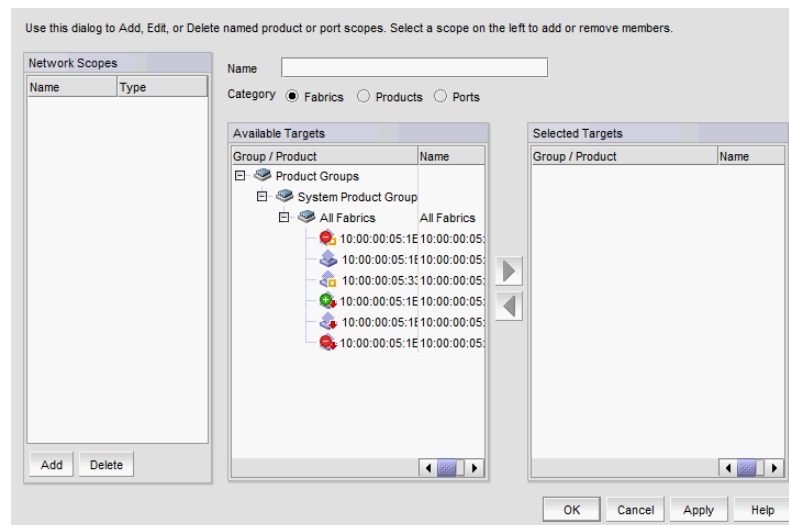


FIGURE 57 Edit Scopes dialog box

2. Click **Add**.

A new network scope displays in the **Network Scopes** list.

3. Enter a name for the scope in the **Name** field.

4. Select one of the following 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.

5. Select one or more the objects you want to include in the network from the **Available** list and click the right arrow button.

The objects display in the **Selected** list. To remove an object from the **Selected** list, select it and click the left arrow button.

6. Click **OK** to save your changes and close the **Edit Scope** dialog box.

Editing a user-defined network scope

You can edit any user-defined network scope.

1. Click the **Network Scope** ellipsis button.
The **Edit Scope** dialog box displays with a list of existing user-defined network scopes in the **Network Scopes** list.
2. Select the network scope you want to edit in the **Network Scopes** list.
The network scope details display in the right side fields.
3. Change the name for the scope in the **Name** field, if needed.
4. To add objects, select one or more 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.
5. To remove an object from the **Selected Targets** list, select it and click the left arrow button.
6. Click **OK** to save your changes and close the **Edit Scope** dialog box.

Deleting a user-defined network scope

You can edit any user-defined network scope.

1. Click the **Network Scope** ellipsis button.
The **Edit Scope** dialog box displays with a list of existing user-defined network scopes in the **Named Scopes** list.
2. Select the network you want to delete in the **Named Scopes** list.
3. Remove all objects from the **Selected Targets** list.
To remove an object from the **Selected Targets** list, select it and click the left arrow button.
4. Click **Delete**.
5. Click **OK** to save your changes and close the **Edit Scope** dialog box.

Setting the data display time frame

Setting the time scope in the dashboard toolbar configures the data display time range for the status widgets and performance monitors that include a time range.

NOTE

Time scope does not affect the Events widget. For the Events widget, you set the time scope within the widget.

NOTE

sFlow monitors only display data for up to 1 day.

From the dashboard, select one of the following duration options for which you want to display data from the **Time Scope** list.

- **30 Minutes** – Displays data for the previous half hour.
- **1 Hour** – Displays data for the previous hour.
- **6 Hours** – Displays data for the previous 6 hours.
- **12 Hours** – Displays data for the previous 12 hours.
- **1 Day** – Displays data for the previous day.
- **3 Days** – Displays data for the previous 3 days.
- **1 Week** – Displays data for the previous week.
- **1 Month** – Displays data for the previous month.

The displayed data changes to the new time frame for any status widget or performance monitor affected by time.

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.

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](#)
- [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 Ports Health dashboard

The SAN Ports 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](#)

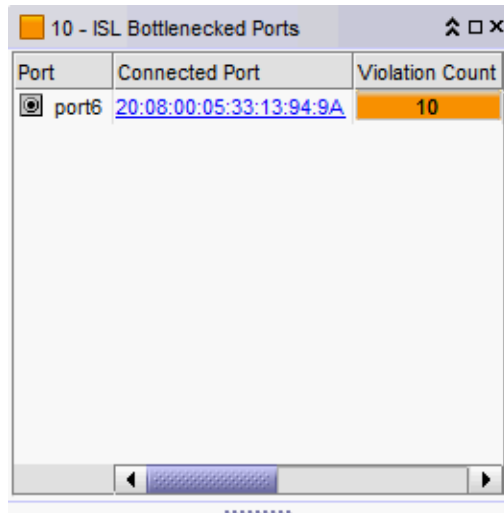
Status widgets

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
- [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 in your SAN
- [Port Health Violations widget](#) — Table view of out of range port health violations. 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 58) 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.



The screenshot shows a window titled "10 - ISL Bottlenecked Ports". It contains a table with the following data:

Port	Connected Port	Violation Count
port6	20:08:00:05:33:13:94:9A	10

FIGURE 58 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 in violation displays 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 967.
- **Product** – The product label, such as product name, IP address, node WWN, domain ID, or zone alias.
- **Type** – The port type.
- **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 [“Setting the network scope”](#) on page 176.
- To display data for a specific duration, refer to [“Setting the data display time frame”](#) on page 179.

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 1207.

Events widget

The **Events** widget (Figure 59) displays the number of events by severity level for a specified time range as a stacked bar graph.

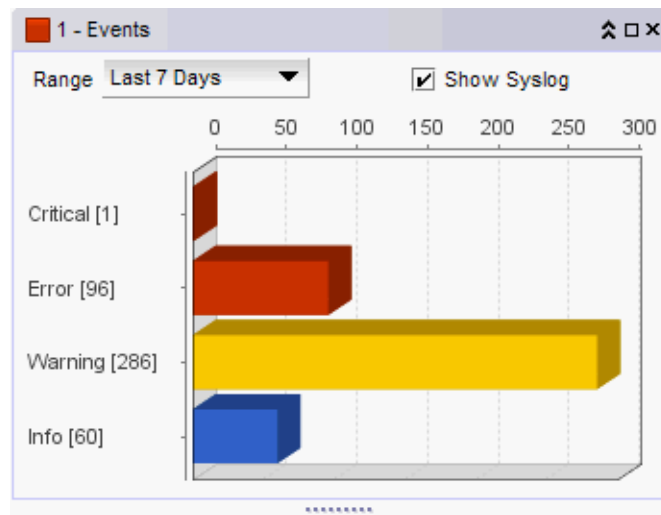









FIGURE 59 Events widget

The **Events** widget includes the following data:

- **Severity icon/widget title/event count** – The color of the worst severity followed by the event count with that severity displays before the widget title.
- **Range list** – Use to customize this widget to display a specific time range. Options include: This Hour, Last Hour, Today, Yesterday, Last 7 Days, and Last 30 Days.
- **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 18](#):

TABLE 18 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 — The network scope does not affect the Events widget. The Events widget always includes all objects in your AOR.
- Time Scope — The time 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 duration and to display Syslog details.

- Display event information for a specific duration by selecting one of the following from the **Range** list:
 - **This Hour** — Displays event information for the current hour beginning when you launch the dashboard.
 - **Last Hour** — Displays event information for the previous hour to when you launch the dashboard.
 - **Today** — Displays event information for the current day beginning at 12:00 AM.
 - **Yesterday** — Displays event information for the previous day beginning at 12:00 AM of the previous day.
 - **Last 7 Days** — Displays event information for the last 7 days, including the current day.
 - **Last 30 Days** — Displays event information for the last 30 days, including the current day.

- 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.

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 event type selected.

For information about report details, refer to “[Fault Management](#)” on page 1063.

Host Adapter Inventory widget

The **Host Adapter Inventory** widget (Figure 60) displays the host adapter products inventory as stacked bar graphs.

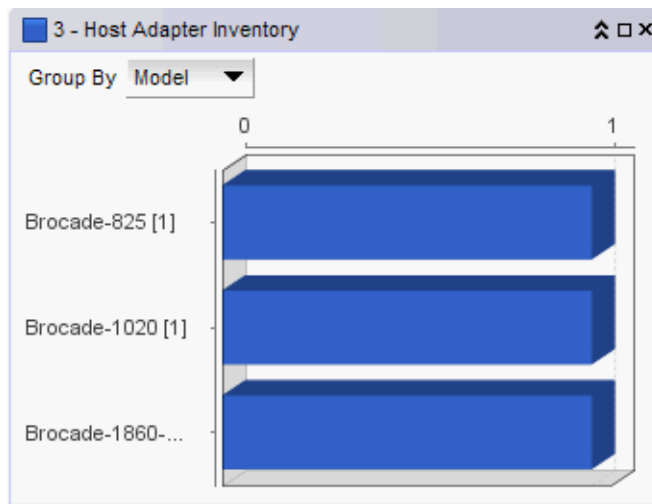


FIGURE 60 Host Adapter Inventory widget

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.
- 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 61) displays the SAN products inventory as stacked bar graphs.




FIGURE 61 SAN 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 to 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 188.

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 62](#)) displays the device status as a pie chart.




FIGURE 62 SAN 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 to 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 188.

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.

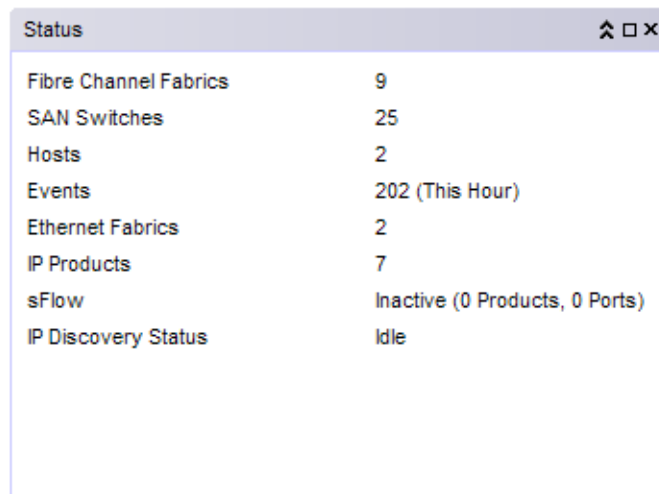
The **SAN Products - Status** dialog box (where *Status* is the section of the widget you selected) displays with the following fields and components:

- **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 1218.
3. Click **Close**.

Status widget

The Status widget (Figure 63) displays the number of products managed and the number of events within the selected event time range.



Status	
Fibre Channel Fabrics	9
SAN Switches	25
Hosts	2
Events	202 (This Hour)
Ethernet Fabrics	2
IP Products	7
sFlow	Inactive (0 Products, 0 Ports)
IP Discovery Status	Idle

FIGURE 63 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.
- Events – The number of events within the last hour.
- 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 [“Setting the network scope”](#) on page 176.
- To display data for a specific duration, refer to [“Setting the data display time frame”](#) on page 179.

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 1218.
- 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 439.

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 alerts you to problems long before they become costly failures.

The MAPS 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 widgets include 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.

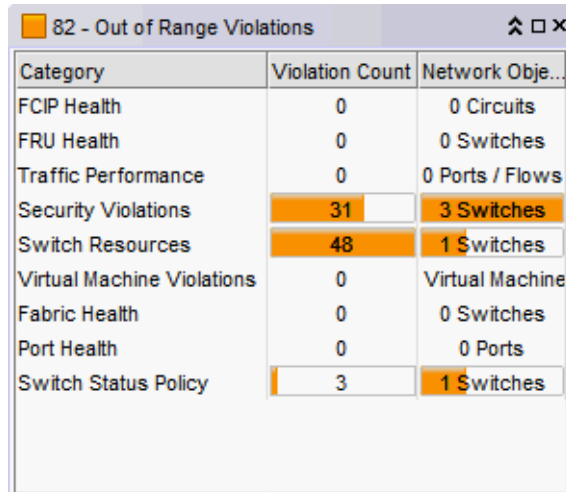
The MAPS widgets display on the main **Dashboard** tab. The Management application provides the following preconfigured MAPS widgets:

- [Out of Range Violations widget](#) — Table view of all out of range threshold violations reported in your SAN.
- [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 64) displays the number of violations for each MAPS category and the number of network objects (such as ports, trunks, switches, and circuits) with that MAPS 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.



Category	Violation Count	Network Objects
FCIP Health	0	0 Circuits
FRU Health	0	0 Switches
Traffic Performance	0	0 Ports / Flows
Security Violations	31	3 Switches
Switch Resources	48	1 Switches
Virtual Machine Violations	0	Virtual Machine
Fabric Health	0	0 Switches
Port Health	0	0 Ports
Switch Status Policy	3	1 Switches

FIGURE 64 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 dashboard categories. Always displays whether or not there is an associated violation. Categories include:
 - Fabric Health
 - FCIP Health
 - FRU Health
 - Port Health
 - Security Violations
 - Switch Resources
 - Switch Status Policy
 - Traffic Performance
 - Virtual Machine Violations
- **Violation Count** – The total number of MAPS 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 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 1 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.

- **Refreshed** — The time of the last update for the widget.

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 [“Setting the network scope”](#) on page 176.
- To display data for a specific duration, refer to [“Setting the data display time frame”](#) on page 179.
- 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 **MAPS > Violations** to navigate to the **Violations** dialog box. For more information, refer to [“Viewing MAPS violations”](#) on page 1179.
- 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 193.
- Double-click the **Virtual Machine Violations** category row to navigate to the **VM Alarms** widget. For more information, refer to the user manual or online help.
- 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 65) 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.

Port	Connected Port	Violation Count	CRC Errors	Invalid Tx Words	Loss of Sync	Link Failures	Loss of Signal
port58	2E:55:00:05:1E:47:16:00	2	0	0	0	2	0

FIGURE 65 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.
- **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 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 Failure** — 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 μ watts, to help determine if the SFP transceiver is in good working condition.
- **SFP Transmit Power** — The amount of outgoing laser, in μ 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.
- **Refreshed** — The time of the last update for the widget.

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 [“Setting the network scope”](#) on page 176.
- To display data for a specific duration, refer to [“Setting the data display time frame”](#) on page 179.
- 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 1218.
- Double-click a row to navigate to the **Violations** dialog box.

Performance monitors

The **Performance Dashboard** provides 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 Dashboard** also provides 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.

Reset the **Performance Dashboard** back to the default size by right-clicking in the white space and selected **Reset to Default**.

The Management application provides the following preconfigured performance monitors:

TABLE 19 Preconfigure performance monitors

Monitor title	Description	Data collectors
Top Port Alignment Errors	Table view of the alignment errors measure	All SAN TE port collector
Top Port C3 Discards	Table view of the C3 discards measure	All SAN FC port collector
Top Port C3 Discards RX TO	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	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	Table view of the discards measure	Port discard count collector
Top Port Encode Error Out	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
Top Port Link Failures	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	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 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

TABLE 19 Preconfigure performance monitors

Monitor title	Description	Data collectors
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

These 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 222.

Displaying monitors on the Performance Dashboard

1. From the **Dashboards** expand navigation bar, double-click **Performance Dashboard**.
The **Performance 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 **Performance Dashboard**.
4. Click **OK**.

Top Port Alignment Errors monitor

The **Top Port Alignment Errors** performance monitor displays the top ports with alignment errors in a table.

The Top 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 221.

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 935.

Top Port C3 Discards monitor

The **Top Port C3 Discards** monitor (Figure 66) displays the top ports with Class 3 frames discarded in a table. There are four port widgets: All, ISL, Initiator, and Target.

Port	Connected Port	C3 Discards	C3 Discards/sec
6e	20:01:00:05:1E:38:A0:1B	8590000000	3372.562
20:02:...	12:82:00:11:0D:00:00:0...	4295000000	1686.281
20:86:...		4295000000	1686.281
20:C3:...	20:02:00:05:1E:53:8A:1A	4295000000	1686.28
20:02:...	20:00:00:05:1E:90:53:7E	27260	0.011
20:00:...	20:00:00:05:1E:90:1B:27	26429	0.01
20:02:...	20:13:00:05:1E:90:48:AD	12209	0.005
20:06:...	10:00:00:05:1E:59:F5:D0	9312	0.004
20:03:...	20:02:00:05:1E:35:9C:86	5001	0.002
20:00:...	20:0A:00:05:1E:90:45:6D	2655	0.001

Refreshed- 12:30 PM

FIGURE 66 Top Port C3 Discards monitor

The **Top Port C3 Discards** monitor includes the following data:

- Severity icon/monitor title – The worst severity of 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.
- **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 221.

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 [“SAN shortcut menus”](#) on page 1218.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Port C3 Discards RX TO monitor

The **Top Port C3 Discards RX TO** monitor (Figure 67) 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.

Port	Connected Port	C3 Discards RX TO	C3 Discards RX TO/sec
TO 150/0/1 NOS ...		35	0
port9	20:1A:00:05:1E:9B:8D:5C	16	0

Refreshed- 2:42 PM

FIGURE 67 Top Port C3 Discards RX TO monitor

The **Top Port C3 Discards RX TO** monitor includes the following data:

- Severity icon/monitor title – The worst severity of 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.
- **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 221.

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 1218.
- 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 388.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Port CRC Errors monitor

The **Top Port CRC Errors** monitor (Figure 68) displays the top ports with frames that contain cyclic redundancy check (CRC) errors in a table.

Port	Connected Port	CRC Errors	CRC Errors/sec
port256789	20:C3:00:05:1E:4B:AA:00	247	0
20:C3:00:...	20:02:00:05:1E:53:8A:1A	1	0

Refreshed- 12:30 PM

FIGURE 68 Top Port CRC Errors monitor

The **Top Port CRC Errors** monitor includes the following data:

- **Severity icon/monitor title** — The worst severity of 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.
- **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 221.

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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Port Encode Error Out monitor

The **Top Port Encode Error Out** monitor (Figure 69) displays the top ports with encoding errors outside of frames in a table.

Port	Target	Encode Error Out	Encode Error Out/sec
test	20:00:00:11:0D:A8:00:00	76.943	0.001

FIGURE 69 Top Port Encode Error Out monitor

The **Top Port Encode Error Out** monitor includes the following data:

- Severity icon/monitor title – The worst severity of 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.
- **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 221.

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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Port Link Failures monitor

The **Top Port Link Failures** monitor ([Figure 70](#)) displays the top ports with link failures in a table.

Port	Connected Port	Link Failures	Link Failures/sec
tt		1	0
20:0...	20:00:00:05:1E:90:53:43	1	0

Refreshed- 12:50 PM

FIGURE 70 Top Port Link Failures monitor

The **Top Port Link Failures** monitor includes the following data:

- Severity icon/monitor title — The worst severity of 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 Failures/sec** — The number (error rate) receive link failure errors per second for the duration specified in the monitor.
- **RX Link Failures** — The number (error count) of receive link failure errors.
- **TX Link Failures/sec** — The number (error rate) of transmit link failure errors for the duration specified in the monitor.
- **TX Link Failures** — The number (error count) of transmit 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 221.

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 1207.
- Double-click a row to navigate to the SAN **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Port Link Resets monitor

The **Top Port Link Resets** monitor ([Figure 71](#)) displays the top ports with link resets in a table.

Port	Connected Port	RX Link Resets	TX Link Resets
port032	20:01:00:05:1E:C1:76:08	8	0
20:02:00...	12:82:00:11:0D:00:00:0...	7	0
tt		6	2
20:01:00...	20:01:00:05:1E:53:8A:1A	5	4
20:01:00...	20:00:00:05:1E:53:6B:69	4	4
6e	20:01:00:05:1E:38:A0:1B	4	3
port256...	20:C3:00:05:1E:4B:AA:00	3	3
port062	20:06:00:05:1E:C5:9E:06	2	0
20:C3:0...	20:02:00:05:1E:53:8A:1A	2	0
slot11 p...	20:83:00:05:1E:4B:AA:00	2	0

Refreshed- 12:50 PM

FIGURE 71 Top Port Link Resets monitor

The **Top Port Link Resets** monitor includes the following data:

- Severity icon/monitor title – The worst severity of 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 221.

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 1207.
- Double-click a row to navigate to the SAN **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Port Overflow Errors monitor

The **Top Port Overflow Errors** performance monitor (Figure 72) displays the top ports with overflow errors in a table.

Port	Connected Port	Overflow Errors	Overflow Errors/sec
Te 0/16		818461369	318.239

Refreshed- 7:24 PM

FIGURE 72 Top Port Overflow Errors performance monitor

The Top 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 221.

Top Port Receive EOF monitor

The **Top Port Receive EOF** performance monitor displays the top ports with received end-of-frames in a table.

The Top 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 221.

Top Port Runtime Errors monitor

The **Top Port Runtime Errors** performance monitor displays the top ports with runtime errors in a table.

The Top 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 221.

Top Port Sync Losses monitor

The **Top Port Sync Losses** monitor (Figure 73) displays the top ports with synchronization failures in a table.

Port	Connected Port	Sync Losses	Sync Losses/sec
20:0...		26171	0.01
20:0...		1383	0.001
20:0...		1383	0.001
20:1...		1383	0.001
20:0...	12:82:00:11:0D:00:00:00...	7	0
20:0...	20:00:00:05:1E:53:6B:69	4	0
tt		3	0
20:0...	20:01:00:05:1E:53:8A:1A	3	0
20:0...	22:00:00:04:CF:BD:70:34...	1	0
20:0...	20:00:00:05:1E:90:53:43	1	0

Refreshed- 12:45 PM

FIGURE 73 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.
- **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, 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 221.

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 1207.
- Double-click a row to navigate to the **Custom: Historical Performance Graphs** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Port Too Long Errors monitor

The **Top 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 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 221.

Top Port Traffic monitor

The **Top Port Traffic** monitor (Figure 74) displays the top ports with receive and transmit traffic in a table.

The screenshot shows a window titled "16.241 - Top Port Traffic" with a table of port traffic data. The table has four columns: Port, Connected Port, RX Traffic (MB/s), and TX Traffic (MB/s). The data is as follows:

Port	Connected Port	RX Traffic (MB/s)	TX Traffic (MB/s)
20:05:00:...	22:00:00:04:CF:BD:70:3...	0.013	16.241
20:03:00:...	20:02:00:05:1E:35:9C:86	8.075	0.007
20:06:00:...	10:00:00:05:1E:59:F5:D0	7.116	0.006
20:00:00:...	20:01:00:05:1E:35:9C:86	0.006	7.113
20:02:00:...	20:14:00:05:1E:90:48:AD	0.004	5.085
20:00:00:...	20:0A:00:05:1E:90:45:6D	5.082	0.004
20:00:00:...	20:00:00:05:1E:90:52:FA	0.004	5.081
20:02:00:...	20:00:00:05:1E:90:53:7E	5.077	1.886
20:02:00:...	20:13:00:05:1E:90:48:AD	5.077	0.004
first port	20:02:00:05:1E:90:1B:27	1.887	5.077

Refreshed- 12:55 PM

FIGURE 74 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.
- **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 Traffic (MB/s)** – The top receive traffic in megabits per second.
- **TX Traffic (MB/s)** – The top 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 221.

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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Port Underflow Errors monitor

The **Top Port Underflow Errors** performance monitor displays the top ports with underflow errors in a table.

The Top 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 221.

Top Port Utilization Percentage monitor

The **Top Port Utilization** monitor (Figure 75) displays the top port utilization percentages in a table.

The screenshot shows a window titled "7.734 - Top Port Utilization Percentage". It contains a table with the following columns: Port, Connected Port, RX Port Utilization Percentage, and TX Port Utilization Percentage. The table lists several ports with their respective utilization percentages. The RX column uses blue bars to represent utilization, and the TX column uses white bars. The status bar at the bottom indicates "Refreshed- 1:04 PM".

Port	Connected Port	RX Port Utilization Percentage	TX Port Utilization Percentage
20:05:00...	22:00:00:04:CF:BD:7...	0.006	7.734
20:03:00...	20:02:00:05:1E:35:9...	3.845	0.003
20:02:00...	20:14:00:05:1E:90:4...	0.002	2.421
20:02:00...	20:13:00:05:1E:90:4...	2.418	0.002
20:02:00...	20:00:00:05:1E:90:5...	2.418	0.898
first port	20:02:00:05:1E:90:1...	0.899	2.417
20:06:00...	10:00:00:05:1E:59:F...	1.694	0.001
20:00:00...	20:01:00:05:1E:35:9...	0.001	1.694
20:00:00...	20:0A:00:05:1E:90:4...	1.21	0.001
20:00:00...	20:00:00:05:1E:90:5...	0.001	1.21

FIGURE 75 Top Port Utilization monitor

The **Top Port Utilization** monitor includes the following data:

- Severity icon/monitor title – 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.
- **RX Port Utilization Percentage** – The top receive port utilization percentages.
- **TX Port Utilization Percentage** – The top 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 221.

Accessing additional data from the Top 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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to “Performance Data” on page 935.

Bottom Port Utilization Percentage monitor

The **Bottom Port Utilization Percentage** monitor (Figure 76) displays the bottom port utilization percentages in a table.

The screenshot shows a window titled "0 - Bottom Port Utilization Percentage" with a table of data. The table has four columns: Port, Connected Port, RX Port Utilization Percentage, and TX Port Utilization Percentage. The data is as follows:

Port	Connected Port	RX Port Utilization Percentage	TX Port Utilization Percentage
port9	20:1A:00:05:1...	0	0.001
port20	20:02:00:05:1E...	0.001	0
saa		0.001	0.001
1/1		0.001	0.001
portest		0	0.001
20:04:00...	20:09:00:05:1E...	0.001	0
1/1		0.001	0
port13	10:00:00:06:2B...	0	0.001
wer		0.001	0.001
1/1		0.001	0.001

Refreshed- 2:32 PM

FIGURE 76 Bottom Port Utilization Percentage monitor

The **Top Port Utilization Percentage** monitor includes the following data:

- Severity icon/monitor title – 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.
- **RX Port Utilization Percentage** – The bottom receive port utilization percentages.
- **TX Port Utilization Percentage** – The bottom 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.

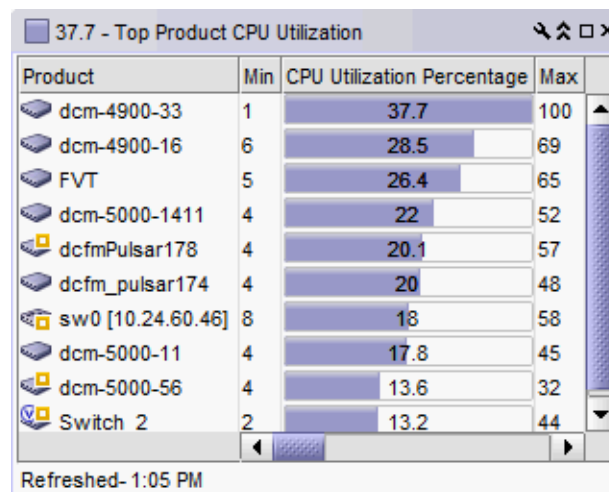
To customize the monitor to display data by a selected time frame as well as customize the display options, refer to [“Editing a preconfigured performance monitor”](#) on page 221.

Accessing additional data from the Top 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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Product CPU Utilization monitor

The **Top Product CPU Utilization** monitor (Figure 77) displays the top product CPU utilization percentages in a table.



The screenshot shows a window titled "37.7 - Top Product CPU Utilization". It contains a table with the following data:

Product	Min	CPU Utilization Percentage	Max
dcm-4900-33	1	37.7	100
dcm-4900-16	6	28.5	69
FVT	5	26.4	65
dcm-5000-1411	4	22	52
dcfmPulsar178	4	20.1	57
dcfm_pulsar174	4	20	48
sw0 [10.24.60.46]	8	18	58
dcm-5000-11	4	17.8	45
dcm-5000-56	4	13.6	32
Switch 2	2	13.2	44

Below the table, it says "Refreshed- 1:05 PM".

FIGURE 77 Top Product CPU Utilization monitor

The **Top Product CPU Utilization** monitor includes the following data:

- Severity icon/monitor title — The worst severity of 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 221.

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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Product Memory Utilization monitor

The **Top Product Memory Utilization** monitor ([Figure 78](#)) displays the top product memory utilization percentages in a table.

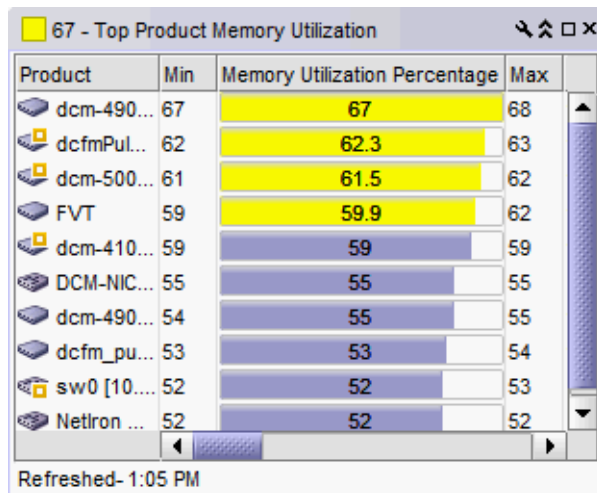


FIGURE 78 Top Product Memory Utilization monitor

The **Top Product Memory Utilization** monitor includes the following data:

- Severity icon/monitor title – The worst severity of 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 221.

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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Product Response Time monitor

The **Top Product Response Time** monitor ([Figure 79](#)) displays the top product response time in a table.

Product	Min	Response Time (ms)	Max
FWS648 Switch...	0	44.4	398
sw0 [10.24.60...	0	9.1	14
TestElkhound [1...	1	3.7	12
sw0 [10.24.60...	0	3.5	11
FWS648 Switch...	0	3.3	24
DCM-NICES202...	0	2.8	21
DCM-CES-76 [1...	0	2.6	20
FGS648P Switc...	0	1	6
Elkhound [10.24...	1	1	1
FCX624 Switch ...	0	0.9	2

Refreshed- 7:54 PM

FIGURE 79 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.

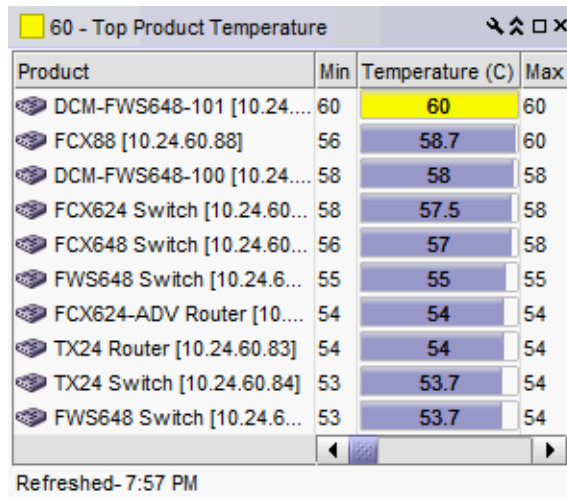
To customize the monitor to display data by a selected time frame as well as customize the display options, refer to [“Editing a preconfigured performance monitor”](#) on page 221.

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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

Top Product Temperature monitor

The **Top Product Temperature** monitor (Figure 80) displays the top product temperature in a table.



The screenshot shows a window titled "60 - Top Product Temperature" with a table of product temperatures. The table has four columns: Product, Min, Temperature (C), and Max. The data is as follows:

Product	Min	Temperature (C)	Max
DCM-FWS648-101 [10.24....	60	60	60
FCX88 [10.24.60.88]	56	58.7	60
DCM-FWS648-100 [10.24....	58	58	58
FCX624 Switch [10.24.60...	58	57.5	58
FCX648 Switch [10.24.60...	56	57	58
FWS648 Switch [10.24.6...	55	55	55
FCX624-ADV Router [10....	54	54	54
TX24 Router [10.24.60.83]	54	54	54
TX24 Switch [10.24.60.84]	53	53.7	54
FWS648 Switch [10.24.6...	53	53.7	54

Refreshed- 7:57 PM

FIGURE 80 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 221.

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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to “Performance Data” on page 935.

Top Products with Unused Ports monitor

The **Top Products with Unused Ports** monitor (Figure 77) displays the top products with ports not in use in a table.

Product	Min	Ports Not In Use	Max
test92	13	331	391
sw01	124	35	144
Reaper1 [10.24.60.36]	68	68	68
dcm-4900-33	63	63	64
dcm-4900-16	43	62	64
sw0 [10.24.60.49]	38	58	60
sw0 [10.24.60.46]	0	53	60
FCX648 Switch [10.24.60....]	48	48	48
FWS648 Switch [10.24.60...]	47	47	48
FWS648 Switch [10.24.60...]	47	47	48

Refreshed- 1:05 PM

FIGURE 81 Top Product CPU Utilization monitor

The **Top Products with Unused Ports** monitor includes the following data:

- Severity icon/monitor title – The worst severity of 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 221.

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 1207.
- Double-click a row to navigate to the **Historical Graphs/Tables** dialog box. For more information, refer to [“Performance Data”](#) on page 935.

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

The **Performance Dashboard** makes it easy for you to 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.
 - FCIP
 - Compression Ratio – The 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 acknowledgements 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.
 - IP (SAN TE ports only)
 - 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.
- 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 82) display the top or bottom number of products (for example, top 10 products) for the selected measure in a table.

Product	Min	Ports Not In Use	Max	Fabric
Reaper1 [10.24.60...]	46	68	68	
dcm-4900-33	42	63	64	10:00:0
dcm-4900-16	43	62	64	10:00:0
sw0 [10.24.60.49]	38	58	60	
DCM-FWS648-101...	47	47	48	
FWS648 Switch [...]	47	47	48	
FWS648 Switch [...]	47	47	48	
DCM-FGS648P-20...	47	47	48	
FWS648 Switch [...]	26	47	48	
FWS648 Switch [...]	26	46	48	

Refreshed- 12:07 PM

FIGURE 82 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 229.

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 935.

Top or bottom port performance monitors

The top or bottom port performance monitors ([Figure 83](#)) display the top or bottom number of ports (for example, bottom 10 ports) for the selected measure in a table.

Port	Connected Port	Sync Losses	Sync Losses/sec
<input checked="" type="checkbox"/> 20:0A:00:05:1E:90:45:72		12460	0.005
<input checked="" type="checkbox"/> 20:04:00:05:1E:07:6A:F8		658	0
<input checked="" type="checkbox"/> 20:05:00:05:1E:07:6A:F8		658	0
<input checked="" type="checkbox"/> 20:11:00:05:1E:90:45:72		658	0

FIGURE 83 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.
- **Severity icon/monitor title** – The worst severity of 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 233.

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 948.

Distribution performance monitors

The distribution performance monitor ([Figure 84](#)) 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 84 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 20 Product measures types

- | | |
|--|---|
| <ul style="list-style-type: none"> • Memory Utilization Percentage • CPU Utilization Percentage • Temperature (C) • Fan Speed (rpm) • Response Time (s) | <ul style="list-style-type: none"> • System Up Time (days) • Ports Not In Use • Ping Packet Loss Percentage • AP Client Count |
|--|---|

TABLE 21 Port measures types

- | | |
|--|--|
| <p>Common</p> <ul style="list-style-type: none"> • Port Utilization Percentage • Traffic • CRC Errors <p>FC</p> <ul style="list-style-type: none"> • 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 | <p>FCIP</p> <ul style="list-style-type: none"> • 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 <p>IP (SAN TE ports only)</p> <ul style="list-style-type: none"> • Receive EOF • Underflow Errors • Overflow Errors • Alignment Errors • Runtime Errors • Too Long Errors • |
|--|--|

- **Refreshed** — The time of the last update for the monitor.

To configure a distribution performance monitor, refer to [“Configuring a user-defined product performance monitor”](#) on page 229 or [“Configuring a user-defined port performance monitor”](#) on page 233.

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 235 or [“Viewing port distribution data details”](#) on page 236.

Time series performance monitors

The time series performance monitors (Figure 85) display the selected measures in a chart.

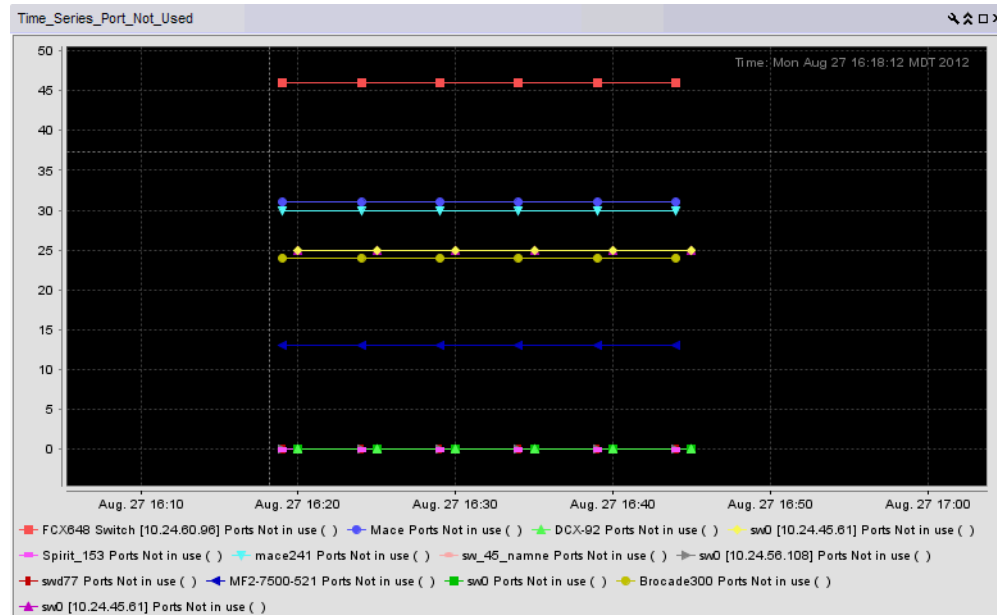


FIGURE 85 Time series performance monitor example

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 display” on page 982).

To configure a time series performance monitor, refer to “Configuring a user-defined product performance monitor” on page 229 or “Configuring a user-defined port performance monitor” on page 233.

Configuring a user-defined product performance monitor

1. From the **Dashboards** expand navigation bar, double-click **Performance Dashboard**.
The **Performance Dashboard** displays.
2. Click the **Customize Dashboard** icon.
The **Customize Dashboard** dialog box displays.

3. Click the **Performance** tab.
4. Click **Add**.

The **Add Performance Dashboard Monitor** dialog box displays.

5. Enter a unique title for the monitor.
The title can be up to 256 characters in length.
6. 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.
7. 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)
8. (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.
9. (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.
10. (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**.

11. (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 232.

Remove targets from the monitor by selecting one or more targets in the **Targets** list and clicking **Remove**.

12. 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.

13. Click **OK** on the **Customize Dashboard** dialog box.

The **Performance Dashboard** dialog box displays with the new monitors 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 235 or [“Viewing port distribution data details”](#) on page 236.
- 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 935.

Adding targets to a user-defined performance monitor

You can only add targets for Time Series monitors.

1. From the **Dashboards** expand navigation bar, double-click **Performance Dashboard**.

The **Performance Dashboard** displays.

2. Click the **Customize Dashboard** icon.

The **Customize Dashboard** dialog box displays.

3. Click the **Performance** tab.

4. Click **Add**.

The **Add Performance Dashboard Monitor** dialog box displays.

5. Select **Time Series** from the **Monitor Type - Product** or **Port** area.

6. Select the port measure for the monitor in the **Measure** area

7. Display data for a specific duration from the **Duration** options.

8. 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/ports and FCIP tunnels to the list of targets.

If you selected a product measure, continue with [step 9](#).

If you selected SAN port measure, continue with [step 9](#).

If you selected a FC IP port measure, go to [step 12](#).

9. Click the **SAN** tab.

10. Select SAN targets from the **Available SAN Sources** list.

11. Click the right arrow button to move the targets to the **Selected Sources** list.

12. Select FCIP targets from the **Available** list.

13. Click the right arrow button to move the targets to the **Selected Sources** list.

14. 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.

15. 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.

16. Click **OK** on the **Customize Dashboard** dialog box.

The **Performance Dashboard** dialog box displays with the new monitors at the bottom of the dashboard.

Configuring a user-defined port performance monitor

1. From the **Dashboards** expand navigation bar, double-click **Performance Dashboard**.
The **Performance Dashboard** displays.
2. Click the **Customize Dashboard** icon.
The **Customize Dashboard** dialog box displays.
3. Click the **Performance** tab.
4. Click **Add**.
The **Add Performance Dashboard Monitor** dialog box displays.
5. 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.
6. 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

FCIP

- 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

IP (SAN TE ports only)

- Receive EOF
- Underflow Errors
- Overflow Errors
- Alignment Errors
- Runtime Errors
- Too Long Errors
-

7. (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.

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 the steps in ["Adding targets to a user-defined performance monitor"](#) on page 232.

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.

- Click **OK** on the **Customize Dashboard** dialog box.

The **Performance Dashboard** dialog box displays with the new monitors 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 235 or [“Viewing port distribution data details”](#) on page 236.
- 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 948.
- 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 388.

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 227).

- Double-click a bar in the graph.

The *Monitor_Title* **Data Details** dialog box displays.

- 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 227).

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.
 - FCIP
 - Compression Ratio – The 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 acknowledgements 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.
 - IP (SAN TE ports only)
 - 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.
 - **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 [//link 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 or 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 86) the selected measures in a chart.




FIGURE 86 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 display” on page 982).

Top or bottom traffic flow performance monitor

The top or bottom traffic flow performance monitors display (Figure 87) the top or bottom number of flows for the selected measure in a table.

Flow Name	Sub Flow Id	Read Frame Count(frames)	Product
fghfgh	20	0	mace_25_t
vbgfcg	22	0	test92

FIGURE 87 Top traffic flow monitor example

The top or bottom flows 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.
- **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 238

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 264).
- Right-click column head to configure the table (refer to “[Customizing application tables](#)” on page 264).

Time series traffic flow performance monitor

The time series traffic flow performance monitors display ([Figure 88](#)) the selected measure in a chart.




FIGURE 88 Traffic flow performance monitor example

The time series traffic flow 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 display](#)” on page 982).

To configure a time series performance monitor, refer to “[Configuring a user-defined traffic flow performance monitor](#)” on page 242.

Configuring a traffic flows monitor from a performance graph

1. Configure the performance graph.
To configure traffic flows performance graph, refer to [//link to flow vision//](#).
2. Click **Publish** 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

1. From the **Dashboards** expand navigation bar, double-click **Performance Dashboard**.
The **Performance Dashboard** displays.
2. Click the **Customize Dashboard** icon.
The **Customize Dashboard** dialog box displays.
3. Click the **Performance** tab.
4. Click **Add**.
The **Add Performance Dashboard Monitor** dialog box displays.
5. 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.

6. 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)

7. (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.

8. (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.

- Select the **Show values** check box.
- Select **greater than** or **less than** from the list.
- Enter a value in the field.

9. (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.

- (Top N and Bottom N monitors only) Select the check box.
- 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 the steps in [“Adding targets to a traffic flow performance monitor”](#) on page 244.

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 **Performance Dashboard** dialog box displays with the new monitors at the bottom of the dashboard.

Adding targets to a traffic flow performance monitor

You can only add targets for Time Series monitors.

1. From the **Dashboards** expand navigation bar, double-click **Performance Dashboard**.

The **Performance Dashboard** displays.
2. Click the **Customize Dashboard** icon.

The **Customize Dashboard** dialog box displays.
3. Click the **Performance** tab.
4. Click **Add**.

The **Add Performance Dashboard Monitor** dialog box displays.
5. Select **Time Series** from the **Traffic Flows** area.
6. Select the one or more measures for the monitor in the **Measure** area
7. Click **Add** beneath the **Targets** table.

The **Performance Dashboard Monitor Targets** dialog box displays.

8. 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.

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 **Performance Dashboard** dialog box displays with the new monitors at the bottom of the dashboard.

7 Traffic flow dashboard monitors

View Management

In this chapter

- SAN tab overview 247
- Icon legend 258
- Customizing the main window 262
- Product List customization 267
- Search 268
- SAN view management overview 271
- SAN topology layout 276
- Grouping on the topology 281

SAN tab overview

The **SAN** tab (Figure 89) 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**.

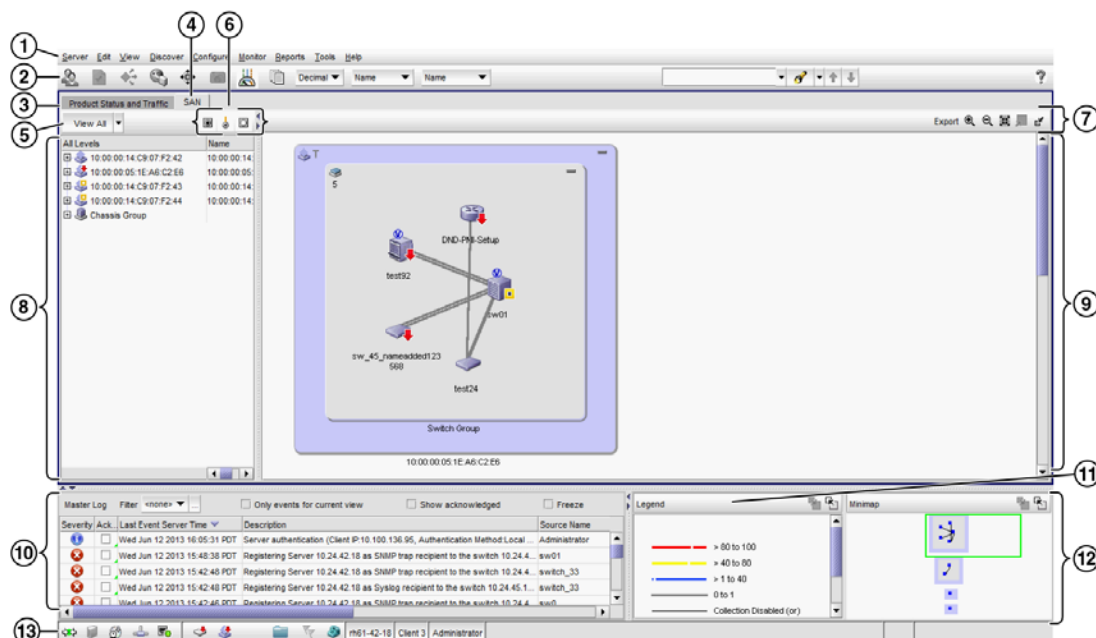


FIGURE 89 Main window - SAN tab

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 1208.
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 249.
3. **Dashboard tab** — Provides a high-level overview of the network managed by Management application server. For more information, refer to “[Dashboard Management](#)” on page 167.
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 to how to view the Product list (All Levels, Products and Ports, Products Only, or Ports Only) and to select which view you want to display in the main window. For more information, refer to “[View All list](#)” on page 250. For step-by-step instruction about creating a view, refer to “[Creating a customized view](#)” on page 271.
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 250.
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 251.
8. **Product List** — Lists the devices discovered in the Management application. For more detailed information, refer to “[Product List](#)” on page 251.
9. **Connectivity Map** — Displays the topology, including discovered and monitored devices and connections. For more information, refer to “[Connectivity Map](#)” on page 253.
10. **Master Log** — Displays all events that have occurred on the Management application. For more information, refer to “[Master Log](#)” on page 255.

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 254.
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 256.
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 257.

SAN main toolbar

The toolbar is located beneath the Menu bar and provides icons to perform various functions.




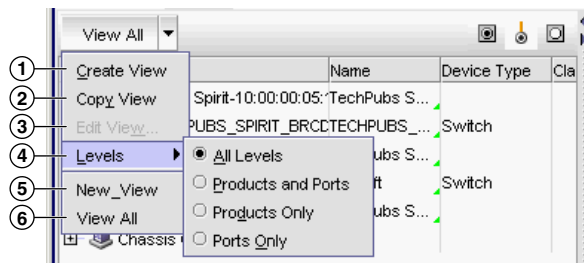
FIGURE 90 SAN main toolbar

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. **Domain ID/Port #** — Use to set the domain ID or port number to display as decimal or hex in the Product List.
10. **Product Label** — Use to set the product label for the devices in the Connectivity Map and Product List.
11. **Port Label** — Use to set the port label for the devices in the Connectivity Map and Product List.
12. **Product List Search** — Use to search for a device in the product list. For detailed instructions, refer to [“Search”](#) on page 268
13. **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 to how to view the Product list (All Levels, Products and Ports, Products Only, or Ports Only) and to select which view you want to display in the main window. Does not display until you discover a fabric. To discover a fabric, refer to “[Discovering fabrics](#)” on page 39.



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 91 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 92 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**.

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 the following tables:
 - “[SAN product icons](#)” on page 258
 - “[SAN group icons](#)” on page 259
 - “[SAN port icons](#)” on page 260
- **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 267.
- **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 264.

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 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 258
- “[SAN group icons](#)” on page 259
- “[Event icons](#)” on page 261

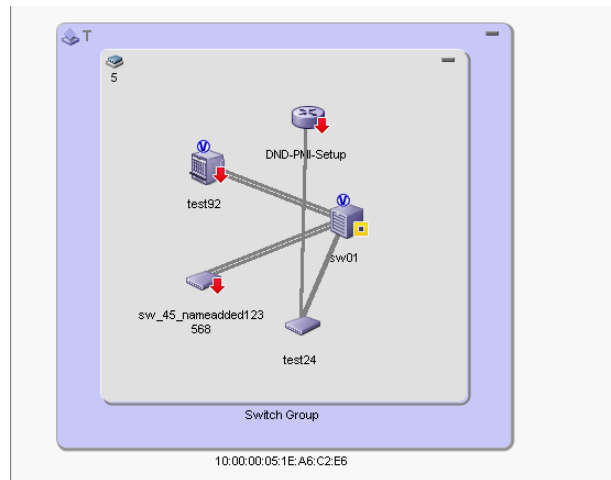


FIGURE 93 Connectivity Map

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 1218.

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.

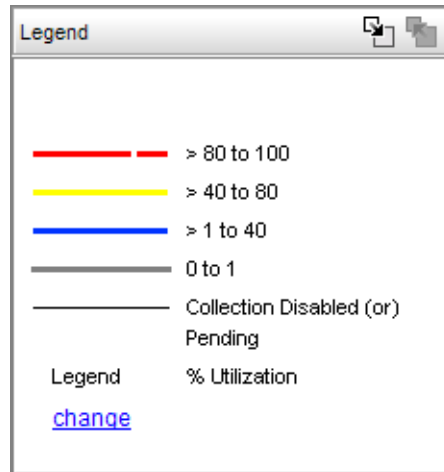


FIGURE 94 Utilization Legend

The colors and their meanings are outlined in the following table.

TABLE 22

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 979.

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 1131. To view event properties, refer to "[Displaying event properties from the Master Log](#)" on page 1129.

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 1063. For a list of the event icons, refer to "[Event icons](#)" on page 261.
- **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.

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 95 SAN 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 (X) 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 96 Status 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.
 - **Network 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 11.
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. **Policy Monitor Status** — Displays whether or not a policy monitor has failed or partially failed. Click to launch the **Policy Monitor** dialog box. For more information about policy monitors, refer to [“Viewing policy monitor status”](#) on page 1048.
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 1088.
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 298.
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 10.
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 9.
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 161.
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.

Icon	Description	Icon	Description
	Fabric		Fabric OS Switch and Blade Switch
	Fabric OS Director		Fabric OS DCB 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)

Icon	Description	Icon	Description
	VC module		Multi-fabric VC module
	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.

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.

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.

Icon	Description	Icon	Description
	Host Group		







SAN port icons




The following table lists the port icons that display in the Product List.

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

SAN product status icons









The following table lists the product status icons that display on the topology.

Icon	Status
No icon	Healthy/Operational
	Attention
	Degraded/Marginal
	Device Added
	Device Removed/Missing
	Down/Failed
	Routed In

Icon	Status
	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 1063.

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 97 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.

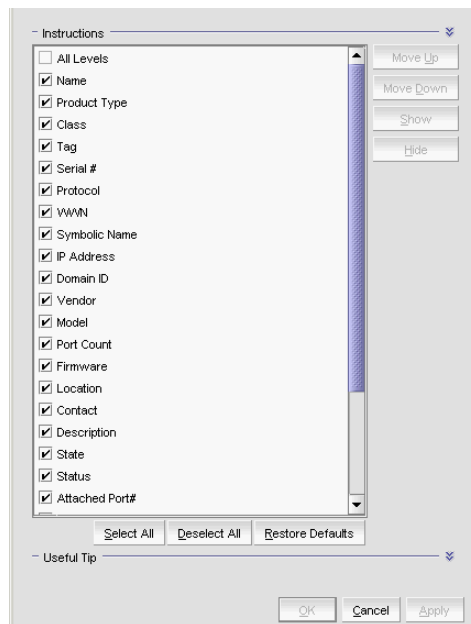


FIGURE 98 Customize Columns dialog box

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 99 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 1283.) 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.

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 270) or by looking for exact matches (refer to [“Searching for an exact match”](#) on page 270).

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

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 1283.) 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

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 1283.) 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 271 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.

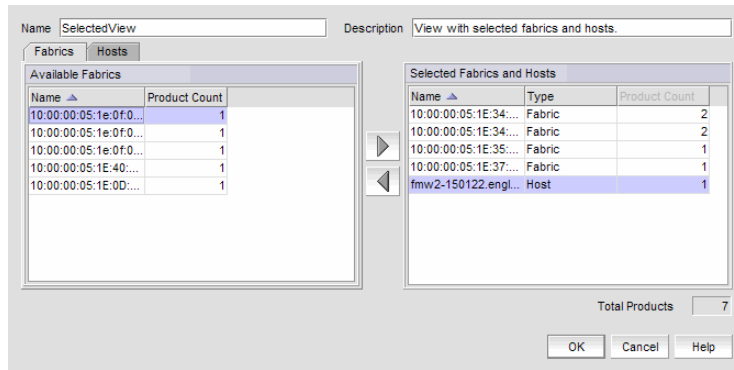


FIGURE 100 Create View dialog box - Fabrics tab

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 101 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.

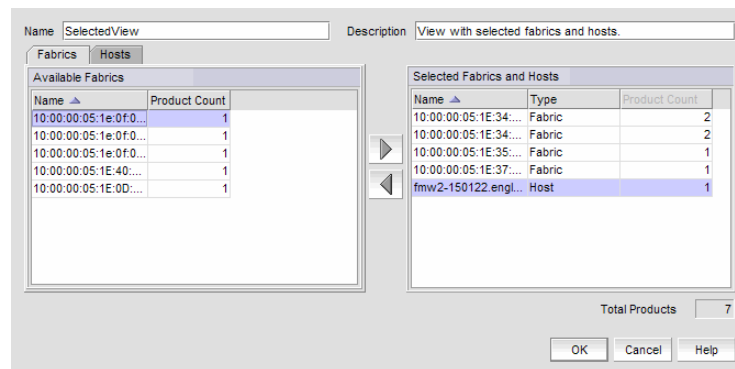


FIGURE 102 Edit View dialog box - Fabrics tab

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 103 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 104 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.

- 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.

- 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).

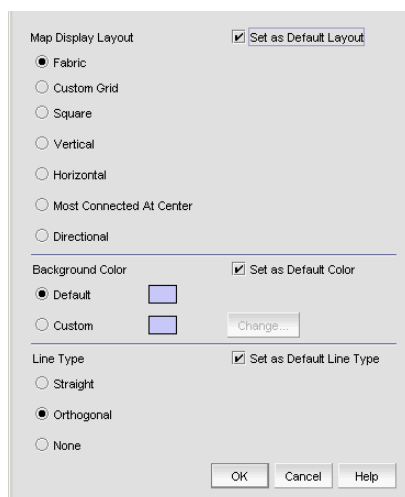


FIGURE 105 Map Display Properties dialog box

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 106](#)).

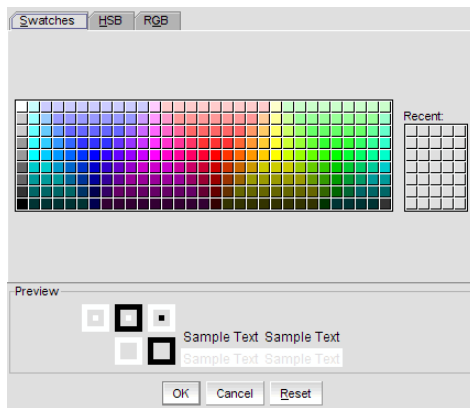


FIGURE 106 Choose a background color dialog box

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

In this chapter

- Call Home overview 284
- Viewing Call Home configurations 285
- Showing a Call Home center 288
- Hiding a Call Home center 288
- Editing a Call Home center 289
- Enabling a Call Home center 296
- Enabling supportSave 296
- Testing the Call Home center connection 297
- Disabling a Call Home center 297
- Viewing Call Home status 298
- Assigning a device to the Call Home center 299
- Removing a device from a Call Home center 299
- Removing all devices and filters from a Call Home center 299
- Defining an event filter 300
- Assigning an event filter to a Call Home center 301
- Assigning an event filter to a device 301
- Overwriting an assigned event filter 302
- Removing all event filter from a Call Home center 302
- Removing an event filter from a device 303
- Removing an event filter from the Call Home Event Filters list 303
- Searching for an assigned event filter 303

Call Home overview

NOTE

Call Home is supported on Windows systems for all modem and 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 only)
- HP LAN (Windows only)
- IBM (Windows only)
- IBM E-mail (Windows and UNIX)
- NetApp E-mail (Windows and UNIX)
- Oracle E-mail (Windows and UNIX)

When configuring modem and HP LAN Call Home centers, you must enter the customer contact information in the device's Element Manager. You may also need to configure the Management application server IP address manually as an SNMP trap recipient for Fabric OS devices.

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.

NOTE

The HP LAN Call Home alert displays the directory separation characters with a double backslash (\\) instead of a single backslash (\).

- 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.
- 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 1233. For more information about events, refer to [“Fault Management”](#) on page 1063.

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.

System requirements

Call Home (except for e-mail and HP LAN) requires the following hardware equipment:

- Any Windows server with an internal or external modem connection
- Analog phone line

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 107).

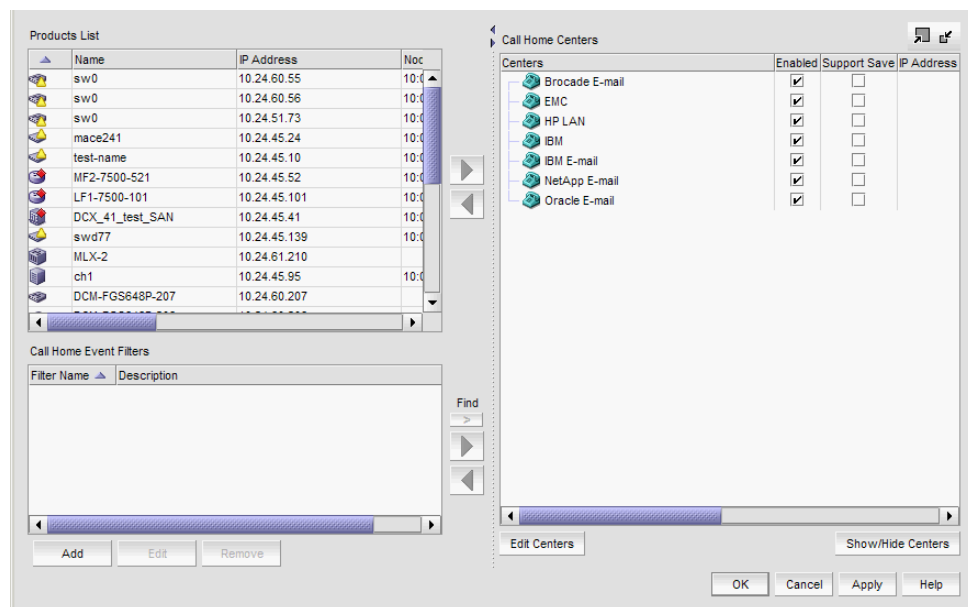


FIGURE 107 Call Home dialog box

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.
 - **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 299). 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 299). 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 300).
- **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 300).
- **Remove** button – Click to remove the event filter (refer to [“Removing an event filter from the Call Home Event Filters list”](#) on page 303) 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 301 or [“Assigning an event filter to a device”](#) on page 301) 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 filter from a Call Home center”](#) on page 302 or [“Removing an event filter from a device”](#) on page 303) 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.
 - **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 and 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 289).
 - **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 288 or [“Hiding a Call Home center”](#) on page 288).
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 (Figure 108).

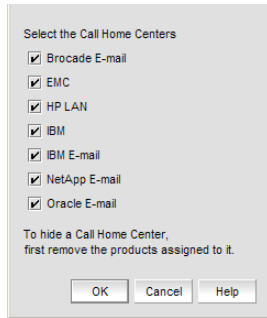


FIGURE 108 Centers dialog box

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 the IBM Call Home center](#) 289
- [Editing an e-mail Call Home center](#) 290
- [Editing the EMC Call Home center](#) 294
- [Editing the HP LAN Call Home center](#) 295

Editing the IBM Call Home center

To edit the IBM Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

The **Call Home** dialog box displays.

2. Select **IBM** 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 109](#)).

FIGURE 109 Configure Call Home Center dialog box (IBM 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. Set the time interval at which to check the Call Home center by selecting the **Set heartbeat interval at ___ days (1-28)** check box and entering the interval in the field.
7. Enter how long you want to wait before timing out the heartbeat interval in the **Time Out** field.
The default is 60 seconds.

8. Enter how often you want to retry the heartbeat interval in the **Retry Interval** field.
The default is 10 seconds.
9. Enter the maximum number of retries in the **Maximum Retries** field.
The default is 3.
10. Enter the primary phone number or extension of the Call Home center in the **Call Home Center - Primary Connection** field.
11. Enter the backup phone number or extension of the Call Home center in the **Call Home Center - Backup Connection** field.
12. Enter the phone number or extension of the local server in the **Local Server - Phone Number** field.
13. Enter the identification number of the local server in the **Local Server - Server ID** field.
14. Click **Send Test** to test the phone number.
The selected Call Home center must be enabled to test the phone number.
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.
15. Click **OK** to close the "Test Event Sent" message.
16. Click **OK**.
The **Call Home** dialog box displays with the Call Home center you edited highlighted in the **Call Home Centers** list.
17. Click **OK** to close the **Call Home** dialog box.

Editing an e-mail Call Home center

E-mail Call Home centers are available for Brocade, 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 110](#)).

FIGURE 110 Configure Call Home Center dialog box (Brocade, 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 292.

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

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 111](#)).

FIGURE 111 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 heartbeat interval at ___ days (1-28)** check box and entering the interval in the field.
7. Enter the path to the ConnectEMC application in the **ConnectEMC** field or browse to the ConnectEMC application location.
8. Enter the phone number or extension of the local server in the **Local Server - Modem #** field.
9. Enter the identification number of the local server in the **Local Server - Cabinet Serial #** field.
10. Enter the site name for the local server in the **Local Server - Site Name** field.
11. Click **Send Test** to test the Connect EMC application.

The selected Call Home center must be enabled to test the ConnectEMC application.

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.

12. Click **OK** to close the “Test Event Sent” message.

13. Click **OK**.

The **Call Home** dialog box displays with the Call Home center you edited highlighted in the **Call Home Centers** list.

14. Click **OK** to close the **Call Home** dialog box.

Editing the HP LAN Call Home center

To edit an HP LAN Call Home center, complete the following steps.

1. Select **Monitor > Event Notification > Call Home**.

The **Call Home** dialog box displays.

2. Select the **HP LAN** 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 112](#)).

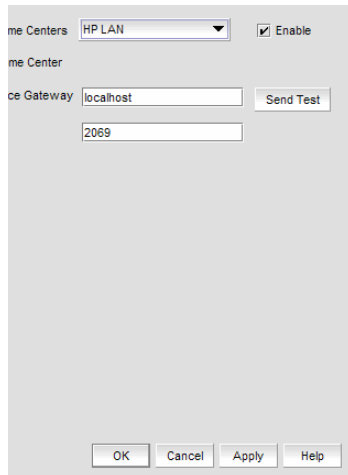


FIGURE 112 Configure Call Home Center dialog box (HP LAN option)

4. Make sure the **HP LAN** 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. Enter the IP address of the Call Home center in the **Service Gateway** field.
The default is 2069.
7. Enter the port number of the Call Home center in the **Port** field.

8. Click **Send Test** to test the address.

The selected Call Home center must be enabled to test the IP address.

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 HP LAN Call Home alert displays the directory separation characters with a double backslash (\\) instead of a single backslash (\).

9. Click **OK** to close the “Test Event Sent” message.
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.

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 5.2 or later or Network OS 2.1.X or later.

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 289.

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.
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 23](#) illustrates and describes the icons that indicate the current status of the Call Home function.

TABLE 23 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: <ul style="list-style-type: none"> • 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 [Appendix B, "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 301 or ["Assigning an event filter to a device"](#) on page 301.
 7. Click **OK** to close the **Call Home** dialog box.

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.

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 [Appendix B, “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 [Appendix B, “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 [Appendix B, “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 filter 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.

9 Searching for an assigned event filter

Third-party tools

In this chapter

- About third-party tools 305
- Starting third-party tools from the application 306
- Launching a Telnet session 306
- Launching an Element Manager 307
- Launching Web Tools 307
- Launching FCR configuration 308
- Launching HCM Agent 310
- Launching Fabric Watch 310
- Single sign on support for IBM 311
- Launch in context support for IBM 312
- Adding a tool 314
- Entering the server IP address of a tool 315
- Adding an option to the Tools menu 316
- Changing an option on the Tools menu 317
- Removing an option from the Tools menu 317
- Changing an option on a device's shortcut menu 319
- Removing an option from a device's shortcut menu 320
- Microsoft System Center Operations Manager (SCOM) plug-in 320

About third-party tools

NOTE

Installing tools is only available with the Trial and Licensed version 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 (such as, Firefox, Windows Explorer, Web Tools, Element Managers, FCR Configuration, HCM Agent and so on) 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 version 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 315.

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 an 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 an 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 a device's Element Manager, complete the following steps.

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 1243.

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 5.0 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 a FC 8 Gbps 16-port Blade (with Integrated Routing license)
- Fabric OS directors configured with a FC 8 Gbps 32-port Blade (with Integrated Routing license)
- Fabric OS directors configured with a 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 will need to enter those credentials to launch Web Tools. For more information about privileges, refer to ["User Privileges"](#) on page 1243.

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 1243.

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 1243.

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

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 Fabric 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 will need to enter those credentials to launch Fabric Watch. For more information about privileges, refer to ["User Privileges"](#) on page 1243.

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
```

10 Launch in context support for IBM

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 140.

4. Make sure any switches you need to manage are discovered by the Management application. Add switch/fabric into the Management application by selecting Discovery > Setup > Add Fabric.

To discover a switch or fabric, refer to [“Discovering fabrics”](#) on page 39.
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_Idf.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 `tpcApplication_Namelfdeployer.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.engliah  
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 `tpcApplication_Nameldfundeployer.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

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

10 Adding a tool

- **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 113](#)).

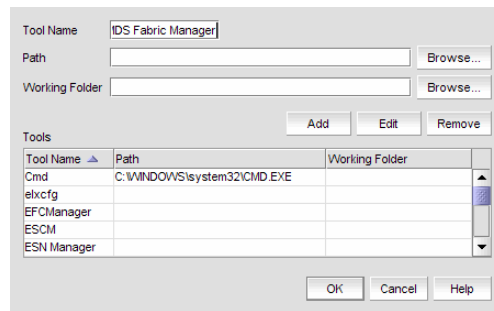


FIGURE 113 Define Tools dialog box

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 316.
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 a 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 114).

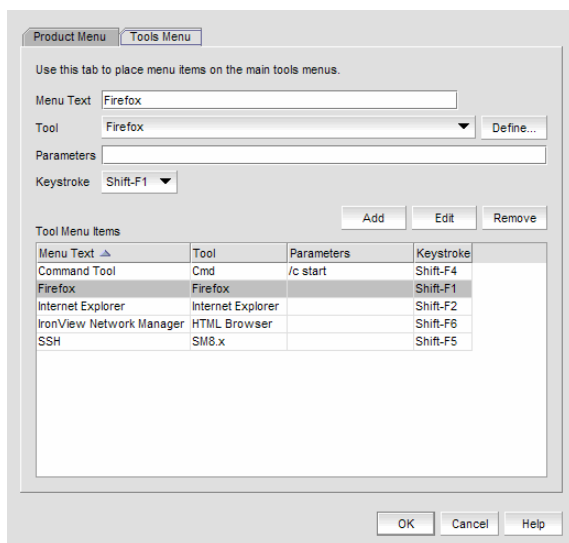


FIGURE 114 Setup Tools dialog box (Tools menu tab)

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 314.
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 a 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.

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 a 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.

3. Select the row of the tool you want to remove in the **Tools Menu Items** table.

10 Adding an option to a device's shortcut menu

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 314.
 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 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. 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 314.

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.

10 Removing an option from a device's shortcut menu

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 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 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 (SCOM) 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.

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 1243.

The 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 when Critical events for switches in the fabric trigger CallHome in the Management application

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 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 monitor statistics. Refer to [“SAN end-to-end monitoring”](#) on page 957.

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 configure the SCOM plug-in, 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**.

Server Management Console

In this chapter

- [Server Management Console overview](#) 323
- [Services tab](#) 324
- [Ports tab](#) 327
- [AAA Settings tab](#) 328
- [Restore tab](#) 342
- [Technical Support Information tab](#) 343
- [HCM Upgrade tab](#) 344
- [SMI Agent Configuration Tool](#) 345

Server Management Console overview

The Server Management Console (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](#) 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 115).

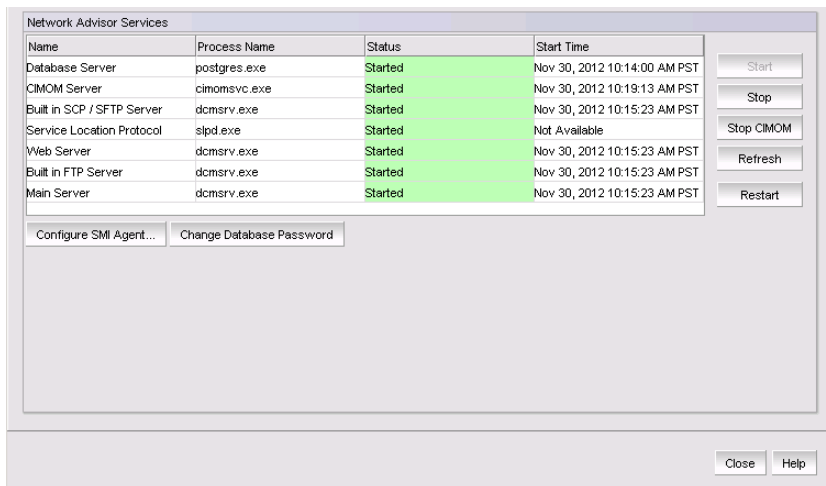


FIGURE 115 Services tab

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 128.

You can also configure the server port settings from the configuration wizard. For instructions, refer to [“Launching the Configuration Wizard”](#) on page 5.

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 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.
- 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 116).

The screenshot shows the AAA Settings configuration interface. At the top, there are four dropdown menus: 'Primary Authentication' (set to 'Radius Server'), 'Secondary Authentication' (set to 'None'), 'Fail Over Option' (set to 'Radius Servers Not Reachable'), and 'Authorization Preference' (set to 'Local Database'). Below these is a table titled 'Radius Servers and Sequence' with columns for 'Network Address', 'TCP Port', 'Timeout(Sec)', 'Attempts', and 'Authentication Type'. The table is currently empty. To the right of the table are buttons for 'Add', 'Edit', 'Delete', 'Up', and 'Down'. At the bottom of the window are buttons for 'Audit Trail', 'Display', 'Test', and 'Apply'.

FIGURE 116 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 330.
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 341.
11. Click **Close** to close the Server Management Console.

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 117).

FIGURE 117 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.

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.

The screenshot shows the configuration interface for LDAP server authentication. At the top, there are dropdown menus for Primary Authentication (set to LDAP Server), Secondary Authentication (set to None), Fail Over Option (set to LDAP Servers Not Reachable), and Authorization Preference (set to Local Database). Below these is a section titled "LDAP Servers and Sequence" which contains a table with the following columns: Network Address, Authentication Type, Security, TCP Port, TimeOut(Sec), and Attempts. To the right of the table are buttons for Add, Edit, Delete, Up, and Down. At the bottom of the interface, there are buttons for Audit Trail, Display, Test, and Apply.

FIGURE 118 AAA Settings tab - LDAP server

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 [“Creating an AD user account”](#) on page 159). 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 159). NmRoles defines the Management application user roles (such as Host Administrator, IP System 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, All IP Products).

3. Add or edit a LDAP server by referring to [“Configuring an LDAP server”](#) on page 333.

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.

4. Rearrange the LDAP servers in the table by selecting a server and click the **Up** or **Down** button to move it.
5. Delete a LDAP server by selecting the server and click **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 137).
 - 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 159).
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 157.
You do not need to create users in the local database.
11. Click **Apply** to save the configuration.

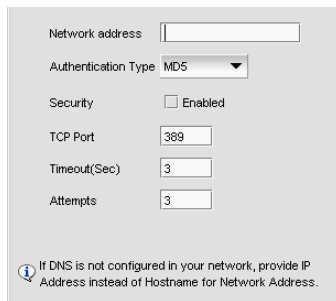
To display the authentication audit trail, refer to [“Displaying the client authentication audit trail”](#) on page 341.
12. Click **Close** to close the Server Management Console.

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 119](#)).

11 AAA Settings tab



Network address

Authentication Type **MDS**

Security Enabled

TCP Port

Timeout(Sec)

Attempts

If DNS is not configured in your network, provide IP Address instead of Hostname for Network Address.

FIGURE 119 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**.

The screenshot shows the configuration interface for TACACS+ servers. At the top, there are four dropdown menus: 'Primary Authentication' set to 'TACACS+ Server', 'Secondary Authentication' set to 'None', 'Fail Over Option' set to 'TACACS+ Servers Not Reachable', and 'Authorization Preference' set to 'Local Database'. Below these is a table titled 'TACACS+ Servers and Sequence' with columns for 'Network Address', 'TCP Port', 'Timeout(Sec)', and 'Attempts'. To the right of the table are buttons for 'Add', 'Edit', 'Delete', 'Up', and 'Down'. At the bottom of the interface are buttons for 'Audit Trail', 'Display', 'Test', and 'Apply'.

FIGURE 120 AAA Settings tab - TACACS+ server

3. Add or edit a TACACS+ server by referring to [“Configuring a TACACS+ server”](#) on page 336.
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 341.
12. Click **Close** to close the Server Management Console.

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 119).

The screenshot shows a dialog box with the following fields and values:

- Network address:
- TCP Port:
- Shared Secret:
- Confirm Secret:
- Timeout(Sec):
- Attempts:

Below the fields, there is a note: "If DNS is not configured in your network, provide IP Address instead of Hostname for Network Address."

FIGURE 121 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 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.

11 AAA Settings tab

The screenshot shows a configuration window for AAA settings. At the top, there are two dropdown menus: 'Primary Authentication' set to 'CAC' and 'Secondary Authentication' set to 'None'. Below these is an 'Authorization Preference' dropdown set to 'Primary Authentication Server'. The main section is titled 'Active Directory Single-Sign-On Account for Network Advisor' and contains four input fields: 'Username', 'Password', 'Confirm Password', and 'Kerberos Service Principal Name'. To the right of the last field is a syntax hint: 'Syntax: - <Service Name>/<Hostname>'. At the bottom left is an 'Audit Trail' section with a 'Display' button. At the bottom right are 'Test' and 'Apply' buttons.

FIGURE 122 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-VNNext-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 encrypts the 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 341.

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 341.
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 341.
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 341.
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.

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 77.

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.

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 123](#)).



FIGURE 123 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**.

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 5.

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.

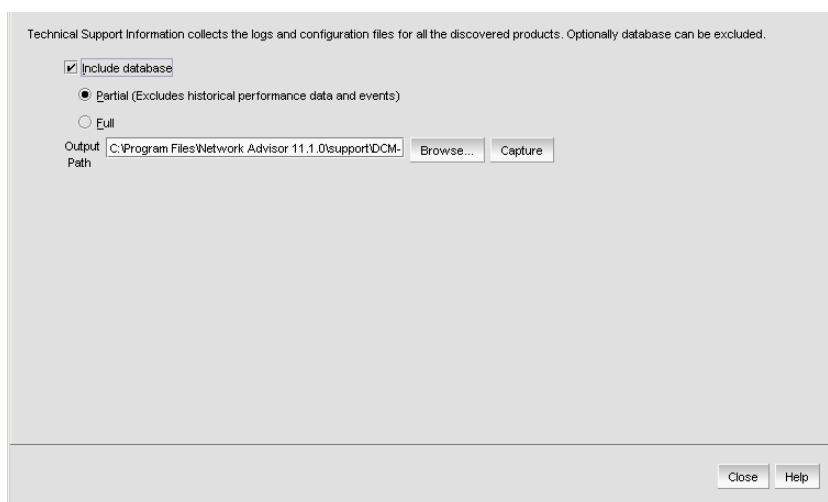


FIGURE 124 Technical Support Information tab

2. Select the **Include database** check box to capture database server support save files and choose one of the following options:
- Select the **Partial** option to exclude historical data and events from the database capture.
 - Select the **Full** option to include historical data and events from the database capture.

NOTE

It is recommended that you only capture the partial database.

NOTE

You should only capture the full database when you need to debug Historical Performance Management or Historical Events issues.

11 HCM Upgrade tab

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 125 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 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 126 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.

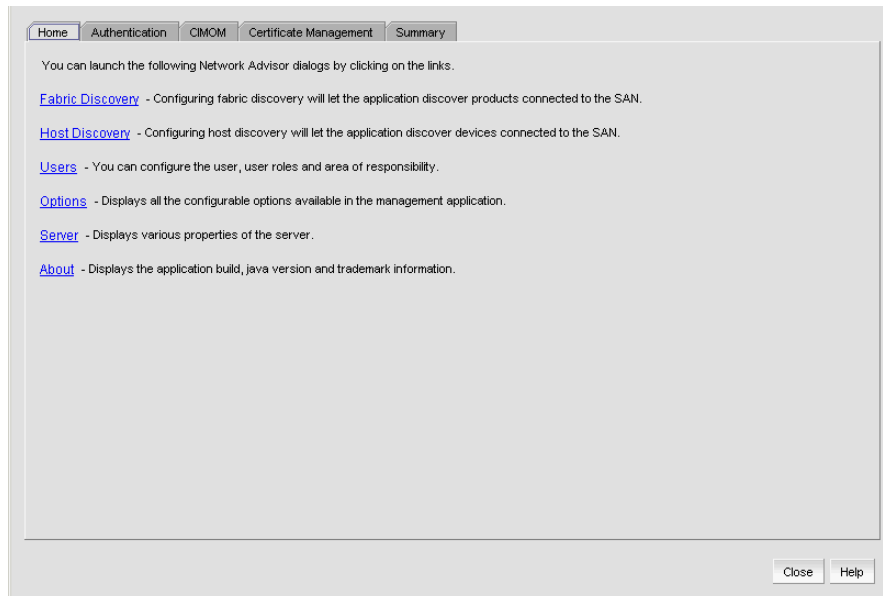


FIGURE 127 SMIA Configuration Tool dialog box

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=WBEM Solutions:f1f65c3b-27f1-4b70-9ced-e412e93a8d5e), (Communication
Mechanism=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),(AuthenticationMechanismDescriptions=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)
```

- `slptool findattr 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 findattr 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=WBEM Solutions: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),(AuthenticationMechanismDescriptions=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 SLP APIs 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 findattr service:wbem:http://IP_Address:Port` and press **Enter**.
 - Type `# Install_Home/cimom/bin/slptool findattr 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 SLP APIs 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 findattr service:wbem:http://IP_Address:Port` and press **Enter**.
 - Type `> slptool findattr 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 “[Discovering fabrics](#)” on page 39.
- **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 58.
- **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 140.
- **Options** – enables you to configure the Management application settings. For step-by-step instructions, refer to “[Application Configuration](#)” on page 75.
- **Server** – enables you to view server properties. For step-by-step instructions, refer to “[Viewing server properties](#)” on page 10.
- **About** – enables you to display information about the Management application, including the build number, Java version, and trademark information.

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 1243.

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 TrustStore 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 TrustStore 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 128 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 1243.

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.

The screenshot shows the 'CIMOM' configuration tab in a web-based management console. At the top, there are navigation tabs: Home, Authentication, CIMOM (selected), Certificate Management, and Summary. Below the tabs, the main content area is titled 'Configure HTTP and HTTPS connections between the CIMOM and CIM Client'. It contains several configuration sections:

- Enable SSL:** A checked checkbox.
- SMI Agent Port#:** A text input field containing '5989'.
- Current Value:** A text input field containing '5989'.
- Default Value:** A text input field containing '5989'.
- IP Configuration:** A section with a dropdown menu for 'Bind Network Address' set to 'TechOPS2008'.
- CIMOM Logs will be written into cimom/server/logs folder:** A section with three settings:
 - Log Level:** A dropdown menu set to 'INFO'.
 - File Size:** A spinner control set to '5' MB.
 - Number of Files:** A spinner control set to '10'.

At the bottom of the configuration area, there is an 'Apply' button. Below the main configuration area, a message states: 'Changes will take effect at the next CIMOM restart. CIMOM server can be restarted from Network Advisor Server Management Console.' At the very bottom right, there are 'Close' and 'Help' buttons.

FIGURE 129 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 1243.

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 1243.

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 1243.

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.

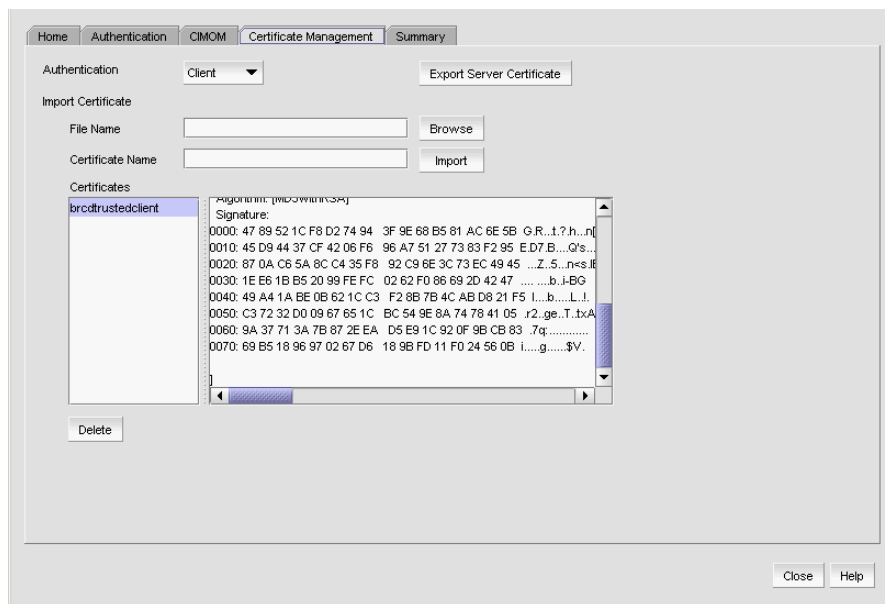


FIGURE 130 Certificate Management tab

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 1243.

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 1243.

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 1243.

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.

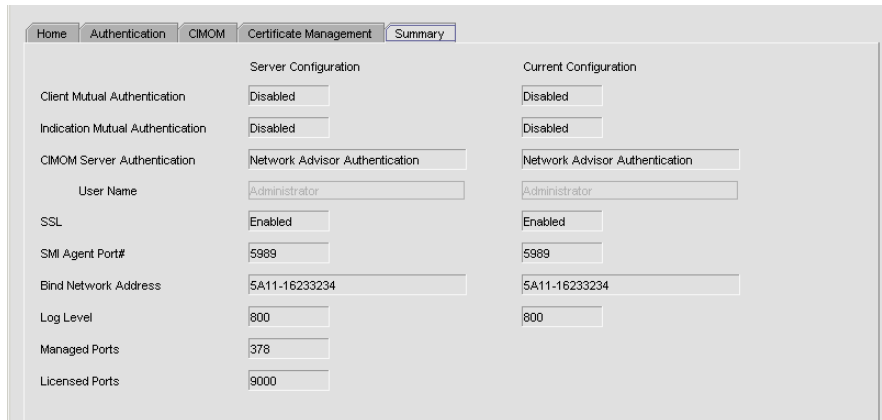


FIGURE 131 Summary tab

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 24

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.

TABLE 24

Field/Component	Description
Log Level	Displays the log level for the Server Configuration and the Current Configuration. Options include the following: <ul style="list-style-type: none">• 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 29.
Licensed Ports	Displays the number of licensed ports.

3. Click **Close** to close the **SMIA Configuration Tool** dialog box.

11 SMI Agent Configuration Tool

SAN Device Configuration

In this chapter

- Configuration repository management 363
- Enhanced group management 380
- Firmware management 380
- Frame viewer 386
- Ports 389
- Port commissioning overview 403
- Port Auto Disable 416

Configuration repository management

(Professional only) Configuration files are run as a DerbyPostgress database as part of the Management application service; however, they are only stored as a flat file. For Windows platforms the default location is `<Install_Home>\data\databases\<Management_Application_Name>.db`

Professional only allows you to back up the configuration repository 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 Windows platforms the default location is
`Install_Home\data\database\Management_Application_Name.db`

For more information about the database fields, refer to [“Database Fields”](#) on page 1307.

Saving switch configurations on demand

NOTE

Save switch configuration is only supported on Fabric OS switches.

NOTE

This feature requires a Trial or Licensed version.

NOTE

To save switch configuration on more than one switch at a time, you must have the Enhanced Group Management license.

NOTE

The Management application enables you to save the same switch configuration to the repository using two methods: on demand (Configure > Configuration > Save) or by defining a schedule (Configuration > Schedule Backup).

Configuration files are uploaded from the selected switches and stored in individual files. Files are named with the convention `cfg_fabricName_switchName_domainID`.

Use this procedure to immediately save switch configurations to the repository. To create a scheduled back up of switch configurations to the repository, refer to [“Scheduling switch configuration backup”](#) on page 366.

1. Select **Configure > Configuration > Save Now**.

The **Save Switch Configurations** dialog box displays.

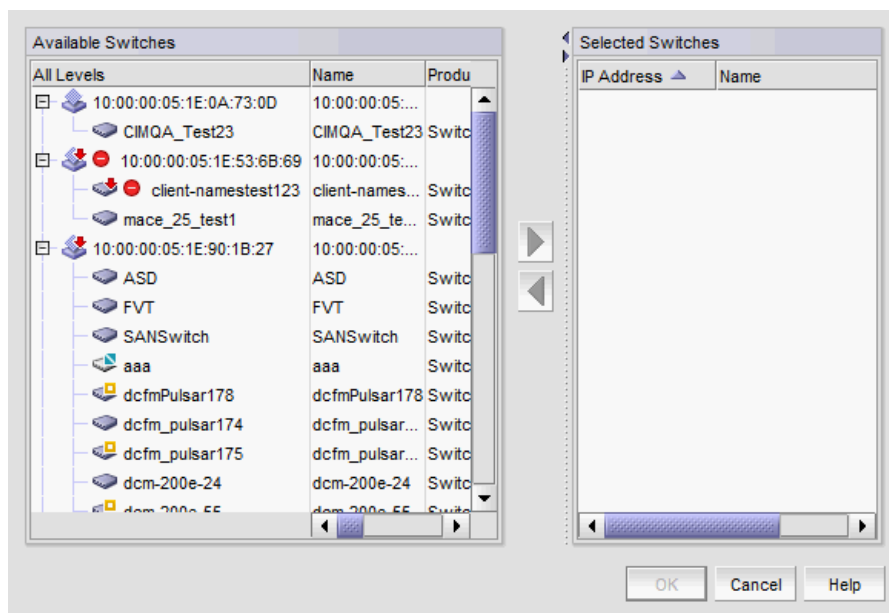


FIGURE 132 Save switch configurations

2. Select the switches for which you want to save configuration files from the **Available Switches** table.
3. Click the right arrow to move the selected switches to the **Selected Switches** table.

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.

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 > Restore**.

The **Restore Switch Configuration** dialog box displays.

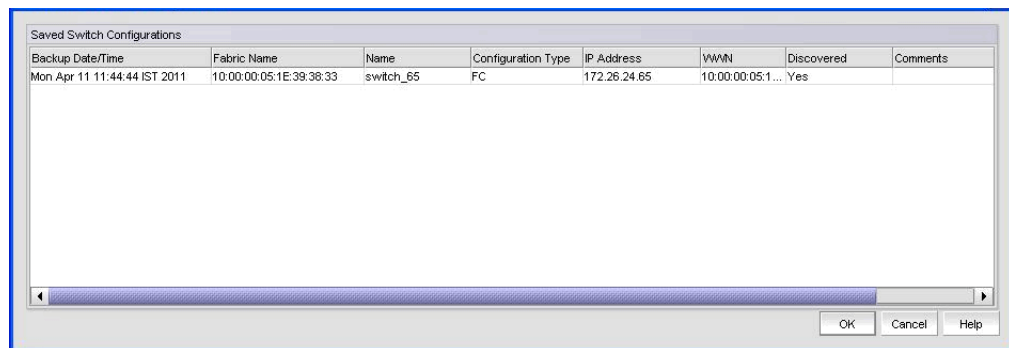


FIGURE 133 Restore Switch Configuration dialog box

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

This feature requires a Trial or Licensed version.

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.

NOTE

The Management application enables you to save the same switch configuration to the repository using two methods: on demand (Configure > Configuration > Save) or by defining a schedule (Configuration > Schedule Backup).

Use this procedure to create a scheduled back up of switch configurations to the repository. To save switch configurations to the repository on demand, refer to [“Saving switch configurations on demand”](#) on page 364.

The configuration files are stored in the Management application database.

1. Select **Configure > Configuration > Save on Schedule**.

The **Schedule Backup of Switch Configurations** dialog box displays.

Enable scheduled backup

Schedule

Frequency **Daily**

Day **Tuesday**

Hour Minute

Time **10** **15**

Purge Backups **30** days and older

Scope - Includes all switches discovered at time of backup

Backup all fabrics

Selected Fabrics

Backup	Fabric Name ▲	Status	# of Switches
<input checked="" type="checkbox"/>	10:00:00:05:1E:34:D...	Down	1
<input checked="" type="checkbox"/>	10:00:00:05:1E:35:3...	Marginal	1
<input checked="" type="checkbox"/>	10:00:00:05:1E:37:B...	Unreachable	1

OK Cancel Help

FIGURE 134 Schedule backup of switch configurations

2. Click 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 back up to run.

If **Frequency** is **Daily**, the Day list is grayed out.

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 back up to run.
- The maximum age allowed before you **Purge Backups**.

The number of purge days (7 through 90) should be at least one day more than the selected backup frequency.

The backup purge thread runs every day at 12:30 PM and deletes all back up configurations that exceed the maximum age allowed.

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.

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 Switch Configuration Repository 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 > Configuration Repository**.

The **Switch Configuration Repository** dialog box displays.

Keep	Backup Date/Time	Fabric Name	Name	Configuration Type	IP Address	WWN	Discovered	Comments
<input type="checkbox"/>	Mon Jun 06 12:01:30 ...	10:00:00:05:1E:53...	Ency_4100	FC	10.24.45.66	10:00:00:05:1E:02...	Yes	

FIGURE 135 Switch Configuration Repository

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
- **Backup Date/Time** — The date and time the last backup occurred. This is the backup that will be restored.
- **Fabric Name** — The name of the fabric to which the selected switch belongs.
- **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.
- **Discovered** — Whether the switch is discovered or not.
 - Yes — The switch is discovered.
 - No — The switch was deleted from discovery.
- **Comments** — Comments regarding the switch.

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 this dialog box:

- [“Keeping a copy past the defined age limit”](#) on page 373
- [“Viewing configuration file content”](#) on page 370

- “Deleting a configuration” on page 372
- “Exporting a configuration” on page 372
- “Importing a configuration” on page 372

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 > Configuration Repository**.

The **Switch Configuration Repository** dialog box displays.

2. Click **View**.

The configuration details display, details include the backed-up switch, including boot parameters, licensing information, and configuration. 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 136 Configuration file content

3. Click **Close** to close the dialog box.
4. Click **Yes** on the message.

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 > Configuration Repository**.

The **Switch Configuration Repository** dialog box displays.

2. Click **View**.

The configuration details display.

3. Enter the information you want to search for in the 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 137 Configuration file content

4. Click **Close** to close the dialog box.
5. Click **Yes** on the message.

Deleting 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 > Configuration Repository**.
The **Switch Configuration Repository** dialog box displays.
2. Select the configuration 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 > Configuration Repository**.
The **Switch Configuration Repository** dialog box displays.
2. Select the configuration you want to export, and click **Export**.
The file chooser appropriate to your operating system displays.
3. Use the file chooser to select the location into which you want to export the configuration.
4. Click **Export**.
The configuration is automatically named (*Device_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 > Configuration Repository**.
The **Switch Configuration Repository** dialog box displays.
2. Click **Import**.
The file chooser appropriate to your operating system displays.
3. Use the file chooser to select the file from which you want to import the configuration, and click **Import**.

Keeping a copy past the defined age limit

NOTE

This feature requires a Trial or Licensed version.

1. Select **Configure > Configuration > Configuration Repository**.
The **Switch Configuration Repository** 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.
3. Click **OK**.

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.

Select **Configure > Configuration > 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:

1. **Overview**, which describes the wizard.
2. **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 25](#) on page 374.
3. **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 26](#) on page 374.
4. **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 27](#) on page 374.
5. **Destination Switches**, which allows you to select the destination switch. For more information about the fields and components of this step, refer to [Table 28](#) on page 376.
6. **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 29](#) on page 377.
7. **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 30](#) on page 377.

To proceed to the next step in the wizard, click **Next**. To return to the previous screen, click **Previous**.

TABLE 25 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: <ul style="list-style-type: none"> • 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 26 Step 3. Source Location

Field/Component	Description
Configuration Repository option	Select to replicate the entire configuration repository 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 27 Step 4. Source Configuration

Field/Component	Description
Saved Switch Configuration table (Configuration Repository only)	Lists the information related to the saved switch, if you selected Configuration Repository on the Source Location screen.
Backup Date/Time (Configuration Repository 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.
Discovered	Whether the switch is discovered or not. Yes – The switch is discovered. No – The switch was deleted from discovery.
Comments	Comments regarding the switch.

TABLE 27 Step 4. Source Configuration (Continued)

Field/Component	Description
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.
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 28 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.
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.

TABLE 28 Step 5. Destination Switches (Continued)

Field/Component	Description
Current Firmware	The current firmware.
Status	The status of the switch .

TABLE 29 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 30 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.

Select **Configure > Configuration > 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:

1. **Overview**, which describes the wizard.
2. **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 31](#) on page 378.
3. **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 32](#) on page 378.
4. **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 33](#) on page 379.
5. **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 34](#) on page 379.
6. **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 35](#) on page 379.

To proceed to the next step in the wizard, click **Next**. To return to the previous screen, click **Previous**.

TABLE 31 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 32 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 33 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 34 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 35 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 380.
- **Security** - enables you to perform Group Security Policy Replication.
For specific instructions for security, refer to [“Configuration repository management”](#) on page 363.
- **Configuration Management** - enables you to perform Group Configuration Upload and Replication.
For specific instructions for configuration management, refer to [“Replicating configurations”](#) on page 373.

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

Firmware download is not supported in pure IPv6 mode.

NOTE

You cannot use Fabric OS firmware download with command line options in the Management application.

Downloading firmware

NOTE

Non-disruptive firmware download (HCL) is not supported when downgrading from Fabric OS version 6.2 to 6.1. 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.

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 138).

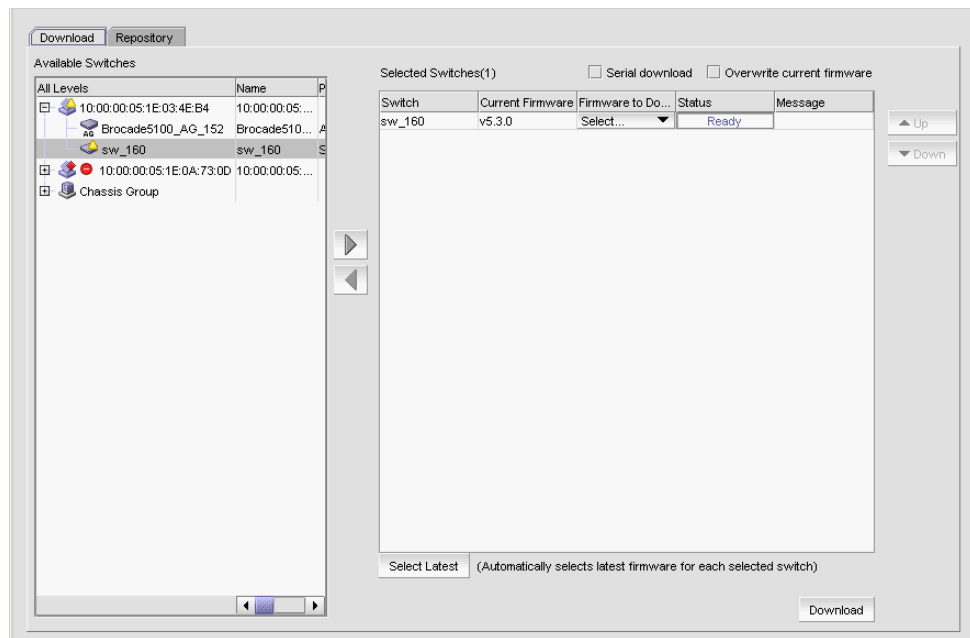


FIGURE 138 Firmware download

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 selected 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:
 - Create user and password.
 - Select the **Shared folders** link and set 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 6.0 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.

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 126.

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 139](#)).

Initially, the repository is empty. You must import firmware files into the repository (refer to [“Importing a firmware file”](#) on page 384). Imported firmware files are then displayed under **Firmware Repository**.

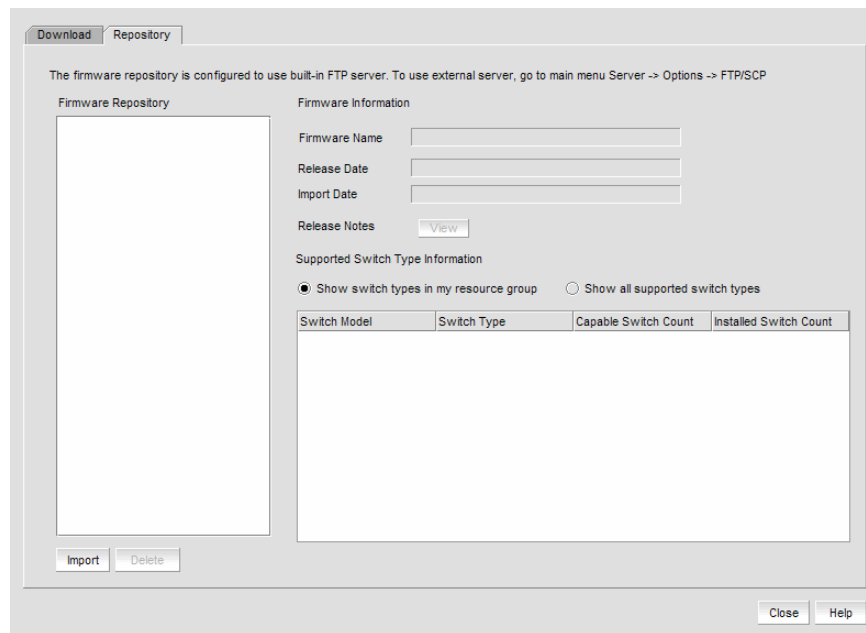


FIGURE 139 Firmware repository

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 384.
 5. Click **Delete** to delete firmware files from the firmware repository. Refer to [“Deleting a firmware file”](#) on page 385.
 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 139](#)).
3. Click **Import**.

The **Import Firmware from File** dialog box displays ([Figure 140](#)).

FIGURE 140 Import firmware

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 102) 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.

Frame viewer

NOTE

Frame viewer is only supported on Fabric OS devices running 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.

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.

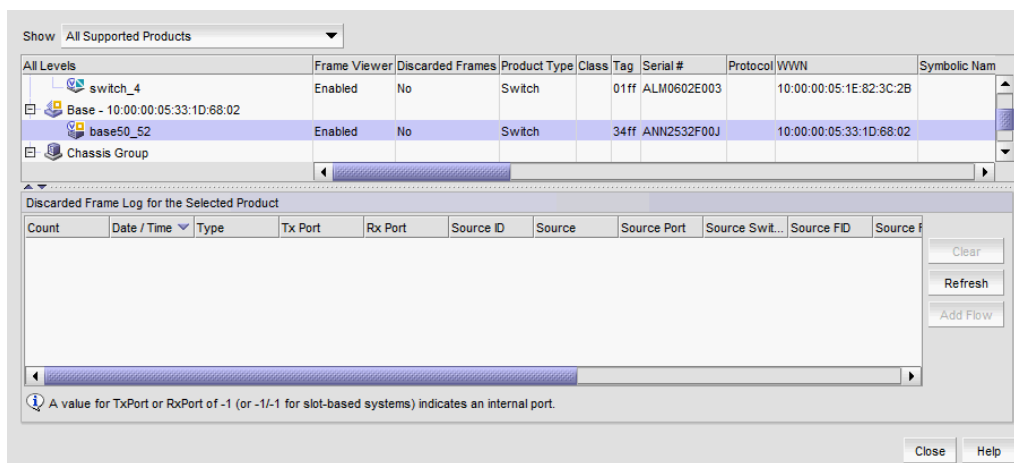


FIGURE 141 Discarded Frames dialog box

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](#)” on page 251) 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**.
 - **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.
 - **Destination ID** – Destination ID in hex PID format.
 - **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 389). 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 389). 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 [“Provisioning flows”](#) on page 996).

NOTE

Flow Vision is supported on platforms running Fabric OS 7.2 and later.

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.
 - **Destination ID** – Destination ID in hex PID format.
 - **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 389). 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 389). 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 “[Provisioning flows](#)” on page 996).

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 386 or [“Viewing discarded frames from a port”](#) on page 388).
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 386 or [“Viewing discarded frames from a port”](#) on page 388).
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 142](#)).

Fabric: 10:00:00:05:1E:90:1B:27

Filter Add Flow All Switches Refresh Help

Port Number	Blade Number	Port Name	Switch	User Port Number	Area ID/Port Index	FC Address	Port WWN	Calculated
0	N/A	hjh	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
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

FIGURE 142 Port Connectivity View dialog box

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 393) in the **Port Connectivity View** dialog box.
- **Add Flow** button — Select a port and click to add a flow definition (refer to “[Provisioning flows](#)” on page 996).

NOTE

Flow Vision is supported on platforms running Fabric OS 7.2 and later.

- **All Switches** 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.
- 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 393 or “[Disabling a port](#)” on page 393.
 - **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).
- **Switch Routing Policy** – Whether a routing policy, for example, port-based routing policy, is enabled.
- **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 143).

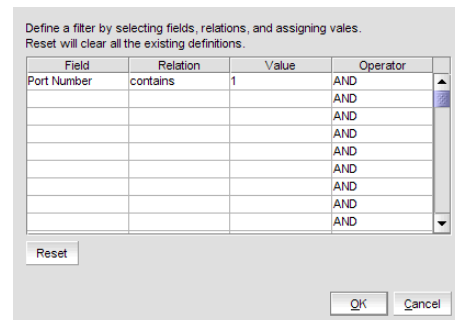


FIGURE 143 Filter dialog box

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 142).

COLUMN	VALUE
Actual Distance	
Area ID (Hex)/Port Index (Hex)	20
Blade Number	N/A
Blocked	
Buffer Limited	N/A
Buffers Needed/Allocated	
COS	
Capability	
Connected Blade Number	N/A
Connected Port Area ID (Hex)/Port Index (Hex) 0 (0x00)	
Connected Port Name	
Connected Port Number	0
Connected Port Speed	
Connected Port State	
Connected Port Status	
Connected Port WWN	
Connected User Port Number (Hex)	
Device Node WWN	
Device Port / Switch Domain Id	
Device Port / Switch Manufacturer	
Device Port / Switch Manufacturing Plant	
Device Port / Switch Name	
Device Port / Switch State	
Device Port / Switch Type Number	
Device Port / Switch WWN	

FIGURE 144 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 1271.

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 36 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 37) may not be available for all products.

TABLE 37 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.
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.

TABLE 37 Port connection properties (Continued)

Field	Description
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 6.2.2d, 6.3.2c, 6.4.1a, or 6.4.2 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.
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.

TABLE 37 Port connection properties (Continued)

Field	Description
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.
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

NOTE

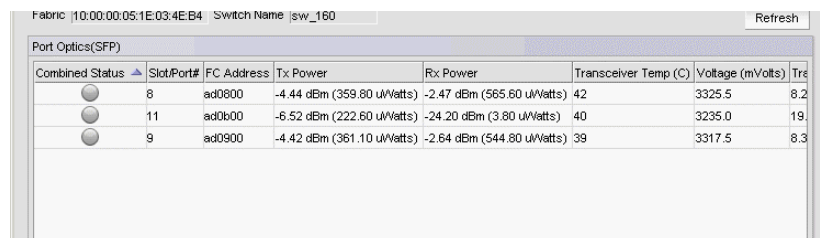
QSFP ports do not display in the **Port Optics** dialog box.

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 [\(Figure 145\)](#).






Combined Status	Slot/Port#	FC Address	Tx Power	Rx Power	Transceiver Temp (C)	Voltage (mVolts)	Tr
	8	ad0800	-4.44 dBm (359.80 uWatts)	-2.47 dBm (565.60 uWatts)	42	3325.5	8.2
	11	ad0b00	-6.52 dBm (222.60 uWatts)	-24.20 dBm (3.80 uWatts)	40	3235.0	19.
	9	ad0900	-4.42 dBm (361.10 uWatts)	-2.64 dBm (544.80 uWatts)	39	3317.5	8.3

FIGURE 145 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 1135.

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.

Port commissioning overview

NOTE

Port commissioning is only supported on Fabric OS devices running Fabric OS 7.1 or later.

Port commissioning provides an automated mechanism to remove an E-Port or F-Port from use (decommission) and to put it back in use (recommission). This feature identifies the target port and communicates the intention to decommission or recommission the port to those systems within the fabric affected by the action. Each affected system can agree or disagree with the action, and these responses are automatically collected before a port is decommissioned or recommissioned.

Note the following restrictions of port commissioning:

- The local switch and the remote switch on the other end of the E-Port or F-Port must both be running Fabric OS 7.1.0 or later.
- Port commissioning is not supported on links configured for encryption or compression.
- Port commissioning is not supported on ports with DWDM, CWDM, or TDM.
- E-Port commissioning requires that the lossless feature is enabled on both the local switch and the remote switch.
- Fabric tracking must be enabled (refer to [“Enabling fabric tracking”](#) on page 132) to maintain the decommissioned port details (such as port type, device port wwn, and so on). Do not accept changes in the Management application client.

Viewing existing CIMOM servers

NOTE

Port commissioning is only supported on Fabric OS devices running Fabric OS 7.1 or later.

Before you can decommission or recommission an F-Port, you must register the CIMOM servers within the fabric affected by the action.

1. Select **Configure > Port Commissioning > Setup**.

The **Port Commissioning Setup** dialog box displays ([Figure 146](#)).

Use this dialog to create a list of systems that will be checked prior to device port (non-E port) decommission and recommission.

Add / Edit Systems and Credentials		Systems List						
Network Address	Description	Network Address	Description	CIMOM Port	Namespace	User ID	Stat...	Last Contacted
				5989	root/cimv2			

Buttons: Import, Export, Change Credentials, Test, OK, Cancel, Help

FIGURE 146 Port Commissioning Setup dialog box

The **Port Commissioning Setup** dialog box has two main areas. The **Add/Edit Systems and Credentials** area enables you to register CIMOM servers (system and credentials) one at a time and contains the following fields and components:

- **Network Address** — Enter the IP address (IPv4 or Ipv6 format) or host name of the CIMOM server in the field.
- **Description** — (Optional) Enter a description of the CIMOM server in the field. The description cannot be over 1024 characters.
- **CIMOM Port** — Enter the CIMOM port number for the CIMOM server in the field. The default port number is 5989.
- **Namespace** — Enter the namespace of the CIM_FCPort in the field. The default namespace is root/cimv2.
- **Credentials - User ID** (Optional) — Enter a user identifier for the CIMOM server in the field. The credentials user identifier cannot be over 128 characters.
- **Credentials - Password** (Optional) — Enter a password in the field. The password cannot be over 512 characters.
- Left arrow button — Select a CIMOM server in the **Systems List** and click to move the defined CIMOM server credentials to the **Add/Edit Systems and Credentials** area for editing or deletion.
- Right arrow button — Click to move the defined CIMOM server credentials from the **Add/Edit Systems and Credentials** area to the **Systems List**.

The **Systems List** details the defined CIMOM server and contains the following data:

- **Network Address** — The IP address (IPv4 or Ipv6 format) or host name of the system.
- **Description** — User-defined description of the system.
- **CIMOM Port** — The CIMOM port number of the system.
- **Namespace** — The namespace of the CIM_FCPort.
- **User ID** — The user identifier for the system.
- **Status** — The system connectivity status. Updates when you test the reachability of the CIMOM server and when you contact the CIMOM server to respond to the F-Port decommission or recommission request. Valid status options include:
 - **OK** — CIMOM server contact successful with current credentials.
 - **Not Contacted Yet** — CIMOM servers configured, connectivity not tested yet.
 - **Credentials Updated** — Credentials changed, connectivity not tested yet.
 - **Credentials Failed** — CIMOM server contact failed with current credentials.
 - **Not Reachable** — CIMOM server not reachable.
 - **Wrong Namespace** — CIMOM server namespace is incorrect.
- **Last Contacted** — The last time you contacted the system. Updates when you test the reachability of the CIMOM server and when you contact the CIMOM server to respond to the F-Port decommission or recommission request.

2. To register a CIMOM server, refer to [“Registering a CIMOM server”](#) on page 405.
3. To edit a CIMOM server, refer to [“Editing CIMOM server credentials”](#) on page 405.
4. To import CIMOM servers, refer to [“Importing CIMOM servers and credentials”](#) on page 406.
5. To export CIMOM servers, refer to [“Exporting CIMOM servers and credentials”](#) on page 406.

6. To edit CIMOM server credentials for one or more CIMOM servers, refer to “[Changing CIMOM server credentials](#)” on page 407.
7. To validate CIMOM server credentials, refer to “[Testing CIMOM server credentials](#)” on page 408.
8. To delete CIMOM servers, refer to “[Deleting CIMOM server credentials](#)” on page 408.
9. Click **OK** to close **Port Commissioning Setup** dialog box.

Registering a CIMOM server

Before you can decommission or recommission an F-Port, you must register the CIMOM servers within the fabric affected by the action.

1. Select **Configure > Port Commissioning > Setup**.
The **Port Commissioning Setup** dialog box displays ([Figure 146](#)).
2. Enter the IP address (IPv4 or Ipv6 format) or host name of the CIMOM server in the **Network Address** field.
3. (Optional) Enter a description of the CIMOM server in the **Description** field.
The description cannot be over 1024 characters.
4. Enter the CIMOM port number for the CIMOM server in the **CIMOM Port** field.
The default port number is 5989.
5. Enter the namespace of the CIM_FCPort in the **Namespace** field.
The default namespace is root/cimv2.
6. (Optional) Enter a user identifier for the CIMOM server in the **Credentials User ID** field.
The credentials user identifier cannot be over 128 characters.
7. (Optional) Enter a password in the **Password** field.
The password cannot be over 512 characters.
8. Click the right arrow button to add the new CIMOM server and credentials to the **Systems List**.
The application validates the mandatory fields.
9. Select the new CIMOM server in the **System List** and click **Test** to check connectivity.
When testing is complete, the updated status displays in the **Status** column of the **Systems List** for the selected CIMOM server.
10. Click **OK** or **Apply** to save your work and save the CIMOM server details in the database.

Editing CIMOM server credentials

1. Select **Configure > Port Commissioning > Setup**.
The **Port Commissioning Setup** dialog box displays ([Figure 146](#)).
2. Select a CIMOM server from the **System List** and click the left arrow button to edit the CIMOM server credentials.

- (Optional) Edit the description of the CIMOM server in the **Description** field.
The description cannot be over 1024 characters.
- Enter the CIMOM port number for the CIMOM server in the **CIMOM Port** field.
The default port number is 5989.
- Enter the namespace of the CIM_FCPort in the **Namespace** field.
The default namespace is root/cimv2.
- (Optional) Enter a user identifier for the CIMOM server in the **Credentials User ID** field.
The credentials user identifier cannot be over 128 characters.
- (Optional) Enter a password in the **Password** field.
The password cannot be over 512 characters.
- Click the right arrow button to update the CIMOM server credentials in the **Systems List**.
- Click **OK** or **Apply** to save your work and save the CIMOM server details in the database.

Importing CIMOM servers and credentials

You can import one or more CIMOM servers (system and credentials) using a CSV formatted file. You can import a maximum of 2,000 CIMOM servers.

- Select **Configure > Port Commissioning > Setup**.
The **Port Commissioning Setup** dialog box displays (Figure 146).
- Click **Import** to import CIMOM server information from a file.
The CSV file must use the following format:
Network Address, User ID,CIMOM Port, Namespace, Description, Password
Example
`10.24.48.100,user,2015,root/cimv2,IBM Host,password`
Network Address is mandatory. If you do not provide values for the User ID,CIMOM Port, Namespace, Description, and Password; the Management application provides default values.
- Browse to the location of the file (.csv format) and click **Open**.
The imported CIMOM servers display in the **Systems List**.
- Click **OK** or **Apply** to save your work and save the CIMOM server details in the database.

Exporting CIMOM servers and credentials

- Select **Configure > Port Commissioning > Setup**.
The **Port Commissioning Setup** dialog box displays (Figure 146).
- Click **Export** to export CIMOM server information to a file.
The **Export Files** dialog box displays.

3. Browse to the location where you want to export the file (.csv format) and click **Save**.

The CSV file uses the following format:

Network Address, User ID,CIMOM Port, Namespace, Description,

Example

```
10.24.48.100,user,2015,root/cimv2,IBM Host,
```

NOTE

Export does not include the password. You can edit the exported file to add the password to the credentials.

4. Click **OK** to close the **Port Commissioning Setup** dialog box.

Changing CIMOM server credentials

You can edit the CIMOM server credentials for one or more CIMOM servers at the same time.

1. Select **Configure > Port Commissioning > Setup**.

The **Port Commissioning Setup** dialog box displays ([Figure 146](#)).

2. Select one or more CIMOM servers from the **System List** table and click **Change Credentials**.

The **Edit Credentials** dialog box displays. If you selected one CIMOM server, the credentials for the selected server display in the dialog box. If you selected more than one CIMOM server, the credential fields are empty.

3. (Optional) Enter a user identifier for the CIMOM server in the **User ID** field.

The user identifier cannot be over 128 characters.

4. (Optional) Enter a password in the **Password** field.

The password cannot be over 512 characters.

5. Click **OK** to close the **Edit Credentials** dialog box.

The **Port Commissioning Setup** dialog box and the status of the selected CIMOM server rows displays “Credentials Updated”.

To validate the credentials, refer to “[Testing CIMOM server credentials](#)” on page 408.

6. Click **OK** to close the **Port Commissioning Setup** dialog box.

Testing CIMOM server credentials

You should validate the CIMOM server credentials before you decommission or recommission ports. During the decommission or recommission of an F-Port, the Management application validates the CIMOM server credentials.

1. Select a device and select **Configure > Port Commissioning > Setup**.

The **Port Commissioning Setup** dialog box displays (Figure 146).

2. Select one or more CIMOM servers from the **System List** table and click **Test**.

The status of the selected CIMOM server rows display “Testing” in the **System List** table. When testing is complete, the CIMOM server connectivity status displays. Valid status options include:

- **OK** – CIMOM server contact successful with current credentials.
- **Not Contacted Yet** – CIMOM servers configured, connectivity not tested yet.
- **Credentials Updated** – Credentials changed, connectivity not tested yet.
- **Credentials Failed** – CIMOM server contact failed with current credentials.
- **Not Reachable** – CIMOM server not reachable.
- **Wrong Namespace** – CIMOM server namespace is incorrect.

3. Click **OK** to close the **Port Commissioning Setup** dialog box.

When the test is complete, an application event displays in the Master Log detailing success or failure.

Deleting CIMOM server credentials

1. Select **Configure > Port Commissioning > Setup**.

The **Port Commissioning Setup** dialog box displays (Figure 146).

2. Select one or more CIMOM server from the **System List** table and click the left arrow button.

The details for the last selected CIMOM server row displays in the **Add/Edit System and Credentials** area.

3. Confirm that this is the CIMOM server you want to delete and click **OK** or **Apply** to delete the CIMOM server from the **Port Commissioning Setup** dialog box.

When the deletion is complete, an application event displays in the Master Log detailing success or failure.

Decommissioning an F-Port

NOTE

You must configure at least one CIMOM server (refer to [“Registering a CIMOM server”](#) on page 405) before you can decommission an F-Port.

NOTE

Fabric tracking must be enabled (refer to [“Enabling fabric tracking”](#) on page 132) to maintain the decommissioned port details (such as port type, device port wwn, and so on). Do not accept changes in the Management application client.

1. Select the F-Port, then select **Configure > Port Commissioning > Decommission > Port**.

The **Port Commission Confirmation** dialog box displays.

2. Choose one of the following options:
 - **Apply Default Settings** (default) — Select to have the Management application contact all registered CIMOM servers within the fabric affected by the action and obtains the status from each CIMOM server. If all CIMOM servers are okay, the Management application sends a CAL Request to decommission the port. If even one CIMOM server is not okay, decommissioning fails.
 - **Force** — Select to force the port decommission. The Management application still contacts all registered CIMOM servers within the fabric affected by the action, but forces the port decommission regardless of the CIMOM server response.

NOTE

If the CIMOM server is not reachable or the credentials fail, F-Port decommission does not occur.

3. Click **OK** on the **Port Commission Confirmation** dialog box.

While decommissioning is in progress, a down arrow icon displays next to the port icon in the Product List. You can view the port commissioning results in the deployment reports (refer to [“Port commissioning deployment report”](#) on page 413).

When the decommission is complete, an application event displays in the Master Log detailing success or failure.

Recommissioning an F-Port

NOTE

You must configure at least one CIMOM server (refer to [“Registering a CIMOM server”](#) on page 405) before you can recommission an F-Port.

Select the F-Port, then select **Configure > Port Commissioning > Recommission > Port**.

While recommissioning is in progress, an up arrow icon displays next to the port icon in the Product List. You can view the port commissioning results in the deployment reports (refer to [“Port commissioning deployment report”](#) on page 413).

When the recommission is complete, an application event displays in the Master Log detailing success or failure.

Decommissioning an E-Port

NOTE

You must enable Lossless DLS on both the source and destination switches before you decommission an E-Port.

NOTE

Fabric tracking must be enabled (refer to [“Enabling fabric tracking”](#) on page 132) to maintain the decommissioned port details (such as port type, device port wwn, and so on). Do not accept changes in the Management application client.

Select the E-Port in the Product List, then select **Configure > Port Commissioning > Decommission > Port**.

While decommissioning is in progress, a down arrow icon displays next to the port icon in the Product List. You can view the port commissioning results in the deployment reports (refer to [“Port commissioning deployment report”](#) on page 413).

When the decommission is complete, an application event displays in the Master Log detailing success or failure.

Recommissioning an E-Port

NOTE

You do not need to enable Lossless DLS before you recommission an E-Port.

Select the E-Port in the Product List, then select **Configure > Port Commissioning > Recommission > Port**.

While recommissioning is in progress, an up arrow icon displays next to the port icon in the Product List. You can view the port commissioning results in the deployment reports (refer to [“Generating a deployment report”](#) on page 918).

When the recommission is complete, an application event displays in the Master Log detailing success or failure.

Decommissioning all ports on a switch

NOTE

Fabric tracking must be enabled (refer to [“Enabling fabric tracking”](#) on page 132) to maintain the decommissioned port details (such as port type, device port wwn, and so on). Do not accept changes in the Management application client.

1. Select the switch or logical switch for which you want to decommission all ports, then select **Configure > Port Commissioning > Decommission > All F-Ports on the Switch**.

NOTE

You can only decommission ports from the logical switch, not the physical chassis.

The **Port Commission Confirmation** dialog box displays.

- Choose one of the following options:

- Apply Default Settings** (default) — Select to have the Management application perform one of the following actions:

The Management application contact all registered CIMOM servers within the fabric affected by the action and obtains the status from each CIMOM server. If all CIMOM servers are okay, the Management application sends a CAL Request to decommission the port. If even one CIMOM server is not okay, decommissioning fails.

- Force** — Select to force the port decommission.

The Management application still contacts all registered CIMOM servers within the fabric affected by the action, but forces the port decommission regardless of the CIMOM server response.

NOTE

If the CIMOM server is not reachable or the credentials fail, port decommission does not occur.

- Click **OK** on the **Port Commission Confirmation** dialog box.

While decommissioning is in progress, a down arrow icon displays next to the port icon in the Product List. You can view the port commissioning results in the deployment reports (refer to [“Port commissioning deployment report”](#) on page 413).

When the decommission is complete, an application event displays in the Master Log detailing success or failure.

Decommissioning all ports on a blade

NOTE

(Virtual Fabrics only) All ports on the blade must be managed by the Management application.

NOTE

Fabric tracking must be enabled (refer to [“Enabling fabric tracking”](#) on page 132) to maintain the decommissioned port details (such as port type, device port wwn, and so on). Do not accept changes in the Management application client.

- Select a port on the blade for which you want to decommission all ports, then select **Configure > Port Commissioning > Decommission > All Ports on the Blade**.

NOTE

You can only decommission ports from the logical switch, not the physical chassis.

The **Port Commission Confirmation** dialog box displays.

- Choose one of the following options:

- Apply Default Settings** (default) — Select to have the Management application perform one of the following actions:

The Management application contact all registered CIMOM servers within the fabric affected by the action and obtains the status from each CIMOM server. If all CIMOM servers are okay, the Management application sends a CAL Request to decommission the port. If even one CIMOM server is not okay, decommissioning fails.

- **Force** – Select to force the port decommission.

The Management application still contacts all registered CIMOM servers within the fabric affected by the action, but forces the port decommission regardless of the CIMOM server response.

NOTE

If the CIMOM server is not reachable or the credentials fail, port decommission does not occur.

3. Click **OK** on the **Port Commission Confirmation** dialog box.

While decommissioning is in progress, a down arrow icon displays next to the port icon in the Product List. You can view the port commissioning results in the deployment reports (refer to [“Port commissioning deployment report”](#) on page 413).

When the decommission is complete, an application event displays in the Master Log detailing success or failure.

Recommissioning all ports on a switch

Select the switch or logical switch for which you want to recommission all ports, then select **Configure > Port Commissioning > Recommission > All F-Ports on the Switch**.

NOTE

You can only recommission ports from the logical switch, not the physical chassis.

While recommissioning is in progress, an up arrow icon displays next to the port icon in the Product List. You can view the port commissioning results in the deployment reports (refer to [“Port commissioning deployment report”](#) on page 413).

When the recommission is complete, an application event displays in the Master Log detailing success or failure.

Recommissioning all ports on a blade

NOTE

All ports on the blade must be managed by the Management application.

Select a port on the blade for which you want to recommission all ports, then select **Configure > Port Commissioning > Recommission > All Ports on the Switch/Blade**.

NOTE

You can only recommission ports from the logical switch, not the physical chassis.

While recommissioning is in progress, an up arrow icon displays next to the port icon in the Product List. You can view the port commissioning results in the deployment reports (refer to [“Port commissioning deployment report”](#) on page 413).

When the recommission is complete, an application event displays in the Master Log detailing success or failure.

Port commissioning deployment report

The **Configuration Deployment** report contains the following parameters:

- **Configuration Name** — Name of the deployment.
For example, Decommission/Recommission - *switch_name*, Decommission/Recommission - *switch_name* - blade, or Decommission/Recommission - *switch_name* - Ports.
- **Description** — A description of the deployment.
- **Module** — The module name. For example, Port Commission.
- **Sub Module** — The sub module name.
- **Deployment Time** — Time when the deployment occurred. Click to launch the detailed deployment results.
- **Status** — Status of the deployment.
- **Creator** — Name of the user that performed the deployment.

The **Deployment Results** contains the following parameters:

- **Configuration Name** — Name of the deployment.
For example, Decommission/Recommission - *switch_name*, Decommission/Recommission - *switch_name* - blade, or Decommission/Recommission - *switch_name* - Ports.
- **Product** — The product name.
- **Status** — The status of the deployment. For example, Allowed or Failed.
- **Reason** — The port level status of the decommission or recommission.

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 discovery. 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 (ADO 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 [“Discovering fabrics”](#) on page 39.

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 (ADO). 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).”
- 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 38](#) details feature support for AD-enabled fabrics in the Management application.

TABLE 38 Feature support for AD-enabled fabrics

Feature	AD context				User interface impact
	ADO	AD255	Not supported	All AD	
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.

TABLE 38 Feature support for AD-enabled fabrics (Continued)

Feature	AD context				User interface impact
	ADO	AD255	Not supported	All AD	
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.
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.

12 Port Auto Disable

TABLE 38 Feature support for AD-enabled fabrics (Continued)

Feature	AD context				User interface impact
	ADO	AD255	Not supported	All AD	
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 6.3 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 6.3 or later.

To view the PAD status, complete the following steps.

1. Select **Monitor > Port Auto Disable**.

The **Port Auto Disable** dialog box displays.

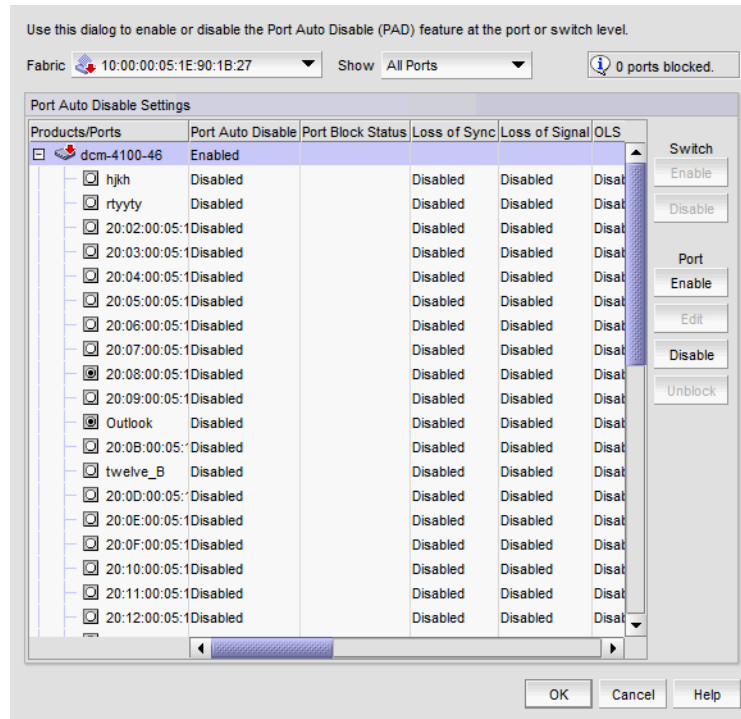


FIGURE 147 Port Auto Disable dialog box

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 6.3 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 418.
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 6.3 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 418.
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 6.3 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 6.3 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 6.3 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

In this chapter

- Host port mapping overview 423
- Creating a new Host 424
- Renaming an HBA Host 425
- Deleting an HBA Host 425
- Viewing Host properties 425
- Associating an HBA with a Host 426
- Importing HBA-to-Host mapping 426
- Removing an HBA from a Host 428
- Exporting Host port mapping 428

Host port mapping overview

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 258. 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).

NOTE

CNA HBAs do not support Host Port Mapping.

The Management application also enables you to discover Hosts directly using Host discovery (for step-by-step instructions, refer to “[Host discovery](#)” on page 58). 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, then you try to discover a Host with the same HBAs using Host discovery, the HBA's 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.

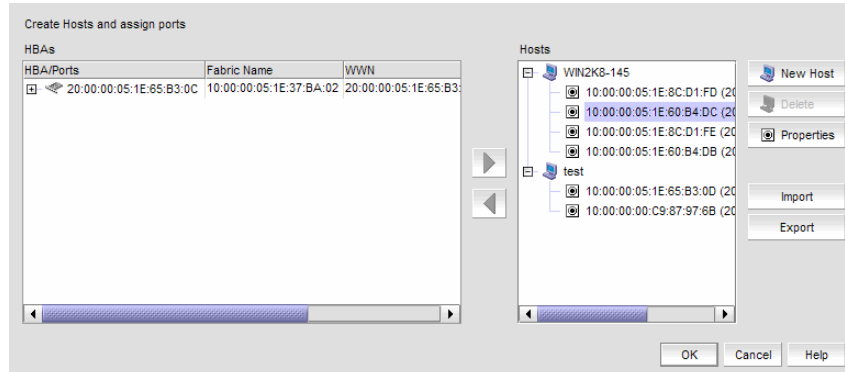


FIGURE 148 Host Port Mapping dialog box

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** table 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** table in alphabetical order. To assign HBAs to this Host, refer to [“Associating an HBA with a Host”](#) on page 426.
 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 426.
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 39](#).

TABLE 39 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.
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 424 and [“Associating an HBA with a Host”](#) on page 426.

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.

13 Exporting Host port mapping

Storage Port Mapping

In this chapter

- [Storage port mapping overview](#) 431
- [Creating a storage array](#) 432
- [Adding storage ports to a storage array](#) 432
- [Unassigning a storage port from a storage array](#) 433
- [Reassigning mapped storage ports](#) 433
- [Editing storage array properties](#) 434
- [Deleting a storage array](#) 434
- [Viewing storage port properties](#) 434
- [Viewing storage array properties](#) 435
- [Importing storage port mapping](#) 435
- [Exporting storage port mapping](#) 437

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 Number** — 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 432.

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.

14 Editing storage array properties

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 436).

TABLE 40 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.

14 Exporting storage port mapping

Host Management

In this chapter

• Host management	439
• Brocade adapters	440
• HCM software	442
• Host adapter discovery	444
• VM Manager	444
• HCM and Management application support on ESXi systems	445
• Connectivity map	447
• View management	447
• Host port mapping	447
• Adapter software	448
• Bulk port configuration	454
• Adapter port WWN virtualization	458
• Role-based access control	463
• Host performance management	464
• Host security authentication	465
• supportSave on adapters	467
• Host fault management	467
• Backup support	469

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*.

Brocade adapters

The following sections describe the three Brocade adapter types:

- “Host Bus Adapters”
- “Converged Network Adapters”
- “Fabric Adapters”

Host Bus Adapters

Brocade offers five models of Fibre Channel Host Bus Adapters (HBAs). These models provide reliable, high-performance host connectivity for mission-critical SAN environments. The Brocade HBAs are listed in [Table 41](#).

TABLE 41 Brocade Fibre Channel HBA models

Model number	Description	Number of ports
825	Dual-port stand-up HBA with a per-port maximum of 8 Gbps using an 8 Gbps SFP. ¹	2
815	Single-port stand-up HBA with a maximum of 8 Gbps using an 8 Gbps SFP. ¹	1
804 ²	Dual-port mezzanine HBA with a per-port maximum of 8 Gbps. This HBA installs in server blades that install in supported blade system enclosures.	2
425	Dual-port stand-up HBA with a per-port maximum of 4 Gbps using a 4 Gbps SFP. ³	2
415	Single-port stand-up HBA with a maximum of 4 Gbps using a 4 Gbps SFP. ³	1

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.

Using Brocade 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 or M-Enterprise operating systems.

Converged Network Adapters

[Table 42](#) describes available Brocade Converged Network Adapters (CNAs) for PCIe x 8 host bus interfaces, hereafter referred to as Brocade CNAs. These adapters provide reliable, high-performance host connectivity for mission-critical SAN environments.

TABLE 42 Brocade Fibre Channel CNA models

Model number	Port speed	Number of ports	Adapter type
1741M-k ¹	10 Gbps maximum	2	Expansion
1020	10 Gbps maximum	2	Stand-up
1010	10 Gbps maximum	1	Stand-up
1007 ²	10 Gbps maximum	2	Expansion

¹The 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.

²The 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 *Brocade 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 43](#) describes the available Brocade 1860 Fabric Adapter model. 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 mode using the Brocade Command Utility (BCU).

TABLE 43 Brocade Fabric Adapter models

Model number	Port speed	Number of ports
1860-1 860-2	16 Gbps FC HBA and 10 Gbps CNA or NIC	1 or 2
1867	16 Gbps FC mezzanine card	2

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 Brocade HBA CIM provider to be installed on the ESXi host.

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 [Chapter 4, “Discovery”](#).

VM Manager

A vCenter server can be discovered by adding a VM Manager to the Management application. Refer to [Chapter 4, “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 149](#).

FIGURE 149 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.

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.

- **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 150](#), displays.

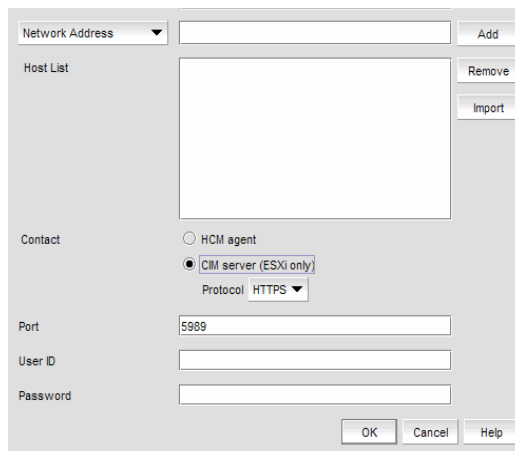


FIGURE 150 Add Host Adapters dialog box

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 259
- [“Host group icons”](#) on page 260
- [“SAN port icons”](#) on page 260

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 [Chapter 8, “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 259. 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 58). 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 [Chapter 13, “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 151](#), displays.

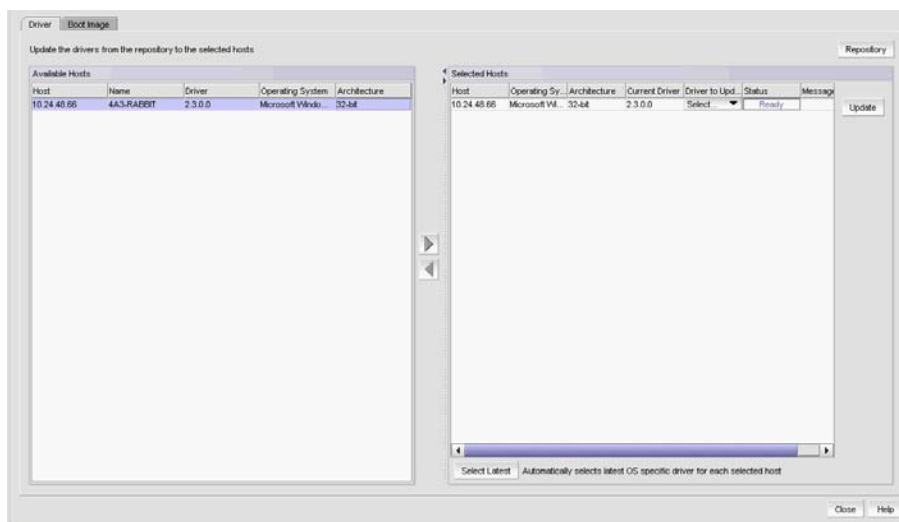


FIGURE 151 Adapter Software dialog box, Driver tab

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 449.

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 152](#), displays.




FIGURE 152 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 153](#), displays.

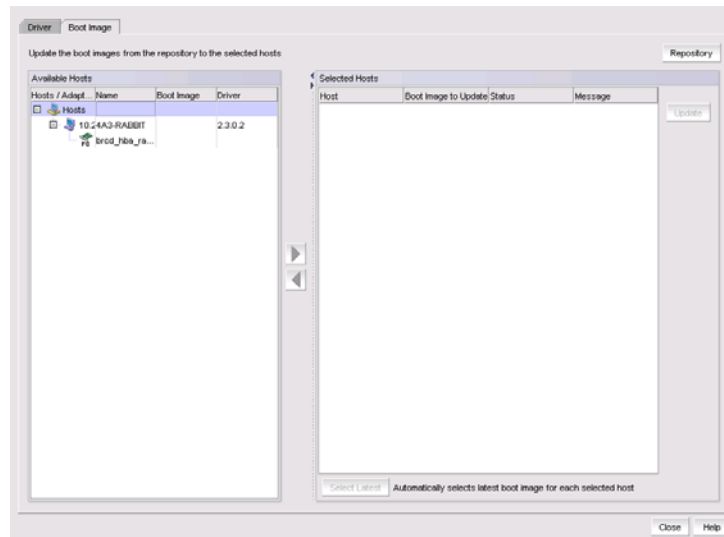


FIGURE 153 Boot Image Management dialog box

3. From the **Boot Image Management** dialog box, click the **Repository** button.

The **Boot Image Repository** dialog box, shown in [Figure 154](#), displays.

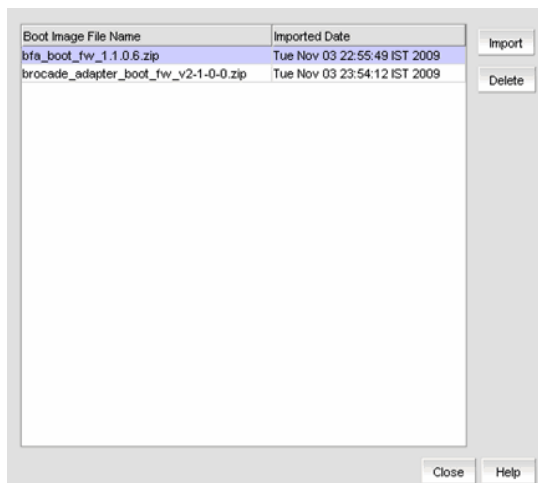


FIGURE 154 Boot Image Repository dialog box

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 469 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.

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 155](#), displays.

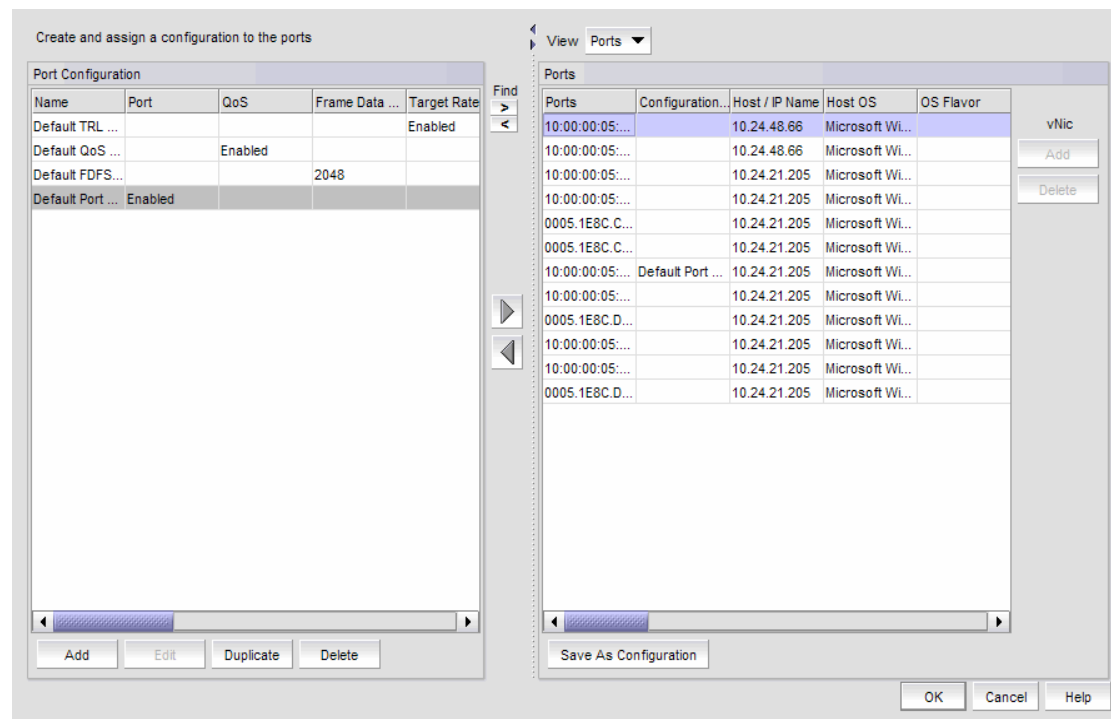


FIGURE 155 Configure Host Adapter Ports dialog box

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 156](#), displays.

FIGURE 156 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 topology 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 455.
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 157](#), 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.

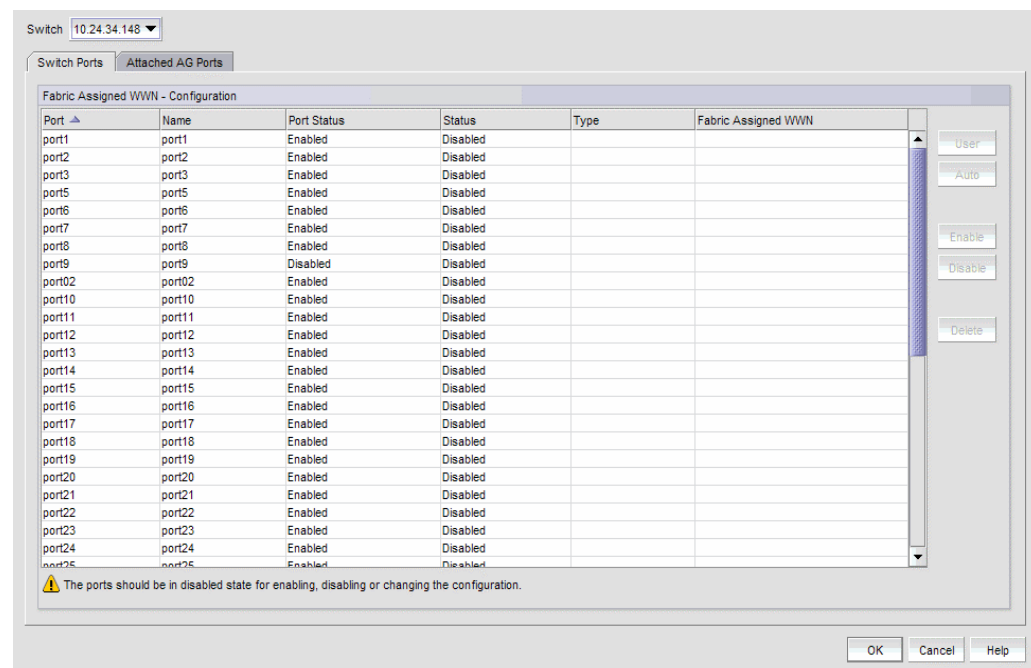


FIGURE 157 Configure Fabric Assigned WWNs dialog box

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 158](#), 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.

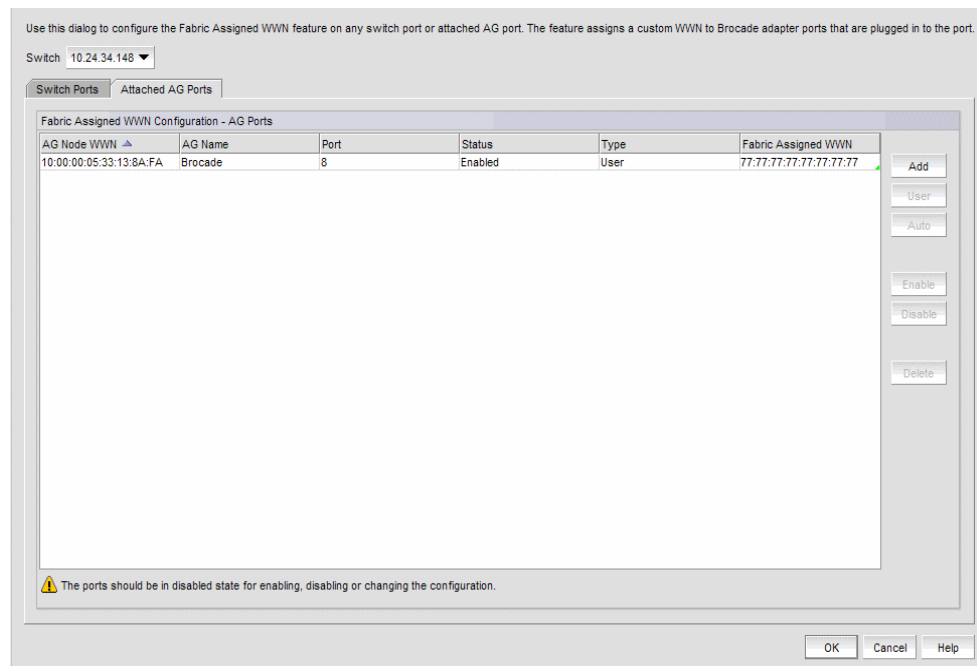


FIGURE 158 Configure Fabric Assigned WWNs dialog box--Attached AG Ports tab

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 [Chapter 6, “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 44](#) lists the counters that are supported for the FC ports and for the HBA and CNA ports.

TABLE 44 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 44 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 942.

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 159](#), displays.




FIGURE 159 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 *Brocade Adapters Administrator's Guide* for information about supportSave on ESXi hosts.

NOTE

You cannot schedule host supportSave information.

Instructions for scheduling and capturing technical support files are detailed in [Chapter 33, "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 1067 for more information.

To configure an e-mail event, use the instructions in [“Configuring e-mail notification”](#) on page 1064.

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 1085.

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

In this chapter

• FCoE overview	471
• Enhanced Ethernet features	472
• FCoE protocols supported	473
• FCoE licensing	474
• Saving running configurations	474
• DCB configuration management	475
• Switch policies	476
• DCB configuration	477
• QoS configuration	490
• FCoE provisioning	496
• VLAN classifier configuration	498
• LLDP-DCBX configuration	502
• 802.1x authentication	506
• Switch, port, and LAG deployment	508
• DCB performance	513
• FCoE login groups	514
• Virtual FCoE port configuration	519

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 Fabric OS DCB switch.

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 6.3x 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 160](#).




FIGURE 160 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 [“Restoring a switch configuration for a selected device”](#) in the “Device Configuration” chapter.

DCB configuration management

Depending on the platform, the DCB switch has one of the configurations shown in [Table 45](#).

TABLE 45 DCB configurations

Device type	Configuration possibilities
IBM blade server	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 16 10-Gbps internal ports • 8 10-Gbps external ports • 4 8-Gbps FC ports
Fabric OS DCB switch	<ul style="list-style-type: none"> • 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 490 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 502 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 506 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 161](#), displays the values only for the external ports of embedded platforms but not for the internal ports.

Products / Ports	Name	Fabric	MAC Address	Interface Mode	Primary IP / Netmask	Status	Date	Protocol Down Reason	Speed	VLAN ID	L2 Mode	LAG ID	LAG Mode	LAG Type
1	INT 01	0005.1ec7.1465	L2	Enabled	Up	Enabled	Up		10	1,4095	Converged			
2	INT 02	0005.1ec7.1466	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
3	INT 03	0005.1ec7.1467	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
4	INT 04	0005.1ec7.1468	L2	Enabled	Up	Enabled	Up		10	1,4095	Converged			
5	INT 05	0005.1ec7.1469	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
6	INT 06	0005.1ec7.146a	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
7	INT 07	0005.1ec7.146b	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
8	INT 08	0005.1ec7.146c	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
9	INT 09	0005.1ec7.146d	L2	Enabled	Up	Enabled	Up		10	1,4095	Converged			
10	INT 010	0005.1ec7.146e	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
11	INT 011	0005.1ec7.146f	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
12	INT 012	0005.1ec7.1460	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
13	INT 013	0005.1ec7.1461	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
14	INT 014	0005.1ec7.1462	L2	Enabled	Down	Enabled	Down		10	1,4095	Converged			
15	EXT 015	0005.1ec7.1463	None	Disabled	Down	admin down	admin down		10	NA				
16	EXT 016	0005.1ec7.1464	None	Disabled	Down	admin down	admin down		10	NA				
17	EXT 017	0005.1ec7.1465	L2	Disabled	Down	admin down	admin down		10	Invalid				
18	EXT 018	0005.1ec7.1466	None	Disabled	Down	admin down	admin down		10	NA				
19	EXT 019	0005.1ec7.1467	None	Disabled	Down	admin down	admin down		10	NA				
20	EXT 020	0005.1ec7.1468	None	Disabled	Down	admin down	admin down		10	NA				
21	EXT 021	0005.1ec7.1469	None	Disabled	Down	admin down	admin down		10	NA				
22	EXT 022	0005.1ec7.146a	None	Disabled	Down	admin down	admin down		10	NA				

FIGURE 161 DCB 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 162](#).




FIGURE 162 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 163](#).

Create LLDP Profiles that will be available for assignment to ports on this switch.

Enable LLDP-DCBX

LLDP Profile

Name:

Description:

Mode: Both Transmit and Receive

Hello(sec): 30 (4-180)

Multiplier: 4 (1-10)

Advertise

Port Description

System Capabilities

System Name

System Description

Management IP Address

Dot1

Dot3

DCBX

FCoE Application

FCoE Logical Link

LLDP Profiles

Name	Description
Global_Configuration	
Lp	34
dafeadf	erw3qrsa
ddf	
sadsadsad	3123
stgsdf	

Details of Selected Profile

Mode	Both Transmit and Rece...
Hello	30
Multiplier	4
Advertise	Port Description
	System Capabilities
	System Name
	Management IP Address
	Dot3
	DCBX
	FCoE Application
	FCoE Logical Link

OK Cancel Help

FIGURE 163 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 166](#).
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 477 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 164](#).

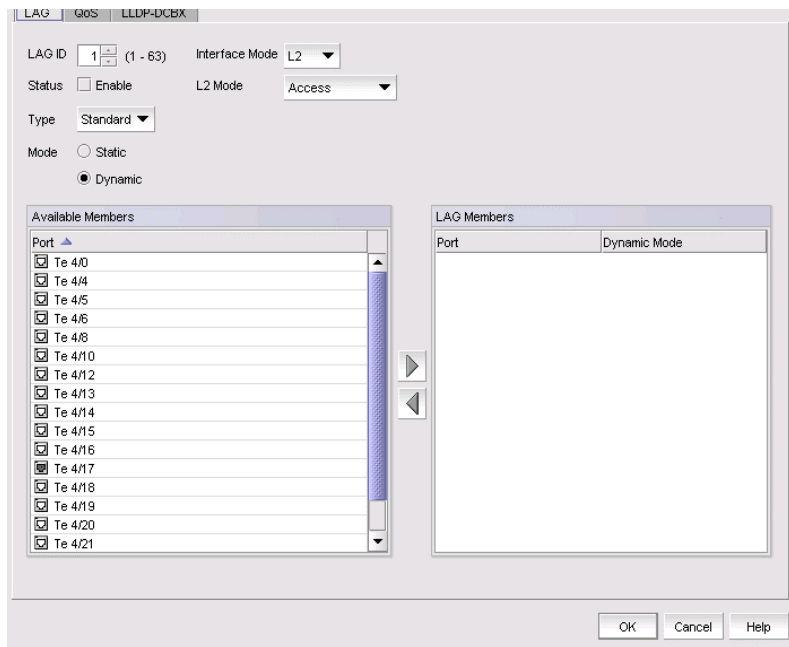


FIGURE 164 Add LAG dialog box

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 165](#)).

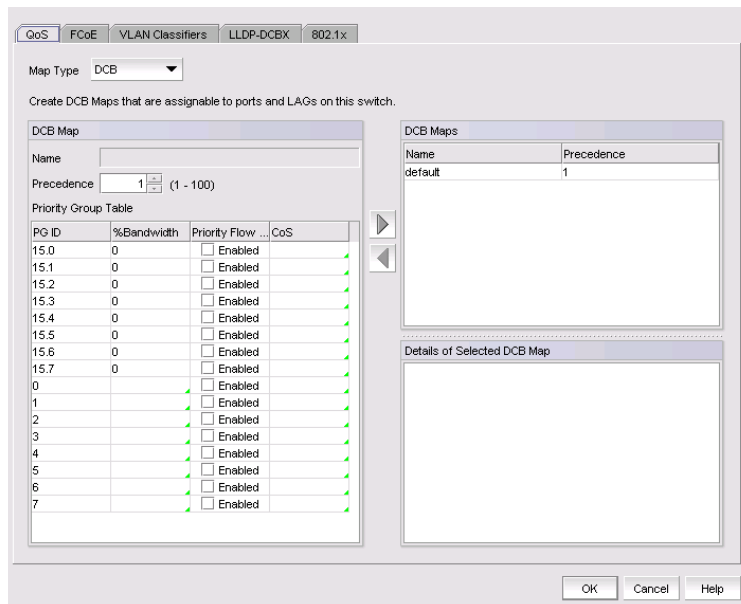


FIGURE 165 Edit Switch dialog box

4. Configure the policies for the **Edit Switch** dialog box tabs, which are described in the following sections:
 - [“QoS configuration”](#) on page 490
 - [“FCoE provisioning”](#) on page 496
 - [“VLAN classifier configuration”](#) on page 498
 - [“LLDP-DCBX configuration”](#) on page 502
 - [“802.1x authentication”](#) on page 506
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 166](#).

FIGURE 166 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.

- 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.

- Click **OK** when you have finished modifying the DCB port parameters.
The **Deploy to Ports** dialog box displays.
- Click **OK** after changing the attributes of the current deployment.
The **Deployment Status** dialog box launches.
- Click **Start** on the **Deployment Status** dialog box to save the changes to the selected port or ports.
- 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).

- Select **Configure > DCB**.

The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

- Select the link aggregation group (LAG) from the **Products/Ports** list.
- Click **Edit**.

The **Edit LAG** dialog box displays, as shown in [Figure 167](#).

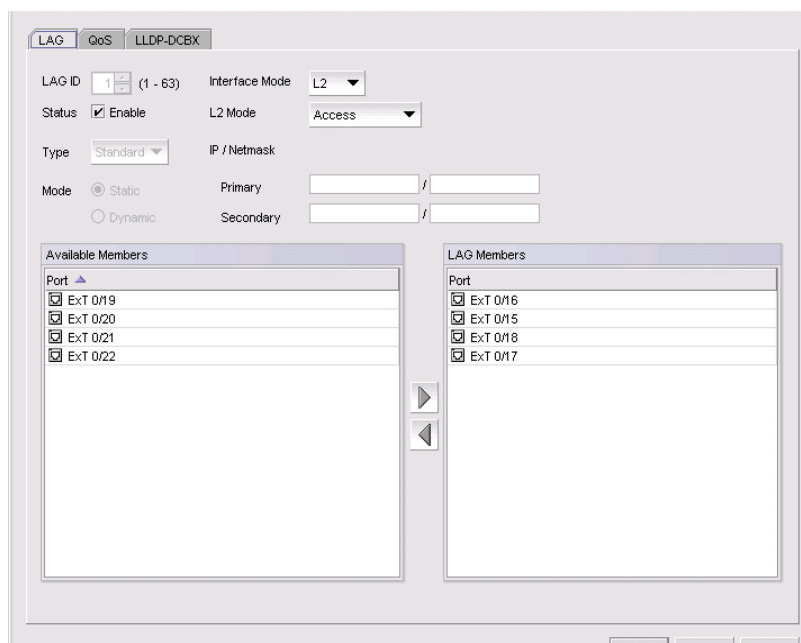


FIGURE 167 Edit LAG dialog box

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.

- 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.

- 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.

- Select **Configure > DCB**.

The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

- 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.

- Click **Enable**.

The **Confirmation and Status** dialog box launches with the selected ports or LAGs.

- 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.

- 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.

- Select **Configure > DCB**.

The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.

- Select one or more LAGs (that can span multiple switches) that you want to delete from the **Products/Ports** list.

- Click **Delete**.

The **Confirmation and Status** dialog box launches with the selected LAGs.

- 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.

- 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 46](#) shows examples of how priority grouping might be allocated in a 15-priority group scenario.

TABLE 46 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	on
No bandwidth % configuration allowed		

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 46](#) 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 168](#).

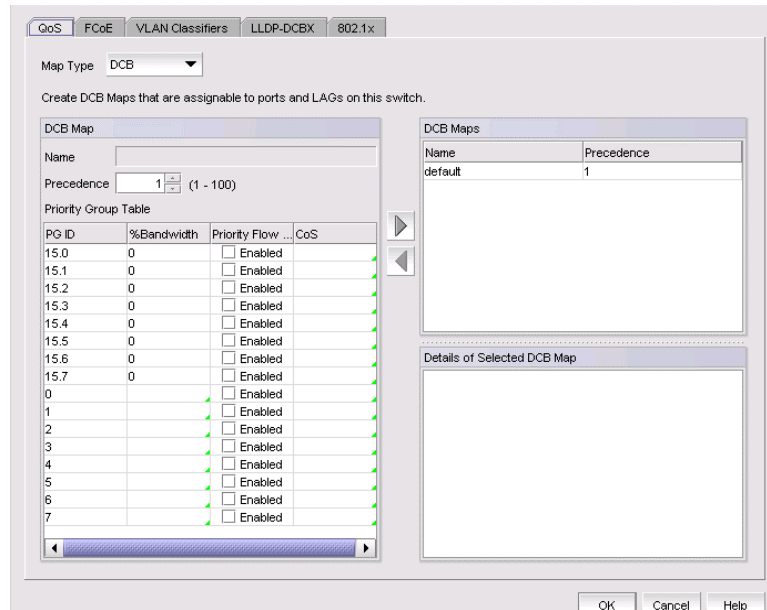


FIGURE 168 QoS, Create DCB Map dialog box

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 46](#) 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.

- 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.

- Select **Configure > DCB**.
The **DCB Configuration** dialog box displays, showing the status of all DCB-related hardware and functions.
- Select a port or LAG, and click **Edit**.
- Click the **QoS** tab on the **Edit Port** or **Edit LAG** dialog box.
The **QoS** dialog box displays.
- Click the **Assign a map** check box.
- Select **Traffic Class** in the **Map Type** list.
- Select a Traffic Class map in the **Traffic Class Map** list.
- 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 508 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 165](#).

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 497 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 169](#). 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.

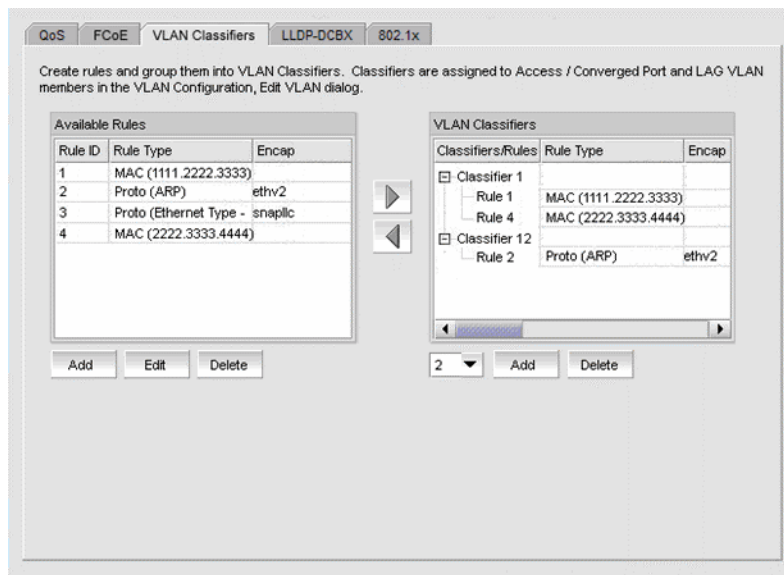


FIGURE 169 Edit Switch dialog box, VLAN Classifiers tab

4. Click the **Add** button under the **Available Rules** list.

The **Add Rules** dialog box displays, as shown in [Figure 170](#).

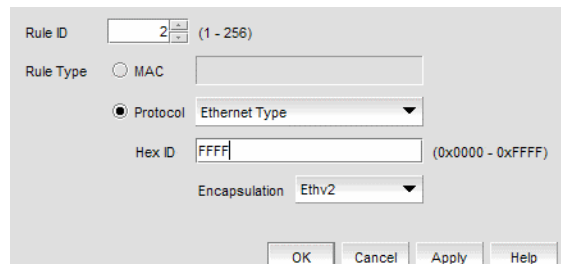


FIGURE 170 Add Rules dialog box

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 Classifier** 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 171](#).
3. Click the **LLDP-DCBX** tab.

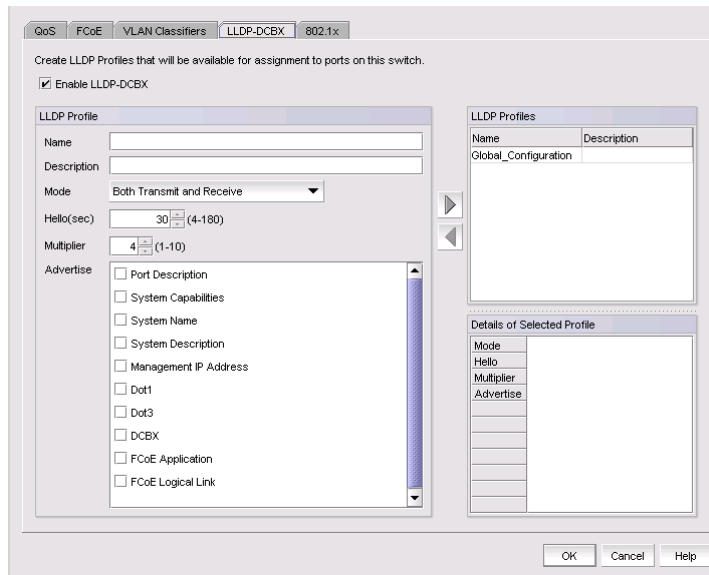


FIGURE 171 Edit Switch dialog box - LLDP-DCBX tab

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 503. 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.
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 508 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 507.
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 172](#).

4. Click the **Enable 802.1x** check box to enable 802.1x authentication.

The 802.1x parameters are enabled for editing.

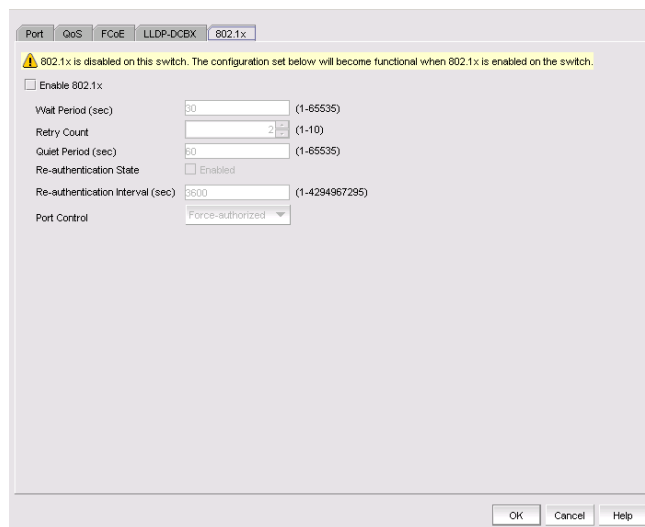


FIGURE 172 802.1x dialog box

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 508 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 173](#), [Figure 174](#), and [Figure 175](#).

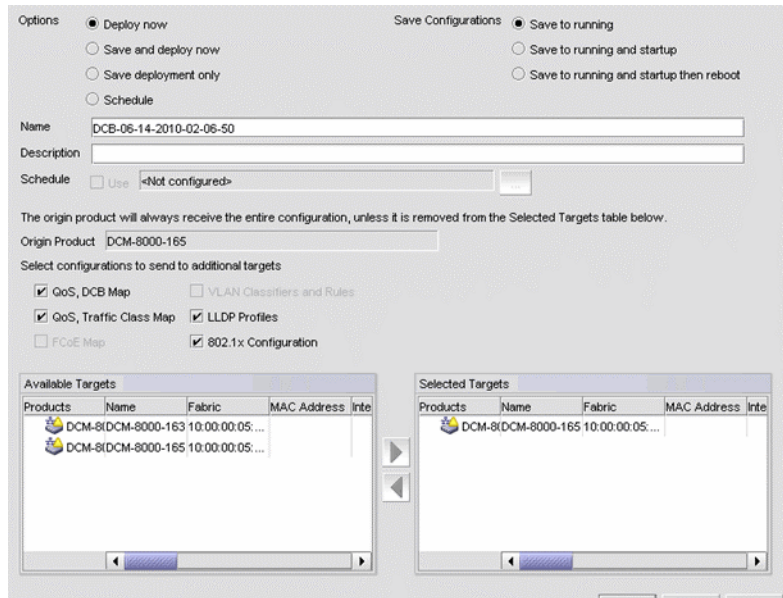


FIGURE 173 Deploy to Products dialog box

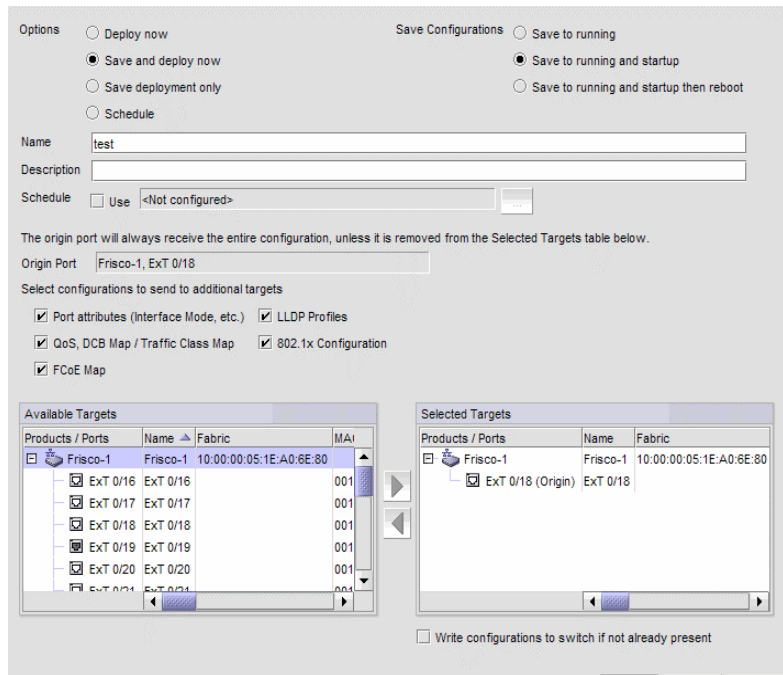


FIGURE 174 Deploy to Ports dialog box

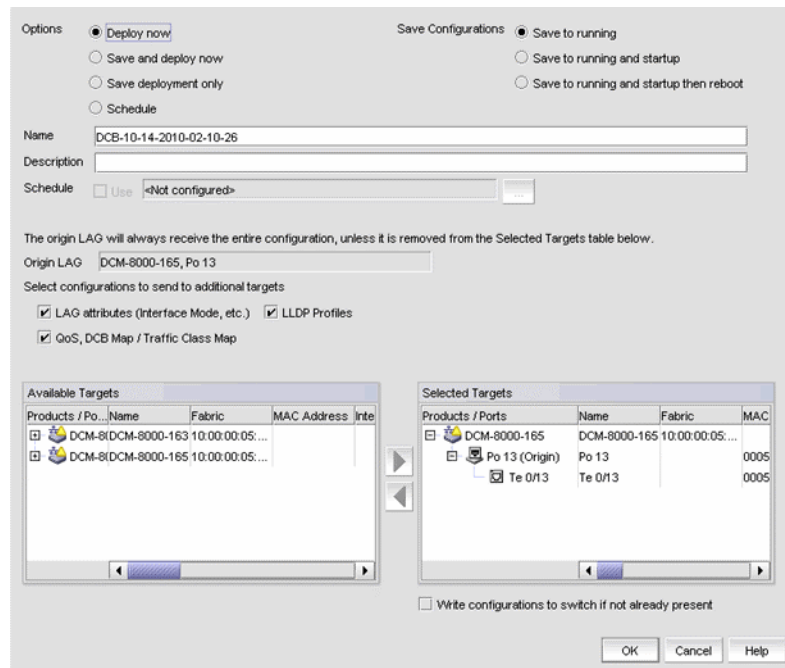


FIGURE 175 Deploy to LAGs dialog box

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 47 lists the restrictions that exist when deploying source switches to target switches.

TABLE 47 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.	<p>Allows Fabric OS DCB switch and FCOE10-24 DCB blade with Fabric OS version 6.4.2 or earlier.</p> <p>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.</p>	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	<p>Allows FCOE10-24 DCB blade with Fabric OS 6.4.1_fcoe or Fabric OS 7.0.0.</p> <p>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.</p> <p>Excludes Fabric OS DCB switch and FCOE10-24 DCB blade with Fabric OS 6.4.2 or earlier.</p>	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.	<p>Allows Fabric OS DCB switch and FCOE10-24 DCB blade with Fabric OS 7.0.0.</p> <p>Excludes all others.</p>	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	<p>Allows Fabric OS Converged 10 Gbe switch module for IBM BladeCenter with Fabric OS 6.3.1_cee, Fabric OS 6.3.1_dcb.</p> <p>Allows Dell M8428-k switch with Fabric OS 6.3.1_dell, Fabric OS 6.3.1_dcb.</p>	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	<p>Allows Fabric OS Converged 10 Gbe switch module for IBM BladeCenter with Fabric OS 6.3.1_cee, Fabric OS 6.3.1_dcb.</p> <p>Allows Dell M8428-k switch with Fabric OS 6.3.1_dell, Fabric OS 6.3.1_dcb.</p>	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 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 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 Fabric OS DCB 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, as shown in [Figure 176](#).

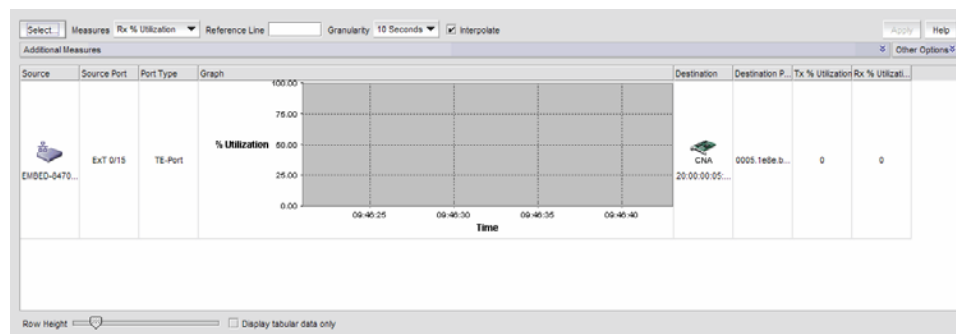


FIGURE 176 Real Time Performance Graphs dialog box - SAN tab

For complete information about Real Time Performance Graphs, refer to [“SAN real-time performance data”](#) on page 942.

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 177](#).

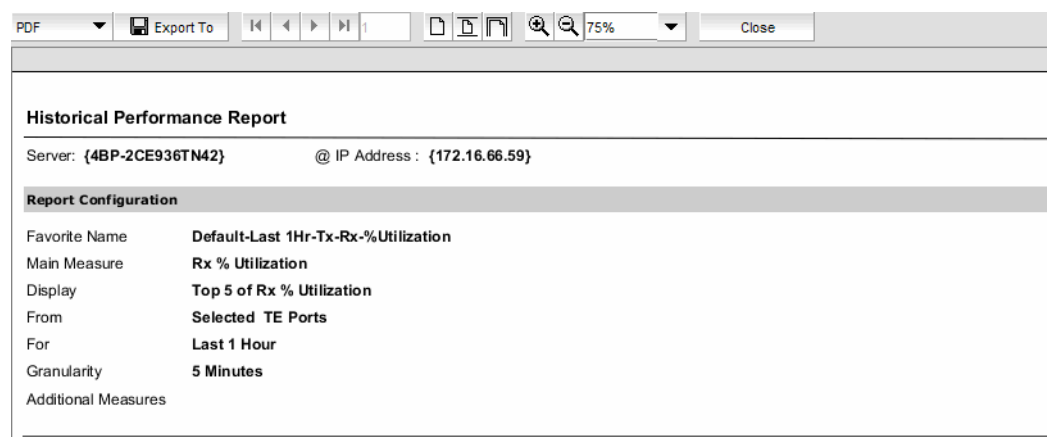


FIGURE 177 Historical Performance Report dialog box

For complete information about Historical Performance Graphs, refer to [“SAN historical performance data”](#) on page 946.

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.

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 178](#).

- 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 178 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 515.
- 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 517.
- Click **Delete** to remove the login group from the list. See [“Deleting one or more FCoE login groups”](#) on page 518.

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 179](#).

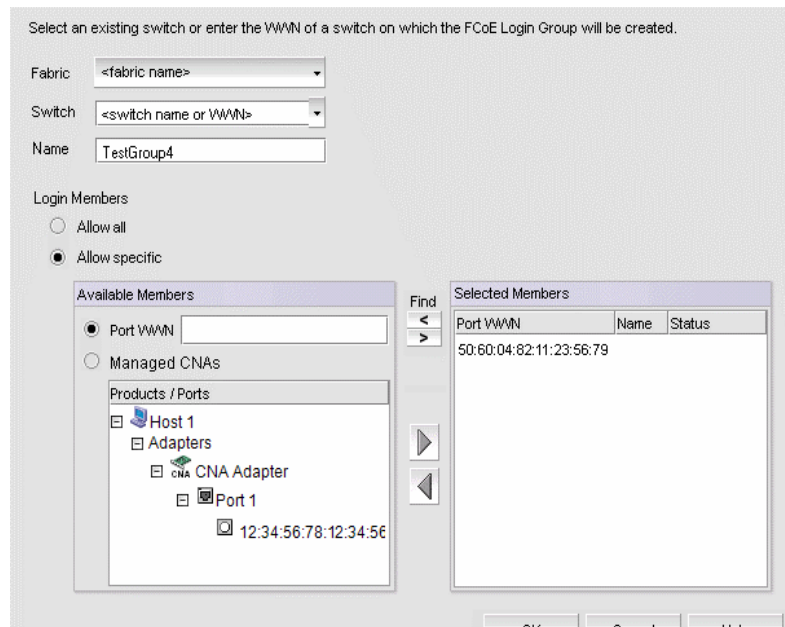


FIGURE 179 Add Login Group dialog box

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.

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 166](#).

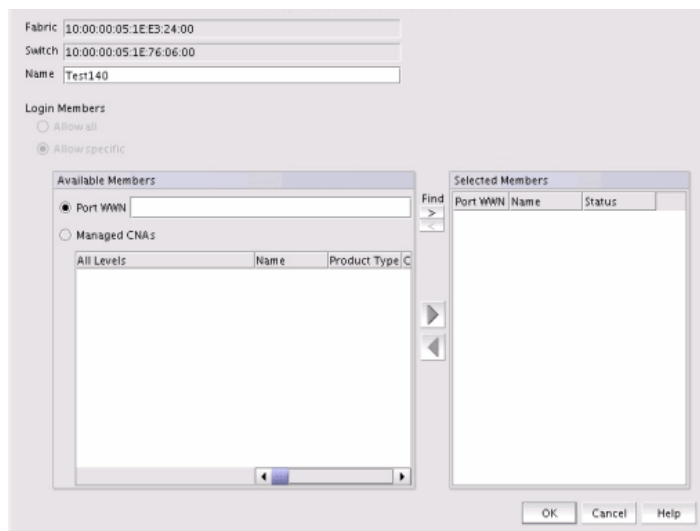


FIGURE 180 Edit Login Group dialog box

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 181](#).

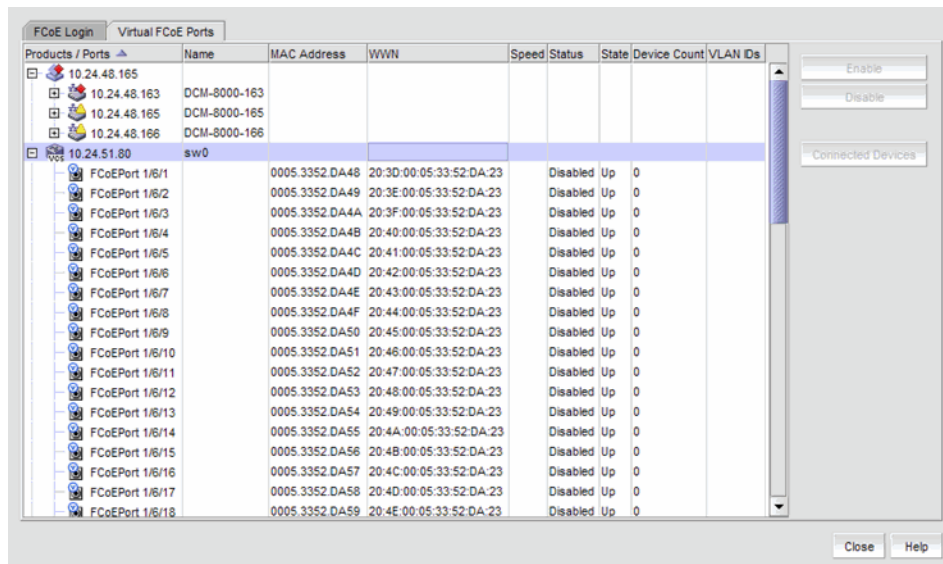


FIGURE 181 Virtual FCoE Ports dialog box

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.

16 Virtual FCoE port configuration

Security Management

In this chapter

- [Layer 2 access control list management](#) 523
- [Security configuration deployment](#) 532

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.

NOTE

Layer 2 ACLs can filter traffic for both Fabric OS and IronWare FCoE devices.

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 182 *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 534

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 524.

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 524.

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 534

Copying a standard Layer 2 ACL configuration (Fabric OS)

To copy 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 duplicate 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 524.
 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 524.
 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 534

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.

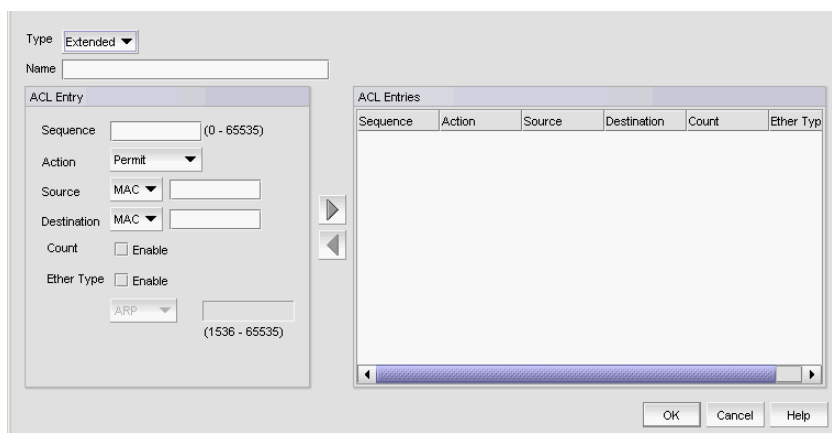


FIGURE 183 *Device_Name* - Layer 2 ACL Configuration (Extended) dialog box

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 534

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 526.
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 526.
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 534

Copying an extended Layer 2 ACL configuration (Fabric OS)

To copy an 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 ACL you want to copy 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 526.
The updated ACL entry displays in the **ACL Entries** list. To edit additional ACL entries, repeat [step 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 526.

The new ACL entry displays in the **ACL Entries** list. To add additional ACL entries, repeat [step 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.
- 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](#).
- 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 534

Assigning a Layer 2 ACL configuration to an interface (Fabric OS)

To assign Layer 2 ACL configuration to a interface, complete the following steps.

- Select **Configure > Security > Layer 2 ACL > Port**.
The **Port Selection - Layer 2 ACL** dialog box displays.
- 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*.
- Click **OK**.

The *Device_Name - Port_Number/LAG LAG_Number- Layer 2 ACL Configuration* dialog box displays.

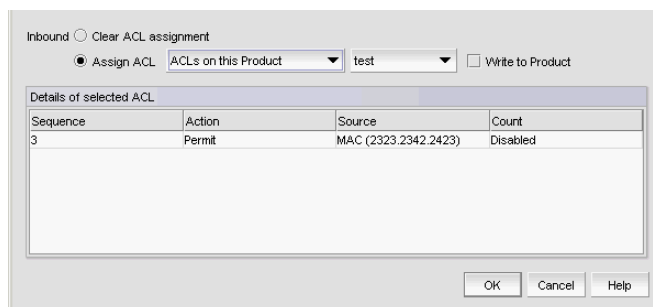


FIGURE 184 *Device_Name - Port_Number- Layer 2 ACL Configuration* dialog box

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 532.

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 532.

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 534

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 534

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 534.

Security configuration deployment

Figure 185 shows the standard interface used to deploy security configurations.

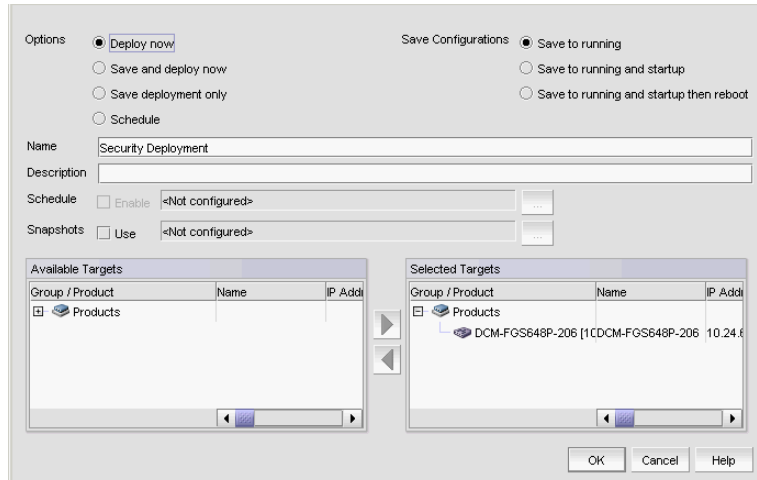


FIGURE 185 Deploy to Product/Ports dialog box

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 915.

Deploying a security configuration on demand

To deploy a security configuration immediately, complete the following steps.




FIGURE 186 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.

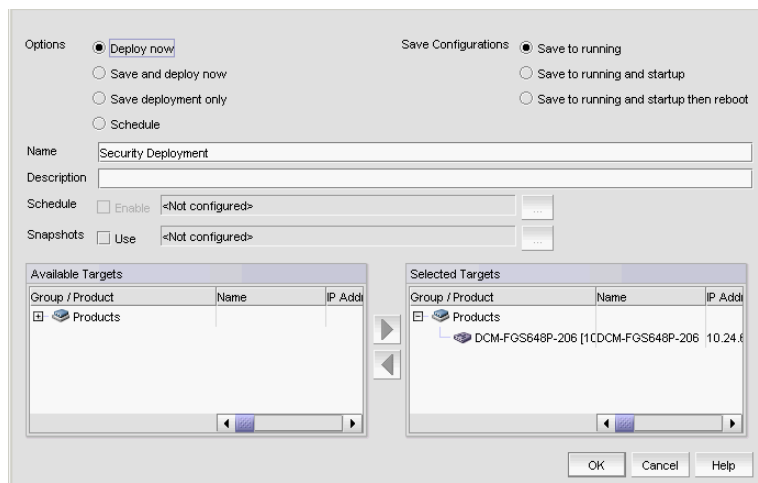


FIGURE 187 Deploy to Product/Ports dialog box

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 188 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 536.
 - To configure hourly deployment, refer to [“Configuring an hourly deployment schedule”](#) on page 536.
 - To configure daily deployment, refer to [“Configuring a daily deployment schedule”](#) on page 537.
 - To configure weekly deployment, refer to [“Configuring a weekly deployment schedule”](#) on page 537.
 - To configure monthly deployment, refer to [“Configuring a monthly deployment schedule”](#) on page 537.
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 535.

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 535.

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 535.

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 535.

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 535.

17 Security configuration deployment

FC-FC Routing Service Management

In this chapter

- [Devices that support Fibre Channel routing](#) 539
- [Fibre Channel routing overview](#) 540
- [Guidelines for setting up Fibre Channel routing](#) 541
- [Connecting edge fabrics to a backbone fabric](#) 542
- [Configuring routing domain IDs](#) 544

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
- 24-port, 16 Gbps Edge Switch
- 48-port, 16 Gbps FC Switch
- 4 Gbps Router, Extension Switch
- 8 Gbps Extension Switch
- Any of the following blades on a Director chassis:
 - 4 Gbps Router, Extension Blade
 - FC 8 GB 16-port Blade
 - FC 8 GB 32-port Blade
 - FC 8 GB 48-port Blade - The shared ports area (ports 16-47) cannot be used as EX_Ports.
 - 8 Gbps Extension Blade

- 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 4-slot or 16 Gbps 4-slot Backbone Chassis only)
 - FC 8 GB 48-port Blade - The shared ports area (ports 16-47) cannot be used as EX_Ports.
 - FC 8 GB 48-port Enhanced Blade (16 Gbps 4-slot or 16 Gbps 4-slot Backbone Chassis only)
 - FC 8 GB 64-port Blade
 - 8 Gbps Extension Blade
 - 16 Gbps 32-port Blade
 - 16 Gbps 48-port 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 <i>backbone fabric</i> 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 189 on page 541 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 189 A 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 (IMO) 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 830 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 190](#)). 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 190 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 `for_fd_domainID`. For example, if the domain ID is 3, the name of the front domain is `for_fd_3`.

9. Configure LSAN zones in each fabric that will share devices.

For specific instructions, refer to “Configuring LSAN zoning” on page 782.

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 191).

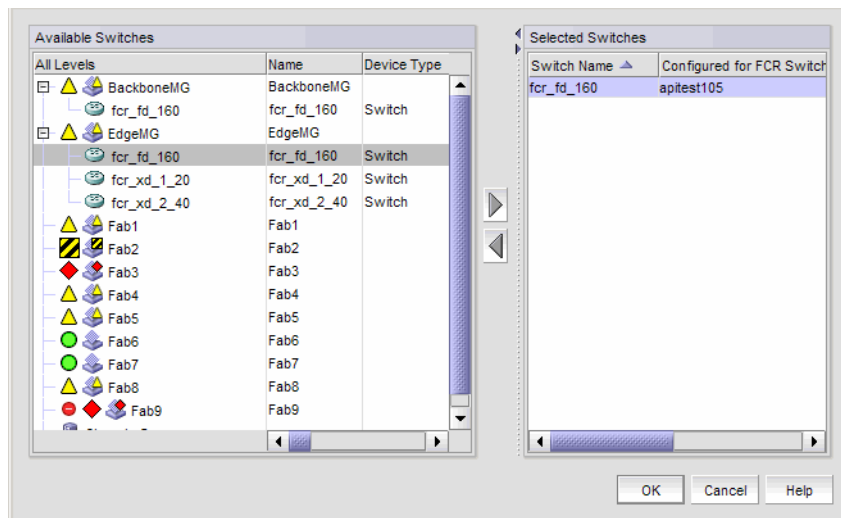


FIGURE 191 Configure Routing Domain IDs dialog box

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

In this chapter

- [Virtual Fabrics overview](#) 545
- [Virtual Fabrics requirements](#) 547
- [Configuring Virtual Fabrics](#) 550

Virtual Fabrics overview


NOTE

Virtual Fabrics requires that you have at least one Virtual Fabrics-enabled physical chassis running Fabric OS 6.2.0 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 192](#), Switch_2 is a logical switch.




FIGURE 192 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 48 lists definitions of Virtual Fabrics terms.

TABLE 48 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.
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: <ul style="list-style-type: none"> • 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 6.2.0 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 6.2.0 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 6.2.0 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 6.2.0 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 551.

[Table 49](#) lists the Virtual Fabrics-capable physical chassis and the number of logical switches allowed for each of those physical chassis.

TABLE 49 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
8 Gbps Extension Switch	4
8-slot Backbone Chassis	8
4-slot Backbone Chassis	8

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 50](#).

TABLE 50 Blade and port types supported on logical switches for backbone chassis

Logical switch type	Ports
Default logical switch	<ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> - FC ports: E_Ports, F_Ports, and VE_Ports - GE ports: VE_Ports • 8-slot and 4-slot Backbone Chassis – ICL ports
Logical switch	<ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> - FC ports: E_Ports, F_Ports, and VE_Ports - GE ports: VE_Ports • 8-slot and 4-slot Backbone Chassis – ICL ports
Base switch	<ul style="list-style-type: none"> • 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 EX_Ports • 8 Gbps Extension Blade <ul style="list-style-type: none"> - FC ports: E_Ports, EX_Ports, VE_Ports, and VEX_Ports - GE ports: VE_Ports • 8-slot and 4-slot Backbone Chassis – ICL Ports

NOTE

In the 8-slot Backbone Chassis, ports 48–63 of the FC 8 GB 64-port Blade 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 550:

- 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.
- **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 coarse wave division multiplexing (CWDM) is being used. Long Distance Fabric requires a license.
- 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.**
 - **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.

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 551 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 552 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 552 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 547.

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.

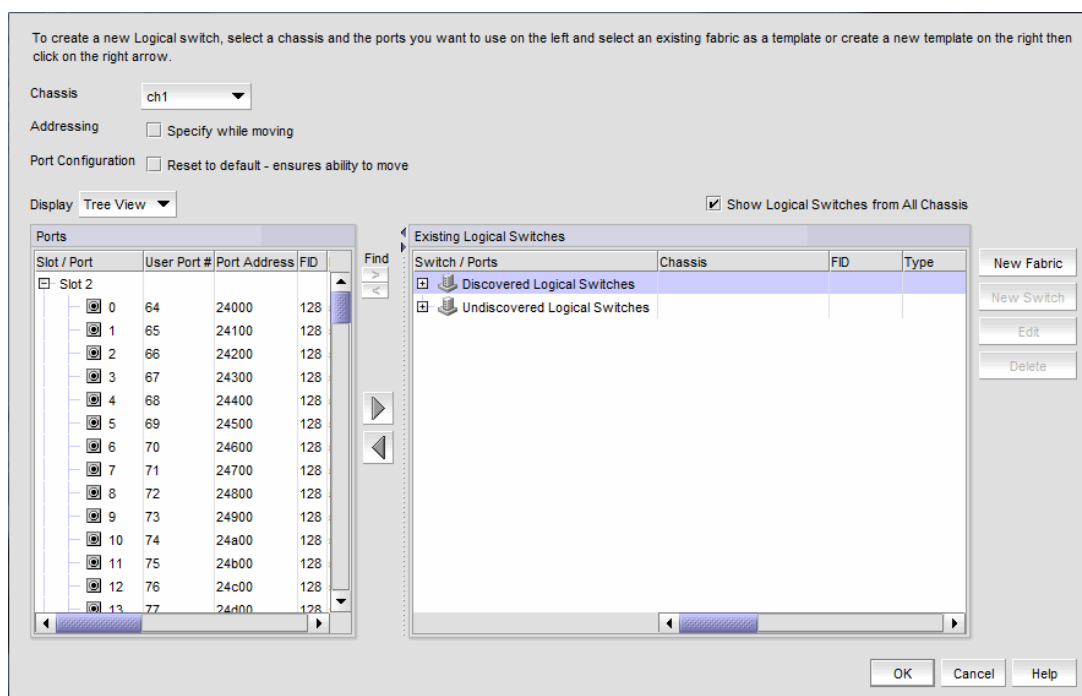


FIGURE 193 Logical Switches dialog box

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 194 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.
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 555.

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 47 for instructions.
- To rediscover the physical chassis, refer to [“Discovering fabrics”](#) on page 39 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.

NOTE

Ports are disabled before moving from one logical switch to another.

9. (Optional) Select the **Unbind Port Addresses while moving them** check box.

- 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.

- When the changes are complete, click **Close**.

Removing ports from a logical switch

- Select **Configure > Virtual Fabric > Logical Switches**.

The **Logical Switches** dialog box displays.

- Select the physical chassis to which the ports belong in the **Chassis** list.
- 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**.

- Click the left arrow button.

A message displays indicating that the ports will be moved to the default logical switch.

- 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.

- (Optional) Perform the following steps to assign the ports to a logical switch other than the default logical switch.

- Select the destination logical switch in the **Existing Logical Switches** list.

- Click the right arrow button.

The ports display in the selected logical switch node in the **Existing Logical Switches** list.

- 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.

NOTE

Ports are disabled before moving from one logical switch to another.

- (Optional) Select the **Unbind Port Addresses while moving them** check box.
- 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.

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 552.

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.

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.

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 47 for instructions.
 - b. Rediscover the physical chassis. Refer to “[Discovering fabrics](#)” on page 39 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.

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

In this chapter

- Encryption Center features 564
- Encryption user privileges 565
- Smart card usage 566
- Network connections 576
- Blade processor links 577
- Encryption node initialization and certificate generation 578
- Key Management Interoperability Protocol 578
- Steps for connecting to a DPM appliance 581
- Steps for connecting to a DPM appliance 581
- Steps for connecting to an LKM/SSKM appliance 587
- Steps for connecting to an ESKM/SKM appliance 592
- Steps for connecting to a TEKA appliance 603
- Steps for connecting to a TKLM appliance 608
- Steps for connecting to a KMIP-compliant SafeNet KeySecure 612
- Steps for connecting to a KMIP-compliant keyAuthority 631
- Encryption preparation 632
- Creating a new encryption group 633
- Adding a switch to an encryption group 670
- Replacing an encryption engine in an encryption group 676
- High availability clusters 677
- Configuring encryption storage targets 680
- Configuring hosts for encryption targets 689
- Adding target disk LUNs for encryption 691
- Adding target tape LUNs for encryption 698
- Moving targets 701
- Configuring encrypted tape storage in a multi-path environment 702
- Tape LUN write early and read ahead 703
- Tape LUN statistics 704
- Encryption engine rebalancing 709
- Master keys 710
- Security settings 719
- Zeroizing an encryption engine 719

- Using the Encryption Targets dialog box 720
- Redirection zones 721
- Disk device decommissioning 722
- Rekeying all disk LUNs manually 725
- Thin provisioned LUNs..... 730
- Viewing time left for auto rekey 731
- Viewing and editing switch encryption properties..... 732
- Viewing and editing encryption group properties 737
- Encryption-related acronyms in log messages 753

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 195](#).) 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**.

Encryption Devices (Group View)	Fabric	Switch / Engine Status	Switch Group Membership Stat...	Target Status	HA Cluster
76ud6					
DCX-4s	FX-824 blade	Healthy	Group Leader		
TEMS					
mace241	10:00:00:05:1E:53:6B:69	Healthy	Group Leader		
Engine		Online		10 OK	
tkim					
Mace26	10:00:00:05:1E:53:6B:69	⚠ Marginal	Group Leader		
Engine		Online		⚠ 1 Offline	
<NO GROUP DEFINED>					
DCX	FX-824 blade	Healthy	ⓘ Not a member		
mace25	10:00:00:05:1E:53:6B:69	⚠ Marginal	ⓘ Not a member		
Engine		⚠ Awaiting initialization		None configured	

FIGURE 195 Encryption Center dialog box

Beginning with Fabric OS 6.4, 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 565 describes the Role-based Access Control privileges that are specific to encryption.
- [“Smart card usage”](#) on page 566 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 576 describes the network connections that must be in place to enable encryption.

- “Blade processor links” on page 577 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 578 lists the security parameters and certificates that are generated when an encryption node is initialized.
- “Steps for connecting to a DPM appliance” on page 581 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 51 lists the associated roles and their read/write access to specific operations. The functions are enabled from the **Encryption Center** dialog box:

TABLE 51 Encryption privileges

Privilege	Read/Write
Storage Encryption Configuration	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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.

TABLE 51 Encryption privileges (Continued)

Privilege	Read/Write
Storage Encryption Security	<ul style="list-style-type: none"> • 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 195](#) on page 564.)
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 196](#).)

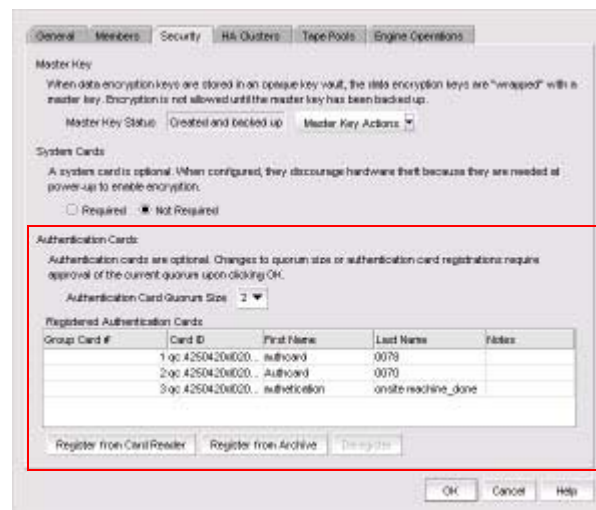


FIGURE 196 Security tab - registering authentication cards

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 197](#).)

The screenshot shows a dialog box titled "Add Authentication Card". It contains the following elements:

- Instruction: "To register an authentication card, you will need a card reader attached to the management station."
- Step 1: "1) Insert a card into the card reader and wait for the card's ID to appear below." Below this is a text field labeled "Card Serial #".
- Step 2: "2) Enter card assignment information. First name and last name are required. Skip this step if the card has previously been registered." Below this are two text fields: "First Name" and "Last Name".
- Field: "Notes:" followed by a large text area.
- Step 3: "3) This card has previously been registered. Enter a card password below and click OK." Below this are two text fields: "Card Password" and "Re-type Password". A small note "Case sensitive, 8-34 characters" is positioned between the two password fields.
- Status bar: "Status: Waiting for card to be inserted..."
- Buttons: "OK", "Cancel", and "Help" at the bottom right.

FIGURE 197 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 195](#) on page 564.)
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 198](#).)

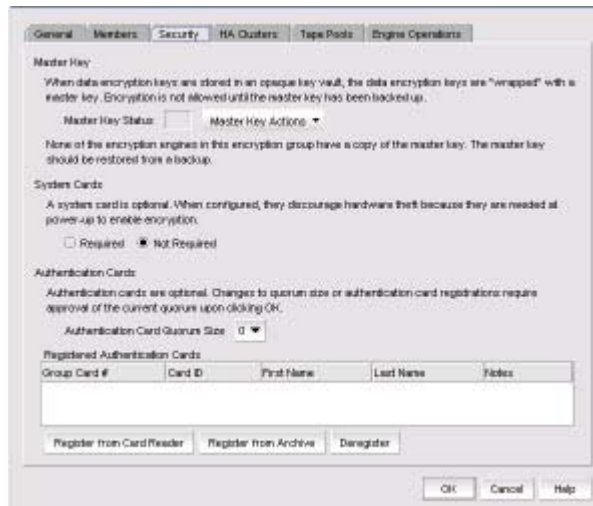


FIGURE 198 Encryption Group Properties dialog box - Security tab

3. Click **Register from Archive**.

The **Authentication Cards** dialog box displays. (Refer to [Figure 199](#).) The table lists the smart cards that are in the database.

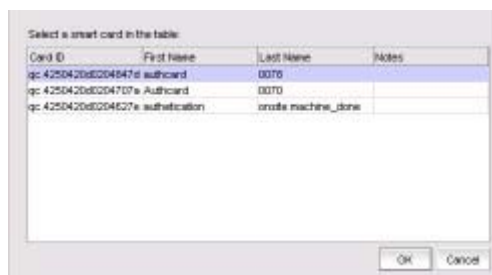


FIGURE 199 Authentication Cards dialog box - Registering smart cards from archive

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 195](#) on page 564.)
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 200](#).)



FIGURE 200 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 195](#) on page 564.)
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 195](#) on page 564.)
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 200](#) on page 571.) 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 195](#) on page 564.)
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 200](#) on page 571.)

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 195](#) on page 564.)
2. Select **Smart Card > Smart Card Tracking** from the menu task bar to display the **Smart Card Asset Tracking** dialog box. (Refer to [Figure 201](#).)

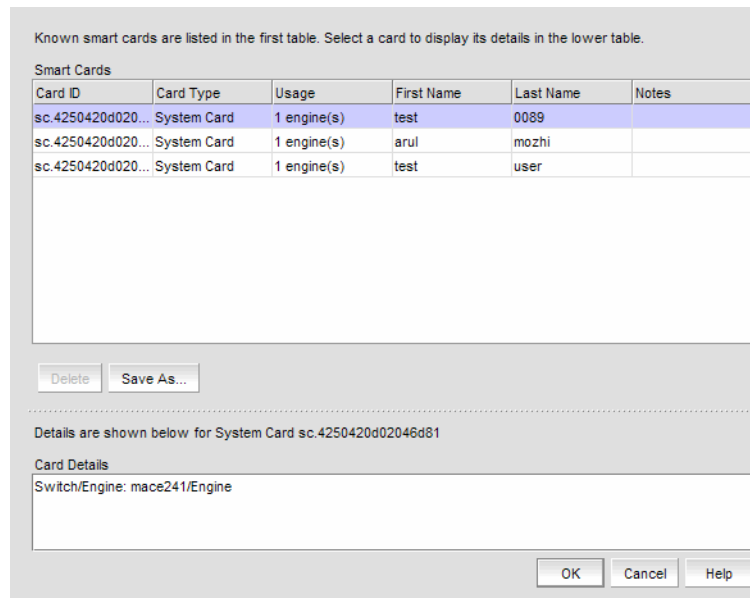


FIGURE 201 Smart Card Asset Tracking dialog box

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 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 does not modify the information that is stored on the card. Notes are optional.
- **Delete** button: Deletes a selected smart card from the 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 database. Deleting smart cards from the 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 202](#).)

To edit a smart card, you will need a card reader attached to the management station.

1) Insert a card into the card reader and wait for the card's ID to appear below. Then enter the card password and click Login button to retrieve card information from the card.

Card ID

Card Password

2) Change card assignment information.

Card Assignment

First Name Last Name

Notes

3) To change the password, select the check box below and enter the new password.

Change password

New Password

Case sensitive, 8-24 characters

Re-type Password

Status: Waiting for card to be inserted ...

FIGURE 202 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 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 195](#) on page 564.)
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 203](#).)

FIGURE 203 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 , 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 . 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 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 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 :

- 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 612.
- TEKA 4.0. The Thales keyAuthority is a KMIP-compliant server that can be configured with the ; 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 .

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 . 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 . 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 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 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 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:		

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 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 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 582.
2. Submit the KAC CSR for signing by a CA. Refer to [“Submitting the CSR to a certificate authority”](#) on page 583.
3. Set the KAC certificate registration expiry. Refer to [“KAC certificate registration expiry”](#) on page 583.
4. Import the signed certificate into the Fabric OS encryption node. Refer to [“Importing the signed KAC certificate”](#) on page 584.
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 584.
 - [“Uploading the KAC certificate onto the DPM appliance \(manual identity enrollment\)”](#) on page 585.
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 586.

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 583.

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 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 204.](#))



FIGURE 204 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 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 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 **SSLCertificateFile** 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 586.

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 205](#).)

If groups are not visible in the **Encryption Devices** table, select **View > Groups** from the menu bar.

DPM	
Encryption Group Name	DPM
Group Status	OK - Converged
Deployment Mode	Transparent
Failback Mode	Automatic
Key Vault Type	RSA Data Protection Manager (DPM)
REPL Support	Disabled
Primary Key Vault IP Address (IPv4 or hostname)	10.38.145.22
Primary Key Vault Connection Status	Connected
Backup Key Vault IP Address (IPv4 or hostname)	None
Backup Key Vault Connection Status	Key Vault Not Configured
High Availability Mode	(Not Applicable)
User Authentication	(Not Applicable)
Certificate Type	(Not Applicable)
Vendor Name	(Not Applicable)

If you specify a key vault IP address above, then you must enter a key vault certificate below.
If a key vault address is not specified above, then entries below are ignored.

Primary Key Vault Certificate

Version: V3
Subject: CN=RSA Key Manager Appliance Demo Root CA - 2012-08-20 15:12:35
Signature Algorithm: SHA1withRSA, OID = 1.2.840.113549.1.1.5
Key: Sun RSA public key, 2048 bits

Load from File...

Backup Key Vault Certificate

<None>

Load from File...

OK Cancel Help

FIGURE 205 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 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 :

- Install and launch the NetApp DataFort Management Console. Refer to [“Launching the NetApp DataFort Management Console”](#) on page 588.
- Establish the trusted link. Refer to [“Establishing the trusted link”](#) on page 588.
- Obtain and import the LKM/SSKM certificate. Refer to [“Obtaining and importing the LKM/SSKM certificate”](#) on page 589.
- Export and register encryption node certificates on LKM/SSKM. Refer to [“Exporting and registering the switch KAC certificates on LKM/SSKM”](#) on page 589.
- If required, create an LKM/SSKM cluster for high availability. Refer to [“LKM/SSKM key vault high availability deployment”](#) on page 590.
- Understanding Data Encryption Keys (DEKs). Refer to [“Data Encryption Keys”](#) on page 591.

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 195](#) on page 564.)
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 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 195](#) on page 564.)
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 206](#).) 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 206 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 :

- Configure a Brocade group on the ESKM/SKM. Refer to [“Configuring a Brocade group on ESKM/SKM”](#) on page 593.
- 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 594.
- Set up a local CA on the ESKM/SKM. Refer to [“Setting up the local Certificate Authority \(CA\) on ESKM/SKM”](#) on page 595.
- Download the CA certificate. Refer to [“Downloading the local CA certificate from ESKM/SKM”](#) on page 596.
- Create and install an ESKM/SKM server certificate. Refer to [“Creating and installing the ESKM/SKM server certificate”](#) on page 596.

- Enable an SSL connection. Refer to [“Enabling SSL on the Key Management System \(KMS\) Server”](#) on page 598.
- Configure a cluster of ESKM/SKM appliances for high availability. Refer to the following sections:
 - [“Creating an ESKM/SKM High Availability cluster”](#) on page 598
 - [“Copying the local CA certificate for a clustered ESKM/SKM appliance”](#) on page 599
 - [“Adding ESKM/SKM appliances to the cluster”](#) on page 599
- Export and sign the encryption node certificate signing requests. Refer to [“Signing the encryption node KAC certificates”](#) on page 600.
- Import the signed certificates into the encryption node. Refer to [“Importing a signed KAC certificate into a switch”](#) on page 601.

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.

NOTE

Fabric OS 6.2.0 uses `brocduser1` as a standard user name when creating a Brocade group on ESKM/SKM. If you downgrade to version 6.2.0, the user name is overwritten to `brocduser1`, and the Brocade group user name must be changed to `brocduser1`.

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 195](#) on page 564.)
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 207](#).)

FIGURE 207 Key Vault Credentials dialog box

The dialog box contains the following information:

- **Primary Key Vault:** Preselected. 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 208](#).)
4. Enter information required by the **Create Local Certificate Authority** 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.

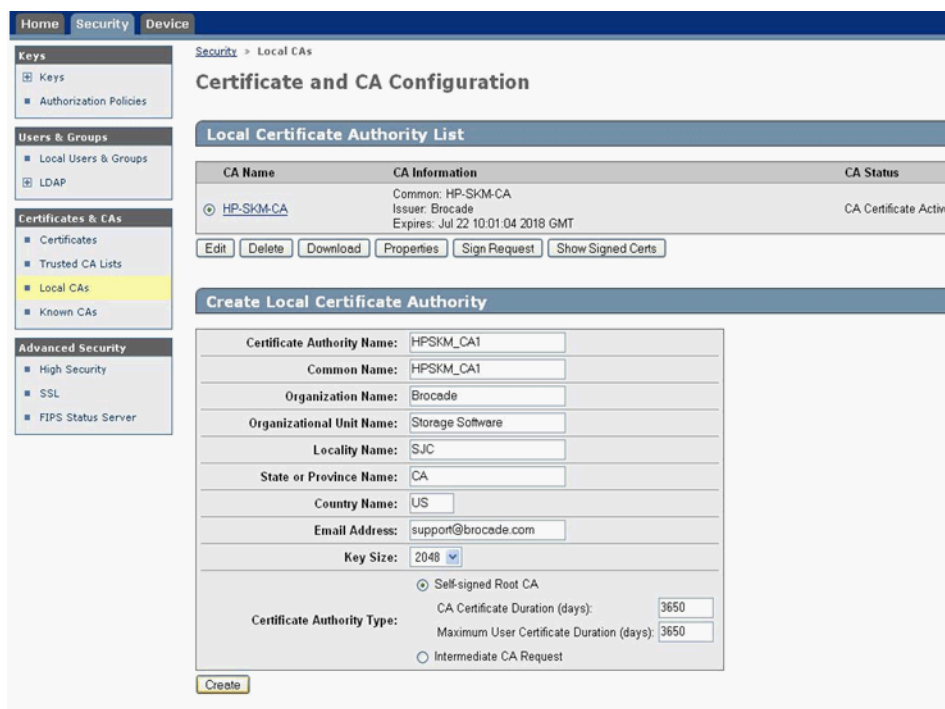


FIGURE 208 Creating an HP ESKM/SKM local CA

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 595 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 595.
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 596.
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 598.

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 596 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 195](#) on page 564.)
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 195](#) on page 564.)
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 209](#).)



FIGURE 209 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. Beginning with Fabric OS 6.3.0, 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/SKMs in the cluster. After the DEK is archived successfully, it gets synchronized with other ESKMs/SKMs in the cluster. If archival is successful, 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 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 :

1. Set up network connections to TEKA. Refer to [“Setting up TEKA network connections”](#) on page 604.
2. Create a TEKA client. Refer to [“Creating a client on TEKA”](#) on page 605.
3. Establish TEKA key vault credentials. Refer to [“Establishing TEKA key vault credentials on the switch”](#) on page 606.
4. Sign encryption node certificate signing requests. Refer to [“Exporting the Fabric OS node self-signed KAC certificates”](#) on page 609.
5. Import the signed requests onto the encryption nodes. Refer to [“Converting the KAC certificate format”](#) on page 609.

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 210.](#))

The screenshot shows the THALES Network Settings page. At the top, there are navigation tabs: Summary, Users, Network (selected), Date & Time, Licensing, and Logs. Below these are sub-tabs: General, SNMP, Remote Syslog, and Email Alerts. The main content area is titled 'Network Settings' and is divided into four sections:

- Management Interface:** Contains input fields for IP address, Subnet mask, and Gateway.
- KM Server Interface:** Contains input fields for IP address, Subnet mask, and Gateway.
- Common Settings:** Contains input fields for HostName, Domain, Primary DNS, and Secondary DNS.
- Service Settings:** Contains input fields for HTTPS Port (443), SSH Port (22), and KM Server Port (9000). It also has checkboxes for 'Enable SSH' and 'Enable KM Server', both of which are checked.

At the bottom of the form are 'Save' and 'Reset' buttons.

FIGURE 210 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 211](#).)

The user name and user group name are applicable only for TEKA /Thales key vault. They are used for creating the client account on the key vault.

User Name

User Group Name

FIGURE 211 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 212](#).)

THALES Help | Logout

Summary Users Groups **Clients** Trusts Keys Logs

Clients

Showing clients 1 to 10 of 18

1 of 2 Page size: 10

<input type="checkbox"/>	Name	Type	Group	Home Directory	Certificate	Details
<input type="checkbox"/>	neptunetop	P1619	brcd1	/neptunetop/		
<input type="checkbox"/>	mace52	P1619	brcd	/mace52/		
<input type="checkbox"/>	mace51	P1619	brcd	/mace51/		
<input type="checkbox"/>	mace190-2	P1619	brcd2	/mace190-2/		
<input type="checkbox"/>	mace160	P1619	brcd1	/mace160/		
<input type="checkbox"/>	Mace158	P1619	brcd	/Mace158/		
<input type="checkbox"/>	Cliff101	P1619	brcd1	/Cliff101/		
<input type="checkbox"/>	Cliff	P1619	brcd1	/Cliff/		
<input type="checkbox"/>	client1	P1619	brcd	/client1/		
<input type="checkbox"/>	brcduser2	P1619	brcd	/brcduser2/		

Delete | Add Client

FIGURE 212 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 . 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 213](#).)

FIGURE 213 Key Vault Credentials dialog box

The dialog box contains the following information:

- **Primary Key Vault** selector: Preselected.
 - **Secondary Key Vault** selector: Active 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 214](#).)



FIGURE 214 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 609.
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 609.
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 609.
7. Create a certificate for the TKLM server. Refer to [“Creating a self-signed certificate for TKLM”](#) on page 610.
8. Import the node KAC certificates. Refer to [“Importing the Fabric OS encryption node KAC certificates to TKLM”](#) on page 610.
9. Export the server CA certificate to a LINUX or Windows host. Refer to [“Exporting the TKLM self-signed server certificate”](#) on page 611.
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 611.
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 the 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 609.

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/tpktklmV2/bin/wsadmin.sh -username TKLMAdmin
-password <password> -lang jython
```

For Windows:

```
<installed directory>\ibm\tivoli\tpktklmV2\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/tpktklmV2/products/tklm/
```

For Windows:

```
<installed directory>\ibm\tivoli\tpktklmV2\products\tklm\
```

4. Transfer the TKLM certificate that was previously exported into the TKLM server file system to the 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 195](#) on page 564.)
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 215](#).)



FIGURE 215 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 . Any KMIP-compliant server can be reregistered as a KMIP key vault on the 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 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 613.)
2. Create a local CA. (Refer to [“Creating a local CA”](#) on page 614.)
3. Create a server certificate. (Refer to [“Creating a server certificate”](#) on page 615.)
4. Create a cluster. (Refer to [“Creating a cluster”](#) on page 620.)
5. Create a Brocade group on the KeySecure appliance. (Refer to [“Configuring a Brocade group on the KeySecure”](#) on page 621.)
6. Register the user name and password. (Refer to [“Registering the KeySecure Brocade group user name and password”](#) on page 622.)
7. Export and sign the encryption node certificate signing requests. (Refer to [“Signing the encryption node KAC CSR on KMIP”](#) on page 623.)
8. Import the signed certificates into the encryption node. (Refer to [“Importing a signed KAC certificate into a switch”](#) on page 625.)
9. Back up the certificates (Refer to [“Backing up the certificates”](#) on page 626.)
10. Configure the KMIP server. (Refer to [“Configuring the KMIP server”](#) on page 628.)
11. Add a secondary node to the cluster. (Refer to [“Adding a node to the cluster”](#) on page 629.)

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. (Refer to [Figure 216](#).)

The screenshot shows the SafeNet KeySecure Management Console interface. The top navigation bar includes 'Home', 'Security', and 'Device'. The left sidebar has a tree view with categories: Keys, Users & Groups, CAs & SSL Certificates, and Advanced Security. The 'Advanced Security' section is expanded to show 'High Security'. The main content area is titled 'High Security Configuration' and contains three main sections:

- FIPS Compliance:** A section with a 'Help' icon and a dropdown menu showing 'Is FIPS Compliant: Yes'.
- High Security Settings:** A section with a 'Help' icon containing two sub-sections:
 - Key Security:**
 - Disable Creation and Use of Global Keys:
 - Disable Non-FIPS Algorithms and Key Sizes:
 - Disable RSA Encryption and Decryption:
 - Device Security:**
 - Disable FTP for Certificate Import, Backup and Restore:
 - Disable Certificate Import through Serial Console Paste:
 - Disable Hotswappable RAID Drives:
- Security Settings Configured Elsewhere:** A section with a 'Help' icon containing a table of settings:

Allow Key and Policy Configuration Operations:	Disabled (FIPS compliant)
Allow Key Export:	Disabled (FIPS compliant)
User Directory:	Local (FIPS compliant)
LDAP Administrator Server Configured:	No (FIPS compliant)
Allowed SSL Protocols:	TLS 1.0 (FIPS compliant)
Enabled SSL Ciphers:	Only FIPS compliant ciphers

FIGURE 216 KeySecure High Security Configuration page

2. Under **FIPS Compliance**, set **FIPS Compliance** to **Yes**.

This ensures that only TLS 1.0 connections are supported between the 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. (Refer to [Figure 217](#).)

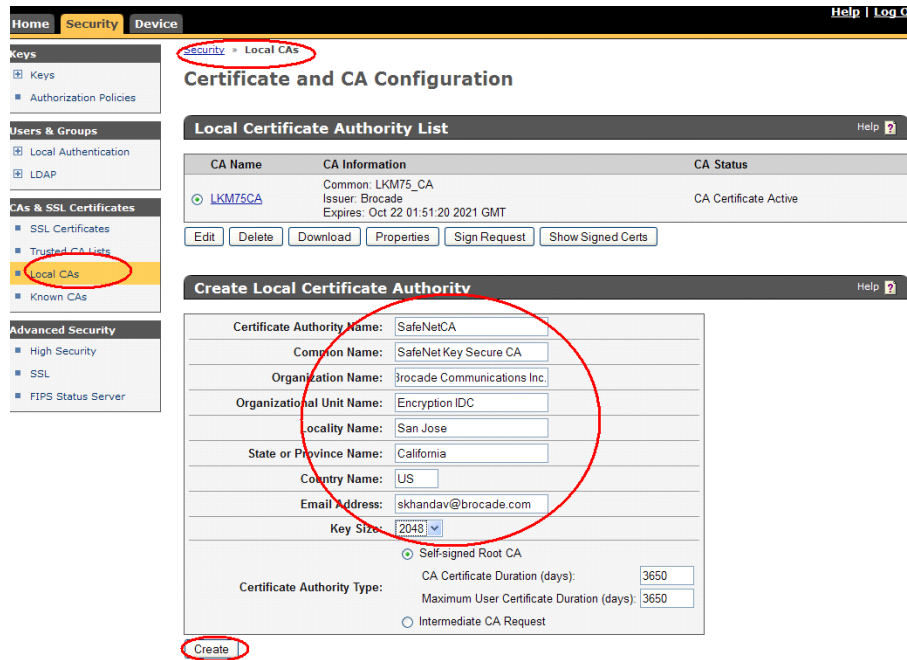


FIGURE 217 KeySecure Certificate and CA Configuration page - Create Local Certificate Authority

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. (Refer to [Figure 218](#).)

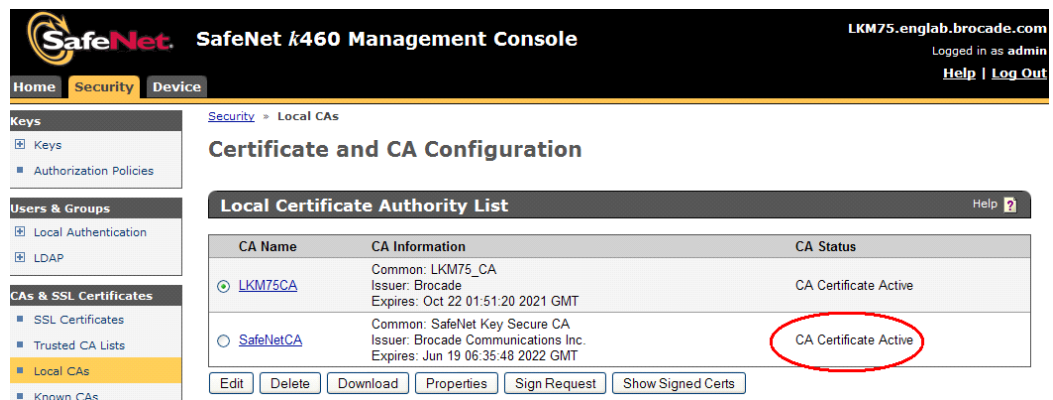


FIGURE 218 KeySecure Certificate and CA Configuration page - Local Certificate Authority List

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. (Refer to [Figure 219](#).)

The screenshot shows the KeySecure web interface. The top navigation bar includes 'Home', 'Security', and 'Device'. The left sidebar has a tree view with 'SSL Certificates' selected. The main content area is titled 'Certificate and CA Configuration'. It features a 'Certificate List' table and a 'Create Certificate Request' form. The table has columns for Certificate Name, Certificate Information, Certificate Purpose, and Certificate Status. The form has fields for Certificate Name, Common Name, Organization Name, Organizational Unit Name, Locality Name, State or Province Name, Country Name, Email Address, and Key Size. A red circle highlights the 'Create Certificate Request' button and the form fields.

Certificate Name	Certificate Information	Certificate Purpose	Certificate Status
LKM75_Cert	Common: LKM75_ServerCert Issuer: Brocade Expires: Oct 21 01:53:47 2021 GMT	Server	Active

FIGURE 219 KeySecure Certificate and CA Configuration page

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. (Refer to [Figure 220](#).)

20 Steps for connecting to a KMIP-compliant SafeNet KeySecure

The screenshot shows the 'Certificate List' section of the SafeNet KeySecure management console. A table lists certificates with columns for Certificate Name, Certificate Information, Certificate Purpose, and Certificate Status. The entry 'Safenet75ServerCert' is circled in red, showing a status of 'Request Pending'. A warning message below the table states: 'Warning: Certificate requests should be backed up for protection'. Below the table is a 'Create Certificate Request' form with fields for Certificate Name, Common Name, Organization Name, Organizational Unit Name, Locality Name, State or Province Name, Country Name (set to US), Email Address, and Key Size (set to 2048).

FIGURE 220 KeySecure Certificate and CA Configuration page - Certificate List

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. (Refer to [Figure 221](#).)

The screenshot shows the 'Certificate Request Information' page for the 'Safenet75ServerCert'. It displays the certificate name and key size (2048). Below this, the 'Subject' field is populated with: CN: SafeNet Key Secure Server Certificate, O: Brocade Communications Inc., OU: Encryption IDC, L: San Jose, ST: California, C: US, and emailAddress: skhandav@brocade.com. The main content of the page is a large block of base64-encoded text representing the certificate request, which is circled in red. At the bottom, there are buttons for 'Download', 'Install Certificate', 'Create Self Sign Certificate', and 'Back'.

FIGURE 221 KeySecure Certificate and CA Configuration page - Certificate Request Information

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**. (Refer to [Figure 222](#).)

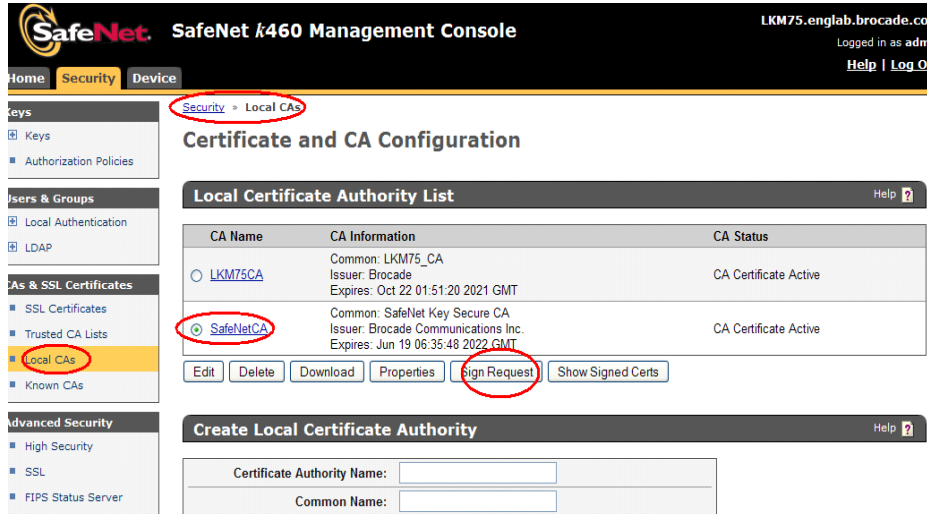


FIGURE 222 KeySecure Certificate and CA Configuration page - Local Certificate Authority List

The **Sign Certificate Request** dialog box displays. (Refer to [Figure 223](#).)

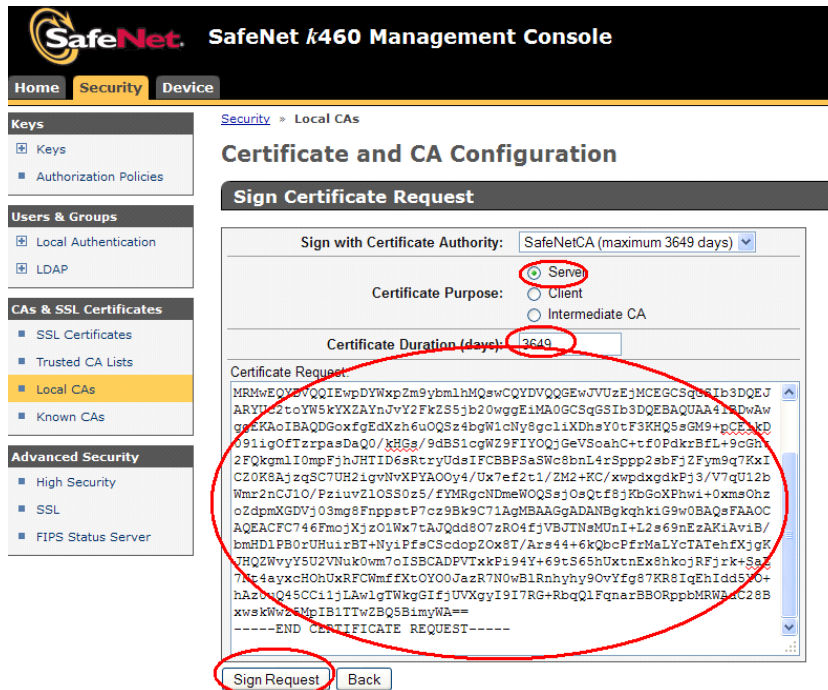


FIGURE 223 KeySecure Certificate and CA Configuration page - Sign Certificate Request

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**. (Refer to [Figure 224](#).)

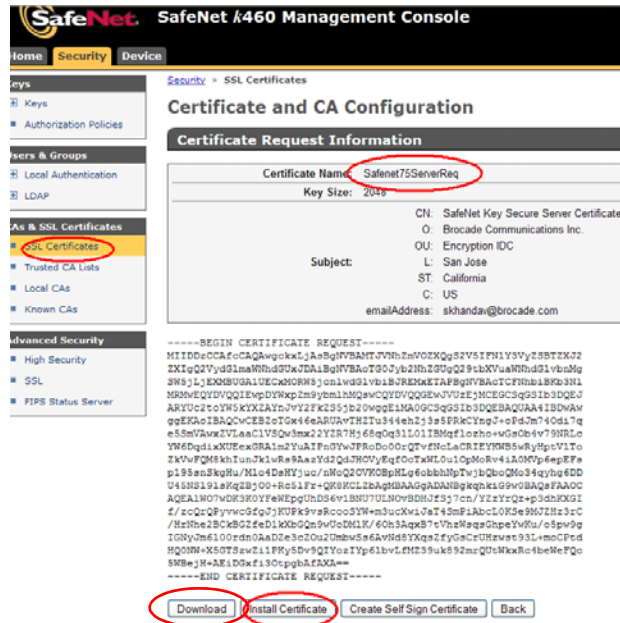


FIGURE 224 KeySecure Certificate and CA Configuration page - Certificate Request Information

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**. (Refer to [Figure 225](#).)

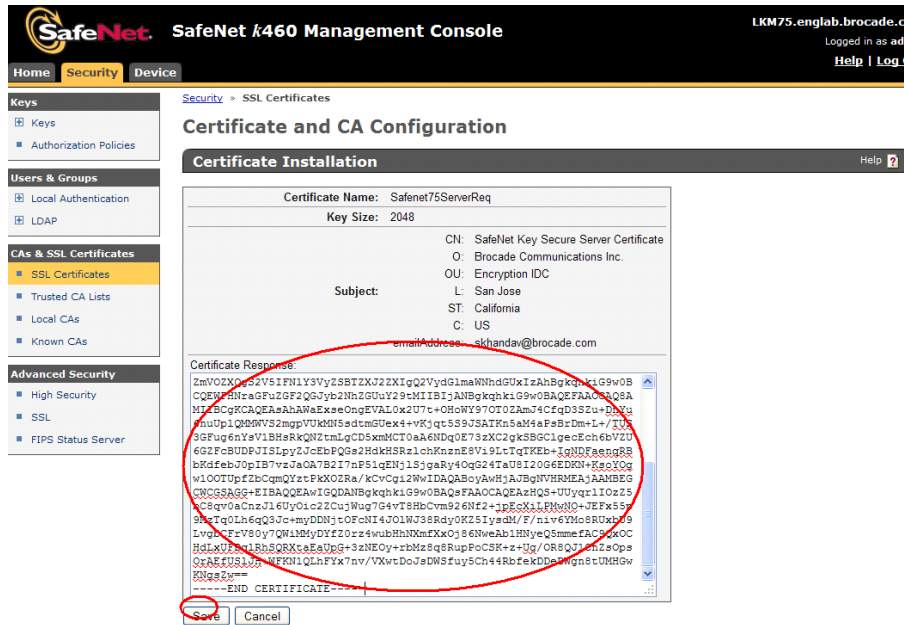


FIGURE 225 KeySecure Certificate and CA Configuration page - Certificate Installation

14. After the page refreshes, the new certificate information is displayed in the **Certificate List** table. (Refer to Figure 226.)

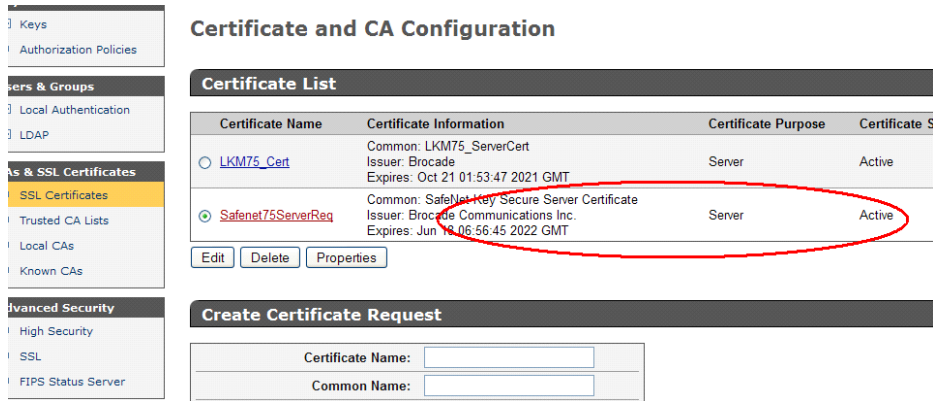


FIGURE 226 KeySecure Certificate and CA Configuration page - Certificate List

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. (Refer to [Figure 227](#).)

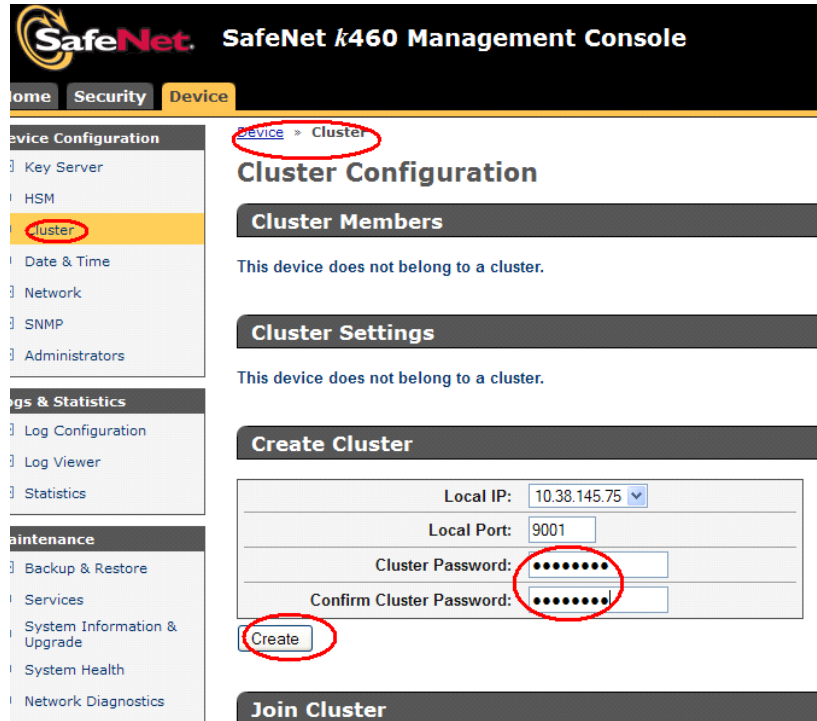


FIGURE 227 KeySecure Cluster Configuration page

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. (Refer to [Figure 228](#).)
3. Verify the cluster status is shown as **Active**.

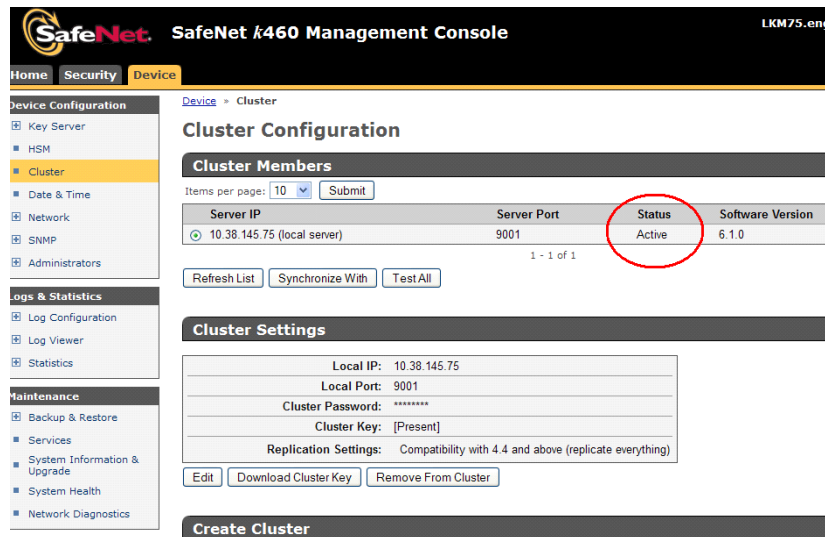


FIGURE 228 KeySecure Cluster Configuration page - Cluster Members

- Under **Cluster Settings**, click **Download Cluster Key**. (Refer to [Figure 229](#).)
You are prompted to enter a local file name.

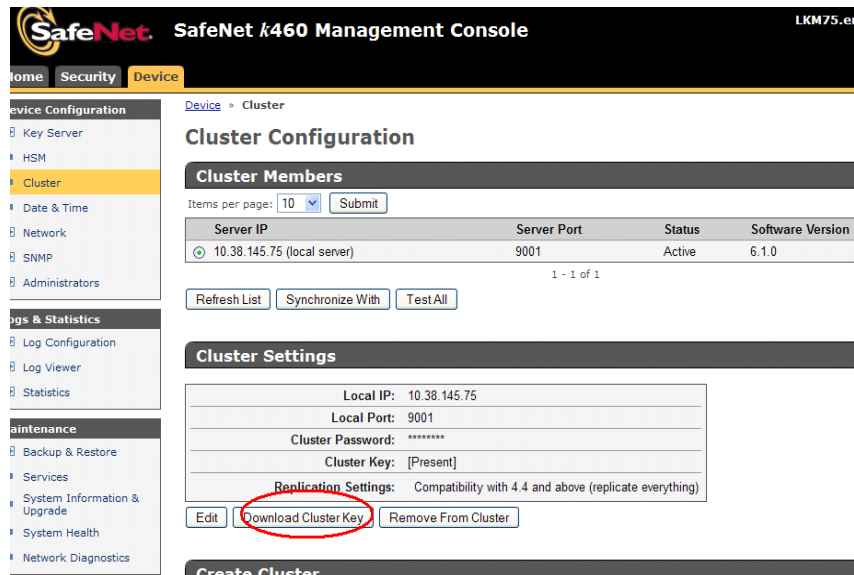


FIGURE 229 KeySecure Cluster Configuration page - Cluster Settings

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.

- Log in to the KeySecure web management console using the admin password.
- 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 230](#).)

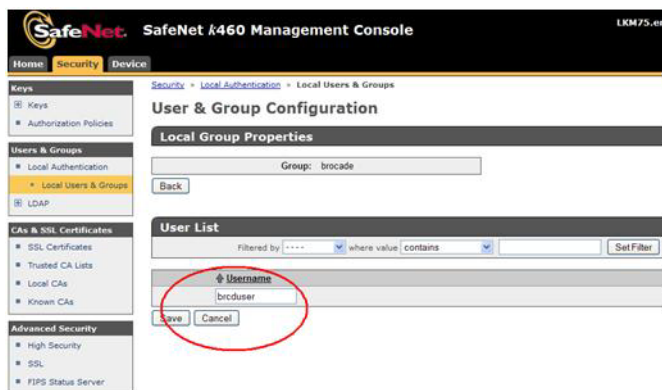


FIGURE 230 User & Group Configuration page - Local Group Properties and User List

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 195](#) on page 564.)
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 231](#)).

To change the existing key vault credentials, select a key vault position, and enter the user name and password. The existing credentials would be overwritten if the operation succeeded. This operation is only applicable for TEKA (Thales), SKM and KMIP key vault.

Primary Key Vault
 Secondary Key Vault

User Name
 User Group Name
 Password
 Re-type Password

FIGURE 231 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 195](#) on page 564.)
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 232](#).)




FIGURE 232 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 195](#) on page 564.)
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 233](#).)



FIGURE 233 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. (Refer to [Figure 234](#).)

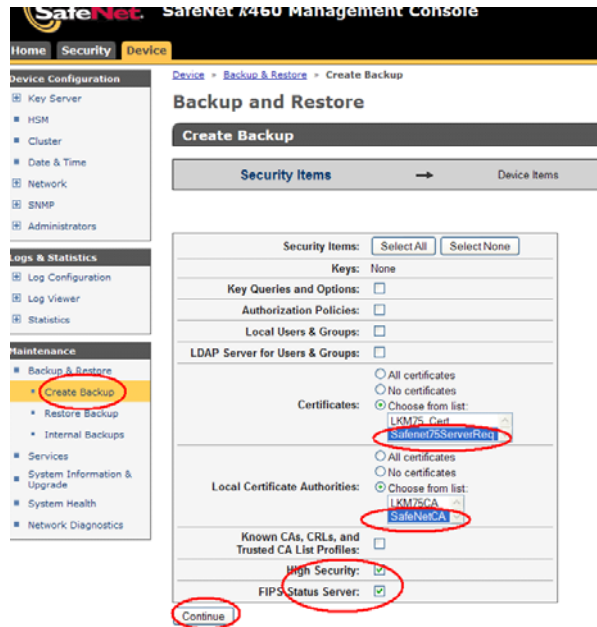


FIGURE 234 Backup and Restore page

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 235](#).)

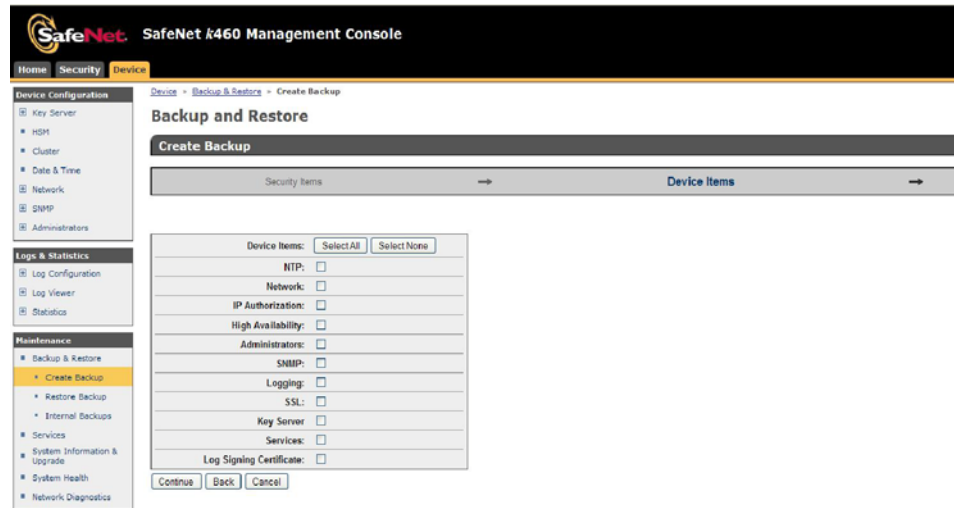


FIGURE 235 Backup and Restore page - Device items

5. Select the items for backup, then click **Continue**.

The **Create Backup** page displays, which is used for setting backup details. (Refer to [Figure 236](#).)

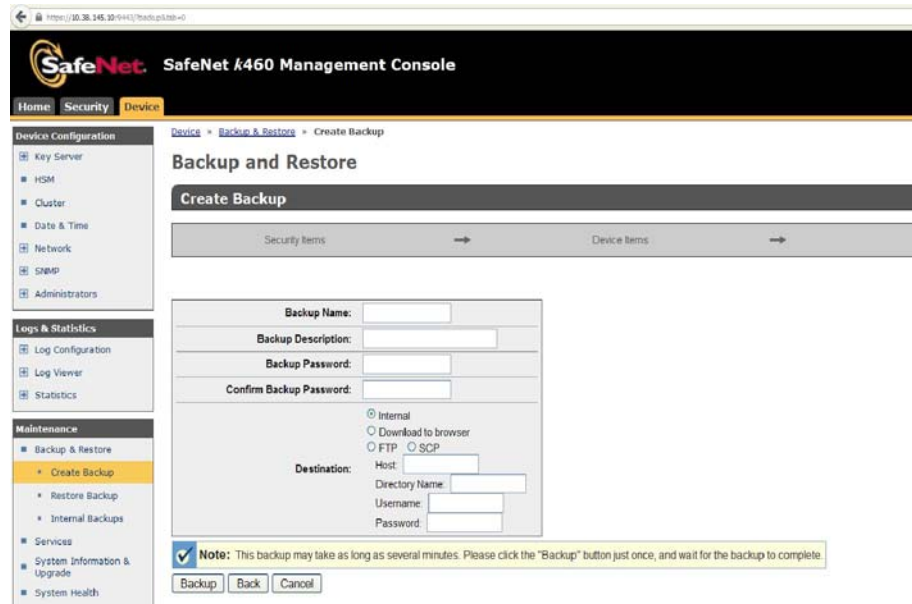


FIGURE 236 Backup and Restore page - Backup details

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. (Refer to [Figure 237](#).)

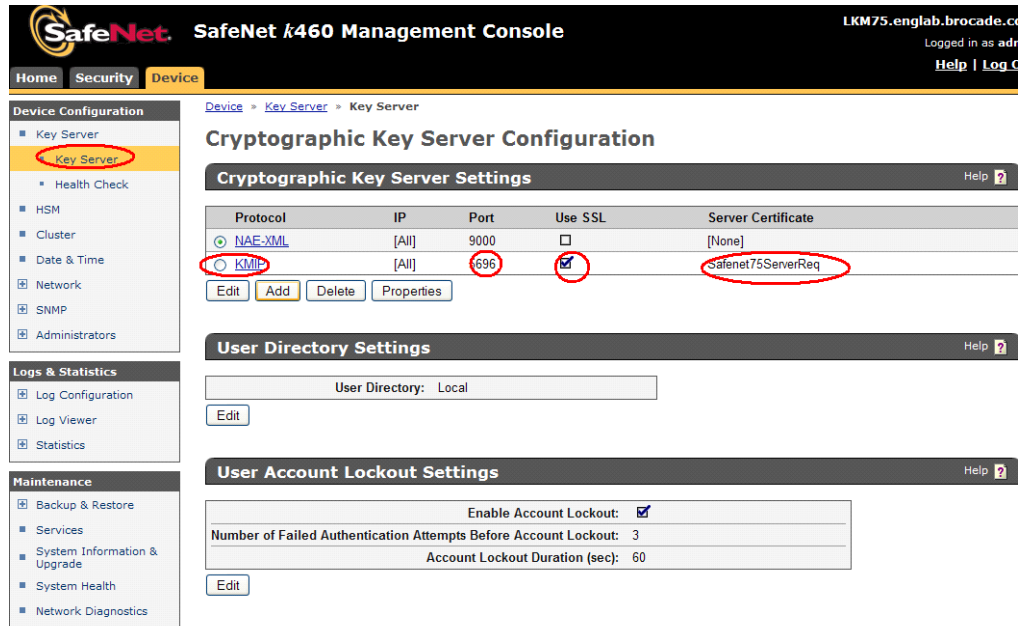


FIGURE 237 KeySecure Cryptographic Key Server Configuration page

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. (Refer to [Figure 238](#).)

The screenshot shows the 'Cluster Configuration' page in the KeySecure Management Console. The page is divided into three main sections: 'Cluster Members', 'Cluster Settings', and 'Create Cluster'. The 'Cluster Members' section shows 'This device does not belong to a cluster.' The 'Cluster Settings' section also shows 'This device does not belong to a cluster.' The 'Create Cluster' section contains a form with fields for 'Local IP' (10.38.145.76), 'Local Port' (9001), 'Cluster Password', and 'Confirm Cluster Password'. Below this is a 'Join Cluster' section with fields for 'Local IP' (10.38.145.76), 'Local Port' (9001), 'Cluster Member IP' (10.38.145.75), 'Cluster Member Port' (9001), 'Cluster Key File' (C:\Documents and Settings\...), and 'Cluster Password'. The 'Join' button and the 'Cluster Password' field are circled in red.

FIGURE 238 KeySecure Cluster Configuration page

2. Under **Join Cluster**, enter the cluster information that you configured for the primary KeySecure node. (Refer to [“Creating a cluster”](#) on page 620.)
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. (Refer to [Figure 239](#).)

20 Steps for connecting to a KMIP-compliant SafeNet KeySecure

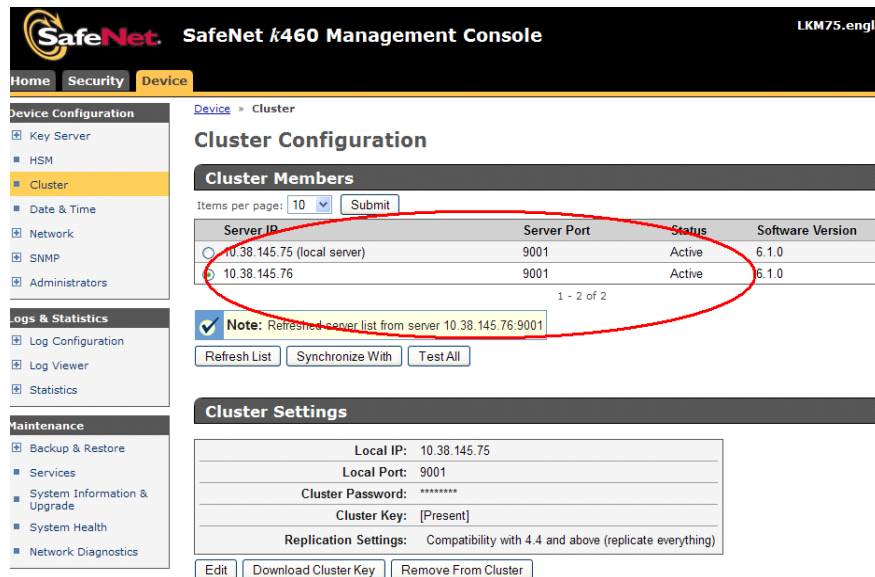


FIGURE 239 KeySecure Cluster Configuration page - Cluster Members

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. (Refer to [Figure 240](#).)

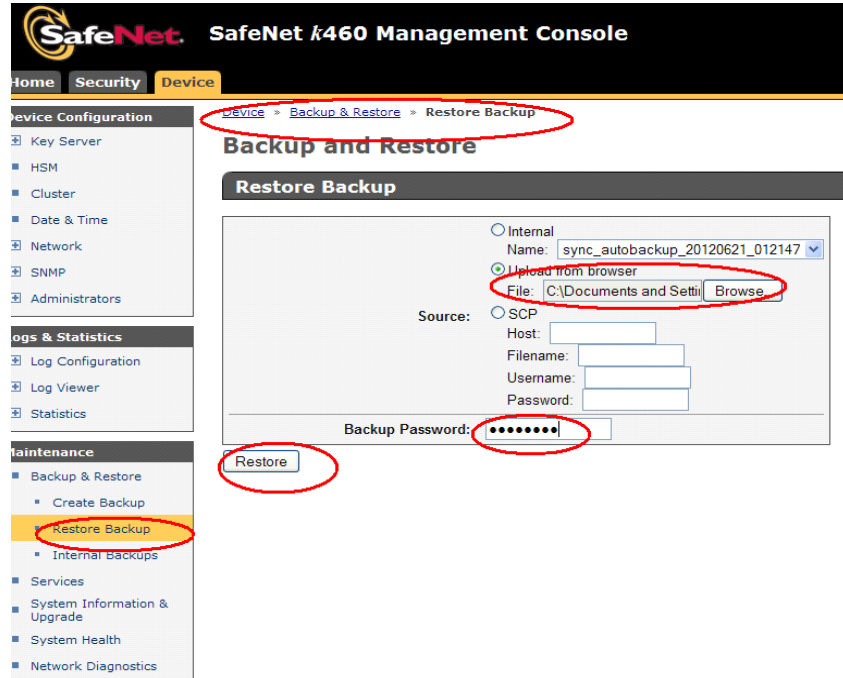


FIGURE 240 KeySecure Backup and Restore page

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. (Refer to [Figure 241.](#))

NOTE

A message displays, advising that the secondary node requires a restart.

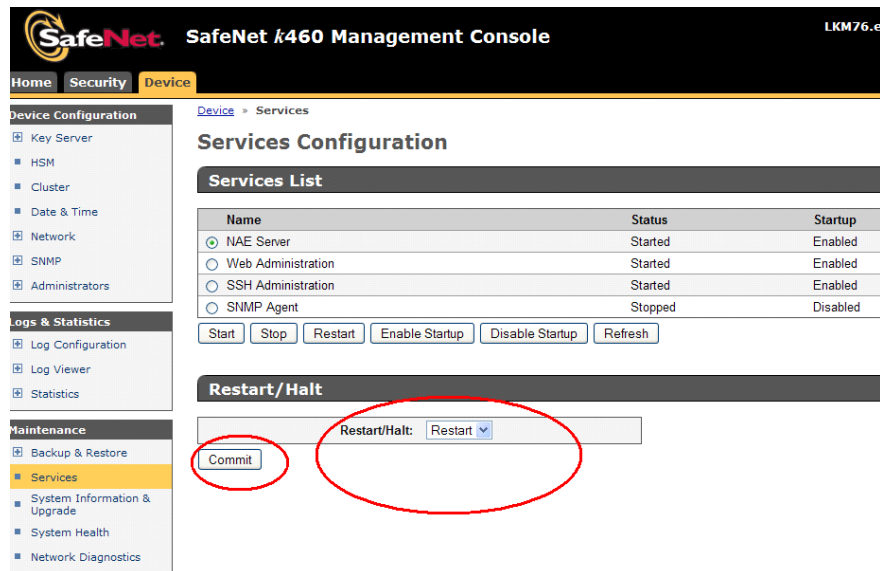


FIGURE 241 KeySecure Services Configuration page

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 242](#).)




FIGURE 242 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 243](#).) 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.

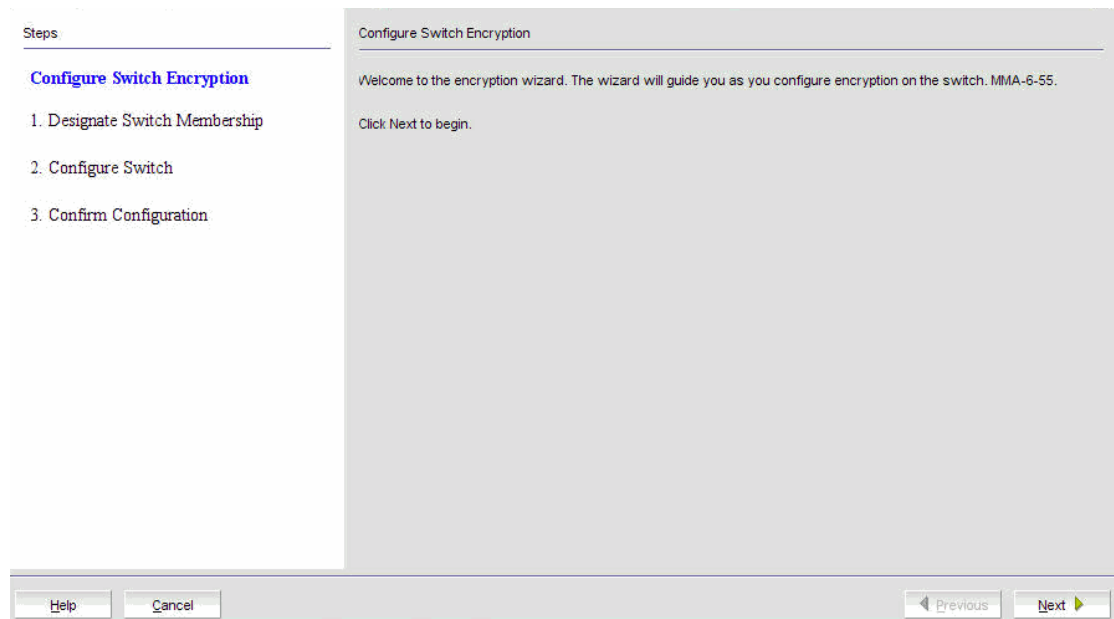


FIGURE 243 Configure Switch Encryption wizard - welcome screen

4. From the **Configure Switch Encryption** welcome screen, click **Next** to begin.

The **Designate Switch Membership** dialog box displays (Figure 244). 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

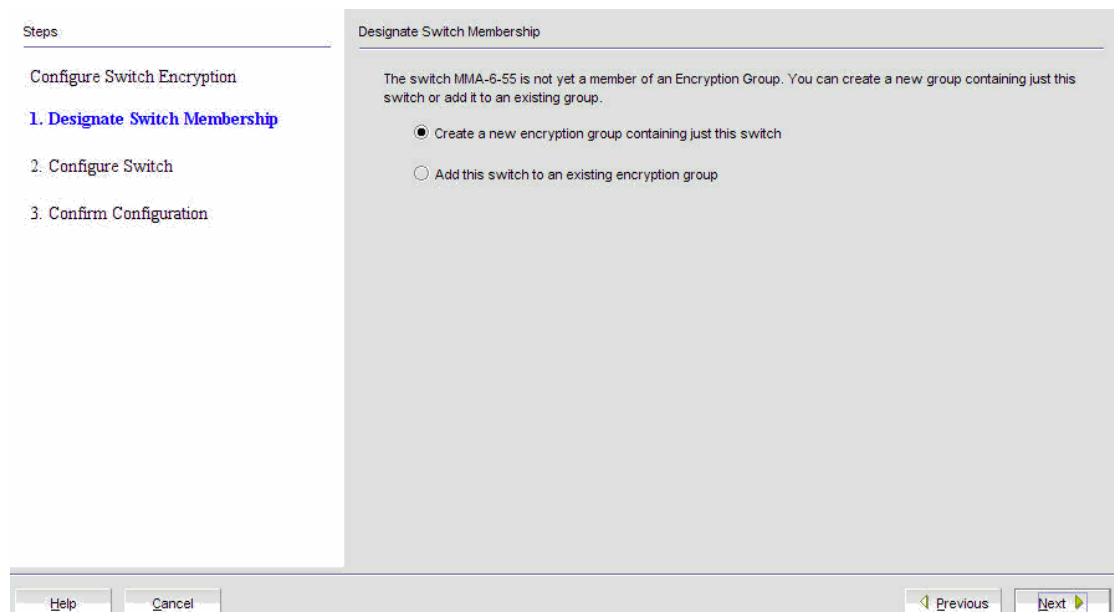


FIGURE 244 Designate Switch Membership dialog box

- 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 670.

The **Create a New Encryption Group** dialog box displays. (Refer to [Figure 245](#).)

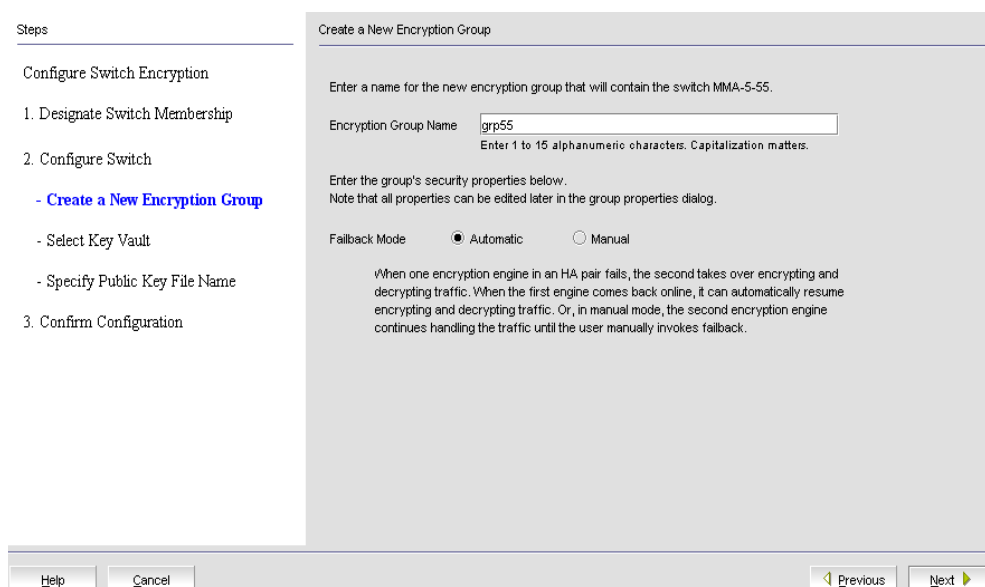


FIGURE 245 Create a New Encryption Group dialog box

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.

- 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 246](#).)

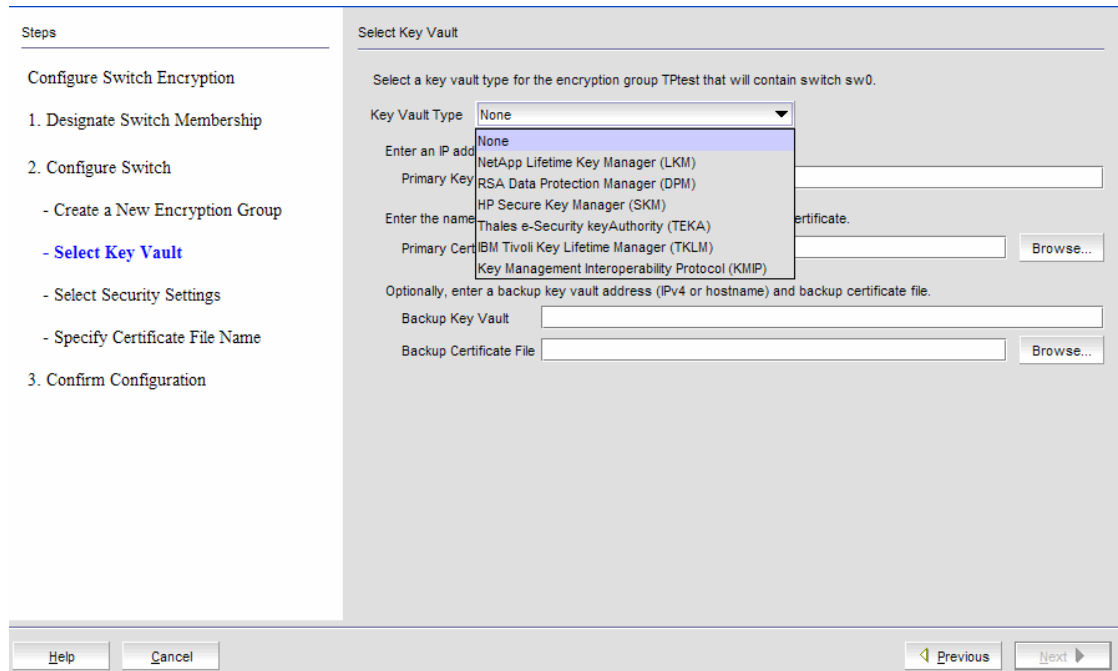


FIGURE 246 Select Key Vault dialog box

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 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.

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 633.

Figure 247 shows the key vault selection dialog box for DPM.

The screenshot shows a 'Select Key Vault' dialog box. On the left, a 'Steps' pane lists the following steps: 'Configure Switch Encryption', '1. Designate Switch Membership', '2. Configure Switch' (with sub-steps: '- Create a New Encryption Group', '- Select Key Vault', '- Select Security Settings', '- Specify Certificate File Name'), and '3. Confirm Configuration'. The 'Select Key Vault' step is highlighted in blue. The main dialog area contains the following fields and options: 'Key Vault Type' is a dropdown menu set to 'RSA Data Protection Manager (DPM)'; 'Primary Key Vault' is a text input field with a placeholder 'Enter an IP address (IPv4 or hostname) for the primary key vault.'; 'Primary Certificate File' is a text input field with a placeholder 'Enter the name of the file holding the primary key vault's CA certificate.' and a 'Browse' button; 'REPL Support' has two radio buttons, 'Enabled' and 'Disabled', with 'Disabled' selected. At the bottom, there are 'Help', 'Cancel', 'Previous', and 'Next' buttons.

FIGURE 247 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 248](#).)



FIGURE 248 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 249](#).)




FIGURE 249 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 250](#).)

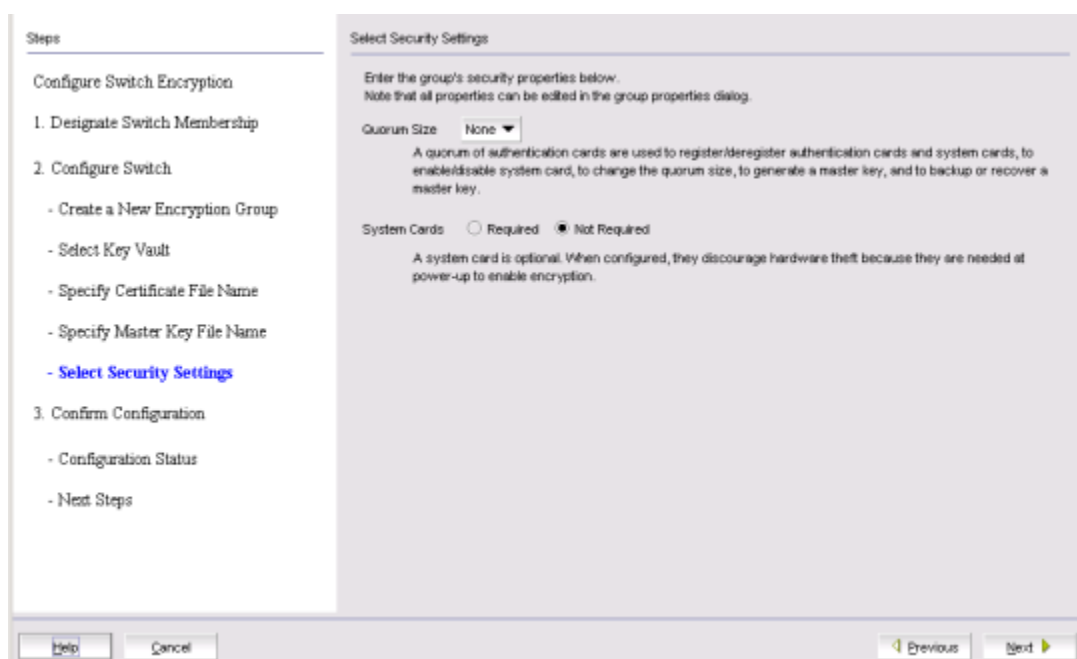


FIGURE 250 Select Security Settings dialog box

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 251](#).) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

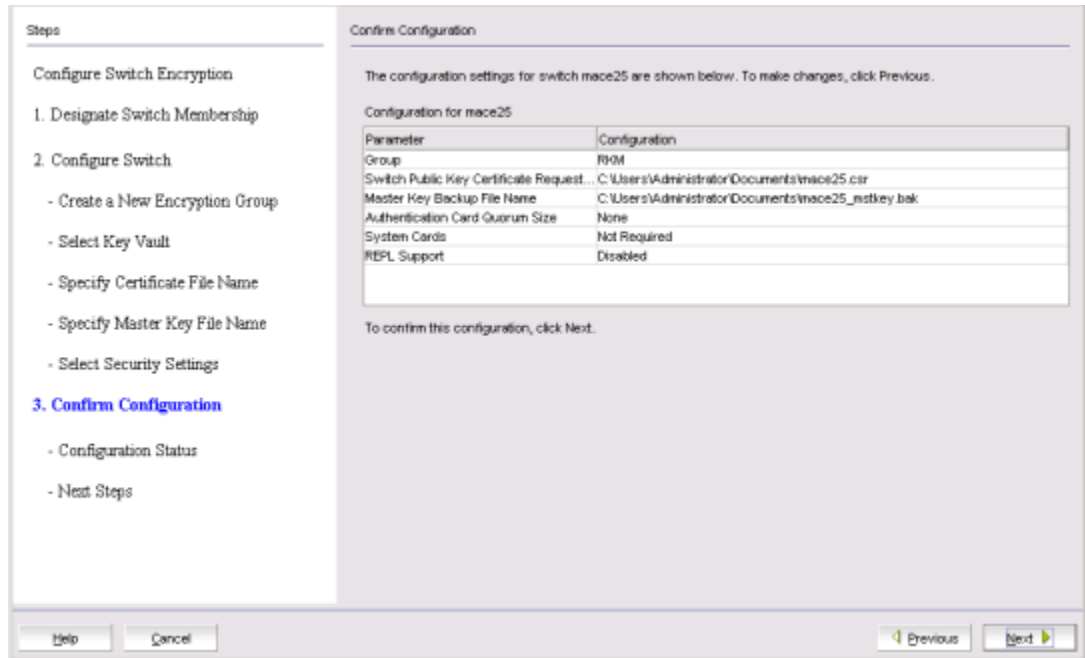


FIGURE 251 Confirm Configuration dialog box

The **Configuration Status** dialog box displays. (Refer to Figure 252.)

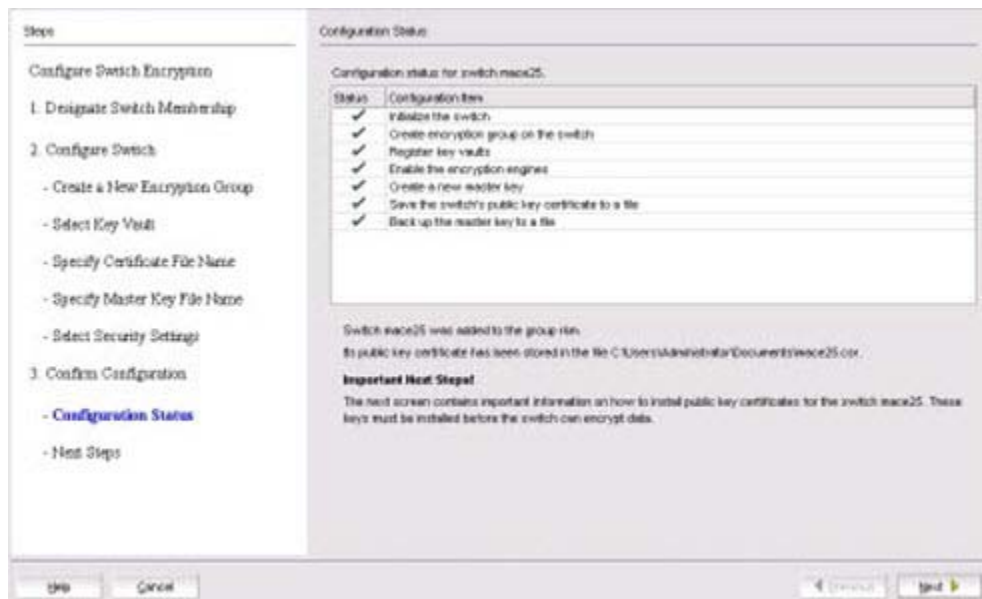


FIGURE 252 Configuration Status dialog box

- 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 253.) Instructions for installing public key certificates for the encryption switch are displayed.

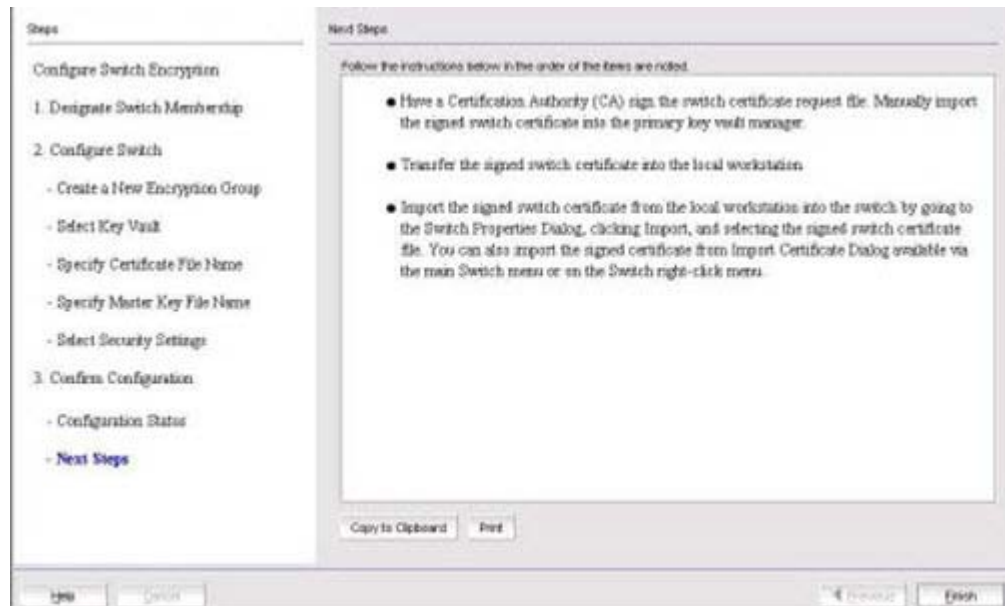


FIGURE 253 Next Steps dialog box

- 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 633.

[Figure 254](#) shows the key vault selection dialog box for LKM/SSKM.

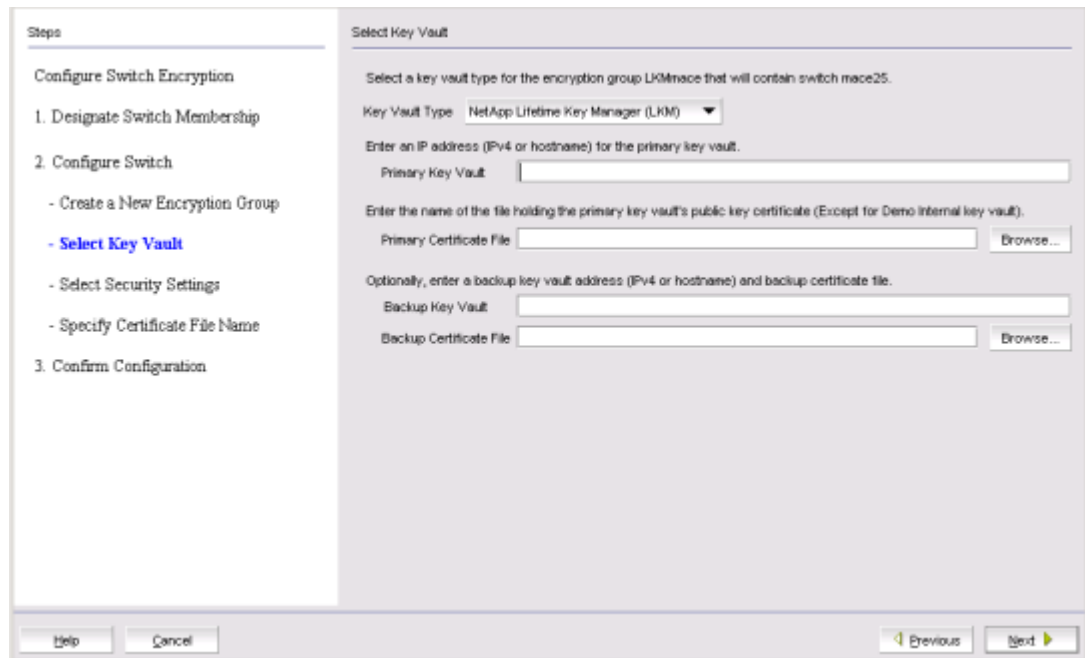


FIGURE 254 Select Key Vault dialog box for LKM/SSKM

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 255.](#))

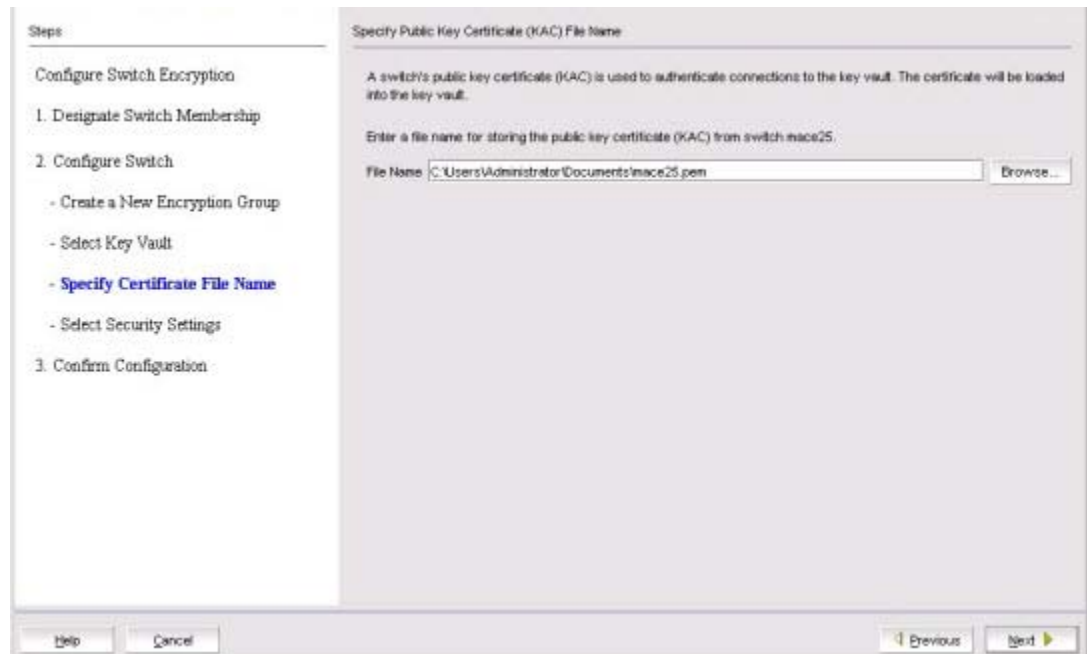


FIGURE 255 Specify Public Key Certificate (KAC) File Name dialog box

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 256](#).)

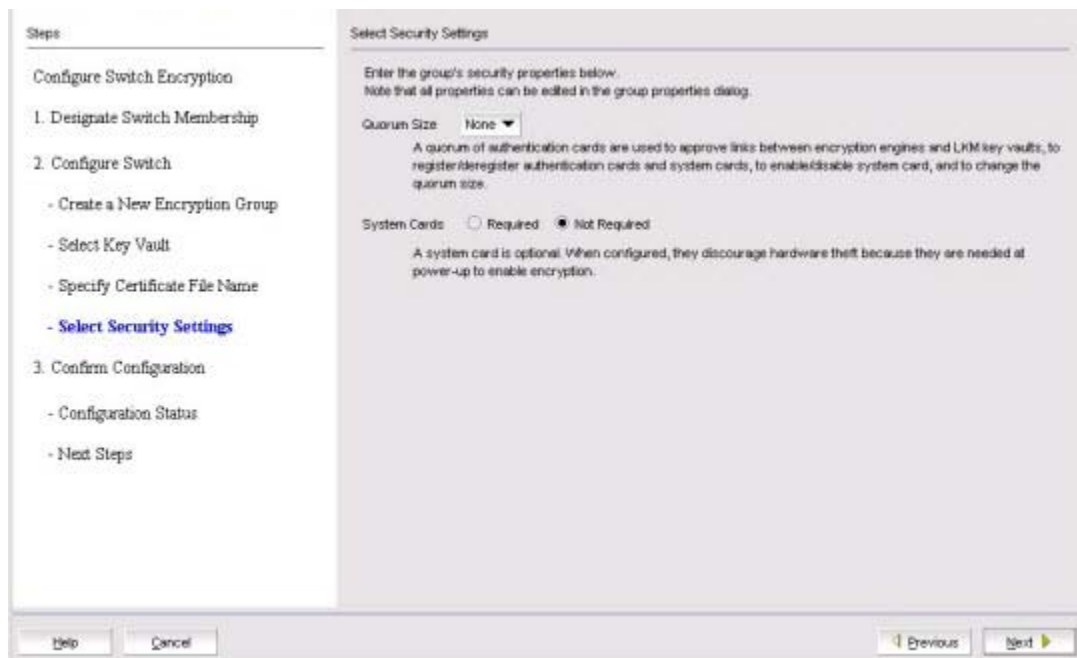


FIGURE 256 Select Security Settings dialog box

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 257](#).) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.




FIGURE 257 Confirm Configuration dialog box

The Configuration Status dialog box displays. (Refer to Figure 258.)




FIGURE 258 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, sends API commands to verify the switch configuration. See [“Understanding configuration status results”](#) on page 669 for more information.

8. Verify the information is correct, then click **Next**.

The **Next Steps** dialog box displays. (Refer to [Figure 259](#).) Instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

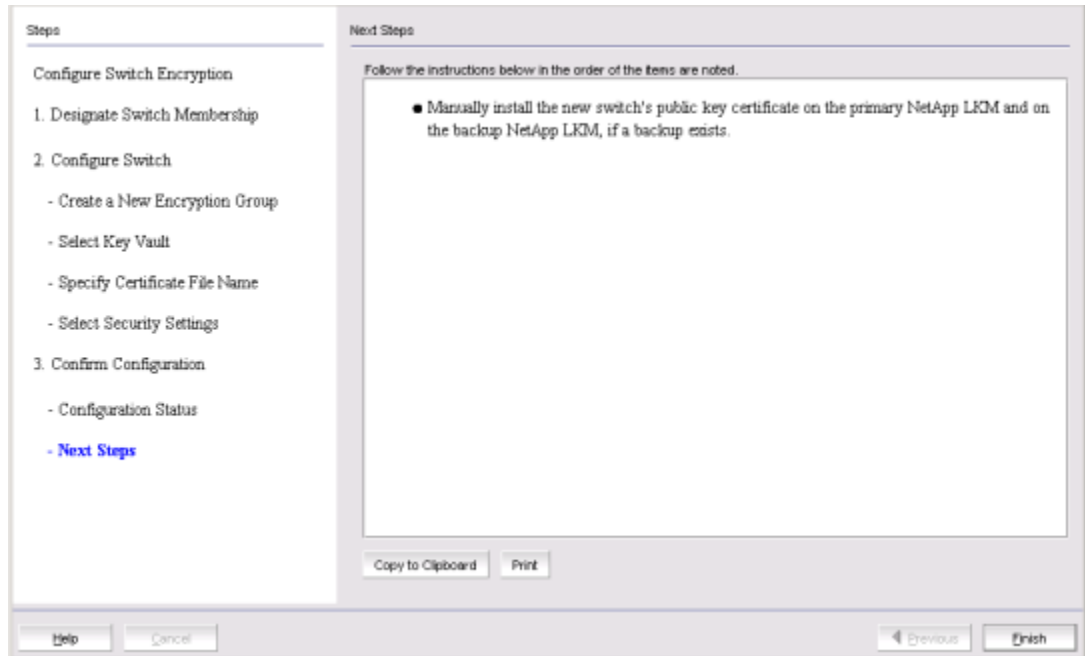


FIGURE 259 Next Steps dialog box

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 669.

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 633.

Figure 260 shows the key vault selection dialog box for ESKM/SKM.

FIGURE 260 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 261](#).)

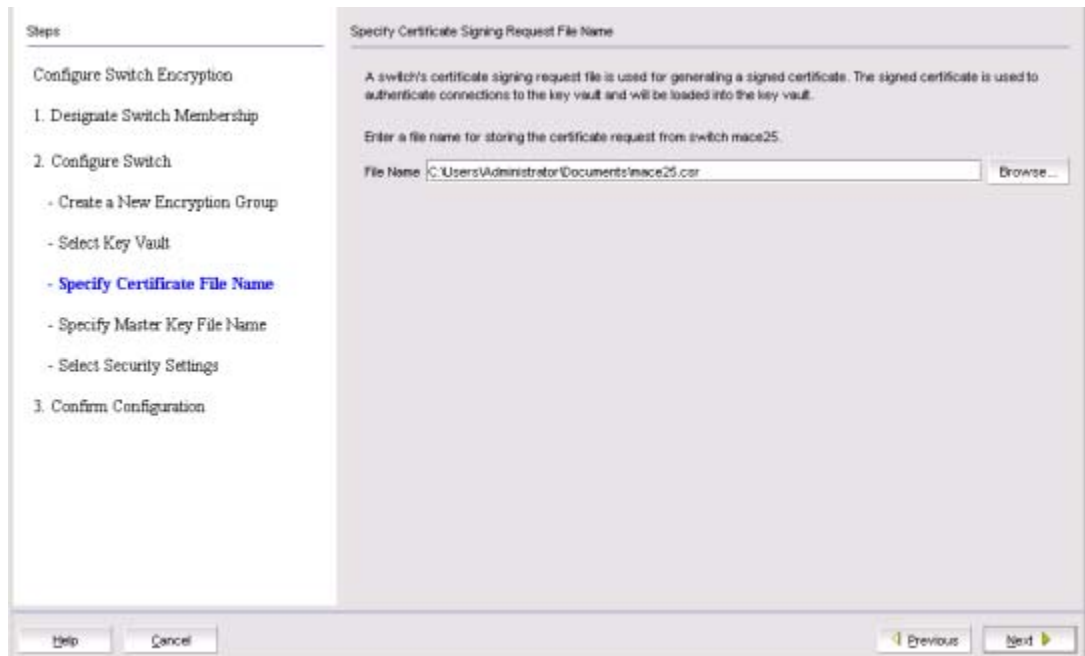


FIGURE 261 Specify Certificate Signing Request File Name dialog box

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 262.](#))

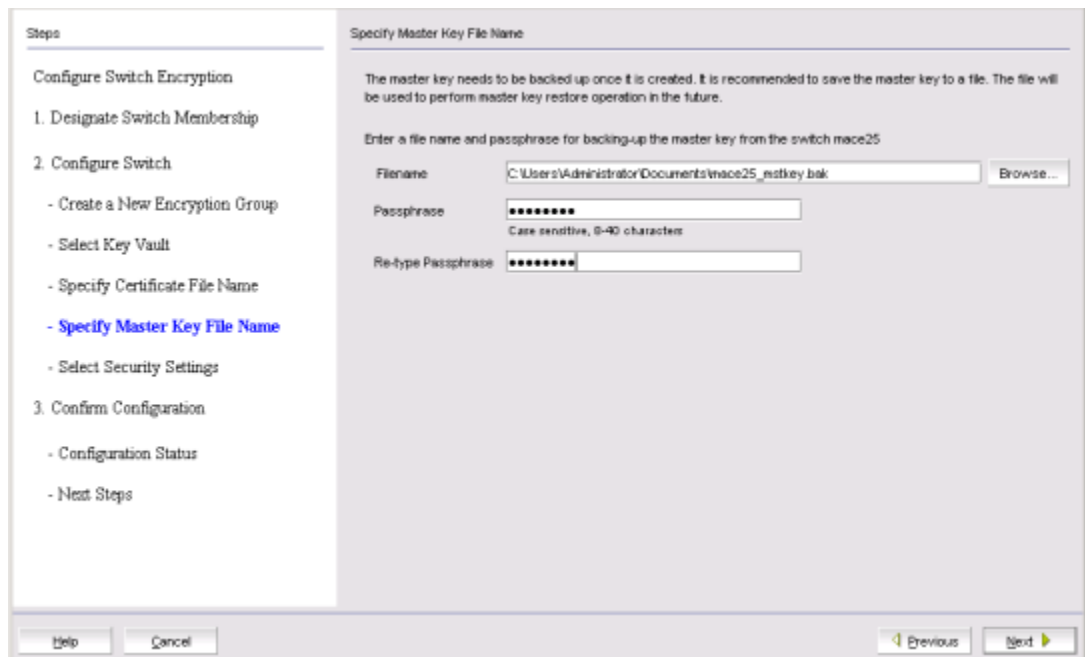


FIGURE 262 Specify Master Key File Name dialog box

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 263](#).)

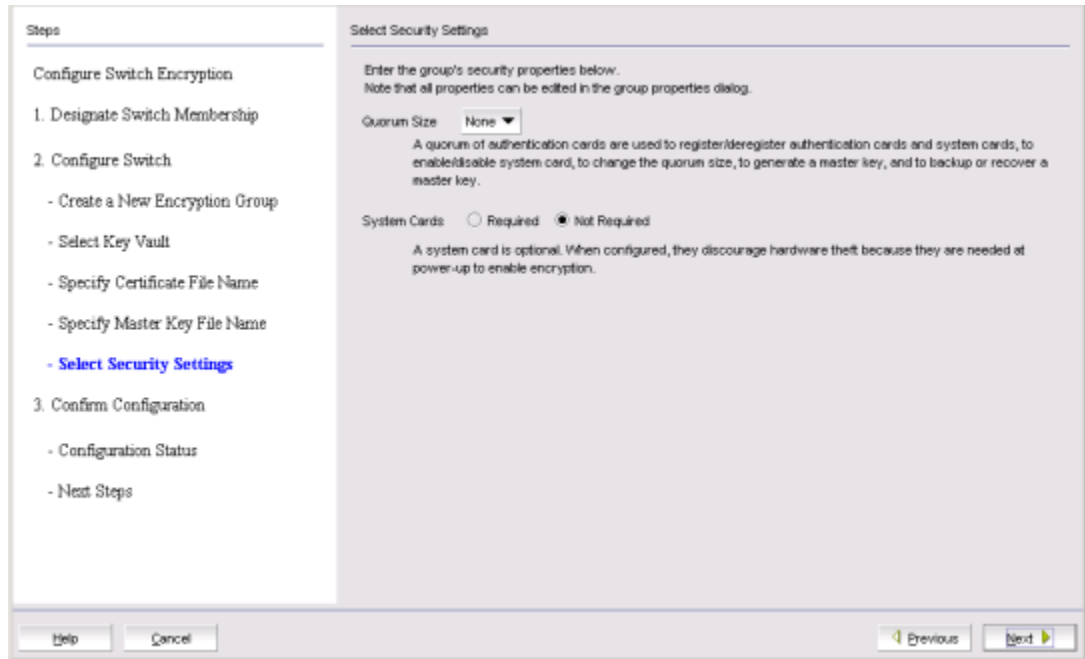


FIGURE 263 Select Security Settings dialog box

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 264](#).) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.




FIGURE 264 Confirm Configuration dialog box

The Configuration Status dialog box displays. (Refer to Figure 265.)




FIGURE 265 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, sends API commands to verify the switch configuration. See [“Understanding configuration status results”](#) on page 669 for more information.

11. Review important messages, then click **Next**.

The **Next Steps** dialog box displays. (Refer to [Figure 266](#).) Instructions for installing public key certificates for the encryption switch are displayed.

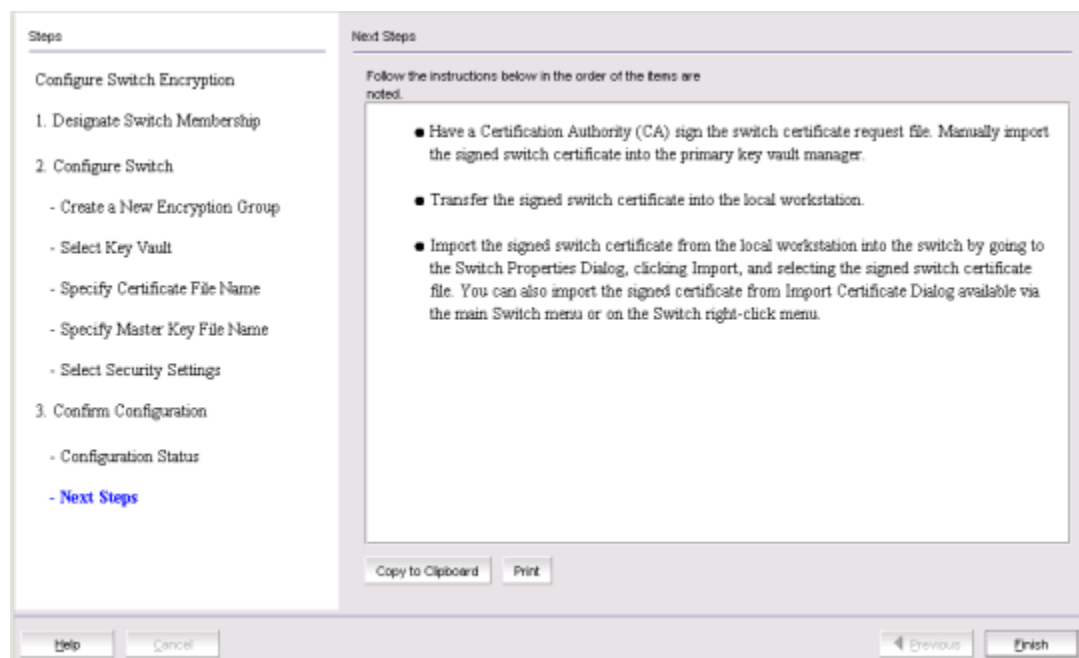


FIGURE 266 Next Steps dialog box

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 669.

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 633.

[Figure 267](#) shows the key vault selection dialog box for TEKA.

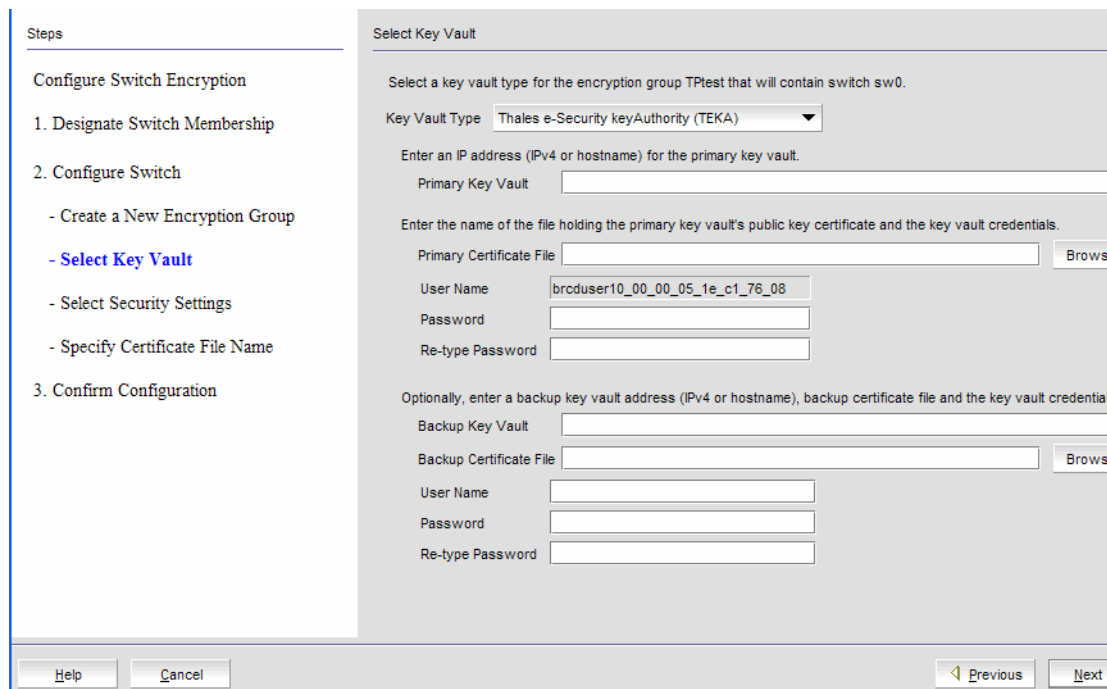


FIGURE 267 Select Key Vault dialog box for TEKA

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 268](#).)

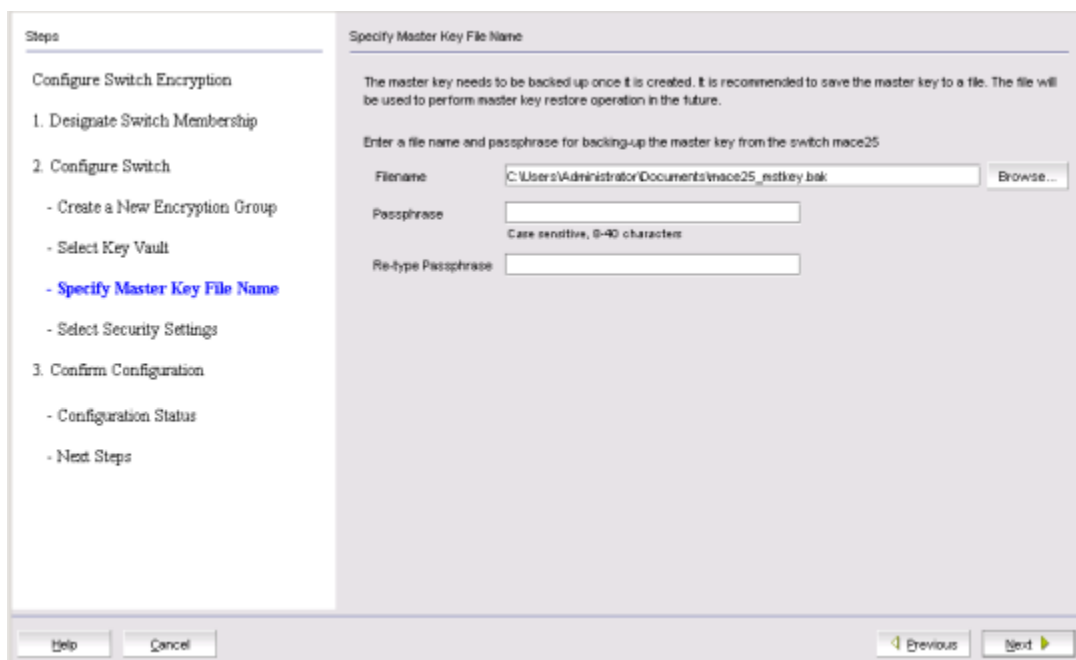


FIGURE 268 Specify Master Key File Name dialog box

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 269](#).)




FIGURE 269 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 270](#).) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

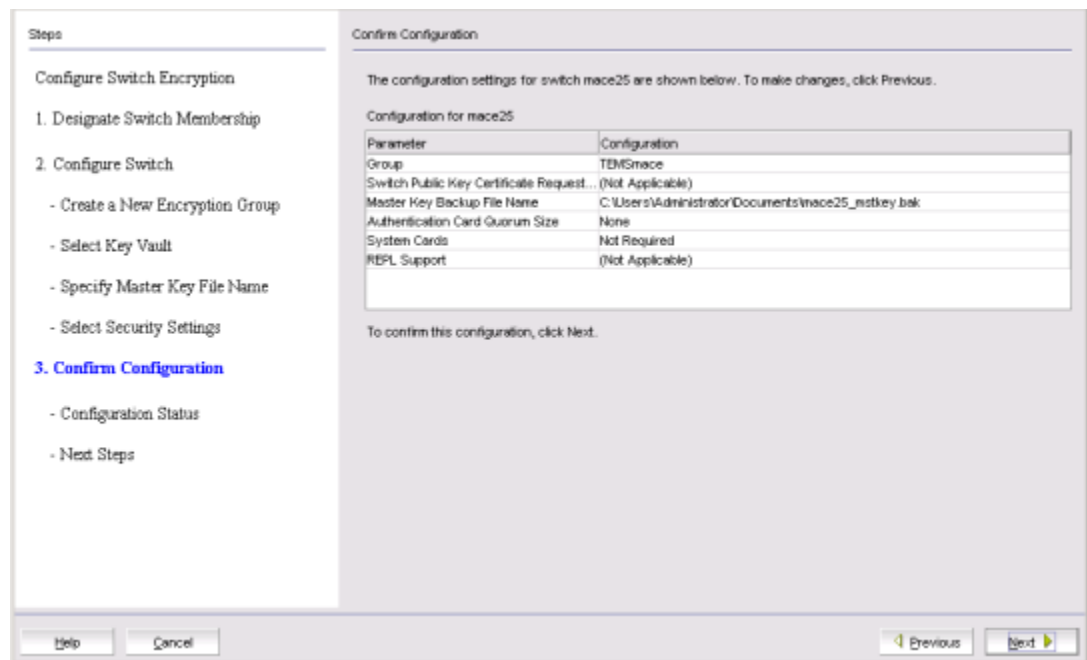


FIGURE 270 Confirm Configuration dialog box

The **Configuration Status** dialog box displays. (Refer to [Figure 271](#).)

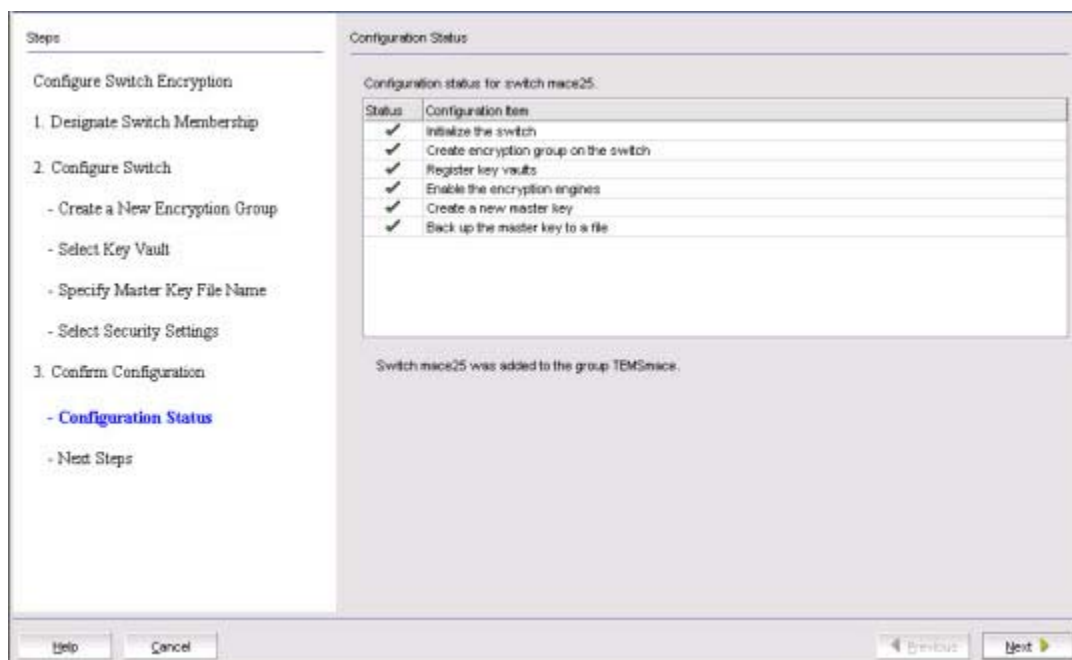


FIGURE 271 Configuration Status dialog box

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, sends API commands to verify the switch configuration. See [“Understanding configuration status results”](#) on page 669 for more information.

11. Click **Next**.

The **Next Steps** dialog box displays. (Refer to [Figure 272](#).) Instructions for installing public key certificates for the encryption switch are displayed.

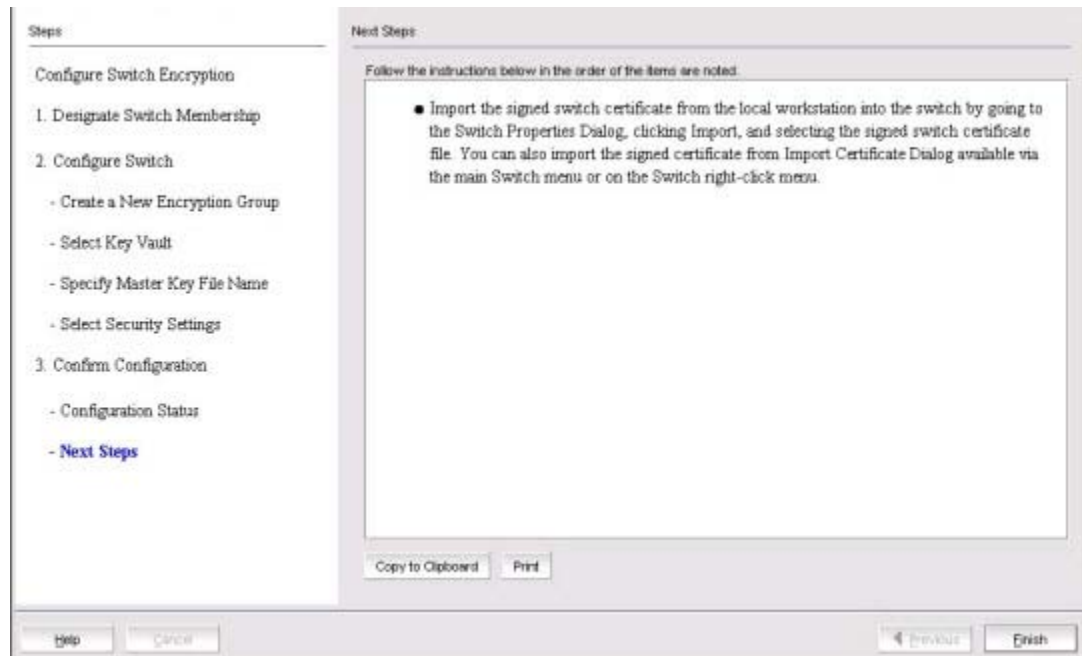


FIGURE 272 Next Steps dialog box

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 669.

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 633.

[Figure 273](#) shows the key vault selection dialog box for TKLM.

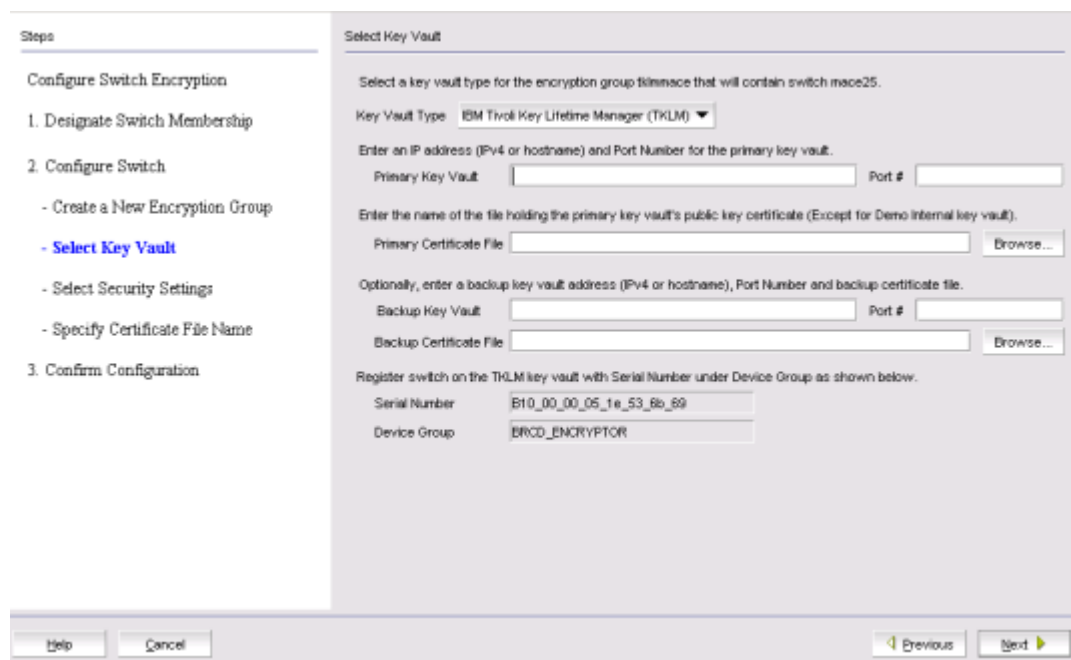


FIGURE 273 Select Key Vault dialog box for TKLM

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 275.](#))

20 Creating a new encryption group

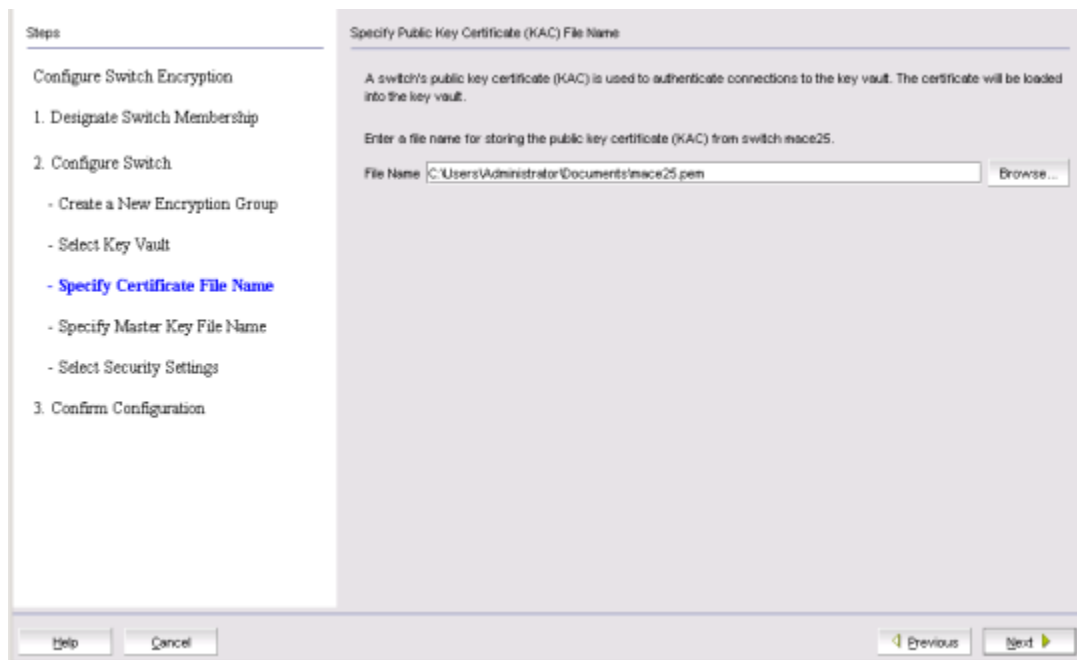


FIGURE 274 Specify Public Key Certificate (KAC) File Name dialog box

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 275](#).)

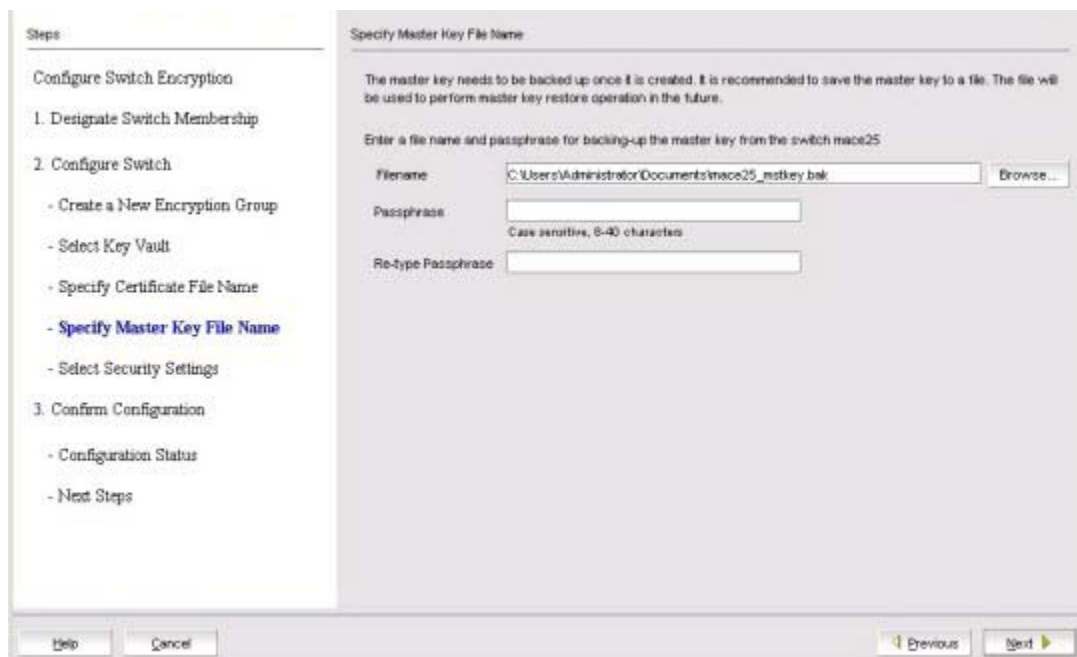


FIGURE 275 Specify Master Key File Name dialog box

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 276](#).)

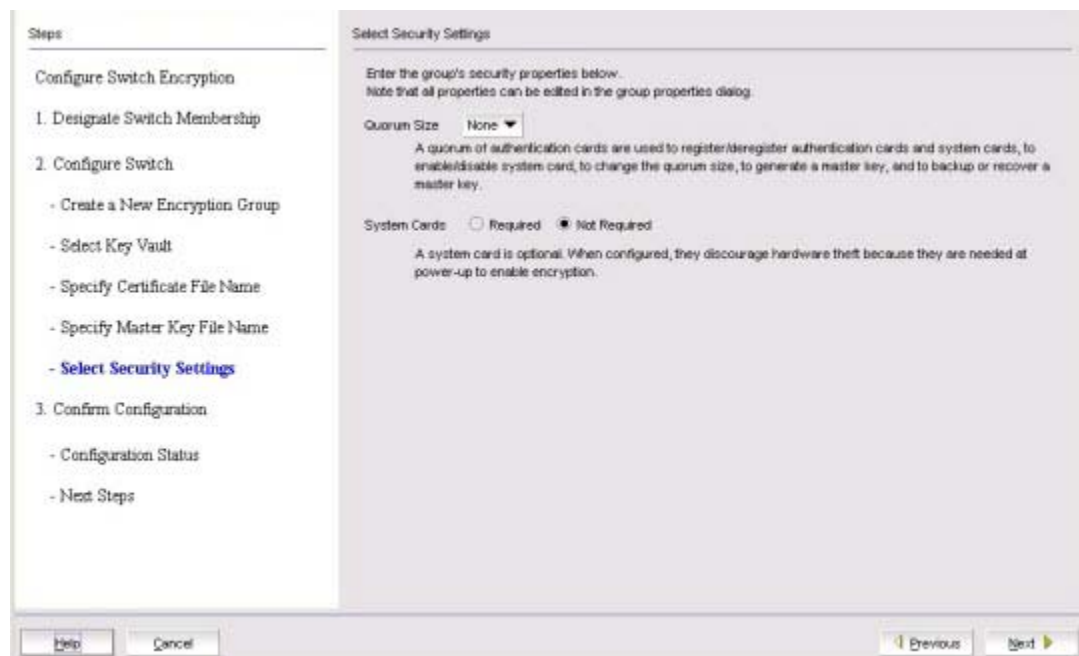


FIGURE 276 Select Security Settings dialog box

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 277](#).) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.




FIGURE 277 Confirm Configuration dialog box

The Configuration Status dialog box displays. (Refer to Figure 278.)

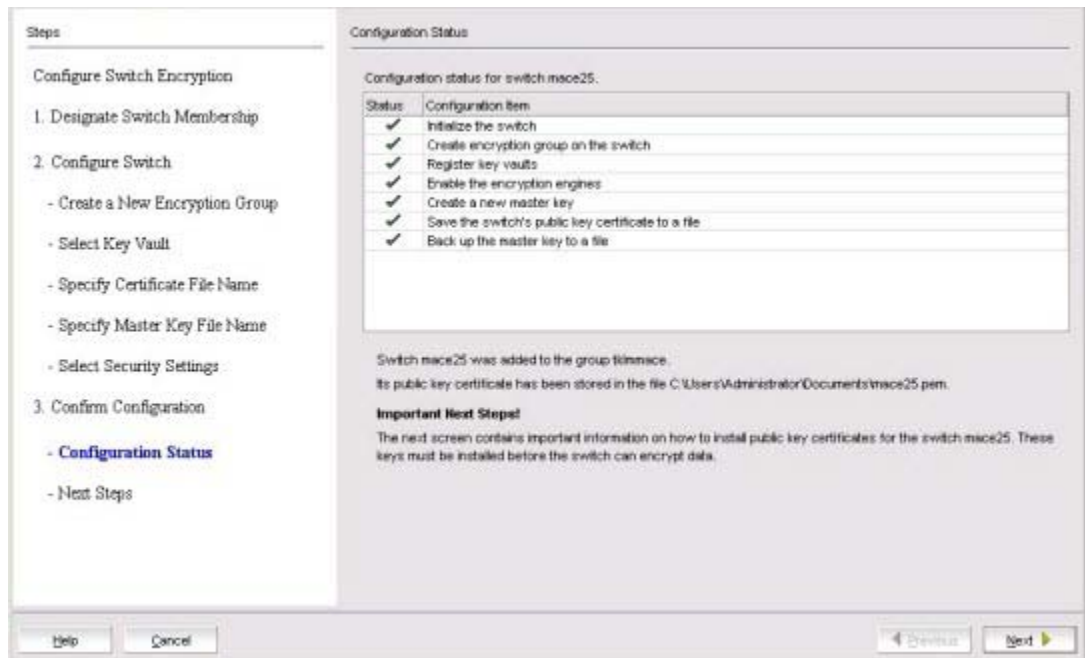


FIGURE 278 Configuration Status dialog box

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, sends API commands to verify the switch configuration.

11. Click **Next**.

The **Next Steps** dialog box displays. (Refer to [Figure 279](#).) Instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

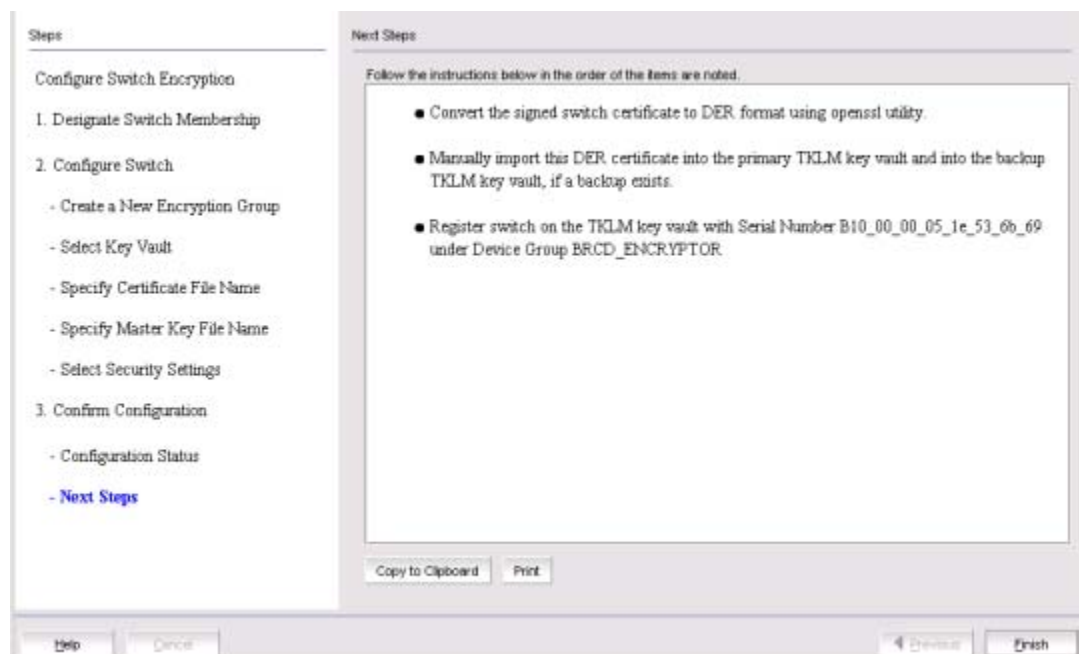


FIGURE 279 Next Steps dialog box

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 669.

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 633.

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 280 shows the key vault selection dialog box for KMIP.

FIGURE 280 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 KAC certificate is signed by a CA, imported back on the and registered as a KAC certificate. The CA will be registered as a key vault certificate on the . If you selected **CA Signed**, the wizard opens the **Specify Public Key Certificate (KAC) File Name** dialog box (Figure 281). 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 and the 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 282.) Go to Step 8.

7. Click **Next**.

The **Specify Public Key Certificate (KAC) File Name** dialog box displays. (Refer to Figure 281.)

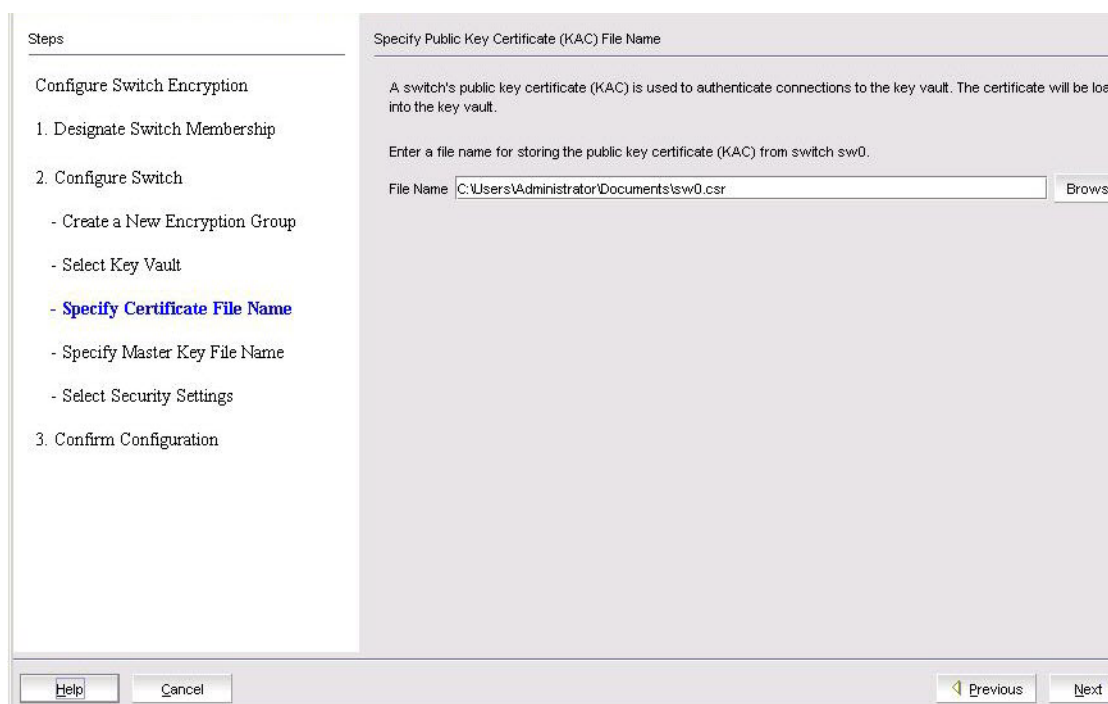


FIGURE 281 Specify Public Key Certificate (KAC) File Name dialog box

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 282.)

20 Creating a new encryption group

Steps

Configure Switch Encryption

1. Designate Switch Membership
2. Configure Switch
 - Create a New Encryption Group
 - Select Key Vault
 - Specify Certificate File Name
 - **Specify Master Key File Name**
 - Select Security Settings
3. Confirm Configuration
 - Configuration Status
 - Next Steps

Specify Master Key File Name

The master key needs to be backed up once it is created. It is recommended to save the master key to a file. The file will be used to perform master key restore operation in the future.

Enter a file name and passphrase for backing-up the master key from the switch sw0

Filename Browse...

Passphrase

Case sensitive, 8-40 characters

Re-type Passphrase

Help Cancel Previous Next

FIGURE 282 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 283](#).)

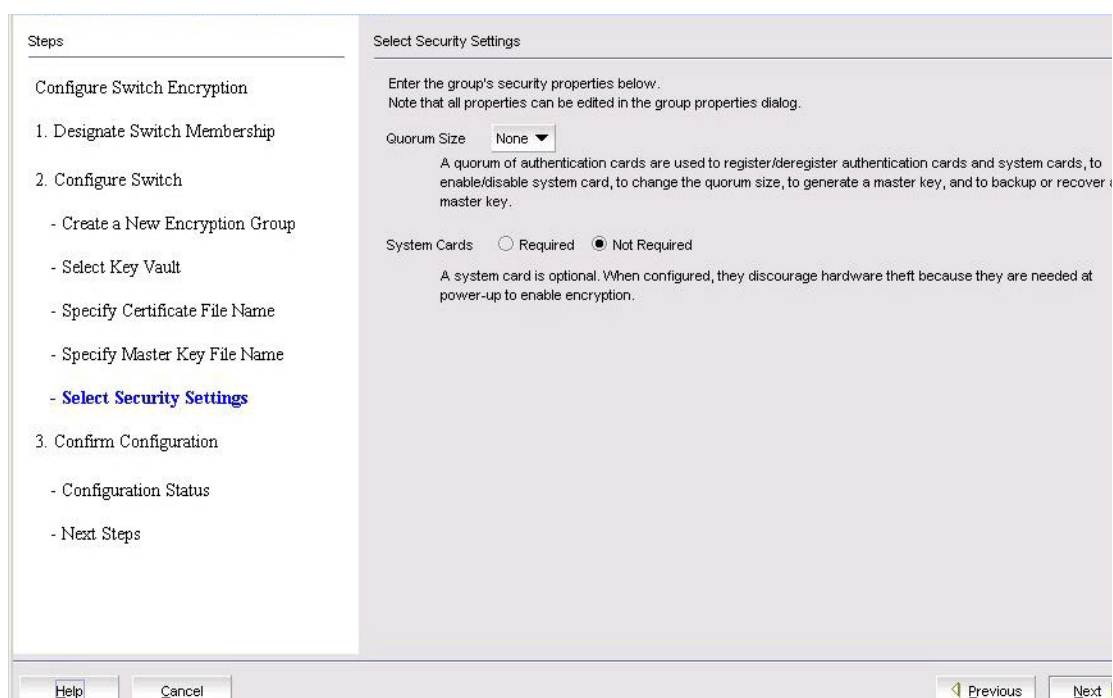


FIGURE 283 Select Security Settings dialog box

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 284](#).)

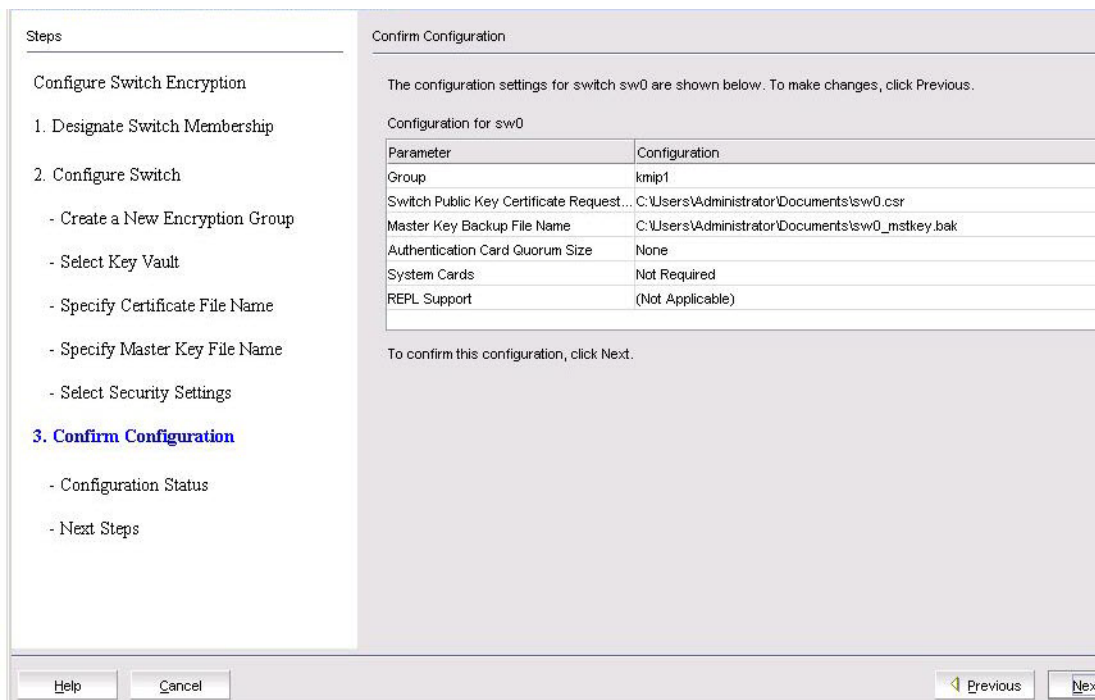


FIGURE 284 Confirm Configuration dialog box

- 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 285](#).)

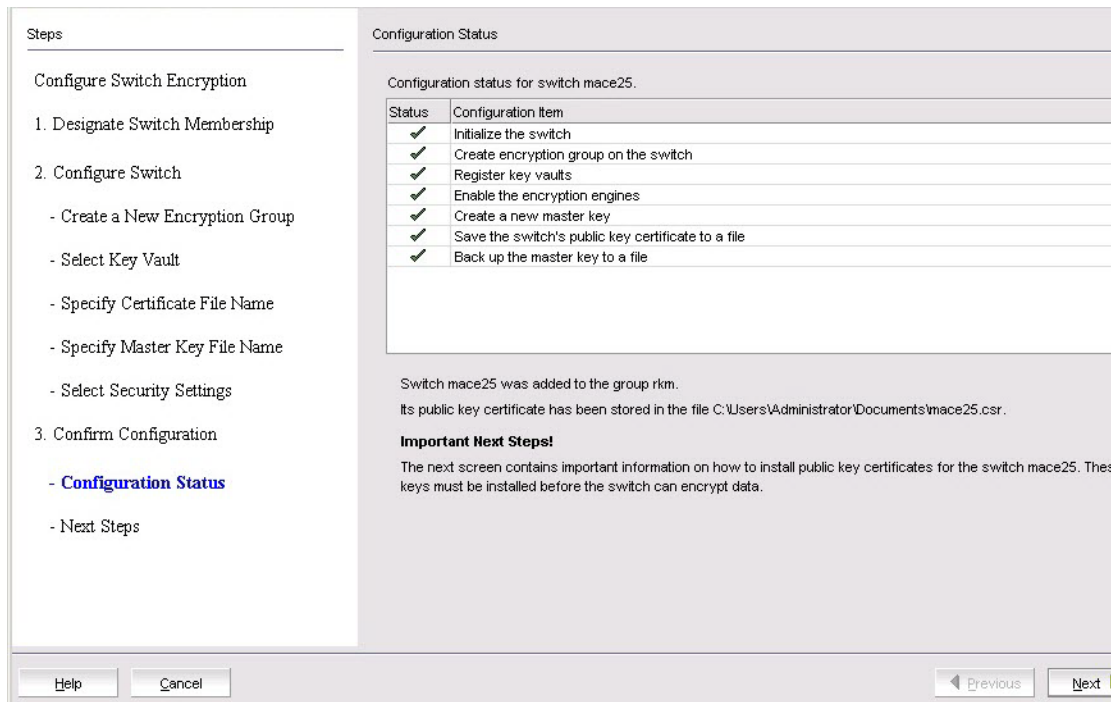


FIGURE 285 Configuration Status dialog box

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, sends API commands to verify the switch configuration.

15. Click **Next**.

The **Next Steps** dialog box displays. (Refer to [Figure 286](#).) Instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

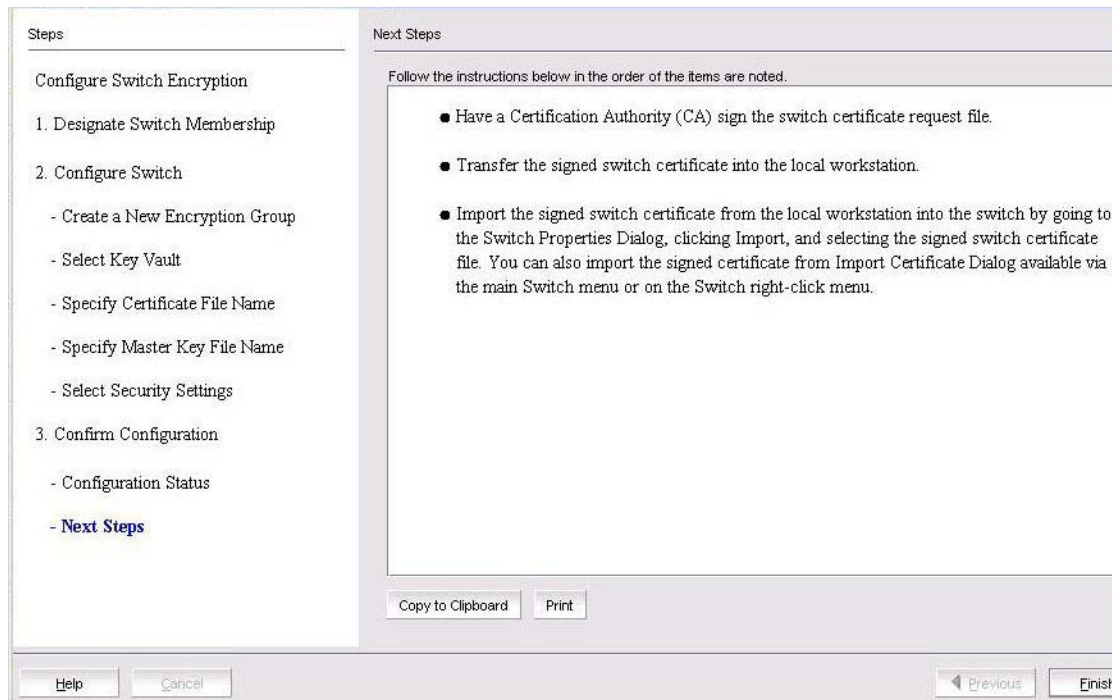


FIGURE 286 Next Steps dialog box

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, 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, performs the **cryptocfg --initnode** command.
2. Create an encryption group on the switch. 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. registers the key vault using the **cryptocfg --reg keyvault** command.
4. Enable the encryption engines. initializes an encryption switch using the **cryptocfg --initEE [<slotnumber>]** and **cryptocfg --regEE [<slotnumber>]** commands.
5. Create a new master key. (Opaque key vaults only). 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. saves the KAC certificate in the specified file.
7. Back up the master key to a file. (Opaque key vaults only). 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 195](#) on page 564.)
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 287](#).)

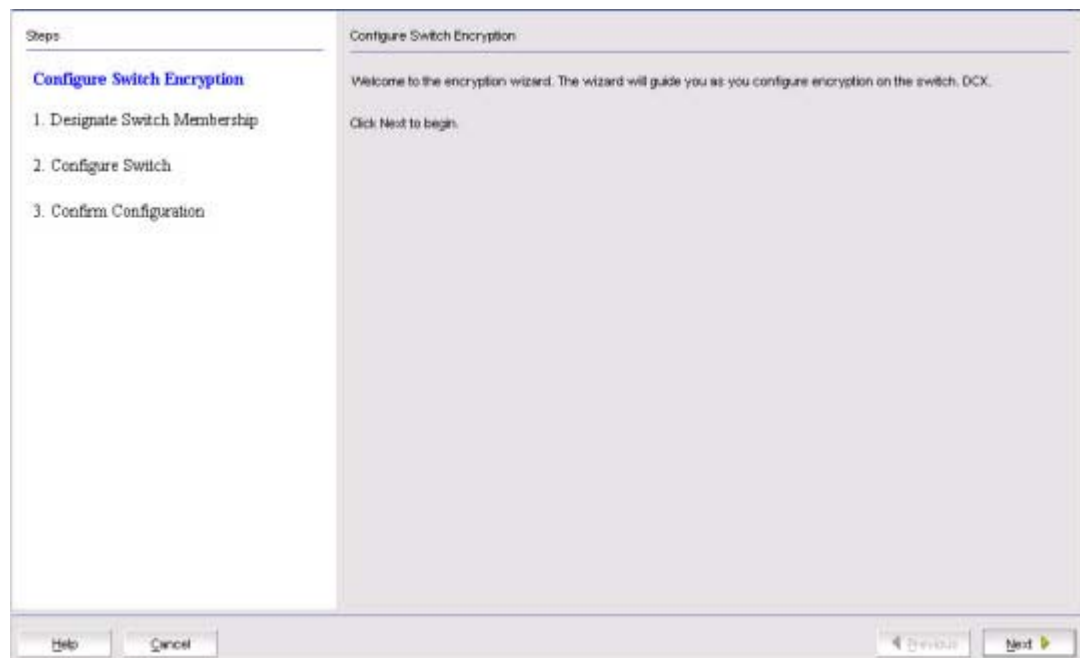


FIGURE 287 Configure Switch Encryption wizard - welcome screen

3. Click **Next**.

The **Designate Switch Membership** dialog box displays. (Refer to [Figure 288](#).)

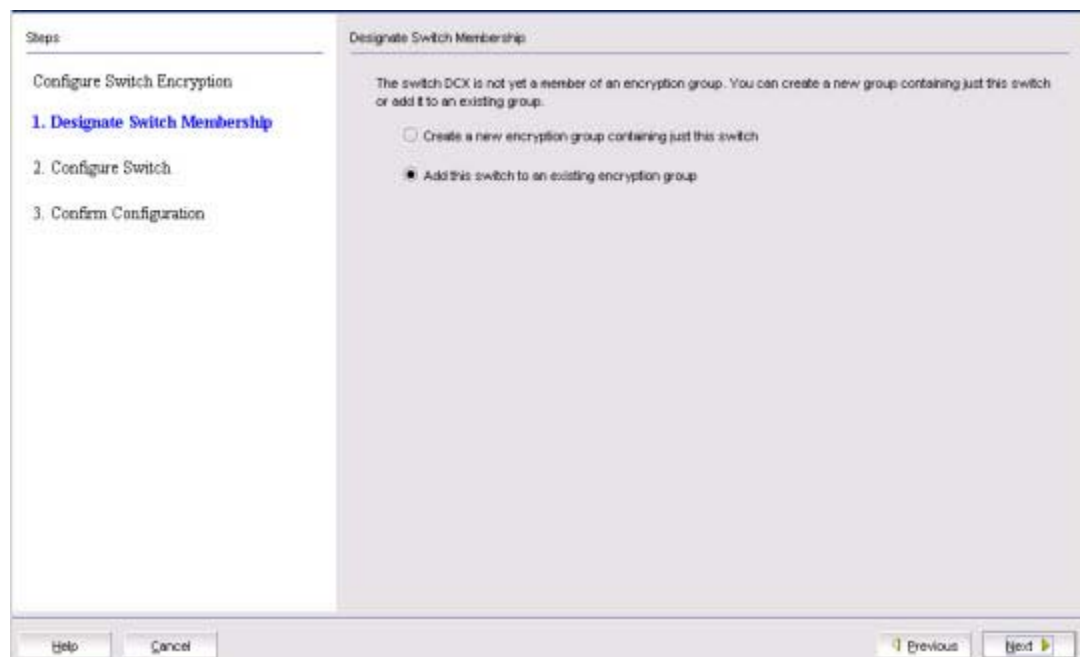


FIGURE 288 Designate Switch Membership dialog box

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 289](#).)

20 Adding a switch to an encryption group

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 633.

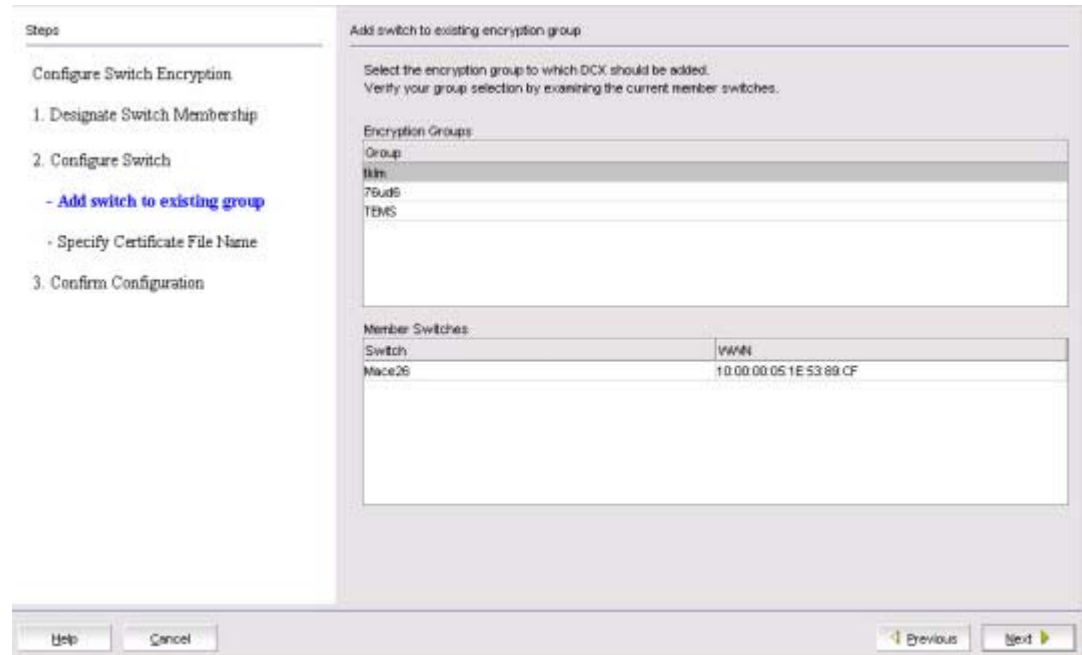


FIGURE 289 Add Switch to Existing Encryption Group dialog box

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 290](#).)

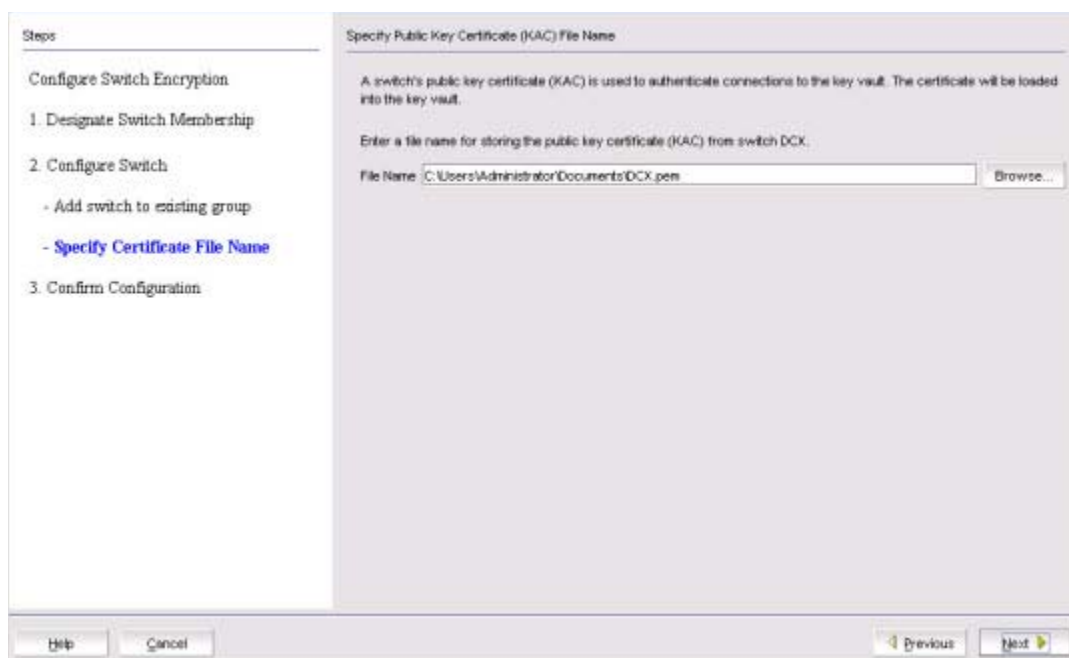


FIGURE 290 Specify Public Key Certificate (KAC) File Name dialog box

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 291](#).) Confirm the encryption group name and switch public key certificate file name you specified are correct, then click **Next**.

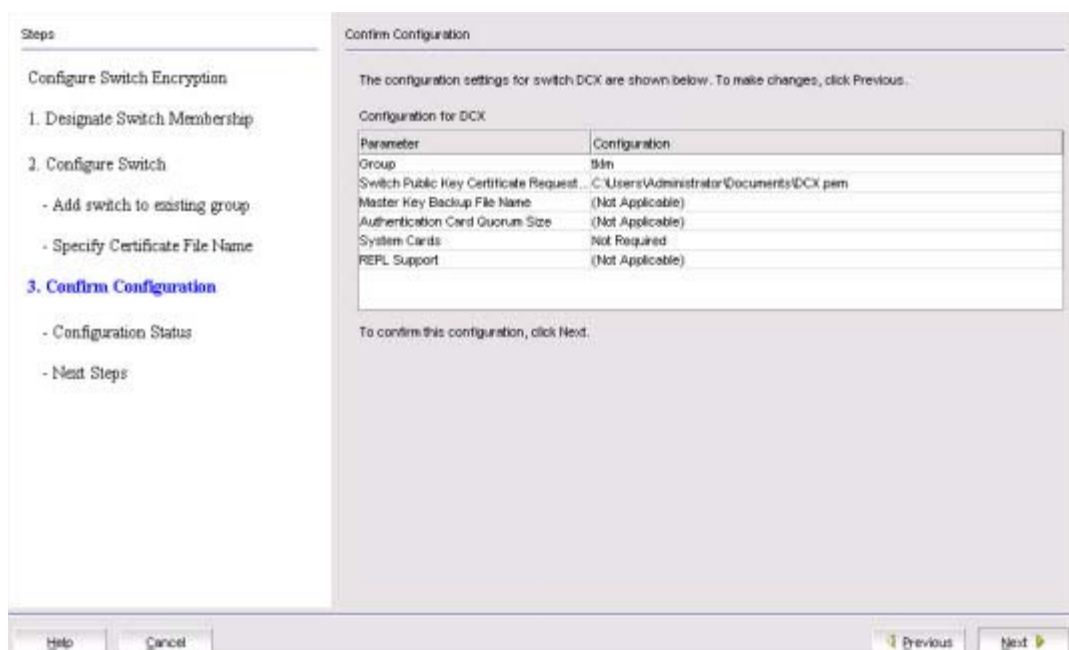


FIGURE 291 Confirm Configuration dialog box

20 Adding a switch to an encryption group

The **Configuration Status** dialog box displays. (Refer to [Figure 292](#).)

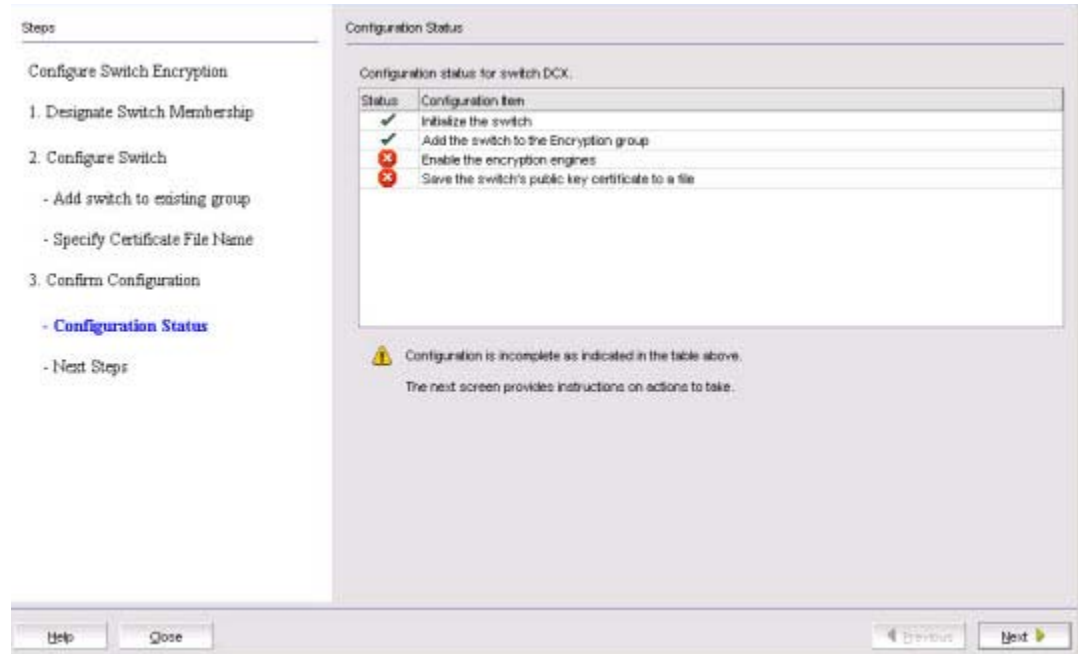


FIGURE 292 Configuration Status dialog box

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 293](#).) Instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

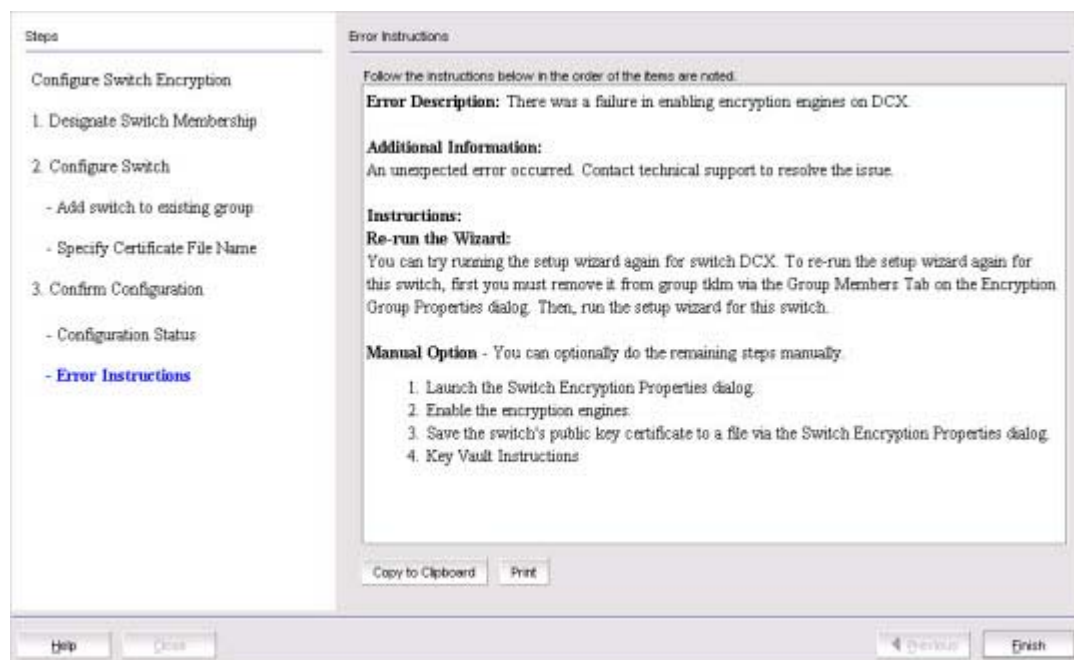


FIGURE 293 Error Instructions dialog box

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 195](#) on page 564.)
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 294](#).)

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 294 Engine 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

In Fabric OS 6.3.0 and later, 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 es, 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 195](#) on page 564.)
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 295](#).)

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 295 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 195](#) on page 564.)
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 295](#).)
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 295](#).)

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 , 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 , complete the following steps:

1. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box. (Refer to [Figure 195](#) on page 564.)
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 295](#).)

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 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 195](#) on page 564.)
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 296](#).)

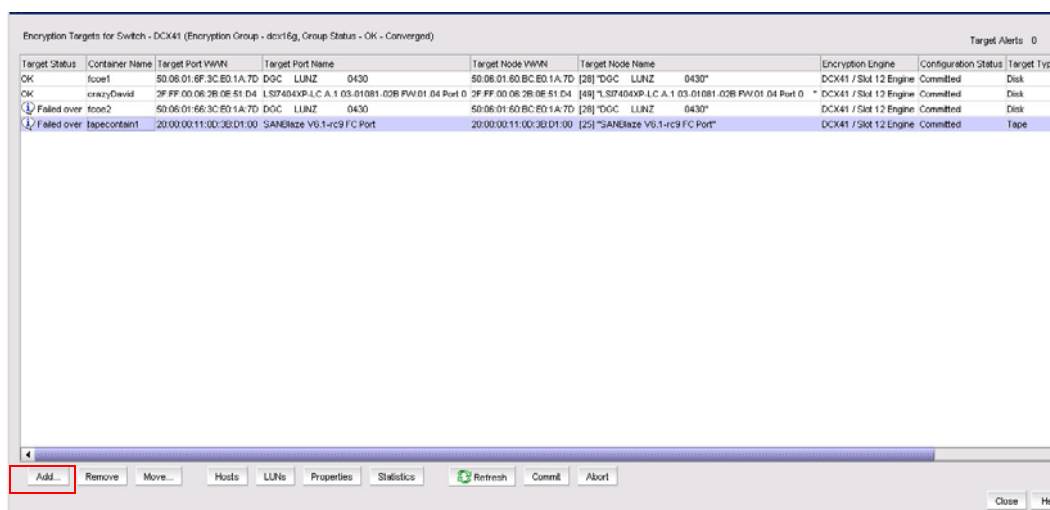


FIGURE 296 Encryption Targets dialog box

3. Click **Add**.

The **Configure Storage Encryption** wizard welcome screen displays. (Refer to [Figure 297](#).)

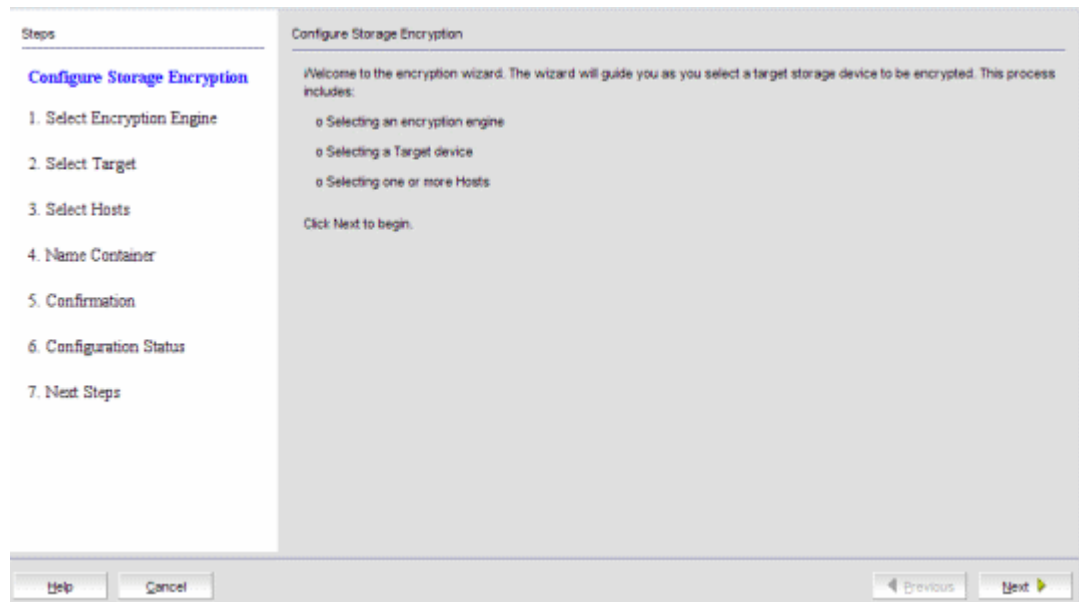


FIGURE 297 Configure Storage Encryption wizard - welcome screen

4. Click **Next**.

The **Select Encryption Engine** dialog box displays. (Refer to [Figure 298](#).)




FIGURE 298 Select 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 299](#).) 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.

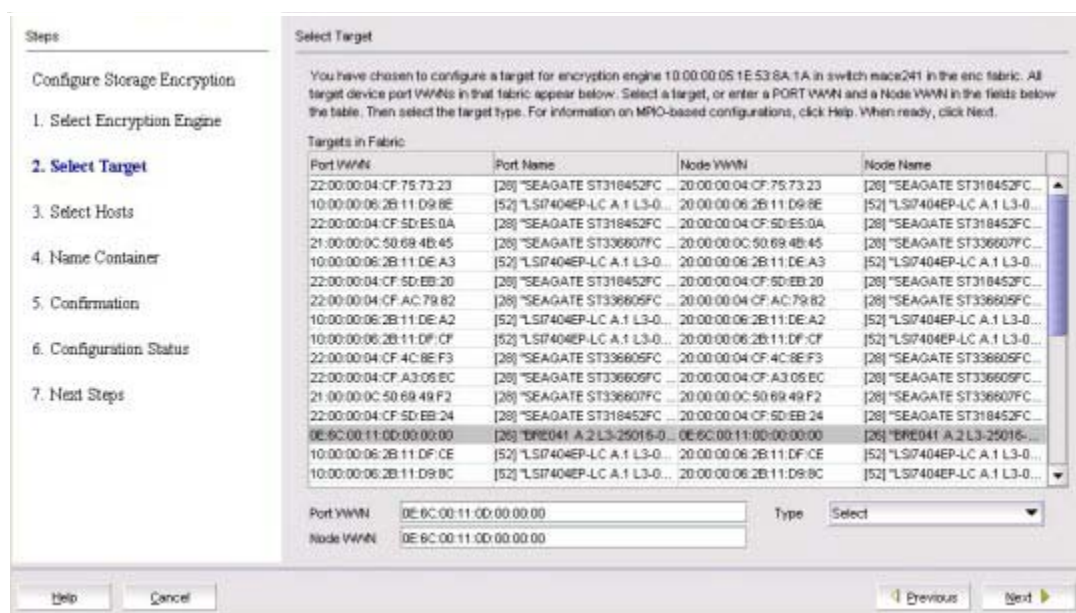


FIGURE 299 Select Target dialog box

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 300](#).) 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.

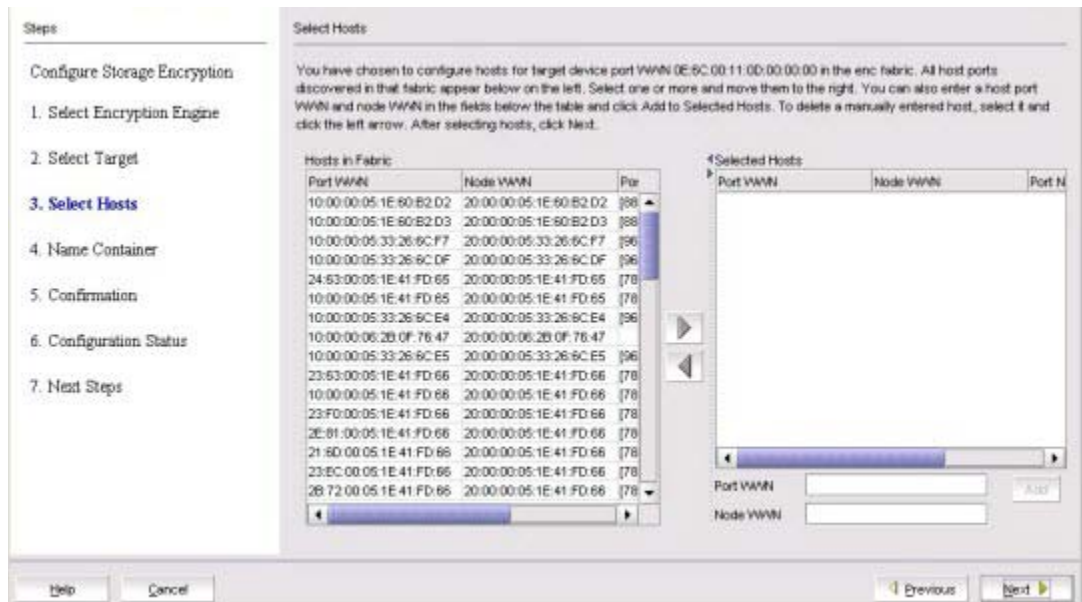


FIGURE 300 Select Hosts dialog box

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 301](#).) 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 301 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 302](#).) The confirmation screen confirms and completes configuration of encryption engines, targets, and hosts.

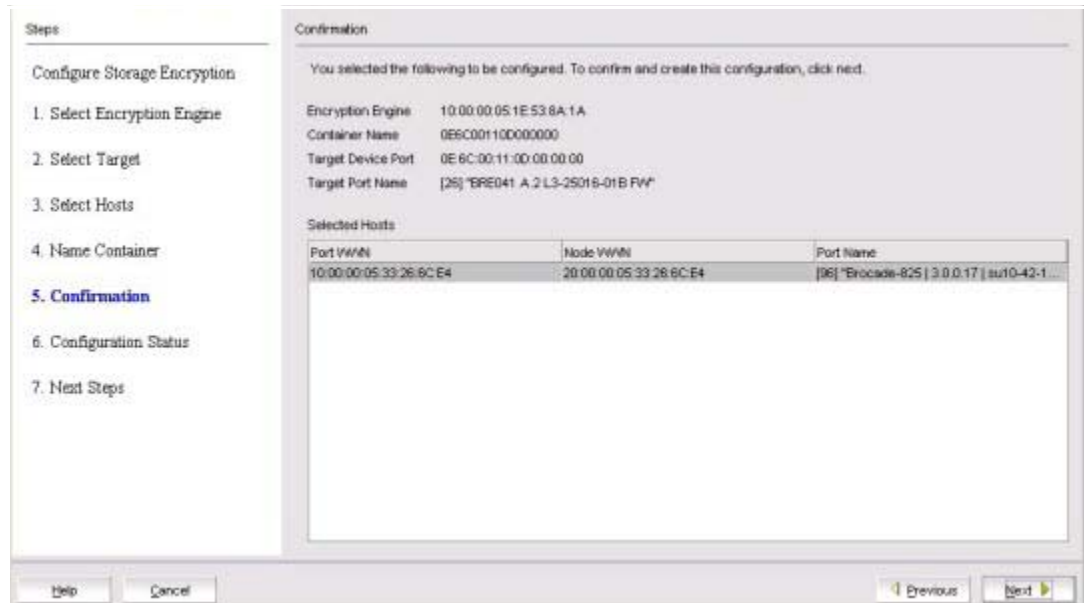


FIGURE 302 Confirmation screen

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 303](#).) 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 303 Configuration 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.

- 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 304](#).) Post-configuration instructions for installing public key certificates for the encryption switch are displayed. These instructions are specific to the key vault type.

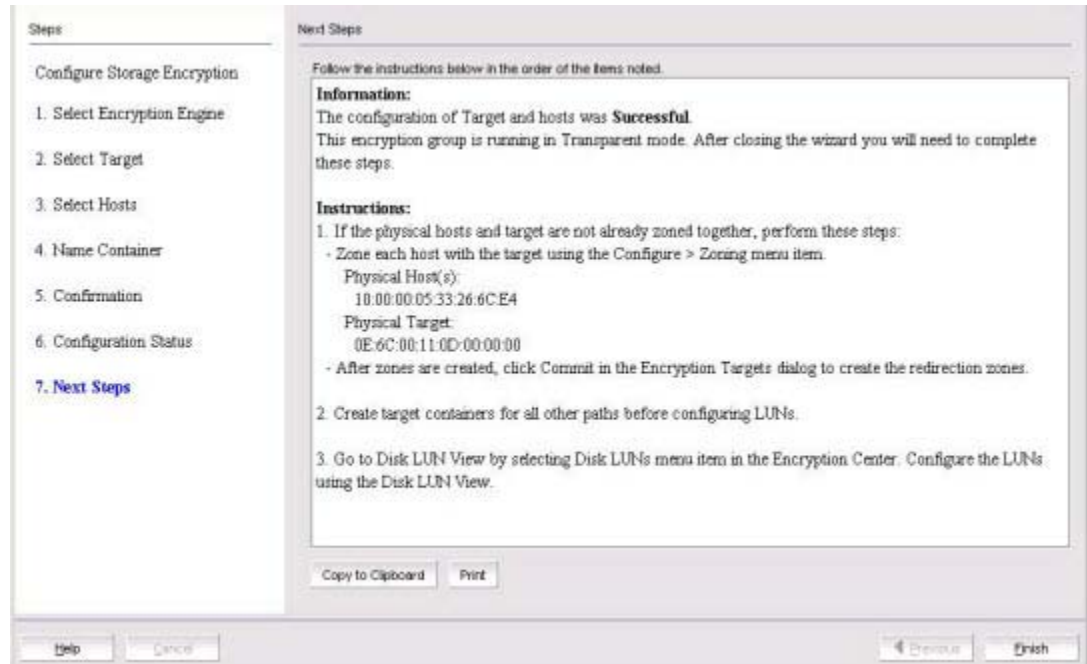


FIGURE 304 Next Steps screen

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.

- 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 195](#) on page 564.)
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 305](#).)

The screenshot shows the 'Encryption Targets for Switch - mso041 (Encryption Group - TENG, Group Status - OK - Converged)' dialog box. It contains a table with the following columns: Target Status, Container Name, Target Port WWN, Target Port Name, Target Type, Target Node WWN, and Target Node Name. The table lists several targets, including disks and tapes, with their respective WWNs and names. At the bottom of the dialog box, there are buttons for 'Add...', 'Remove', 'Move...', 'Hosts', 'LLRP', 'Properties', 'Statistics', 'Refresh', 'Cancel', and 'Abort'. There are also 'Close' and 'Help' buttons at the bottom right.

Target Status	Container Name	Target Port WWN	Target Port Name	Target Type	Target Node WWN	Target Node Name
Offline	500507630000B019	50:05:07:63:00:00:B0:19	IBM_2105750_1.62	Disk	50:05:07:63:00:00:B0:19	[26] IBM_2105750_1.62"
OK	500507630000B019	50:05:07:63:00:00:B0:19	IBM_2105750_1.62	Disk	50:05:07:63:00:00:B0:19	[26] IBM_2105750_1.62"
Offline	220000040F5DEEC1	21:00:00:0C:90:69:4B:29	SEAGATE ST336607FC_0006	Disk	20:00:00:0C:90:69:4B:29	[26] SEAGATE ST336607FC_0006"
Offline	1212001100010001	12:12:00:11:00:01:00:01	BRE041 A.2 L3-25016-01B PW	Disk	12:12:00:11:00:01:00:01	[26] BRE041 A.2 L3-25016-01B PW"
Offline	1212001100010000	12:12:00:11:00:01:00:00	BRE041 A.2 L3-25016-01B PW	Disk	12:12:00:11:00:01:00:00	[26] BRE041 A.2 L3-25016-01B PW"
Offline	11E4001100010002	11:E4:00:11:00:01:00:02	BRE041 A.2 L3-25016-01B PW	Tape	11:E4:00:11:00:01:00:02	[26] BRE041 A.2 L3-25016-01B PW"
Offline	11E4001100010001	11:E4:00:11:00:01:00:01	BRE041 A.2 L3-25016-01B PW	Tape	11:E4:00:11:00:01:00:01	[26] BRE041 A.2 L3-25016-01B PW"
Offline	11E4001100010000	11:E4:00:11:00:01:00:00	BRE041 A.2 L3-25016-01B PW	Tape	11:E4:00:11:00:01:00:00	[26] BRE041 A.2 L3-25016-01B PW"
OK	10E2001100010002	10:E2:00:11:00:01:00:02	BRE041 A.2 L3-25016-01B PW	Disk	10:E2:00:11:00:01:00:02	[26] BRE041 A.2 L3-25016-01B PW"
OK	10E2001100010001	10:E2:00:11:00:01:00:01	BRE041 A.2 L3-25016-01B PW	Disk	10:E2:00:11:00:01:00:01	[26] BRE041 A.2 L3-25016-01B PW"
OK	10E2001100010000	10:E2:00:11:00:01:00:00	BRE041 A.2 L3-25016-01B PW	Disk	10:E2:00:11:00:01:00:00	[26] BRE041 A.2 L3-25016-01B PW"
Offline	100000062B11DFCF	10:00:00:06:2B:11:DF:CF	LSP404EP-LC A.1 L3-01071-0	Disk	20:00:00:06:2B:11:DF:CF	[52] LSP404EP-LC A.1 L3-01071-0"

FIGURE 305 Encryption Targets dialog box

3. Select a target storage device from the list, then click **Hosts**.

The **Encryption Target Hosts** dialog box displays. (Refer to [Figure 306](#).) 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.




FIGURE 306 Encryption Target Hosts dialog box

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.

- **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 308](#).)

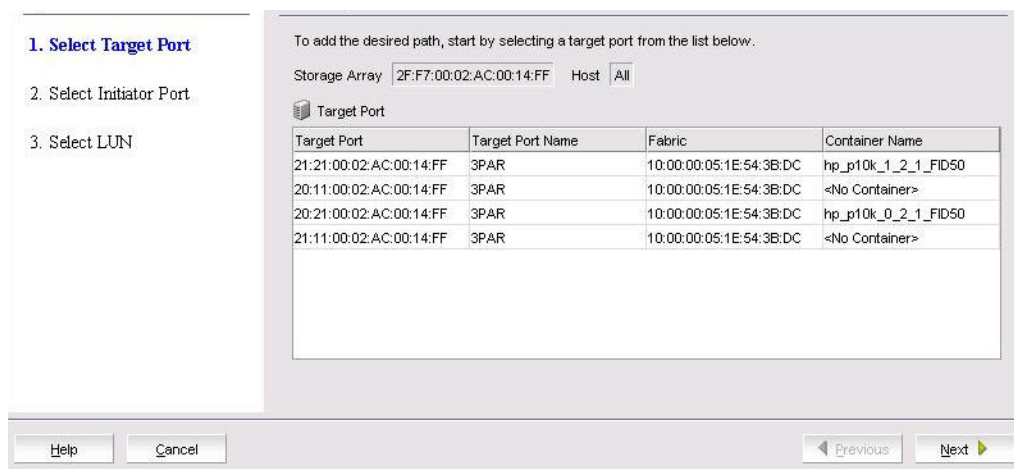


FIGURE 308 Select Target Port dialog box

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 309](#).)

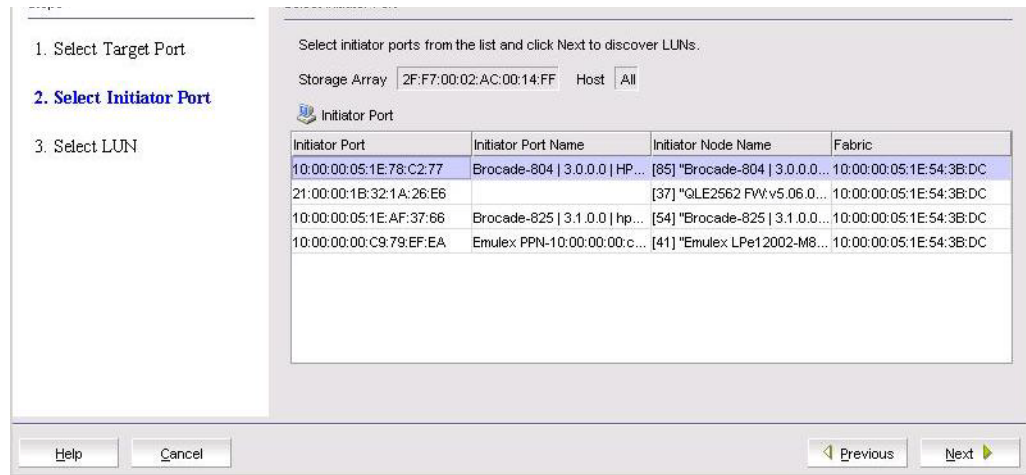


FIGURE 309 Select Initiator Port dialog box

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 310](#).)

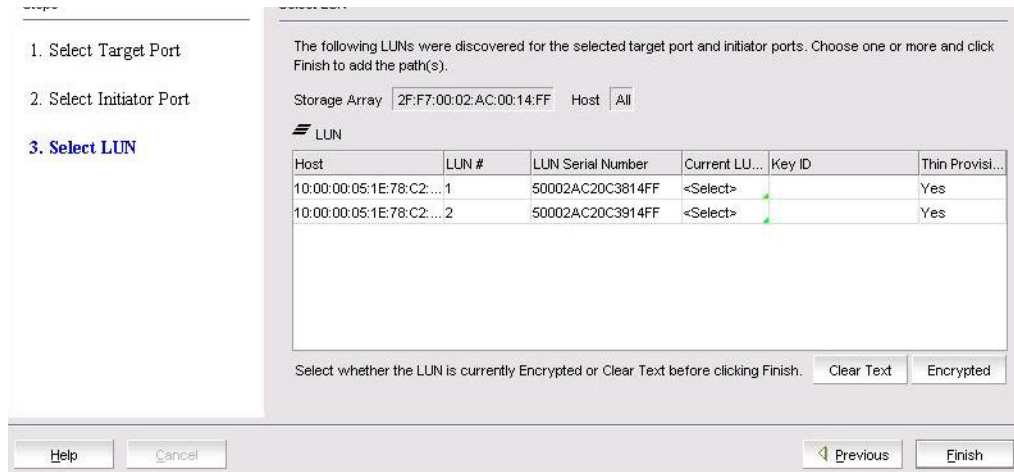


FIGURE 310 Select LUN dialog box

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 **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 311.](#))

to display and configure single-path or multi-path encryption. Start by selecting a storage array to view the current encryption paths configured for LUNs in that array. Click the Add button to create new paths between is. Click OK or Apply button to commit the changes of configuration.
Status: OK - Converged.

#)	Target Port	Initiator Port	Container Name	Switch Name	Fabric	State	Thin Provision LUN	Encryption Mode	Encrypt Exis
20C3B14FF	21:21:00:02:AC:00:14:FF	10:00:00:05:1E:78:C2:77	hp_p10k_1_2_1_FD...	sw062069-FID50	10:00:00:05:1E:54:3B:DC	Encryption enabled	Yes	Native Encryption	Disable
	20:21:00:02:AC:00:14:FF	10:00:00:05:1E:78:C2:77	hp_p10k_0_2_1_FD...	sw062069-FID50	10:00:00:05:1E:54:3B:DC	Encryption enabled	Yes	Native Encryption	Disable
20D1C14FF	21:21:00:02:AC:00:14:FF	10:00:00:05:1E:AF:37:66	hp_p10k_1_2_1_FD...	sw062069-FID50	10:00:00:05:1E:54:3B:DC	Encryption enabled	No	Native Encryption	Mismatch
	20:21:00:02:AC:00:14:FF	10:00:00:05:1E:AF:37:66	hp_p10k_0_2_1_FD...	sw062069-FID50	10:00:00:05:1E:54:3B:DC	Encryption enabled	No	Native Encryption	Enable
FFFFF14FF	21:21:00:02:AC:00:14:FF	10:00:00:05:1E:78:C2:77	hp_p10k_1_2_1_FD...	sw062069-FID50	10:00:00:05:1E:54:3B:DC	Clear text	Unknown	Mismatch	Mismatch
	20:21:00:02:AC:00:14:FF	10:00:00:05:1E:78:C2:77	hp_p10k_0_2_1_FD...	sw062069-FID50	10:00:00:05:1E:54:3B:DC	Disabled (Wrong device type found)	Unknown	Clear Text	Not Applicab
20C3A14FF	21:21:00:02:AC:00:14:FF	10:00:00:05:1E:78:C2:77	hp_p10k_1_2_1_FD...	sw062069-FID50	10:00:00:05:1E:54:3B:DC	Disabled (Found metadata while L...	Unknown	Native Encryption	Enable
	20:21:00:02:AC:00:14:FF	10:00:00:05:1E:78:C2:77	hp_p10k_0_2_1_FD...	sw062069-FID50	10:00:00:05:1E:54:3B:DC	Disabled (Found metadata while L...	Unknown	Mismatch	Mismatch

FIGURE 311 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 312 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 312, 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.

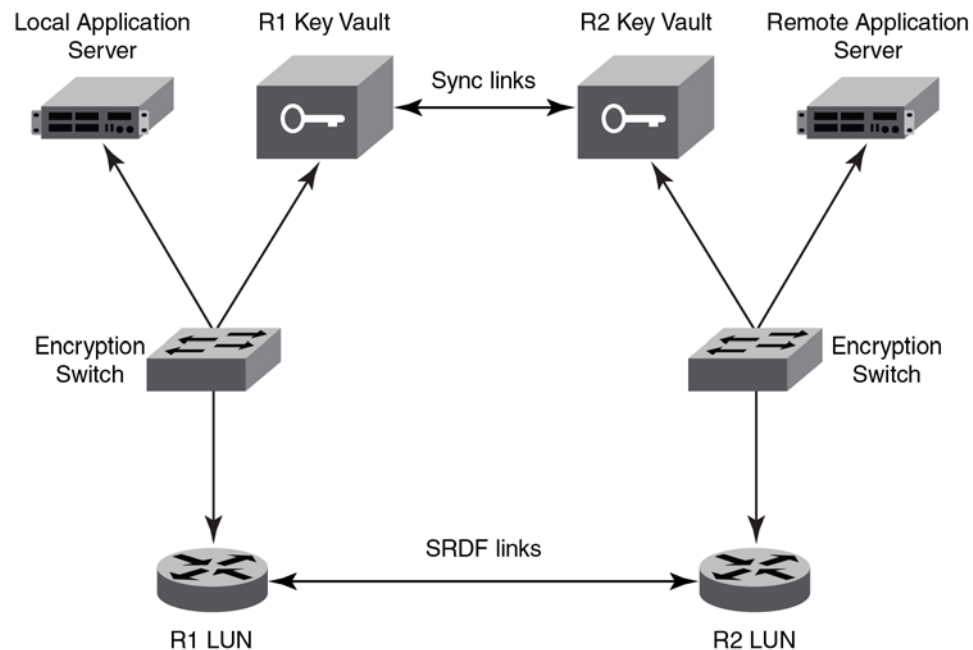


FIGURE 312 Basic SRDF configuration with encryption switches

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 312).

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 195](#) on page 564.)
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 313](#).) Initially, the table is empty. You must add LUNs manually.

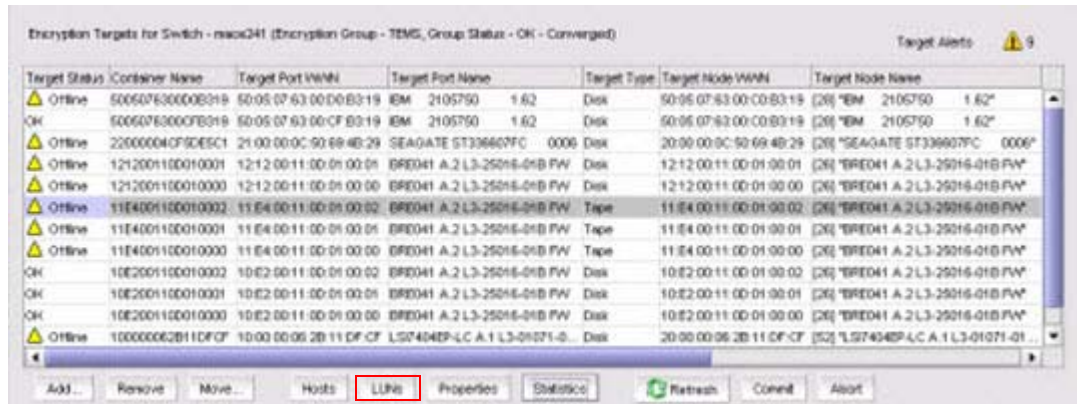


FIGURE 313 Encryption Targets dialog box

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 314](#).)



FIGURE 314 Encryption Target Tape LUNs dialog box

4. Click **Add**.

The **Add Encryption Target Tape LUNs** dialog box displays. (Refer to [Figure 315](#).) 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 315 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 by 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 .

1. Configure Host Port A and Target Port A in the same zone by selecting **Configure > Zoning** from the main menu.
2. Configure Host Port B and Target Port B in the same zone by selecting **Configure > Zoning** from the main menu.
3. Select **Configure > Encryption** from the menu task bar to display the **Encryption Center** dialog box (Refer to [Figure 195](#) on page 564).
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, 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.

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 195](#) on page 564.)
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 316](#).)

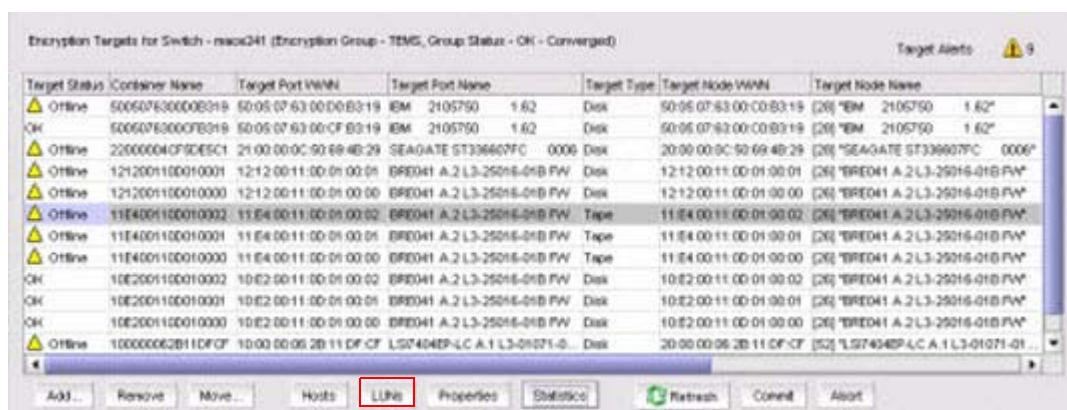


FIGURE 316 Encryption Targets dialog box

3. Select a target tape storage device from the table, then click **LUNs**.

The **Encryption Target Tape LUNs** dialog box displays. (Refer to [Figure 317](#).)

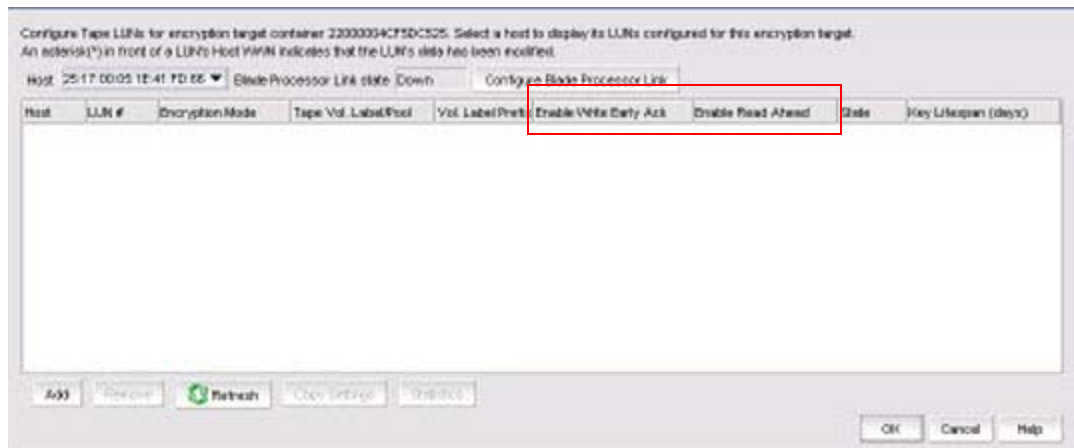


FIGURE 317 Encryption Target Tape LUNs dialog box - Setting tape LUN read ahead and write early

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 195](#) on page 564.)
 - 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 195](#) on page 564.)
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 318](#).) A list of the configured CryptoTarget containers is displayed.

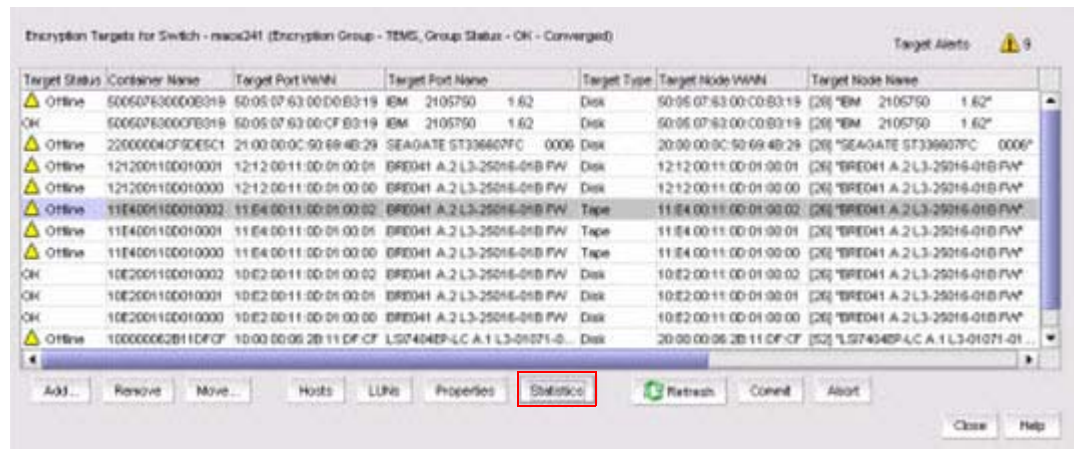


FIGURE 318 Encryption Targets dialog box

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 319](#).) A list of the statistics for all LUNs that are members of the selected tape container is displayed.

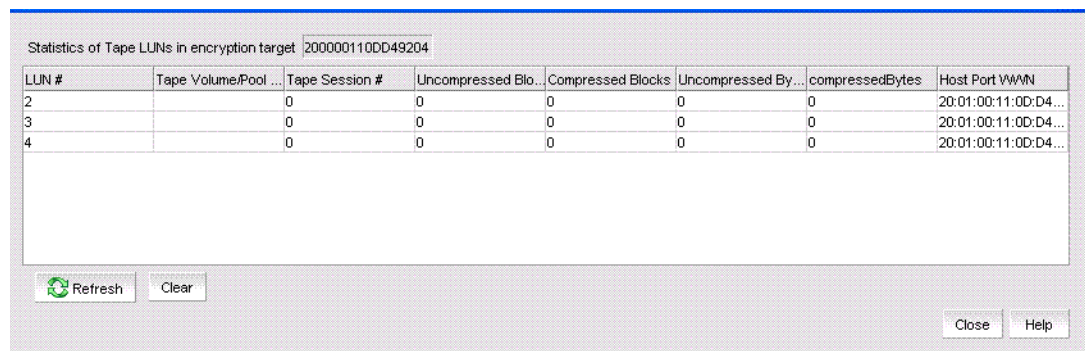


FIGURE 319 Tape LUN Statistics dialog box

The dialog box contains the following information:

- **LUN #:** The number of the logical unit for which statistics 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 195](#) on page 564.)
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 316](#).)

3. Select a tape target storage device, then click **LUNs**.

The **Target Tape LUNs** dialog box displays. (Refer to [Figure 320](#).) A list of the configured tape LUNs is displayed.

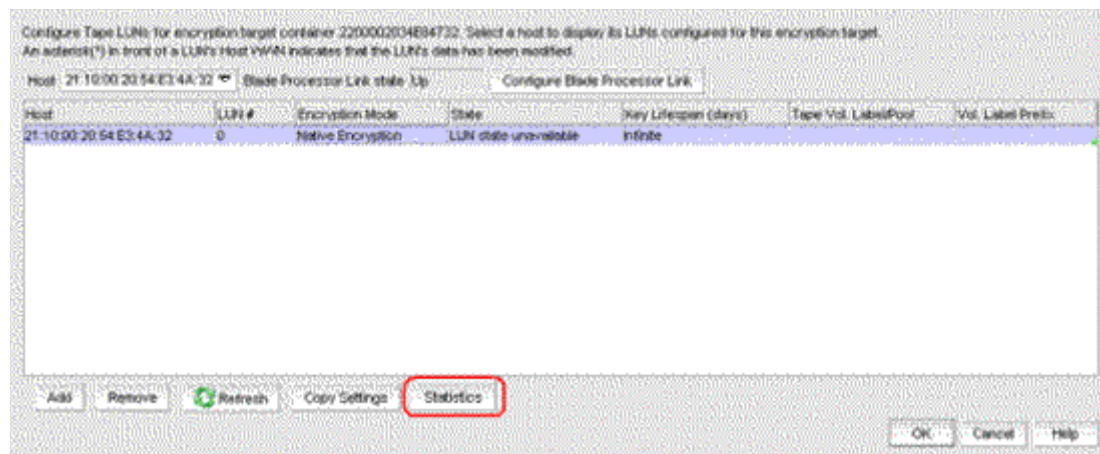


FIGURE 320 Target Tape LUNs dialog box

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 321](#).) The statistic results based on the LUN or LUNs you selected is displayed. Tape LUN statistics are cumulative.

LUN #	Tape Volume/Pool Label	Tape Session #	Uncompressed Blocks	Compressed Blocks	Uncompressed Bytes	Compressed Bytes	Host Port WWN
6		0	0	0	0	0	10:00:00:05:33:2...

FIGURE 321 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 195](#) on page 564.)
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 322](#).) A list of configured CryptoTarget containers is displayed.

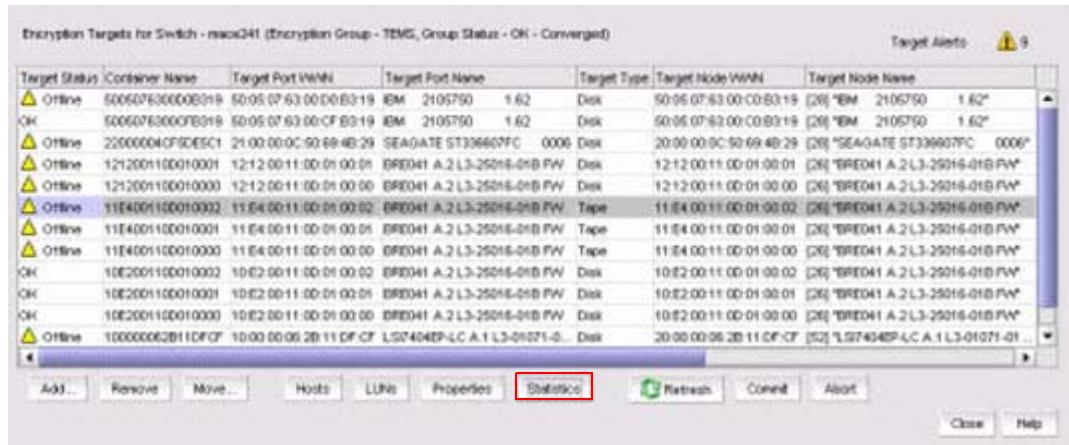


FIGURE 322 Encryption Targets dialog box

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 323](#).) The statistics for all LUNs that are members of the selected tape container are displayed.

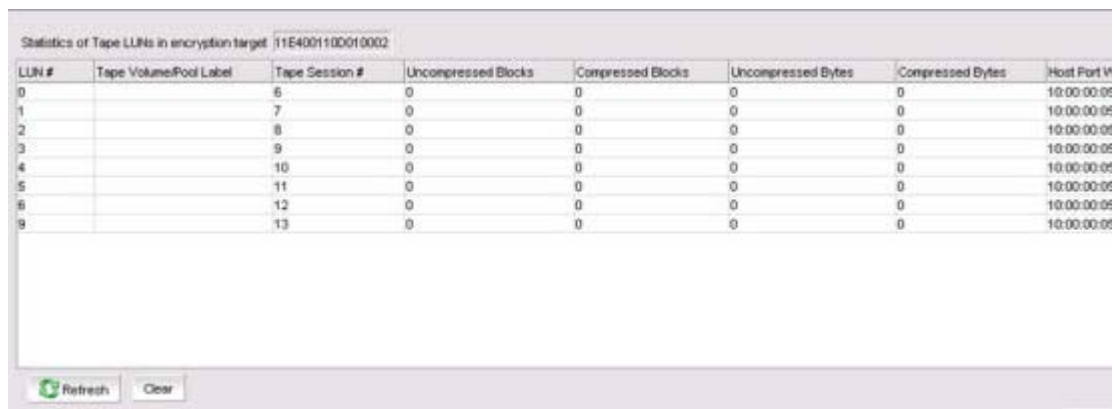


FIGURE 323 Tape LUN Statistics dialog box

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 195](#) on page 564.)
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 722.

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 718 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 .

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 712
- [“Saving a master key to a key vault”](#) on page 713
- [“Saving a master key to a smart card set”](#) on page 714

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 715
- [“Restoring a master key from a key vault”](#) on page 716
- [“Restoring a master key from a smart card set”](#) on page 717

- **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 718.

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 195](#) on page 564.)
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. (Refer to [Figure 324](#).)

FIGURE 324 Master Key Backup dialog box - Backup Destination to file

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 195](#) on page 564.)
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. (Refer to [Figure 325](#).)

FIGURE 325 Backup Master Key for Encryption Group dialog box - Backup Destination to key vault

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 195](#) on page 564.)
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. (Refer to [Figure 326](#).)

FIGURE 326 Backup Master Key for Encryption Group dialog box - Backup Destination to smart cards

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 195](#) on page 564.)
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. (Refer to [Figure 327](#).)

FIGURE 327 Restore Master Key for Encryption Group dialog box - Restore from file

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 195](#) on page 564.)
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. (Refer to [Figure 328](#).)

Select a Master Key to Restore

Active Master Key - The resulting key will be used for all new data encryption.

Alternate Master Key - The resulting key can be used for reading old tapes.

Restore From: Key Vault

Key ID:

Enter a passphrase to decrypt the master key

Passphrase:

Capitalization matters, 8-40 characters.

OK Cancel Help

FIGURE 328 Restore Master Key for Encryption Group dialog box - Restore from key vault

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 195](#) on page 564.)
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. (Refer to [Figure 329](#).)

FIGURE 329 Restore Master Key for Encryption Group dialog box - Restore from smart cards

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 195](#) on page 564.)
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 195](#) on page 564.)
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 195](#) on page 564.)
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 195](#) on page 564.)
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 330](#).) 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.

Target Status	Container Name	Target Port WWN	Target Port Name	Target Type	Target Node WWN	Target Node Name
Offline	500507630000B019	50:05:07:63:00:00:B0:19	EM_2105750_1.62	Disk	50:05:07:63:00:00:B0:19	[OK] "EM_2105750_1.62"
OK	500507630000B019	50:05:07:63:00:00:CF:B3:19	EM_2105750_1.62	Disk	50:05:07:63:00:00:B0:19	[OK] "EM_2105750_1.62"
Offline	22000004CF9DE5C1	21:00:00:0C:90:69:4B:29	SEAGATE ST336667FC_0006	Disk	20:00:00:0C:50:69:4B:29	[OK] "SEAGATE ST336667FC_0006"
Offline	1212001100010001	12:12:00:11:00:01:00:01	BRE041 A.2 L3-25016-01B FW	Disk	12:12:00:11:00:01:00:01	[OK] "BRE041 A.2 L3-25016-01B FW"
Offline	1212001100010000	12:12:00:11:00:01:00:00	BRE041 A.2 L3-25016-01B FW	Disk	12:12:00:11:00:01:00:00	[OK] "BRE041 A.2 L3-25016-01B FW"
Offline	11E4001100010002	11:E4:00:11:00:01:00:02	BRE041 A.2 L3-25016-01B FW	Tape	11:E4:00:11:00:01:00:02	[OK] "BRE041 A.2 L3-25016-01B FW"
Offline	11E4001100010001	11:E4:00:11:00:01:00:01	BRE041 A.2 L3-25016-01B FW	Tape	11:E4:00:11:00:01:00:01	[OK] "BRE041 A.2 L3-25016-01B FW"
Offline	11E4001100010000	11:E4:00:11:00:01:00:00	BRE041 A.2 L3-25016-01B FW	Tape	11:E4:00:11:00:01:00:00	[OK] "BRE041 A.2 L3-25016-01B FW"
OK	10E2001100010002	10:E2:00:11:00:01:00:02	BRE041 A.2 L3-25016-01B FW	Disk	10:E2:00:11:00:01:00:02	[OK] "BRE041 A.2 L3-25016-01B FW"
OK	10E2001100010001	10:E2:00:11:00:01:00:01	BRE041 A.2 L3-25016-01B FW	Disk	10:E2:00:11:00:01:00:01	[OK] "BRE041 A.2 L3-25016-01B FW"
OK	10E2001100010000	10:E2:00:11:00:01:00:00	BRE041 A.2 L3-25016-01B FW	Disk	10:E2:00:11:00:01:00:00	[OK] "BRE041 A.2 L3-25016-01B FW"
Offline	100000062B11DFCF	10:00:00:06:2B:11:DF:CF	L57404P-LC A.1 L3-01071-B	Disk	20:00:00:06:2B:11:DF:CF	[OK] "L57404P-LC A.1 L3-01071-B"

FIGURE 330 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 .
- 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 331](#).)

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 195](#) on page 564.)
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 318](#) on page 705.)

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 195](#) on page 564.)
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 331](#).)

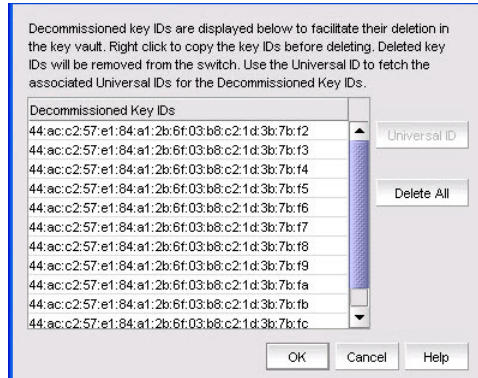


FIGURE 331 Decommissioned Key IDs dialog box

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 725.

- **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 195](#) on page 564.)
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 331](#).)

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 332](#).)




FIGURE 332 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 195](#) on page 564.)
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 333](#).)

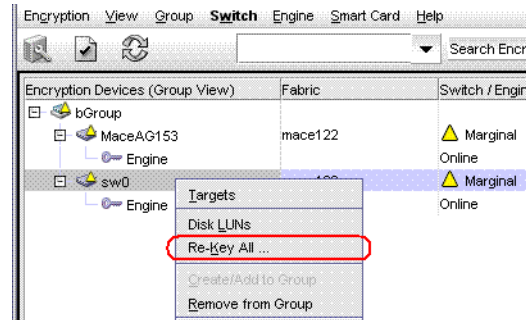


FIGURE 333 Selecting the Re-Key All operation

If REPL support is enabled on the encryption group, a confirmation dialog box displays, asking whether to rekey mirror LUNs.

3. Click **Yes** to include 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 195](#) on page 564.)
 - 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 305](#).)

6. Select a disk LUN device from the table, then click **LUNs**.

The **Encryption Targets Disk LUNs** dialog box displays. (Refer to [Figure 334](#).) The dialog box lists the status of the rekey operation.

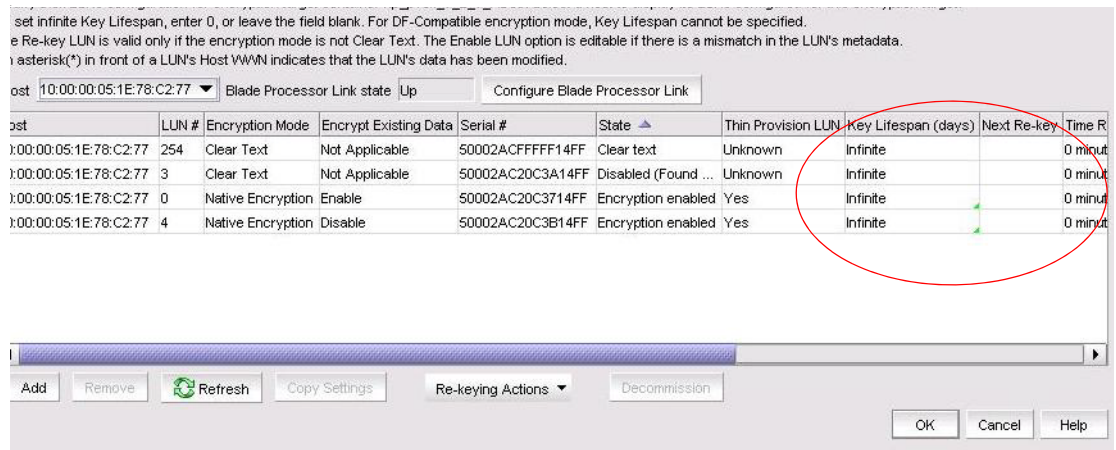


FIGURE 334 Pending manual rekey operations

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 195](#) on page 564.)
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 335](#).) Initially the list is empty. You must add LUNs manually.

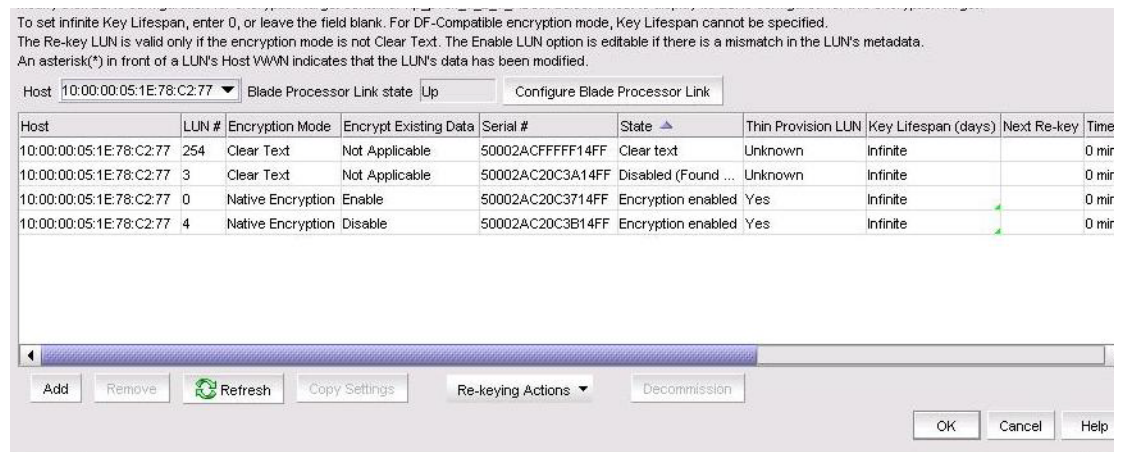


FIGURE 335 Encryption Target Disk LUNs dialog box

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 the **Encryption Center** dialog box. (Refer to [Figure 195](#) on page 564.)
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 336](#).)

Re-Key Sessions Status for Switch - mace_25_test

LUN #	LUN Serial #	Re-Key Session #	Percent Complete	Re-Key State	Re-Key Role	Block Size	Container Name	Host Port WWN	Current LBA	Number of Blocks	Thin Provision LUN
5	600110D01...	1	0	Waiting for ...	Primary/Act...	512	149E00110D0...	10:00:00:05:3...	1	2048	No
5	600110D01...	1	0	Waiting for ...	Primary/Re...	512	149E00110D0...	10:00:00:05:1...	1	2048	No

Refresh

Close Help

FIGURE 336 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 es 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 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 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 or 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 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 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 195](#) on page 564.)
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 305](#).)

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 337](#).)

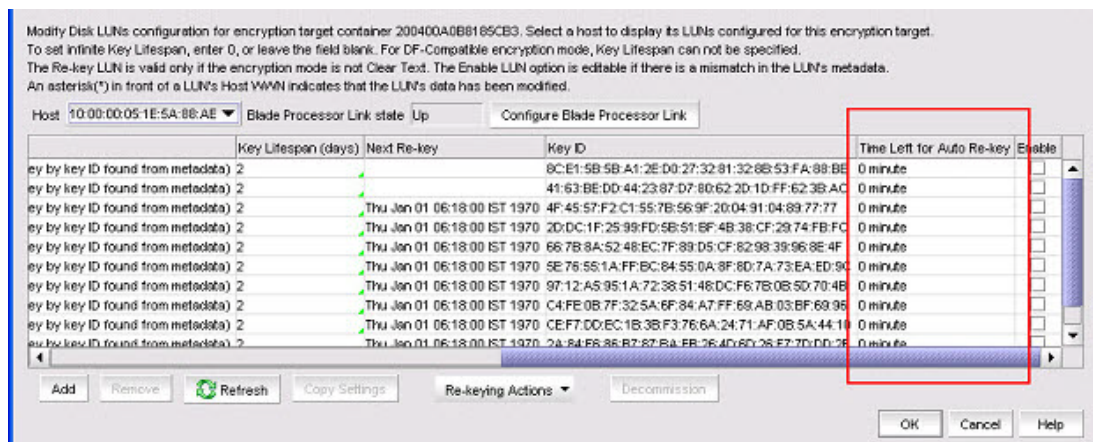


FIGURE 337 Encryption Targets Disk LUNs dialog box - Time left for auto rekey

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 195](#) on page 564.)
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 338](#).)

- **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:**
 - **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 Fabric OS 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 Fabric OS prior to v7.1.0, Key Vault Type is displayed as “Thales Key Manager (TEMS)”.
 - **Tivoli Key Lifetime Manager (TKLM):** No other key vault reference is used.
 - **Key Management Interoperability Protocol (KMIP):** Any KMIP-compliant server can be registered as a key vault on the 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*.

- **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 195](#) on page 564.)
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 339](#).)

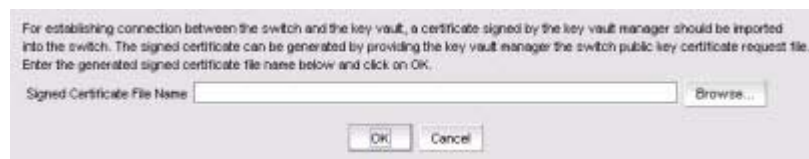


FIGURE 339 Import Signed Certificate dialog box

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 195](#) on page 564.)
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 195](#) on page 564.)
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 340](#).)

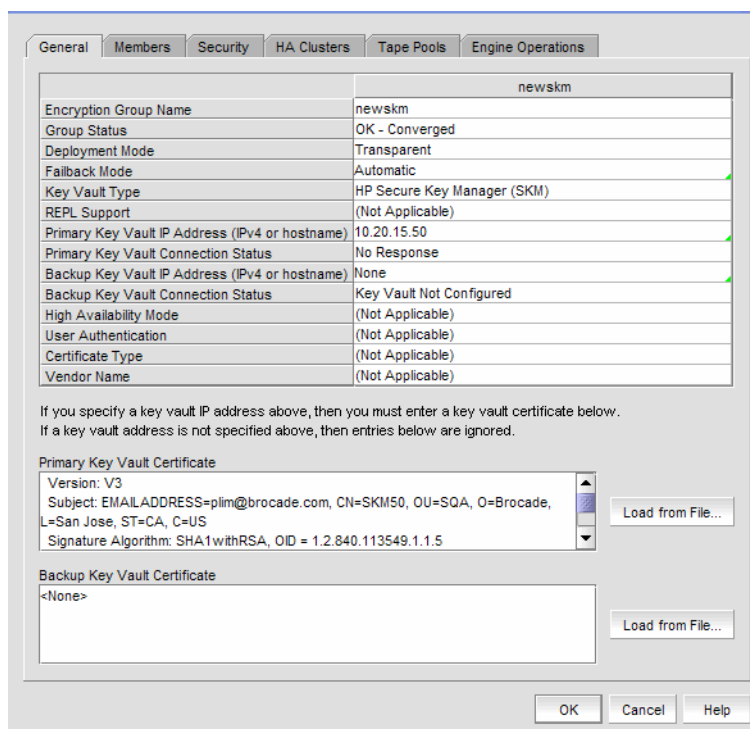


FIGURE 340 Encryption Group Properties dialog box

The dialog box contains the following information:

- **General** tab: For a description of the dialog box, refer to “[General tab](#)” on page 738.
- **Members** tab: For a description of the dialog box, refer to “[Members tab](#)” on page 742.
- **Security** tab: For a description of the dialog box, refer to “[Security tab](#)” on page 744.
- **HA Clusters** tab: For a description of the dialog box, refer to “[HA Clusters tab](#)” on page 746.
- **Link Keys** tab: (*Visible for LKM/SSKM only.*) For a description of the dialog box, refer to “[Link Keys tab](#)” on page 748.
- **Tape Pools** tab: For a description of the dialog box, refer to “[Tape Pools tab](#)” on page 749.
- **Engine Operations** tab: For a description of the dialog box, refer to “[Engine Operations tab](#)” on page 752.

General tab

The **General** tab is viewed from the **Encryption Group Properties** dialog box. (Refer to [Figure 341](#).) 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.

newskm	
Encryption Group Name	newskm
Group Status	OK - Converged
Deployment Mode	Transparent
Failback Mode	Automatic
Key Vault Type	HP Secure Key Manager (SKM)
REPL Support	(Not Applicable)
Primary Key Vault IP Address (IPv4 or hostname)	10.20.15.50
Primary Key Vault Connection Status	No Response
Backup Key Vault IP Address (IPv4 or hostname)	None
Backup Key Vault Connection Status	Key Vault Not Configured
High Availability Mode	(Not Applicable)
User Authentication	(Not Applicable)
Certificate Type	(Not Applicable)
Vendor Name	(Not Applicable)

If you specify a key vault IP address above, then you must enter a key vault certificate below.
If a key vault address is not specified above, then entries below are ignored.

Primary Key Vault Certificate

Version: V3
Subject: EMAILADDRESS=plim@brocade.com, CN=SKM50, OU=SQA, O=Brocade, L=San Jose, ST=CA, C=US
Signature Algorithm: SHA1withRSA, OID = 1.2.840.113549.1.1.5

Backup Key Vault Certificate

<None>

OK Cancel Help

FIGURE 341 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:**
 - **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 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 . 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 . 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 . 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 and registered as a KAC certificate. The CA will be registered as a key vault certificate on the .
 - **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 342](#).) 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 342 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 52 explains the impact of removing switches.

TABLE 52 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.	<ul style="list-style-type: none"> • 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 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.



CAUTION

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.

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 343.) 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.

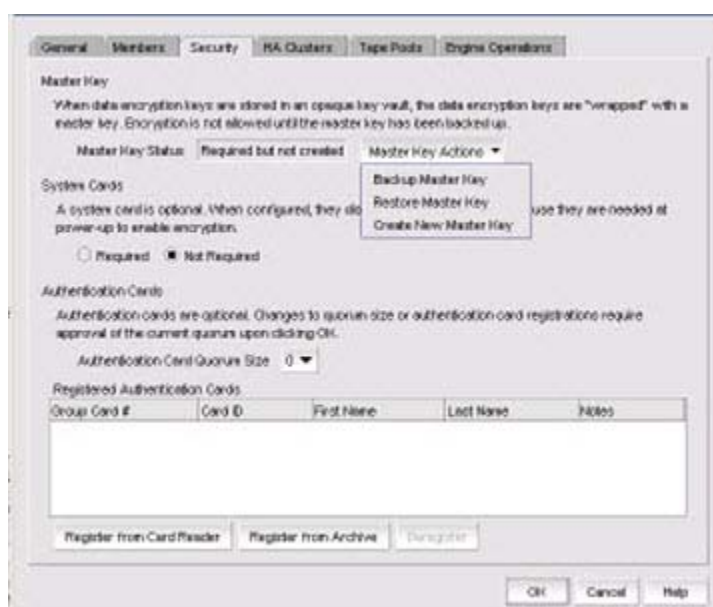


FIGURE 343 Encryption Group Properties dialog box - Security tab

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 677.

The **HA Clusters** tab is viewed from the **Encryption Group Properties** dialog box. (Refer to [Figure 344](#).) 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 577.
- **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 680.

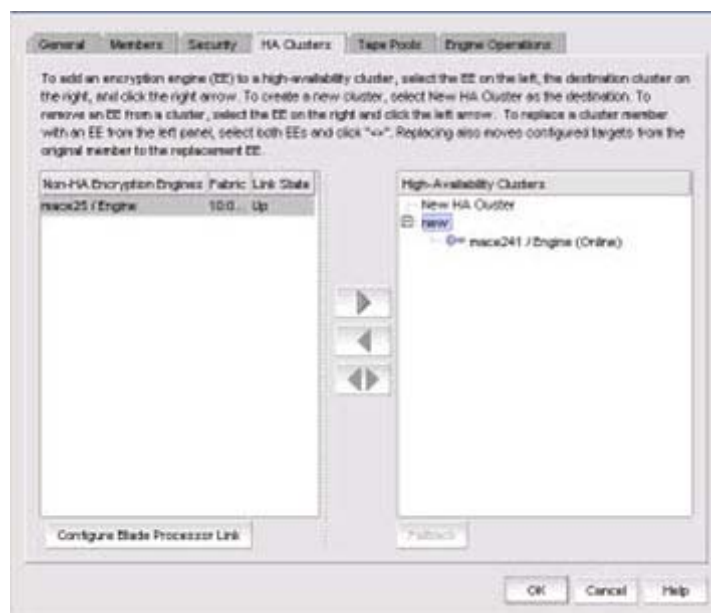


FIGURE 344 Encryption Group Properties dialog box - HA Clusters tab

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 345](#).) 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 345 Encryption 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 346](#).) 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 346 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 195](#) on page 564.)
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. (Refer to [Figure 347](#).) The **Name** tape pool label type is the default; however, you can change the tape pool label type to **Number** (Refer to [Figure 348](#).)

FIGURE 347 Add Tape Pool by name dialog box

FIGURE 348 Add Tape Pool by number dialog box

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 676.

The **Engine Operations** tab is viewed from the **Encryption Group Properties** dialog box. (Refer to [Figure 349](#).) 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 349 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 53](#) lists some of those acronyms.

TABLE 53 Encryption acronyms

Acronym	Name
EE	Encryption Engine
EG	Encryption Group
HAC	High Availability Cluster

20 Encryption-related acronyms in log messages

Zoning

In this chapter

- [Zoning overview](#) 755
- [Zone database size](#) 758
- [Zoning configuration](#) 759
- [LSAN zones](#) 781
- [LSAN tagging](#) 786
- [Traffic Isolation zones](#) 786
- [Boot LUN zones](#) 792
- [Zoning administration](#) 794

Zoning overview

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 350](#), 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 350 Zoning

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 721 for more information.
- **LSAN zones**
Provide device connectivity between fabrics without merging the fabrics. Refer to “[LSAN zones](#)” on page 781 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.
- Traffic Isolation zones (TI zones)
Isolate inter-switch traffic to a specific, dedicated path through the fabric. Refer to [“Traffic Isolation zones”](#) on page 786 for more information.

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.

NOTE

Online zoning is supported in Fabric OS and in Network OS.

Offline zoning

NOTE

Fabric OS 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 794.

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:

- Names must start with an alphabetic character and may contain alphanumeric characters and the underscore (_) character.
- 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 numeric character or a special character.
- 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.

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 (1 MB).

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.
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 760.
6. Add members to each zone.
For specific instructions, refer to [“Adding members to a zone”](#) on page 761 and [“Creating a new member in an LSAN zone”](#) on page 785.
7. Create a zone configuration.
For specific instructions, refer to [“Creating a zone configuration”](#) on page 769.
8. Activate the zone configuration.
For specific instructions, refer to [“Activating a zone configuration”](#) on page 771.
9. Set zoning policies, if necessary.
For specific instructions, refer to [“Enabling or disabling the default zone for fabrics”](#) on page 765.
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 758.
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 QoSX_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 761 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 762.

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 761.

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 758.

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.

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 758.)
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.
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 click **OK**.The **Zoning Policies** dialog box closes and the **Zone DB** tab displays.
8. 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 758 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.
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 758 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 768.
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 758.
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 758.)
7. Add zones to the zone configuration.
For step-by-step instructions, refer to [“Adding zones to a zone configuration”](#) on page 770.
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.

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 758.

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 758.
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 760.

8. Add members to each zone.

For specific instructions, refer to [“Adding members to a zone”](#) on page 761 and [“Creating a member in a zone”](#) on page 762.

9. Create a zone configuration.

For specific instructions, refer to [“Creating a zone configuration”](#) on page 769.

10. Activate the zone configuration.

For specific instructions, refer to [“Activating a zone configuration”](#) on page 771.

11. Set zoning policies, if necessary.

For specific instructions, refer to “[Enabling or disabling the default zone for fabrics](#)” on page 765.

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 777.
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 773.
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 that were marked for addition or removal as well as all members not otherwise marked.

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 351](#).




FIGURE 351 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 54) 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, or aliases, but do not mix element types.
 - b. Select the same type of element in the **Editable Zone DB** area.
If you selected a zone configuration in the **Reference Zone DB** area, you must select a zone configuration in the **Editable Zone DB** area.
 - c. Click **Merge**.
If the **Merge** button is inactivated, check that 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. You cannot merge a zone with a zone configuration, for example.
6. (Optional) Merge all elements by clicking **Merge All**.

7. (Optional) Add elements (aliases, zones, and zone configurations) 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.
These are the elements that 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. You can add zones to a zone configuration. And you can add zone configurations to the zone database.
 - c. Click **Add**.
If the **Add** button is inactivated, check that 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 (one that you have added) from the **Editable Zone DB** area and clicking **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 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 777 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 351](#).
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 771 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 783.
6. Add members to each zone.
For specific instructions, refer to [“Adding members to the LSAN zone”](#) on page 784.

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 758.
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.

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 784.

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

In Fabric OS 6.4.0 or later, 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 (Brocade 300)
- 40-port, 8 Gbps FC Switch (Brocade 5100)
- 80-port, 8 Gbps FC Switch (Brocade 5300)
- 48-port, 16 Gbps FC Switch (Brocade 6510)
- 8 Gbps 12-port Embedded Switch (Brocade 5410)
- 8 Gbps 24-port Embedded Switch (Brocade 5424, 5450, 5460)
- 8 Gbps 16-port Embedded Switch (Brocade 5470)
- 8 Gbps 24-port Embedded Switch (Brocade 5480)
- 8 Gbps Extension Switch (Brocade 7800)
- 8 Gbps 8-FC port, 10 GbE 24-DCB port Switch (Brocade 8000)
- 8 Gbps 40-port Switch (Brocade VA-40FC)
- 16 Gbps 4-slot Backbone Chassis (Brocade DCX 8510-4)
- 16 Gbps 8-slot Backbone Chassis (Brocade DCX 8510-8)
- 8-slot Backbone Chassis (Brocade DCX)
- 4-slot Backbone Chassis (Brocade DCX-4S)
- 8 Gbps Encryption Switch (Brocade 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 6.4.0 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 788.

7. Add members to each zone.

For specific instructions, refer to [“Adding members to a Traffic Isolation zone”](#) on page 789.

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 6.1.1 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 758.

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.

For Network OS fabrics, Boot LUN zones are not supported.

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.
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.
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 764 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 54](#).

TABLE 54 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 352](#).




FIGURE 352 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 54) 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 777.

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 **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.
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 2](#) 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.

If the seed switch is running Fabric OS 6.3.0 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.
5. 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.

21 Zoning administration

Fibre Channel over IP

In this chapter

- FCIP services licensing 806
- FCIP Concepts 806
- IP network considerations. 806
- FCIP platforms and supported features. 807
- FCIP trunking 808
- IPsec and IKE implementation over FCIP. 816
- Open systems tape pipelining. 820
- FICON emulation features. 821
- Connecting cascaded FICON fabrics over FCIP 823
- FCIP configuration guidelines. 829
- Configuring an FCIP tunnel 830
- Adding an FCIP circuit 833
- Use TCP/IP DSCP or L2CoS to prioritize FC traffic. 835
- Configuring FCIP tunnel advanced settings. 837
- Viewing FCIP connection properties. 843
- Viewing General FCIP properties 844
- Viewing FCIP port properties. 845
- Editing FCIP circuits. 847
- Disabling FCIP tunnels 848
- Enabling FCIP tunnels 848
- Deleting FCIP tunnels 849
- Displaying FCIP performance graphs. 850
- Displaying FCIP performance graphs for Ethernet ports 850
- Displaying tunnel properties from the FCIP tunnels dialog box 851
- Displaying FCIP circuit properties from the FCIP tunnels dialog box 852
- Displaying switch properties from the FCIP Tunnels dialog box 853
- Displaying fabric properties from the FCIP Tunnels dialog box 854
- Troubleshooting FCIP Ethernet connections 854

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 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:

- The 8 Gbps Extension Switch.
- The 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,

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 55](#) summarizes FCIP capabilities per platform.

TABLE 55 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)	Yes	Yes
• FCIP Fastwrite		
• Tape Acceleration		
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 56](#).

TABLE 56 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

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 353](#). Each circuit represents a portion of the available Ethernet bandwidth provided by the GbE ports that are connected to the WAN.




FIGURE 353 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 354](#), 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 354 Link loss and retransmission over peer lowest metric circuit

In [Figure 355](#), 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 355 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 833 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 811.

Examples of circuit failover in groups

Tables [Table 57](#) through [Table 59](#) provide examples of how failover occurs on circuits with different bandwidths configured in failover groups.

[Table 57](#) 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 57 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.
Circuit 3 Metric 1	1	500Mb	Only when circuit 1 fails.
Circuit 4 Metric 1	2	1000Mb	Only when circuit 2 fails.

Table 58 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 58 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 59 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 circuit 1 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 59 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 356](#). 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 815. 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.

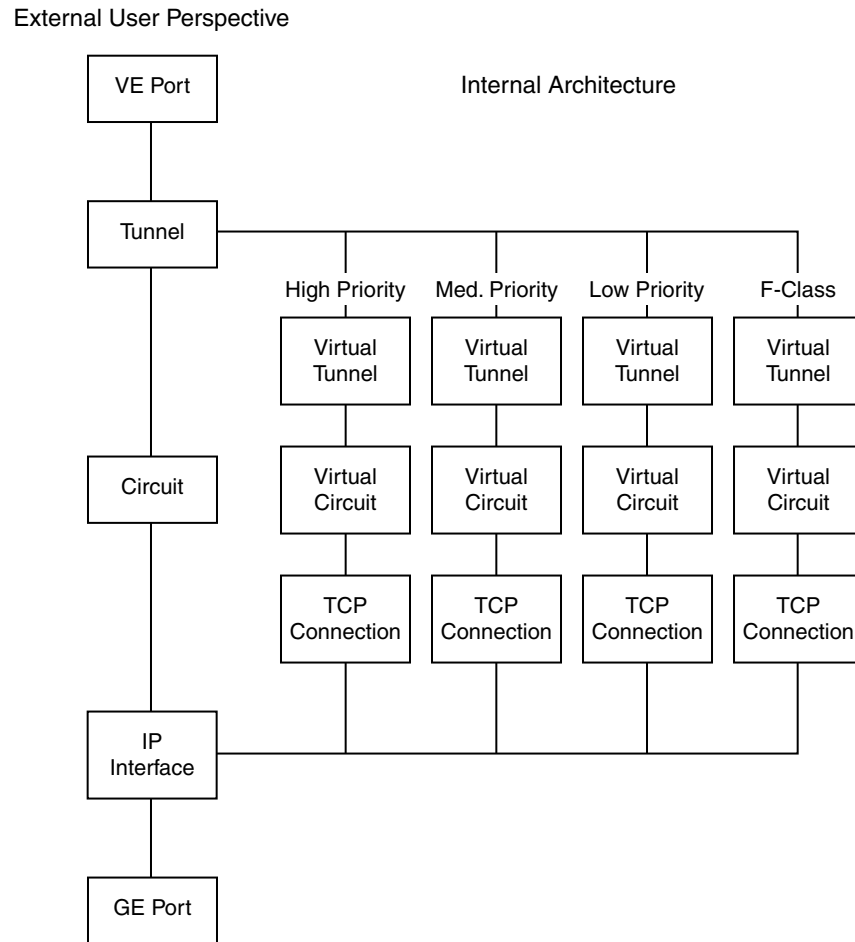


FIGURE 356 TCP connections for handling QoS SID/DID-based FC traffic prioritization

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 357).

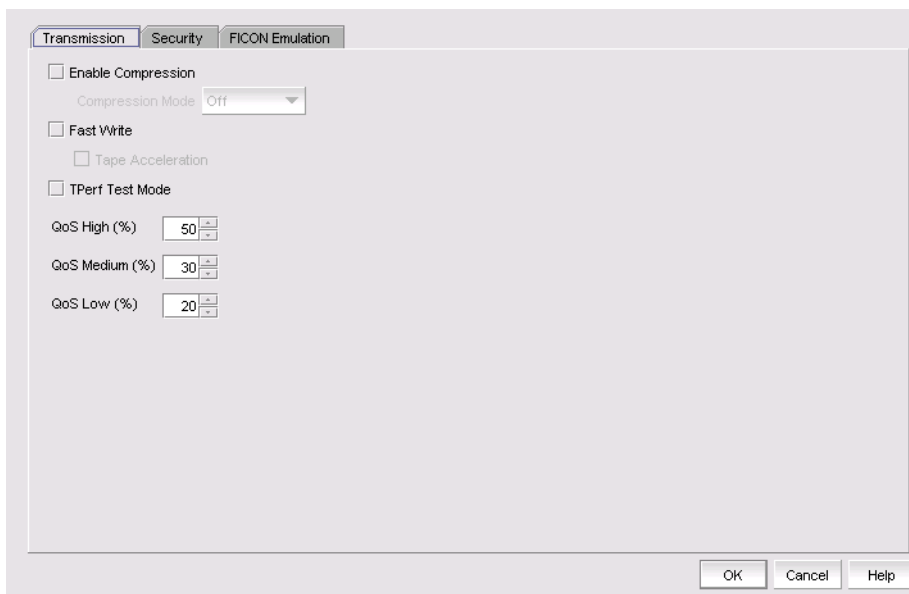


FIGURE 357 Advanced Settings Transmission Tab

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.

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.
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 60 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
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

TABLE 60 Default Mapping of DSCP priorities to L2CoS Priorities (Continued)

DSCP priority/bits	L2CoS priority/bits	Assigned to:
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 61](#) when configuring tunnels to use OSTP.

TABLE 61 OSTP constraints

FCIP Fastwrite	Tape Acceleration
Each GbE port supports up to 2048 simultaneous accelerated exchanges, which means a <i>total of 2048 simultaneous exchanges combined</i> for Fastwrite and Tape Acceleration.	Each GbE port supports up to 2048 simultaneous accelerated exchanges, which means a <i>total of 2048 simultaneous exchanges combined</i> 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 61 OSTP constraints

FCIP Fastwrite	Tape Acceleration
Class 3 traffic is accelerated with Fastwrite.	Class 3 traffic is accelerated between host and sequential device.
	<p data-bbox="930 394 1461 562">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.</p> <p data-bbox="930 573 1461 741">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.</p>
	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 6.4.1 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 are 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 358](#).



FIGURE 358 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 [Figure 360](#).




FIGURE 359 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” [Chapter 25, “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 827.

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 829.
2. Configure tunnels circuits between the switches by following steps under [“Configuring an FCIP tunnel”](#) on page 830
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 827. 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 833
 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” [Chapter 25, “FICON Environments”](#).
 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 decussated FICON fabric merge” [Chapter 25, “FICON Environments”](#).
 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 824 and the “Traffic Isolation zones” section of [Chapter 21, “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 62 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 62 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 **FCIP Tunnel Advanced Settings** dialog box **Transmission** tab (access by selecting **Advanced Settings** 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 that 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 55](#) 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.

1. Select **Configure > FCIP Tunnels**.

The **FCIP Tunnels** dialog box is displayed ([Figure 360](#)).

The screenshot shows the FCIP Tunnels dialog box. The top section displays a tree structure of discovered fabrics, Extension Switches, and configured tunnels. The bottom section shows a detailed view of the selected fabric.

Products	Switch One	Switch Two	Total Circuits	Tunnel Operational Status	Administrative Status	Description
7800 fabric						
switch202FVT						
Tunnel 16 (VE)	switch202FVT		1	Disabled	Disabled	
Tunnel 17 (VE)	switch202FVT		1	Up	Enabled	
Tunnel 18 (VE)	switch202FVT		1	Disabled	Disabled	
Tunnel 19 (VE)	switch202FVT		1	Disabled	Disabled	
Tunnel 20 (VE)	switch202FVT		2	Up	Enabled	
Tunnel 21 (VE)	switch202FVT		1	In Progress	Enabled	
10.00.00.05:1E:53:6B:69						
MF2-7500-521						
Tunnel 0 (VE 16)	MF2-7500-521		1	Active	Enabled	
Tunnel 1 (VE 17)	MF2-7500-521		1	Active	Enabled	
FX8-24 blade						
DCX_FVT_128						

Fabric	
Name	7800 fabric
FOS Name	SJFabric
Seed Switch	10.24.49.202
AD Enabled	No
Status	Marginal
Switch and AG Count	2
Description	
Principal Switch	10.24.49.202
Active Zone Configuration	
Last Discovery	Tue Feb 15 13:18:02 PST 2011
Tracked	Yes
Location	
Contact	

FIGURE 360 FCIP Tunnels dialog box (fabric selected from Product tree)

The dialog box displays a tree structure of all discovered fabrics, Extension Switches, and configured tunnels. Details such as 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 is displayed ([Figure 361](#)). The name of the switch you selected is displayed in the **Switch** field under **Switch One Settings**. This dialog 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** circuits using the function 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.

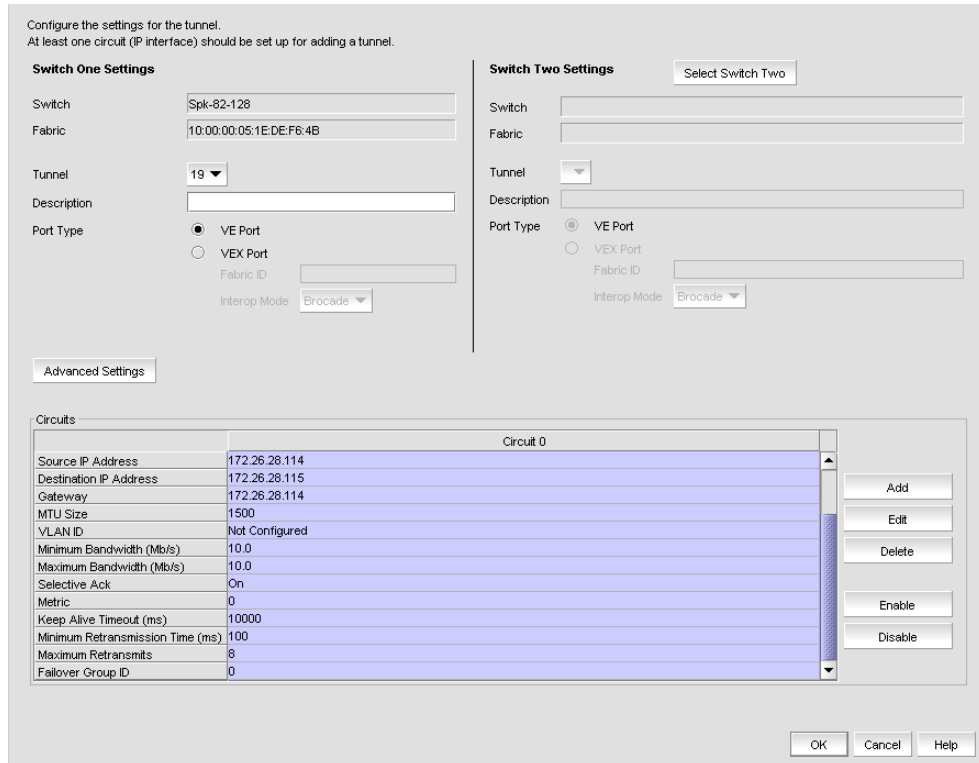


FIGURE 361 Add FCIP Tunnel dialog box

- c. Click **Select Switch Two** under **Switch Two Settings** on the **Add FCIP Tunnels** dialog box to display the **Select Switch** dialog box.
The **Select Switch** dialog box displays discovered Extension Switches.
- d. Select the switch you want to connect to switch one and click **OK**.
- e. The switch and fabric names display in the **Switch Two Settings** area of the **Add FCIP Tunnel** dialog box.
- f. Enter a description of the tunnel in the **Description** field.

NOTE

You cannot assign a **Tunnel ID** until after at least one circuit is configured. The **Add Circuit** dialog returns you to the **Add FCIP Tunnel** dialog to allow you to select the **Tunnel ID**.

- g. Skip to step [step 4](#) and continue configuration.

- To edit the configuration for an existing FCIP tunnel and circuits between two switches, follow these steps:

NOTE

You cannot edit an active tunnel; disable the tunnel before making changes.

- From the **FCIP Tunnels** dialog box (refer to [step 1](#)), select the FCIP tunnel that you want to configure under the **Products** tree.
- 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** circuits using the function 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.

- Change configuration settings as required using the following steps.
- 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 is displayed. Continue with [“Adding an FCIP circuit”](#).

Logical switch function on FCIP Tunnels dialog box

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 default and user-user-configured logical switch are discovered, all tunnels and circuits created on logical switch display, including circuits with shared GbE 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*.
- [Chapter 19, “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 (Figure 361 on page 831). For 8 Gbps platforms, you can add multiple FCIP circuits to the tunnel with this button.

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 830.

FIGURE 362 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 geo through ge9
- 10 Gbps mode - Ports xge0 and xge1
- Dual mode - Ports geo through ge9 and xge0

2. Select **Crossport circuit** to configure the 10 GbE 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).
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 Route**, and select a **Gateway address**.
6. Enter the **MTU Size**.

For SAN traffic, the largest possible MTU (Maximum Transmission Unit) size is generally the most efficient. Enter a value between 1260 and 2348 for the 4 Gbps platforms and between 1260 and 1500 for the 8 Gbps platforms. MTU rates must match on both ends of the tunnel.

If you have an active connection between switch one and switch two, click **Verify 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. 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

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.

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 VLAN ID to both switches. You can assign the same or different VLAN IDs to each switch.

The VLAN ID is an integer value between 1 and 4094 which sets the VLAN tag value in the header assigning the traffic to that specific VLAN. Layer two class of service (L2CoS) values may be assigned to establish traffic priorities over a VLAN. This is done as an **Advanced Setting**.

8. The **Metric** option 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 will be 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** for the circuit from 0 to 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 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. For more information on Circuit Failover Grouping, refer to [“Circuit Failover Grouping”](#) on page 811.

NOTE

Failover Group 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 814 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.

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 acknowledgement if your system cannot support selective acknowledgement.
 - 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.

If you select **Advanced Settings**, the **Transmission tab** of the **FCIP Circuit Advanced Settings** dialog box displays ([Figure 363](#)).

FIGURE 363 FCIP Circuit Advanced Settings

- Selective the **Ack** check box to disable selective acknowledgement. This should not be done unless your system cannot support selective acknowledgement.
 - Use the **Keep Alive Time Out (ms)** option to override the default value of 10000 ms. As shown, the range is from 500 to 7200000.
 - Use the **Max. Retransmission Time (ms)** option to override the default value of 100 ms.
 - Use the **Max. Retransmits** option to override the default value of 8. As shown, the range is 1 to 8.
 - Select **L2CoS** and **DSCP** priorities. Refer to “[QOS, DSCP, and VLANs](#)” on page 818 for more information.
 - Select **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.
 14. Click **OK** to close the **Add FCIP Tunnel** dialog box.

Logical switch function in FCIP Add Circuit dialog box

The display and function of circuits created on logical switches and with shared GbE ports varies in the **FCIP Add Circuit** dialog box according default switch and user-configured logical switch discovery as follows. For details on configuring FCIP with logical switches, refer to the “Using FCIP with logical switches” section in the *Fabric OS FCIP Administrator’s Guide*.

- If both the default logical switch and user-configured logical switch are discovered:
 - The **GigE Port** drop list will display all GigE ports in all logical switches, including ports from default switch 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 default logical switch, since 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 adding a circuit, only the GigE ports present in the logical switch with undiscovered 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*.
- [Chapter 19, “Virtual Fabrics”](#).

Circuit configuration failure

When a tunnel cannot be created because the process for adding a new circuit configuration fails, a **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, **Security** tab, and **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 and disabling compression

Data compression can improve performance on long distance connections.

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 **Enable Compression** check box to enable compression.
This enables the **Compression Mode** selector ([Figure 364](#)).

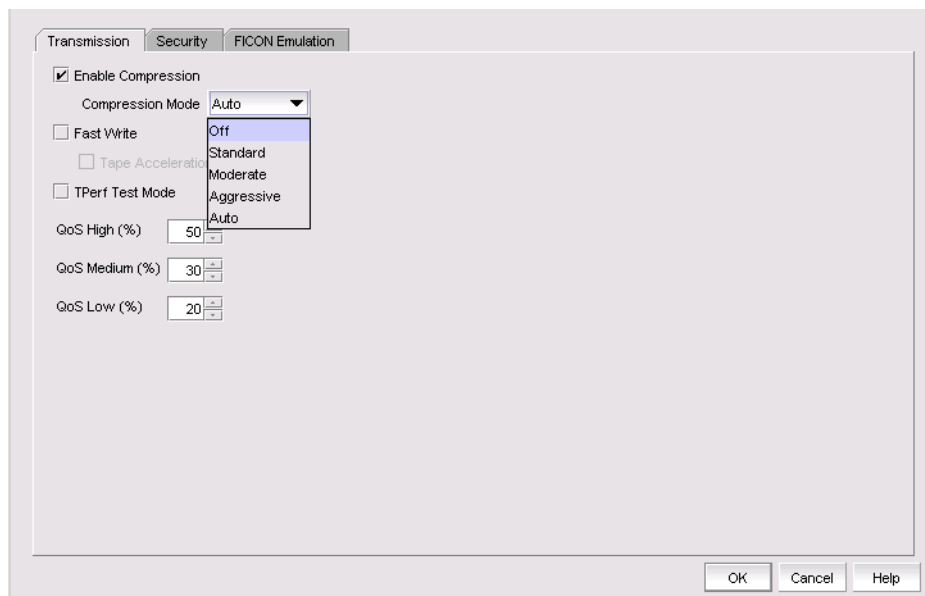


FIGURE 364 Selecting a compression mode

3. Select the desired compression mode.

A **Standard** option provides hardware compression and is available on all platforms. The 8 Gbps Extension Switch and the 8 Gbps Extension Blade provide three additional options for compression. The **Moderate** option enables 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. The **Aggressive** option is a software-only compression option that provides a more aggressive algorithm. This option supports up to 2.5 Gbps of FC traffic. The **Auto** option allows the system to set the best compression mode based on the tunnel's configured bandwidth and the bandwidth of all tunnels in the system.

4. Click **OK** to commit your selection.

To disable compression, click the **Enable Compression** to clear the check mark, and click **OK**.

Enabling Open Systems Tape Pipelining (OSTP)

Latency introduced by a long distance IP connection can negatively impact tape I/O performance. 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. OSTP options are available on the **Transmission** tab.

To enable OSTP, do the following:

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.

3. Select the **Tape Acceleration** check box.
4. Click **OK**.

Enabling Tperf test mode

To enable Tperf test mode, do the following:

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 QoS (Quality of Service) priority percentages from the preset default values of 50% (High), 30% (Medium), and 20% (low). Values for the three priority levels must equal 100%. A minimum of 10% is required for each level. You can adjust percentages in increments of 1%. To configure QoS percentages, do the following:

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 IPsec and IKE policies

IPsec and IKE policies are configured from the **Security** tab. The screens and procedures are platform-dependent. [Figure 365](#) on page 840 shows the screen 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 365 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 drop-down 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 grayed 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 will need to 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 between 12 and 32 alphanumeric characters. The length depends on the chosen IKE policy.
 - For the 8Gbps 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.

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 Fabric OS earlier than 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 366). 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.

Both the ends of the tunnel will be configured identically. If there is a discrepancy, the FICON emulation feature will be disabled.
 Enabling XRC emulation feature requires XRC license to be installed on the switch.
 Enabling Tape Write Pipelining and Read Pipelining feature requires Tape license to be installed on the switch.

Populate Default Values

Enable FICON XRC Emulation Enable FICON Tape Write Emulation

Enable FICON Tape Read Emulation Enable FICON Tape Read Block ID

Enable FICON Tin/Tir Emulation Enable FICON Device Level Ack Emulation

Enable FICON Teradata Read Pipelining Enable FICON Teradata Write Pipelining

FICON Tape Write Max Pipe (1-100) FICON Tape Read Max Pipe (1-100)

FICON Tape Write Max Ops (1-32) FICON Tape Read Max Ops (1-32)

FICON Tape Write Timer (100-1500ms) FICON Tape Max Write Chain (1000000-5000000ms)

FICON Oxid Base (0x0000-0xF000)

OK Cancel Help

FIGURE 366 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 6.4.1 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 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.

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 367).

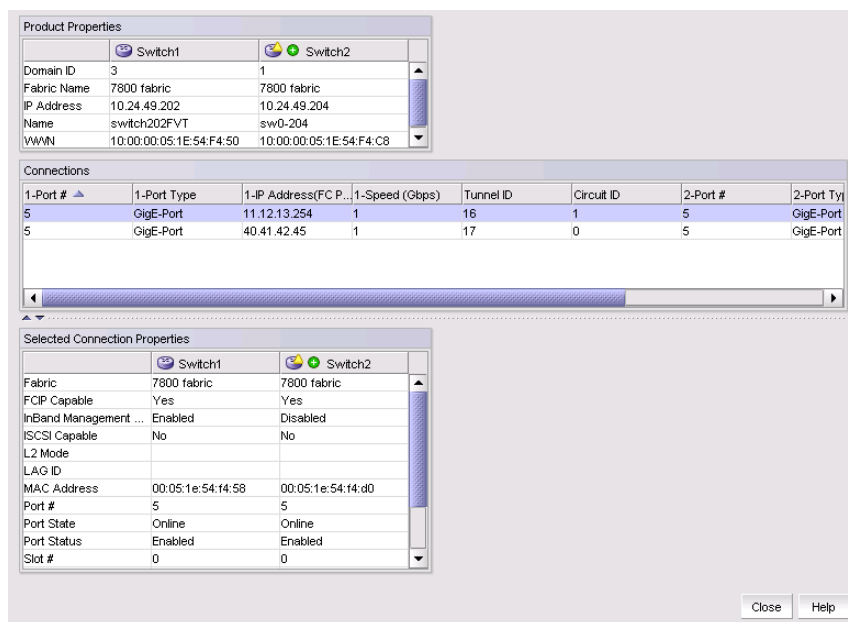


FIGURE 367 FCIP connection properties

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 367.

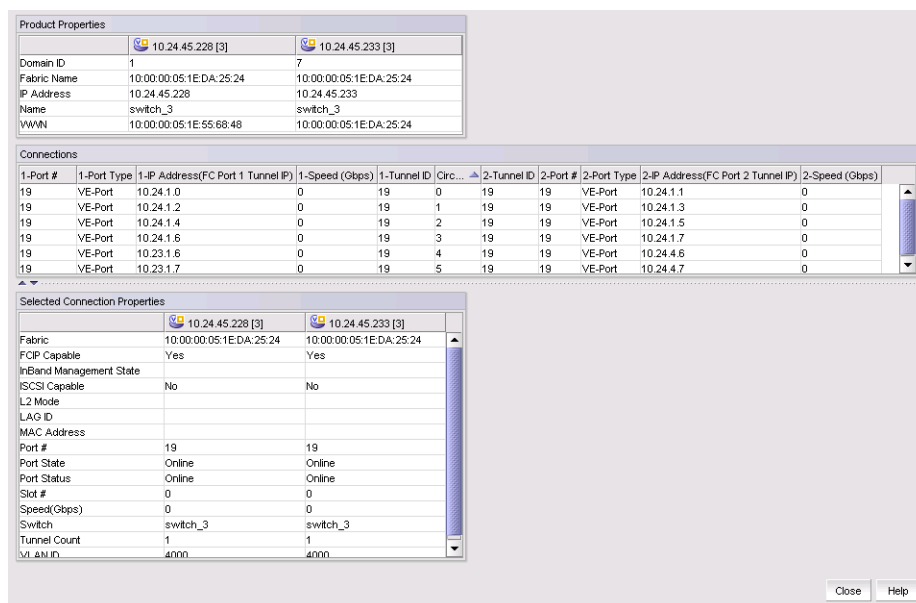


FIGURE 368 FCIP connection properties (default switch not discovered)

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.

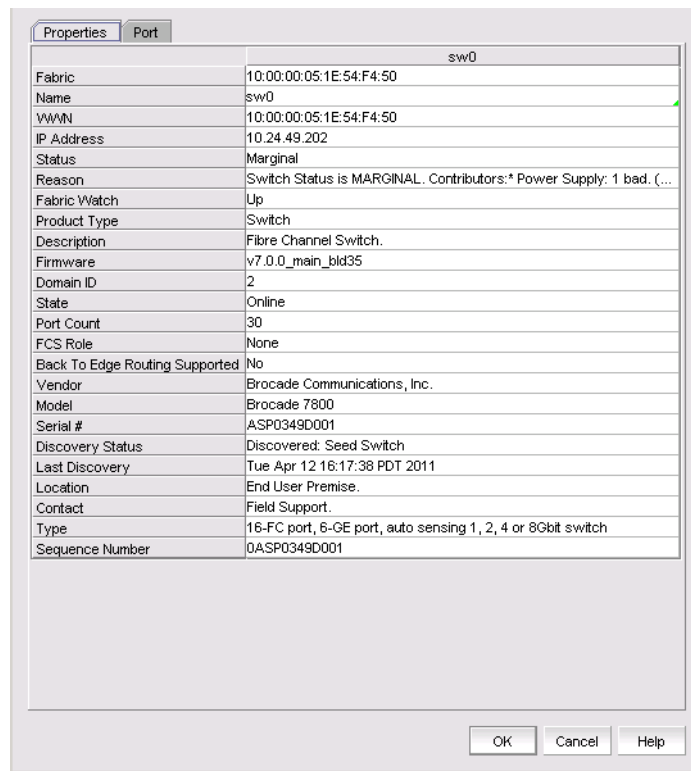


FIGURE 369 General FCIP properties tab (Extension Switch or Blade)

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.

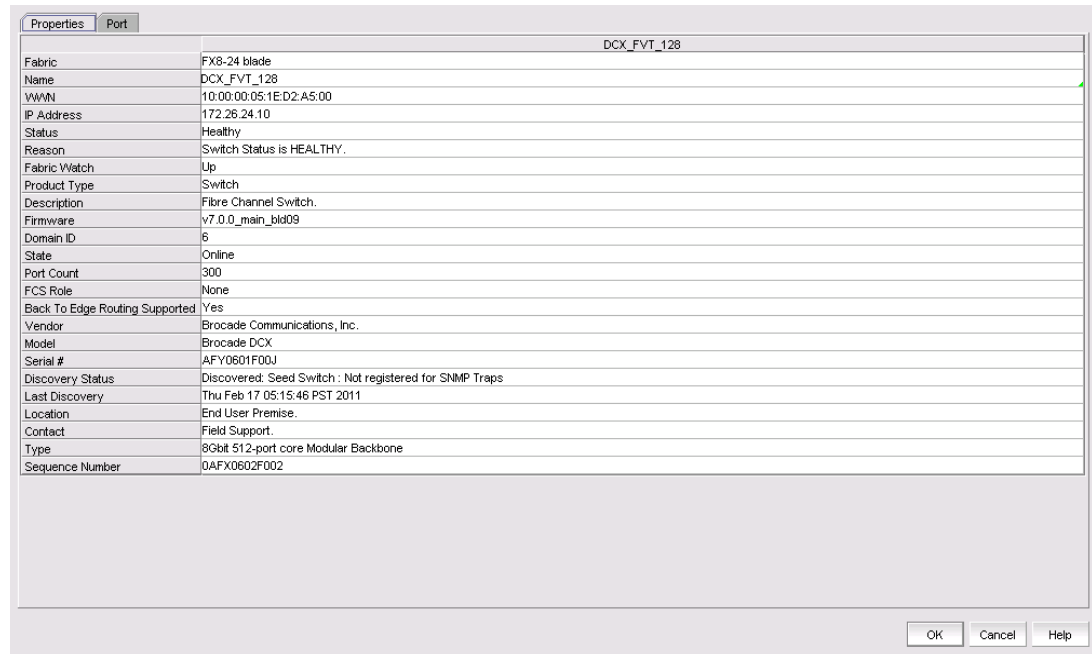


FIGURE 370 General FCIP properties tab (blade chassis)

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 371).

22 Viewing FCIP port properties

Properties Port

Port Count 24 Type FC

	port0	port 1	port2	3rd port	port4	port5	7800
Fabric	7800 fabric	7800 fabric	7800 fabric	7800 fabric	7800 fabric	7800 fabric	7800
Switch	switch202FVT	switch202FVT	switch202FVT	switch202FVT	switch202FVT	switch202FVT	switch202FVT
Name	port0	port 1	port2	3rd port	port4	port5	port6
Slot #	0	0	0	0	0	0	0
Port #	0	1	2	3	4	5	6
User Port #	0	1	2	3	4	5	6
Area ID /Port Index	0/ 0	1/ 1	2/ 2	3/ 3	4/ 4	5/ 5	6/ 6
FC Address	030000	030100	030200	030300	030400	030500	030600
Status	Mod_Inrv	No_Module	No_Light	Online	No_Module	No_Module	No_Module
Additional Port Info	Speed Mismatch / Incom...						
State	Offline	Offline	Offline	Online	Offline	Offline	Offline
Type	U-Port	U-Port	U-Port	F-Port	U-Port	U-Port	U-Port
Port Speed (Gb/s)	11	8	8	4	8	8	8
Port Module	sw		sw				
Port WWN	20:00:00:05:1E:54:F4:50	20:01:00:05:1E:54:F4:50	20:02:00:05:1E:54:F4:50	20:03:00:05:1E:54:F4:50	20:04:00:05:1E:54:F4:50	20:05:00:05:1E:54:F4:50	20:06:00:05:1E:54:F4:50
Protocol	FC	FC	FC	FC	FC	FC	FC
Buffers Desired	0	0	0	0	0	0	0
Buffers Allocated	0	0	0	8	0	0	0
Distance Actual (km)	0	0	0	0	0	0	0
Distance Estimated (km)	0	0	0	0	0	0	0
Long Distance Setting	L0.Normal	L0.Normal	L0.Normal	L0.Normal	L0.Normal	L0.Normal	L0.Normal
Physical/Logical	Physical	Physical	Physical	Physical	Physical	Physical	Physical
Locked Port Type	U-port	EX-Port	U-port	U-port	U-port	U-port	U-port
NPIV Enabled	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Connected Switch	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Performance

OK Cancel Help

FIGURE 371 FC ports properties

- To view VE and VEX port information, select the **VE/VEx** from the **Type** drop-down list (Figure 372).

Properties Port

Port Count 20 Type VE / VEx

Performance

	slot8 port19	slot8 port26	slot8 port27	slot8 port31	slot8 port14	slot8 port29
Fabric	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00
Switch	pluto105	pluto105	pluto105	pluto105	pluto105	pluto105
Name	slot8 port19	slot8 port26	slot8 port27	slot8 port31	slot8 port14	slot8 port29
Slot #	8	8	8	8	8	8
Port #	19	26	27	31	14	29
User Port #	211	218	219	223	206	221
Area ID /Port Index	211/ 211	218/ 218	219/ 219	223/ 223	206/ 206	221/ 221
FC Address	69d300	69da00	69db00	69df00	69ce00	69dd00
Status	UNKNOWN	Disabled - Persistently di...	Disabled - Persistently di...	Disabled - Persistently di...	UNKNOWN	Disabled - Persistently di...
Additional Port Info		Persistently disabled port	Persistently disabled port	Persistently disabled port		Persistently disabled port
State	Offline	Offline	Offline	Offline	Offline	Offline
Type	U-Port	U-Port	U-Port	U-Port	U-Port	U-Port
Port Speed (Gb/s)	0	0	0	0	0	0
Port Module						
Port WWN	20:D3:00:05:33:1D:7B:00	20:DA:00:05:33:1D:7B:00	20:DB:00:05:33:1D:7B:00	20:DF:00:05:33:1D:7B:00	20:CE:00:05:33:1D:7B:00	20:DD:00:05:33:1D:7B:00
Protocol	FCIP	FCIP	FCIP	FCIP	FCIP	FCIP
Buffers Desired						
Buffers Allocated						
Distance Actual (km)						
Distance Estimated (km)						
Long Distance Setting						
Physical/Logical	Logical	Logical	Logical	Logical	Logical	Logical

Add Edit Delete

OK Cancel Help

FIGURE 372 VE/VEx port properties

- To view GbE (Ethernet) port information, select the **GigE** from the **Type** drop-down list (Figure 373).

	8/ge0	8/ge1	8/ge2	8/ge3	8/ge4	8/ge5	8/ge6
Port #	ge0	ge1	ge2	ge3	ge4	ge5	ge6
Fabric	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00	10:00:00:05:33:1D:68:00
Switch	pluto105	pluto105	pluto105	pluto105	pluto105	pluto105	pluto105
Slot #	8	8	8	8	8	8	8
MAC Address	00:05:33:41:CD:B0	00:05:33:41:CD:B1	00:05:33:41:CD:B2	00:05:33:41:CD:B3	00:05:33:41:CD:B4	00:05:33:41:CD:B5	00:05:33:41:CD:B6
Port Status	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Port State	Online	Offline	Offline	Offline	Online	Online	Online
Speed (Gb/s)	1	1	1	1	1	1	1
Tunnel Count	1	1	1	1	1	1	1
ISCSI Capable	No	No	No	No	No	No	No
FCIP Capable	Yes	Yes	Yes	Yes	Yes	Yes	Yes
InBand Management State	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
LAG ID							
L2 Mode							
VLAN ID							
IPsec Supported	No	No	No	No	No	No	No

FIGURE 373 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 374).

FIGURE 374 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 enable 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.

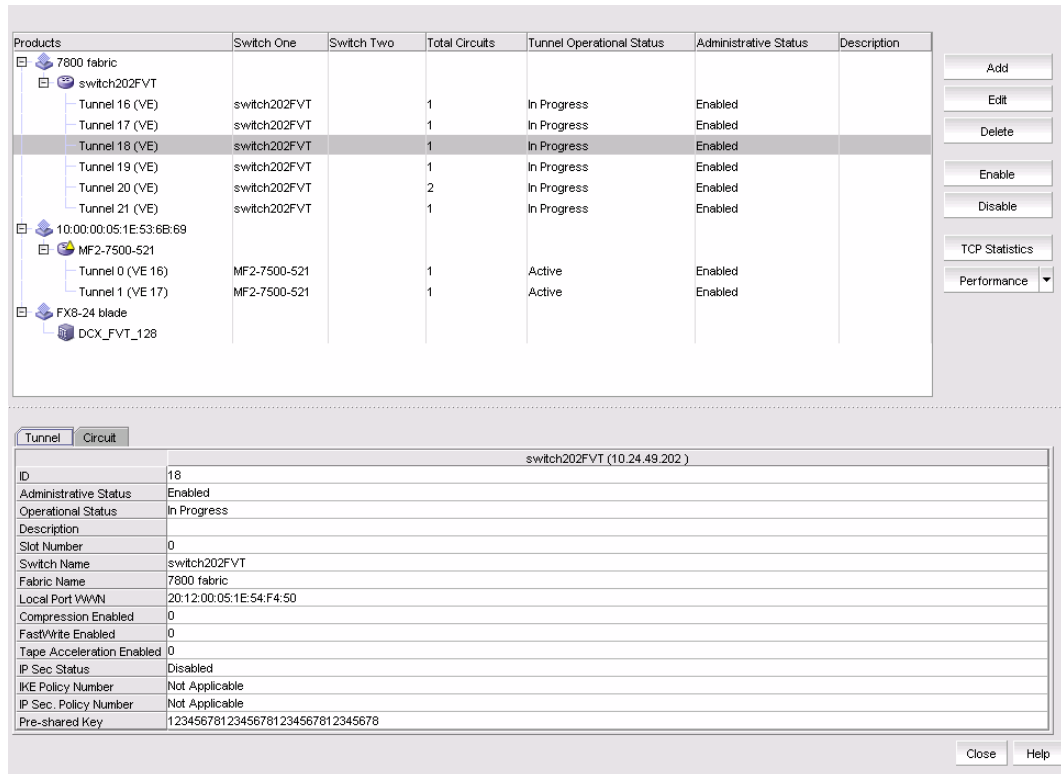


FIGURE 375 Tunnel properties on the FCIP Tunnels dialog box

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 376).

Products	Switch One	Switch Two	Total Circuits	Tunnel Operational Status	Administrative Status	Description
7800 fabric						
switch202FVT						
Tunnel 16 (VE)	switch202FVT		1	In Progress	Enabled	
Tunnel 17 (VE)	switch202FVT		1	In Progress	Enabled	
Tunnel 18 (VE)	switch202FVT		1	In Progress	Enabled	
Tunnel 19 (VE)	switch202FVT		1	In Progress	Enabled	
Tunnel 20 (VE)	switch202FVT		2	In Progress	Enabled	
Tunnel 21 (VE)	switch202FVT		1	In Progress	Enabled	
10:00:00:05:1E:53:6B:69						
MF2-7500-521						
Tunnel 0 (VE 16)	MF2-7500-521		1	Active	Enabled	
Tunnel 1 (VE 17)	MF2-7500-521		1	Active	Enabled	
FX8-24 blade						
DCX_FVT_128						

	Circuit 0	Circuit 2
Switch Name	switch202FVT	switch202FVT
Administrative Status	Enabled	Enabled
Operational Status	In Progress	In Progress
GigE Port	ge4	ge4
Source IP Address	44.44.44.44	40.50.60.71
Destination IP Address	44.44.44.45	40.50.60.70
Gateway	0.0.0.0	0.0.0.0
MTU Size	1500	1500
VLAN ID	22	Not Configured
Minimum Bandwidth (Mb/s)	10.0	10.0
Maximum Bandwidth (Mb/s)	10.0	10.0
Selective Ack	On	On
Metric	0	0
Keep Alive Timeout (ms)	1000	1000
Minimum Retransmission Time (ms)	100	100

FIGURE 376 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 377).

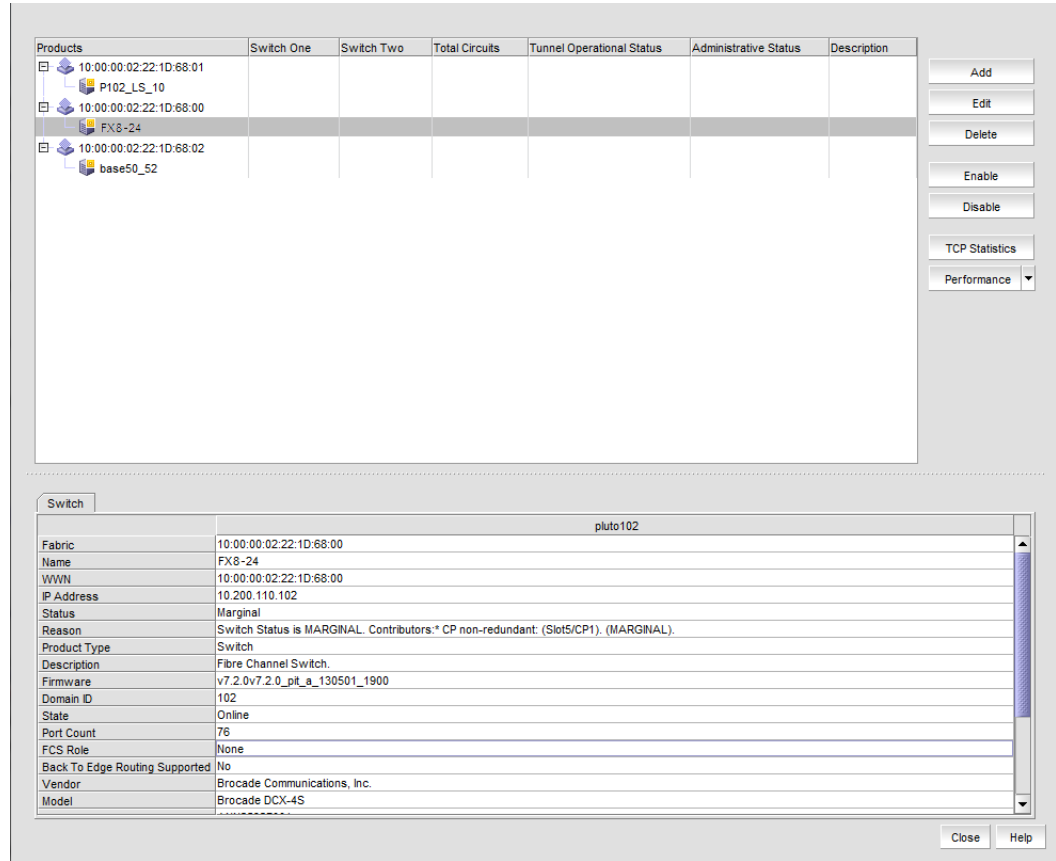


FIGURE 377 Switch properties on the FCIP Tunnels dialog box

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 378).

The screenshot displays the FCIP Tunnels dialog box. The top section shows a tree view of products under '7800 fabric', including 'switch202FVT' and 'sw0-204'. Below this is a table listing tunnels with columns for 'Switch One', 'Switch Two', 'Total Circuits', 'Tunnel Operational Status', 'Administrative Status', and 'Description'. The bottom section, titled 'Fabric', shows properties for '7800 fabric' in a table format.

Products	Switch One	Switch Two	Total Circuits	Tunnel Operational Status	Administrative Status	Description
7800 fabric						
switch202FVT						
Tunnel 16 (VE)	switch202FVT	sw0-204	1	In Progress	Enabled	
Tunnel 17 (VE)	switch202FVT	sw0-204	1	In Progress	Enabled	
Tunnel 18 (VE)	switch202FVT	sw0-204	1	In Progress	Enabled	
Tunnel 19 (VE)	switch202FVT	sw0-204	1	In Progress	Enabled	
Tunnel 20 (VE)	switch202FVT	sw0-204	2	In Progress	Enabled	
Tunnel 21 (VE)	switch202FVT	sw0-204	1	In Progress	Enabled	
sw0-204						
Tunnel 16 (VE)	sw0-204	switch202FVT	1	Up	Enabled	
Tunnel 17 (VE)	sw0-204	switch202FVT	1	Up	Enabled	
Tunnel 18 (VE)	sw0-204	switch202FVT	1	Up	Enabled	
Tunnel 19 (VE)	sw0-204	switch202FVT	1	Up	Enabled	
Tunnel 20 (VE)	sw0-204	switch202FVT	2	Up	Enabled	
Tunnel 22 (VE)	sw0-204	switch202FVT	1	Up	Enabled	

Fabric	
Name	7800 fabric
FOS Name	SJFabric
Seed Switch	10.24.49.202
AD Enabled	No
Status	Marginal
Switch and AG Count	2
Description	
Principal Switch	10.24.49.202
Active Zone Configuration	
Last Discovery	Tue Feb 15 13:18:02 PST 2011
Tracked	Yes
Location	
Contact	

FIGURE 378 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

In this chapter

- [Fabric Binding overview](#) 855
- [High integrity fabrics overview](#) 860

Fabric Binding overview

NOTE

Fabric Binding is supported on Fabric OS 5.2 or later.

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 379](#)).




FIGURE 379 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 857 and [“Disabling fabric binding”](#) on page 858.
 - **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 859.
 - **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 860.

- **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 859.

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

Fabric Binding is only supported on Fabric OS 5.2 or later.

1. Select **Configure > Fabric Binding**.
The **Fabric Binding** dialog box displays ([Figure 379](#)).
2. In the **Fabric List** table, click 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 859 and [“Removing switches from fabric binding membership”](#) on page 860.

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 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.

NOTE

Fabric Binding is only supported on Fabric OS 5.2 or later.

1. Select **Configure > Fabric Binding**.

The **Fabric Binding** dialog box displays ([Figure 379](#)).

2. In the **Fabric List** table, clear the **Enable/Disable** check box for fabrics for which you want to disable fabric binding.
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 857), you can add switches to the fabric binding membership list.

NOTE

Fabric Binding is only supported on Fabric OS 5.2 or later.

To add a switch to the fabric, complete the following steps.

1. Select **Configure > Fabric Binding**.
The **Fabric Binding** dialog box displays ([Figure 379](#)).
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 379](#)).
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 nodeworld wide name (WWN) of the switch in the **Node WWN** field.
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.
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 857), you can remove switches that are not part of the fabric from the membership list.

NOTE

Fabric Binding is only supported on Fabric OS 5.2 or later.

1. Select **Configure > Fabric Binding**.
The **Fabric Binding** dialog box displays ([Figure 379](#)).
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 fabrics 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 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

The term 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

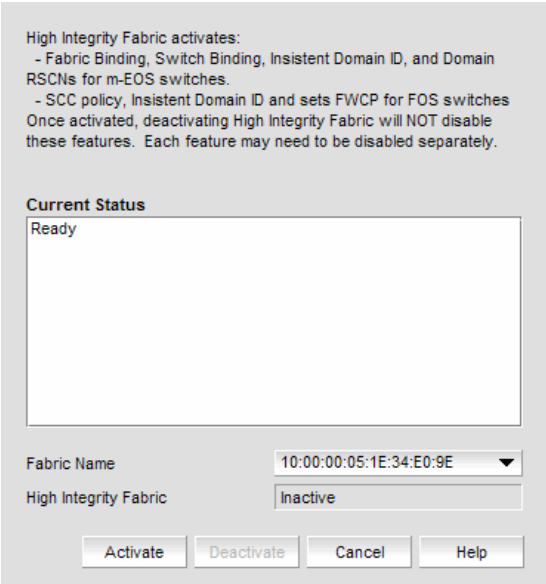
Port binding is a security method for restricting devices that connect to particular switch ports. Port binding should never be used in FICON environments. The FICON channel cannot be added to the port binding list.

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 380).



The screenshot shows the 'High Integrity Fabric' dialog box. At the top, it lists features activated: Fabric Binding, Switch Binding, Insistent Domain ID, and Domain RSCNs for m-EOS switches; and SCC policy, Insistent Domain ID, and FWCP for FOS switches. It notes that deactivating HIF will not disable these features. Below this is a 'Current Status' section with a text area displaying 'Ready'. At the bottom, there are two fields: 'Fabric Name' with a dropdown menu showing '10:00:00:05:1E:34:E0:9E' and 'High Integrity Fabric' with a dropdown menu showing 'Inactive'. At the very bottom are four buttons: 'Activate', 'Deactivate', 'Cancel', and 'Help'.

FIGURE 380 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.

To deactivate a HIF, complete the following steps.

1. Select **Configure > High Integrity Fabric**.

The **High Integrity Fabric** dialog box displays (Figure 380).

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

In this chapter

- [About port fencing](#) 863
- [Thresholds](#) 866
- [Adding thresholds](#) 869
- [Editing thresholds](#) 878
- [Removing thresholds](#) 884

About port fencing

NOTE

This feature is only available for Fabric OS devices.

NOTE

All Fabric OS devices must have Fabric Watch and must be running firmware Fabric OS 6.2 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 6.2 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 382).

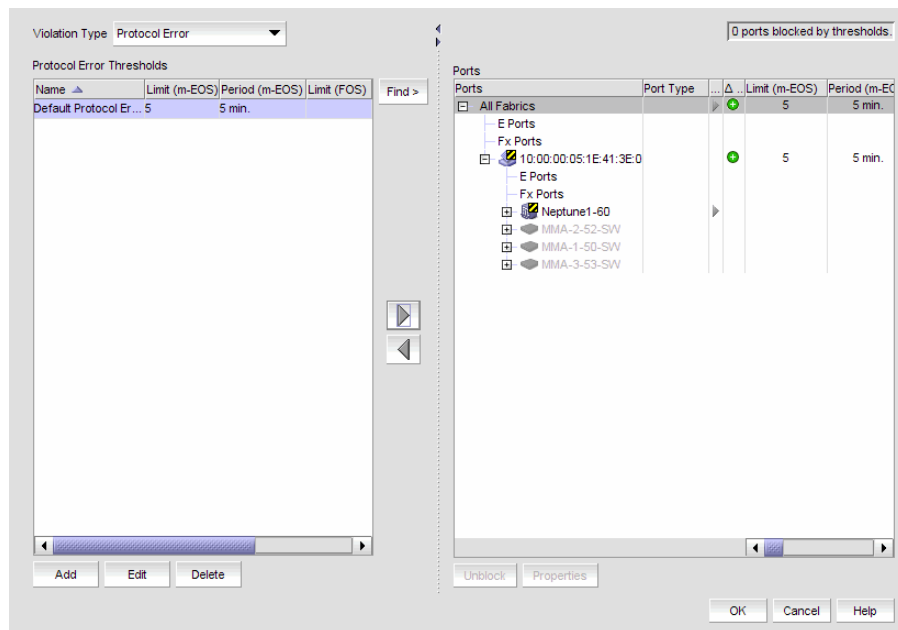


FIGURE 381 Port Fencing dialog box

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 877).
- 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 884).
- **Add** button — Click to add an ISL protocol threshold (refer to [“Adding thresholds”](#) on page 869).
- **Edit** button — Click to edit an ISL protocol threshold (refer to [“Editing thresholds”](#) on page 878).
- **Delete** button — Click to delete an ISL protocol threshold (refer to [“Removing thresholds from the thresholds table”](#) on page 884).
- **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 1265.
 - **Unblock** button – Click to unblock a blocked port after a warning message displays (refer to [“Unblocking a port”](#) on page 877). 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 6.3 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 directors, switches, and blades with a 4 Gbps, 8 Gbps, or 16 Gbps ASIC.

- 32-port, 4 Gbps FC Switch
- 64-port, 4 Gbps FC Switch
- 32-port, 4 Gbps FC Interop Switch
- 4 Gbps Router, Extension Switch
- 4 Gbps Extension Switch
- 4 Gbps 32-port Switch
- 8 Gbps 32-port Switch
- 8 Gbps 40-port Switch
- 8 Gbps 24-port Embedded Switch
- 8 Gbps 16-port Embedded Switch
- 8 Gbps 8-FC port, 10 GbE 24-DCB port Switch
- 16 Gbps 24-port Edge switch
- 16 Gbps 48-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
- 8 Gbps Encryption Switch
- Encryption Blade
- 10 Gbps FCoE Port Router Blade
- FC 8 GB 64-port 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

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 6.3 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 6.3 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 382).

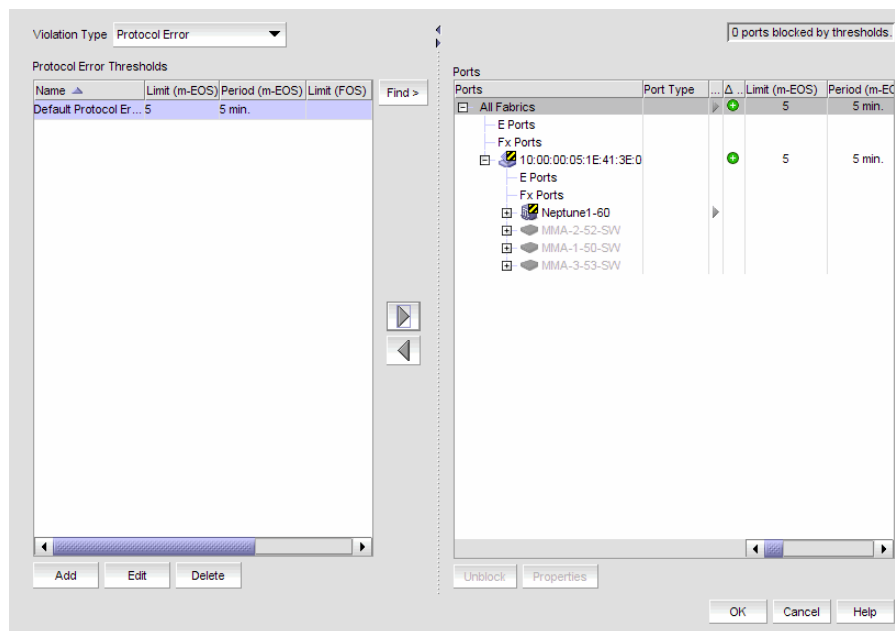


FIGURE 382 Port Fencing dialog box

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.

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 877.
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.

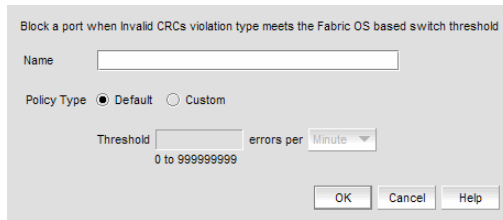


FIGURE 383 Add Invalid CRCs Threshold dialog box

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 877.
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 384 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 877.

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 385 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 877.
 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.

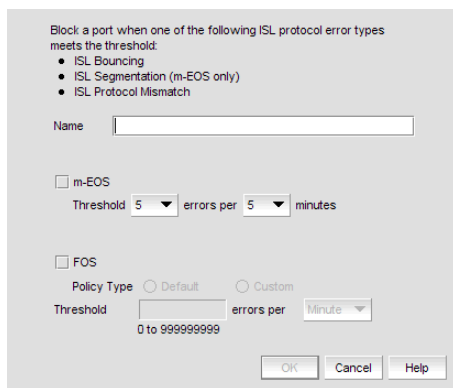


FIGURE 386 Add Protocol Error Threshold dialog box

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 877.
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 6.3 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 387 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 877.
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 869.
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 877.
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 870.
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 877.
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 872.

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 877.

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 873.

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 877.

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 874.
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 877.

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 6.3 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 875.

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 877.

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 inherit 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.

24 Removing thresholds

FICON Environments

In this chapter

- FICON configurations 887
- Configuring a switch for FICON operation 888
- Configuring an Allow/Prohibit Matrix 895
- Configuring an Allow/Prohibit Matrix manually 896
- Saving or copying Allow/Prohibit Matrix configurations to another device 898
- Activating an Allow/Prohibit Matrix configuration 900
- Deleting an Allow/Prohibit Matrix configuration 901
- Changing the Allow/Prohibit Matrix display 901
- Cascaded FICON fabric 902
- Cascaded FICON fabric merge 905
- Port groups 910
- Swapping blades 913

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 388](#) and [Figure 389](#) 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 388 Cascaded configuration, two domains




FIGURE 389 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 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 902.

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 to 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 to 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 to 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 905.

- If you are considering connecting cascaded switches over IP networks, refer to the planning considerations in the “Connecting cascaded FICON fabrics over FCIP” in [Chapter 22, “Fibre Channel over IP”](#).

Configuring the switch

Perform the following steps to configure a switch for FICON operation.

- 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.

- Configure the Management application display for FICON. Refer to the “Setting your FICON display” section of [Chapter 5, “Application Configuration”](#).
- 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.
- 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](#).

- Select **Add** on the **Discover Fabrics** dialog box.

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

- 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 390 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 391 Add Fabric Discovery dialog box (SNMP tab)

Refer to the “SAN discovery overview” section in [Chapter 4, “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 392](#) on page 892.

25 Configuring a switch for FICON operation

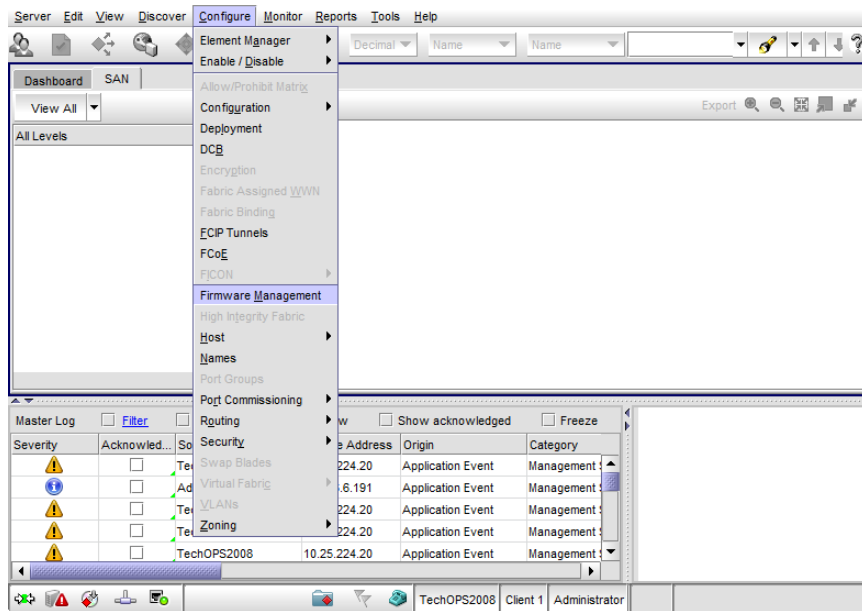


FIGURE 392 Selecting Firmware Management from Configure menu

The Firmware Management dialog box displays.

10. Select the **Download** tab (Figure 393).




FIGURE 393 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 [Chapter 12, “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 548.

15. To configure the switch as part of a fabric, follow procedures under “[Configuring a cascaded FICON fabric](#)” on page 903, 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 [Chapter 24, “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 [Chapter 24, “Port Fencing”](#).

19. Set the zoning policy for the switch by referring to steps under “Enabling or disabling the default zone for fabrics” in [Chapter 21, “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 [Chapter 21, “Zoning”](#).

Be sure to reference the “Zoning and FICON” section of [Chapter 21, “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 895.

22. Configure Call Home by referring to procedures in [Chapter 9, “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 **bottleneckmon** 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 **bottleneckmon** 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 394](#) on page 895.

FC Address	Port Name	Blocked	91	92	93	94	95	96	97	98	99	9A	9E
12		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
13		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
14		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
15		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
16		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
17		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
18		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
19		<input type="checkbox"/>	●	●	●	●	●	⊘	●	●	●	●	●
1A		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
1B		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
1C		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
1D		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
1E		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●
1F		<input type="checkbox"/>	●	●	●	●	●	●	●	●	●	●	●

FIGURE 394 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 896.

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 395](#) on page 896.

FIGURE 395 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 396](#) on page 898.

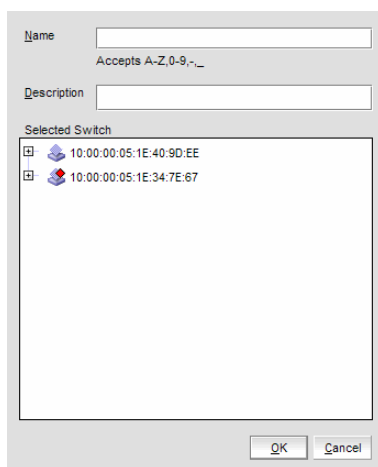


FIGURE 396 Save As/Duplicate dialog box

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 895 or [“Configuring an Allow/Prohibit Matrix manually”](#) on page 896.

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 397](#) on page 899.




FIGURE 397 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 895 or “[Configuring an Allow/Prohibit Matrix manually](#)” on page 896.

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 398](#) on page 900.




FIGURE 398 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.

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 port-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 6.4 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 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.

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** screen of the **FICON Configuration** dialog displays, as shown in [Figure 399](#) on page 903.

FIGURE 399 Configure Cascaded FICON Fabric /Switch dialog box

2. Use the **Fabric** list to select the fabric you want to configure.

NOTE

(Fabric OS switches only) All switches in a fabric must be running Fabric OS version 5.3 or later. If a Fabric OS version earlier than version 5.3 is present in the topology, the fabric is not listed.

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 **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 904. You must enable DLS to select routing policies.

5. 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 6.4 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

Either 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.

6. 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.

7. 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 4](#):

- 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 6.4 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 an 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 6.4 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 6.3 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 distances 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

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 5.3 or later and must be reachable. If a Fabric OS version earlier than version 5.3 is present in the fabric, the fabric is not listed.

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.

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 a secondary fabric where AD, ACL, and zone databases, if defined, will be cleared.
7. Read the bulleted list of actions so you understand the actions that are taken to avoid conflicts when the fabrics are merged.

8. 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 blue check mark indicates successful completion of that **Merge check**. A red stop sign indicates a failed step. If the configuration is successful, all configuration items have blue check marks.

9. If the merge fails, but is recoverable, click **Resolve**.
10. If desired, click **Check Merge Again** to run the merge check test again.
11. Click **Next** to continue.

The **Configure long distance (optional)** dialog box displays. If the distance between the merged fabrics is 10 km or greater, you must configure the connection as a long distance connection. Selecting a distance invokes an algorithm to compute the required number of BB credits available to the port. The longer the link, the greater latency, resulting in the potential for more outstanding frames in the link, and the need for more BB credits. FICON may require more BB credits than the algorithm provides, and it is a good practice to specify a distance that is longer than the actual distance to be sure enough BB credits are allocated.

12. Perform the appropriate following action based on whether the connection is a long distance connection:
 - If it is not a long distance connection, click **Next** to view the **Configure merge** screen. Proceed to [step 13](#).
 - If it is a long distance connection, expand the fabrics under **Selected Fabrics** to the switch port level.
 - a. Select the E_Ports used for the connection on the local switch and on the remote switch, and click the right arrow.
The selected E_Ports are moved to **Selected Ports**.
If there is no E_Port in the selected fabrics, a warning message displays.
 - b. Specify the **Cable length between switch ports**.
The range is from 10 through 500 km. The default is 50 km.
 - c. Select **ARBs** or **IDLEs** to configure the **Fibre Channel Primitive Signal Fill Words**.
For Fabric OS version 6.1.0b or earlier, the setting is always **ARBs**. You cannot change to **IDLEs**.
For Fabric OS version 6.1.0c or later, the default setting is **IDLEs**, however, you can change it to **ARBs**.
 - d. Click **Next**.
The **Configure merge** screen displays.
13. 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.
14. 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 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 11](#) through [step 14](#) of the procedure “[Merging two cascaded FICON fabrics](#)” on page 907 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 895.

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 400](#) on page 911.

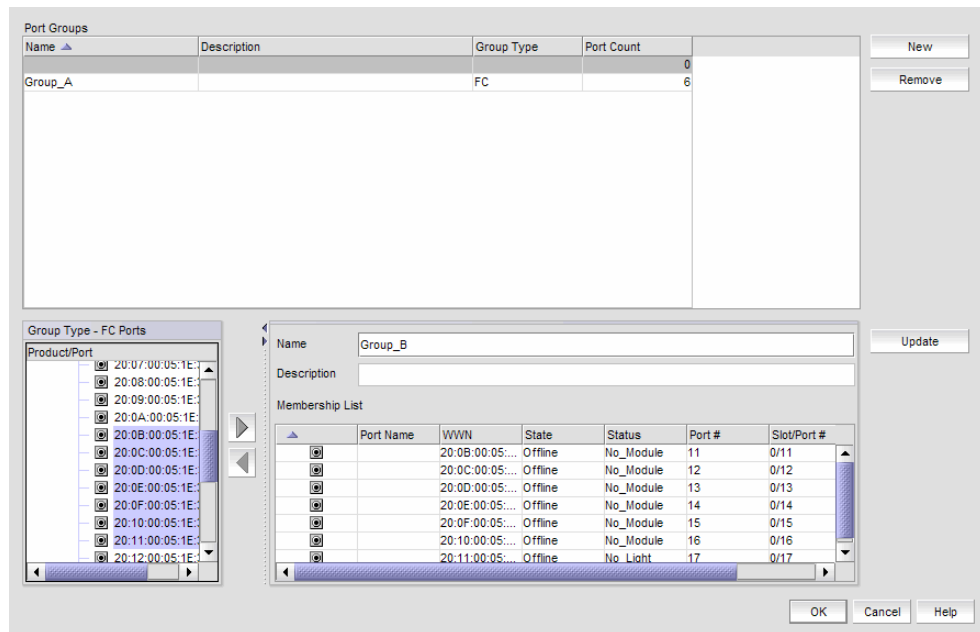


FIGURE 400 Port Groups dialog box

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 51.

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 912.

3. Delete the port group, as needed.

To delete a port group, refer to [“Deleting a port group”](#) on page 913.

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 6.3 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

In this chapter

- [Introduction to the Deployment Manager](#) 915
- [Editing a deployment configuration](#) 916
- [Duplicating a deployment configuration](#) 916
- [Deleting a deployment configuration](#) 917
- [Deploying a configuration](#) 917
- [Viewing deployment logs](#) 917
- [Generating a deployment report](#) 918
- [Generating a deployment configuration snapshot report](#) 918
- [Searching the configuration snapshots](#) 918

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 471 (for DCB configurations)
- [“VLAN Management”](#) on page 925 (for VLAN and STP configurations)
- [“Security Management”](#) on page 523

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 401](#).

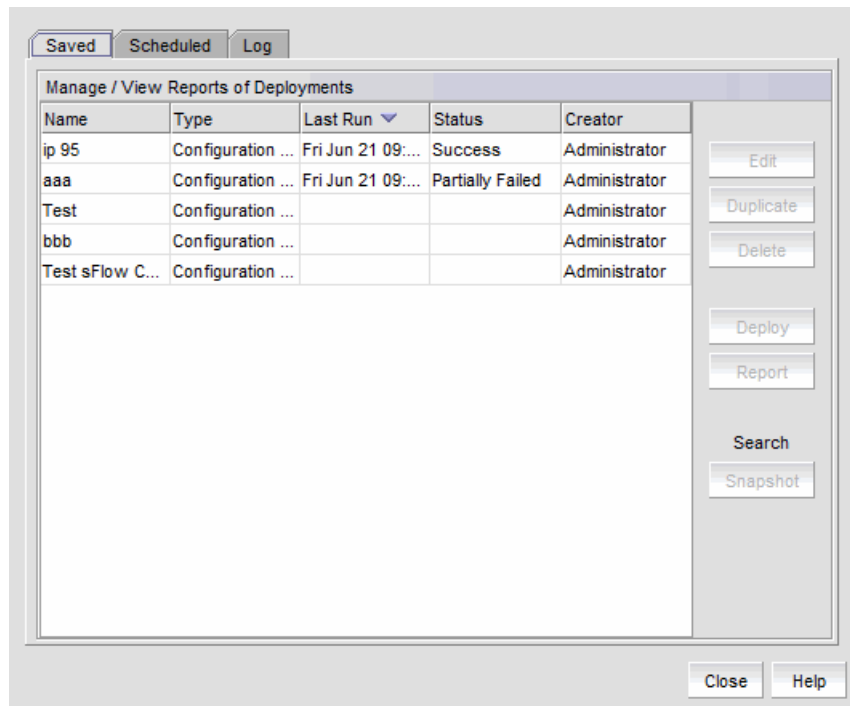


FIGURE 401 Task Scheduler dialog box

2. Select a deployment configuration in the **Saved** or **Scheduled** tab.
3. Click **Edit**.

Policy-based routing configurations cannot be edited.

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 copyn*". 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,

25 Searching the configuration snapshots

Fibre Channel Troubleshooting

In this chapter

- [FC troubleshooting](#) 921
- [FCIP troubleshooting](#)..... 930

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 and 16 Gbps-capable ports only)** – Use to run the following diagnostic port tests on the 10 Gbps-capable and 16 Gbps-capable ports: electrical (16 Gbps only), optical (16 Gbps only), measure link distance, and link traffic.

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 5.2 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 6.2 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 991.

6. Click **Close** on the **Trace Route Summary** dialog box.
7. Click **Cancel** on the **Trace Route** dialog box.

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 924.
 - 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 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.
- Only supported on devices with 16 Gbps-capable F-ports, ICL-ports, and AG N-ports running Fabric OS 7.1 or later.
- 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.

TABLE 63 D-Port test support matrix

D-Ports Tests		Fabric OS 7.0	Fabric OS 7.1				HBA driver 3.2
		E-Port	E-Port	F-Port	AG N-Port	ICL-Port	
Electrical Test		Supported	Supported	Supported	Supported	Not supported	Supported
Optical Test		Supported	Supported	Supported	Supported	Not supported	Not supported
Link Traffic Test	Configuration	Not supported	Supported	Supported	Supported	Supported	Not supported
	Test	Supported	Supported	Supported	Supported	Supported	Not supported
Link Distance		Supported	Supported	Supported	Supported	Supported	Not supported
Link Measurement		Not supported	Not supported	Supported	Supported	Supported	Not 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.

2. Select the ports 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 **Link Traffic Test**.

Refer to [“Configuring link traffic test parameters”](#) on page 929.

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:

The following tests are performed on the selected ports:

- Electrical loopback (16 Gbps only)
- Optical loopback (16 Gbps only)

- Link traffic
- Latency measurement
- Measure link distance

TABLE 64 Supported link distance measurements

SFP speed	Accuracy	Precision
10 Gbps	124 meters	+ or - 50meters
16 Gbps	5 meters	+ or - 5 meters

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.

6. Select a port row in the Selected Ports table to display the detailed status in the **Status Details of the Selected Row** table.

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 65 Status Detail messages

Operation/Test	Possible message
Disable the source or destination port	Disabled the port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> .
	Failed to disable the port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> . Reason: <i>CAL_error_message</i>
Enable the diagnostic mode on source or destination E ports	Enabled diagnostic mode on port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> .
	Failed to enable diagnostic mode on port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> . Reason: <i>CAL_error_message</i>
Enable the source or destination port	Enabled the port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> .

TABLE 65 Status Detail messages

Operation/Test	Possible message
	Failed to enable the port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> . Reason: <i>CAL_error_message</i>
Disable the diagnostic mode on source or destination D ports	Disabled diagnostic mode on port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> . Failed to disable diagnostic mode on port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> . Reason: <i>CAL_error_message</i>
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 <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> . Electrical Loopback Test failed on port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> . Reason: <i>CAL_error_message</i> Electrical Loopback Test skipped on port <i>slot_number/port_number</i> of the switch <i>switch_IP_address</i> . Reason: <i>CAL_error_message</i>
Optical Loopback Test	Successfully completed Optical Loopback Test. Optical Loopback Test failed Optical Loopback Test skipped. Reason: <i>CAL_error_message</i>
Link Traffic Test	Successfully completed Link Traffic Test. Link Traffic Test failed.
Distance between ports	Approximate distance between the ports is <i>numerical_value</i> meters.
Reverse Optical Loopback Test	Successfully completed Reverse Optical Loopback Test. Reverse Optical Loopback Test failed.
Roundtrip link latency	Roundtrip link latency: <i>numerical_value</i> nano-seconds.
Buffers required	Buffers required: <i>numerical_value</i>
Link Traffic Test Configuration used	No. of test frames: <i>numerical_value</i> Million Duration of test (HH:MM): hours:minutes Test frame size: <i>numerical_value</i> Bytes Payload Pattern: <i>pattern_name</i> or <i>frame_data</i> 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.

TABLE 65 Status Detail messages

Operation/Test	Possible message
If any test fails, that test displays as failed and a Failure report displays.	<p>Sample failure report :</p> <p>Errors detected (local): CRC, Bad_EOF, Enc_out</p> <p>Errors detected (remote): CRC, Bad_EO</p> <p>Run portstatssh and porterrshow for more detail on the errors.</p>
HBA Electrical test successful	Successfully completed Electrical Loopback Test on port <i>HBA_port_number</i> of the HBA <i>HBA_node</i>
HBA Electrical test failed	Electrical Loopback Test failed on port <i>HBA_port_number</i> of the HBA <i>HBA_node</i>

7. Click **Close** on the **Diagnostic Port Test** dialog box.

Configuring link traffic test 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 pateren 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).

- Click **OK** on the **Link Traffic Test Configuration** dialog box.

The **Diagnostic Port Test** dialog box displays. Return to [step 5](#) of “[Troubleshooting port diagnostics](#)” on page 925.

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 running Fabric OS 5.2 or later.

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.

- Select **Monitor > Troubleshooting > FCIP > Ping**.
The **IP Ping** dialog box displays.
- Select a switch from the **Available Switches** table.
- Select a port from the **GigE Port** list.
- Select an IP address switch from the **IP Interface** list.
- Enter the remote IP address in the **Remote IP Address** field.
- 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 66 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 67 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.

NOTE

Trace route is only supported on Fabric OS devices running Fabric OS 5.2 or later.

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 68 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 running Fabric OS 5.2 or later.

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:

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 acknowledgements.
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.

Field/Component	Description
DELAY	The average round trip time to send a packet of data and receive the acknowledgement.
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

In this chapter

- [SAN performance overview](#) 935
- [SAN real-time performance data](#) 942
- [SAN historical performance data](#) 946
- [SAN end-to-end monitoring](#)..... 957
- [SAN Top Talker monitoring](#) 961
- [Bottleneck detection](#) 967
- [Thresholds and event notification](#) 974
- [SAN connection utilization](#) 979

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, device Ports, managed HBA Ports, and managed CNA 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, device Ports, managed HBA Ports, managed CNA 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 14 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 6.4.1 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, managed HBA ports, managed CNA ports, E_port trunks, 10GE ports, and FCIP tunnels.
- Rx % Utilization — Available for FC, GE, managed HBA ports, managed CNA ports, 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.
- Dropped Packets – Available for FCIP tunnels only.
- 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 and managed CNA ports.
- Non Frame Coding Error – Available for managed HBA ports and managed CNA ports.

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: 1001:0:0:0:0:0:172
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 37.
 - 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 Fabric-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, managed HBA Ports, managed CNA 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 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.

NOTE

To make sure that statistics for a switch does not fail, you must configure SNMP credentials for the switch. For step-by-step instructions, refer to [“Discovery”](#) on page 37.

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
2. Select **Monitor > Performance > Real-Time Graph**.

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 944.

If you selected a fabric or device, the **Realtime Port Selector** dialog box displays, as shown in [Figure 402](#) on page 943. Continue with [step 3](#).

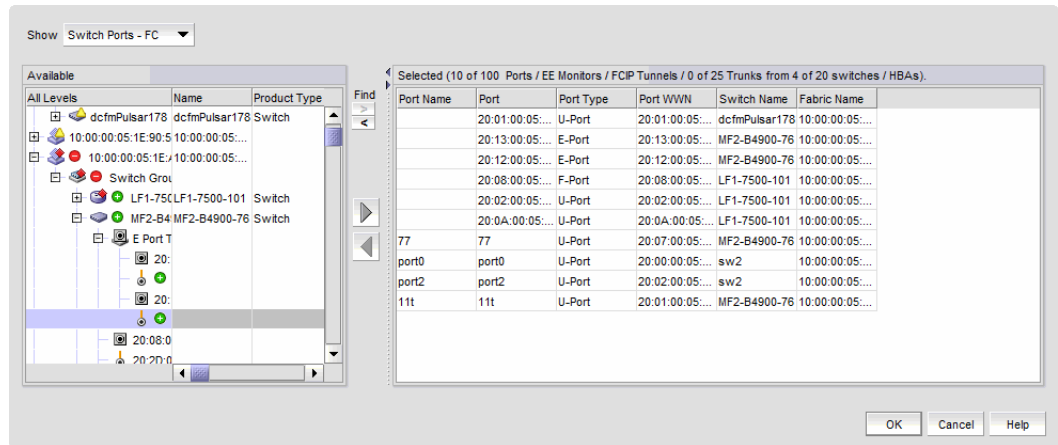


FIGURE 402 Realtime Port Selector dialog box

NOTE

You can set columns in 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.

3. Select the object type from the **Show** list by which you want to graph performance.

NOTE

Devices with 10GE ports must be running Fabric OS 6.4.1 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. Right-click anywhere in the **Available** list and select **Expand All** from the menu.
5. Select the ports or trunks you want to include in the performance graph in the **Available** list. Press **Ctrl** or **Shift** and then click to select more than one port.
6. Click the right arrow to move the selected ports to the **Selected** list.
7. 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 942.
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 936.
3. To select more than one measure, click the **Additional Measures** expand arrows and select the check box for each additional measure. If **Additional Measures** area is not shown, click the down arrow.
For a list of possible performance measures, refer to “[SAN performance measures](#)” on page 936.
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:

- Source Fabric - The source fabric being monitored.
- Source - The source device being monitored.
- Source Port - The source port being monitored.
- Tunnel ID - The ID of the FCIP tunnel being monitored.
- Destination Fabric - The destination fabric.
- Description - Description of the FCIP tunnel.

- 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.
- Destination Tunnel ID - The ID of the destination FCIP tunnel.
- Destination Port Type - The port type through which the selected device is connected to the destination device.
- Additional Measures columns - Displays each measure selected in the **Measures** list and **Additional Measures** area.
- Measures columns - A column for each selected measure in the **Measures** list or **Additional Measures** area.

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 942.
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 statistic counters to zero on a selected device, complete the following steps.

1. Right-click a device on the Connectivity Map or Product List and select **Monitor > Performance > Clear Counters**.
2. Click **Yes** on the message.
A **Port Stats Counter Reset** message displays. If any of the counters do not clear, the message displays a list of the associated ports.
3. Click **OK** on the **Port Stats Counter Reset** message.

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 14 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 1203.
- Configure the graph display using right-click menu options. For more information refer to [“Configuring the graph display”](#) on page 950.

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 403](#) on page 947.

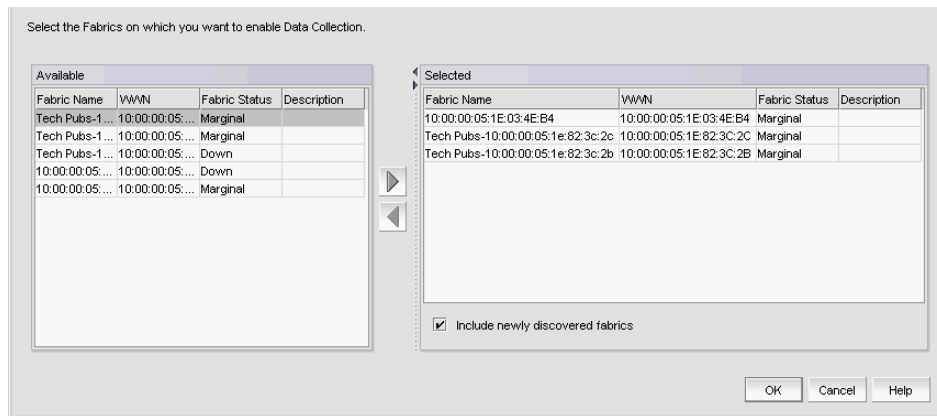


FIGURE 403 Historical Data Collection dialog box

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 6.4.1 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 [“Configuring the performance graph display”](#) on page 982.

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 404](#) on page 948.




FIGURE 404 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 6.4.1 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 951.

- 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 952.

- 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 14 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. 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 948.

To delete user-defined graph, refer to [“Deleting a favorite graph configuration”](#) on page 953.

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 950.

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 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 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 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 405](#) on page 951.

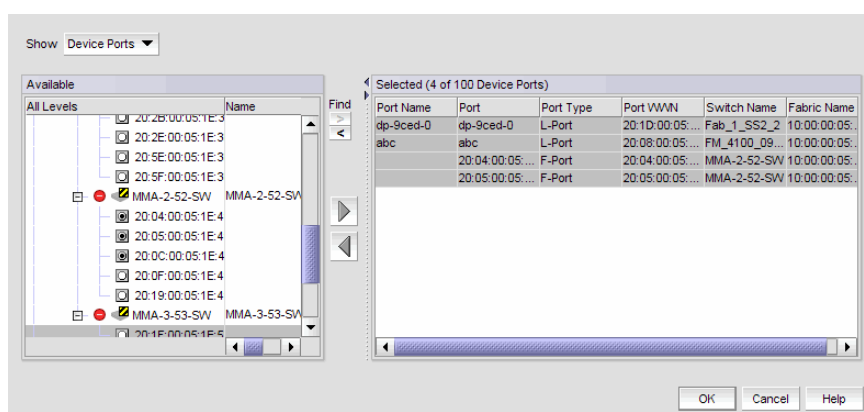


FIGURE 405 Custom Port Selector dialog box

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

Devices with 10GE ports must be running Fabric OS 6.4.1 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 406](#) on page 952. 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 406 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 948.

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 *<Management_Application install-home> \conf \schema \dcm-postgres-schema.sql* location
 - Open the PostgreSQL user interface by double clicking on *<Management_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 [Appendix H, "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_2HOURL 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 69](#) 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 69](#)).

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 69 Performance statistic counters

Counter name	Type	Protocol	Source OID value	Formula
TX	FC	SNMP	.1.3.6.1.3.94.4.5.1.6	$\text{TX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
RX	FC	SNMP	.1.3.6.1.3.94.4.5.1.7	$\text{RX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
TX	GE	SNMP	.1.3.6.1.2.1.31.1.1.1.10	$\text{TX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
RX	GE	SNMP	.1.3.6.1.2.1.31.1.1.1.16	$\text{RX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
TX	FCIP	SNMP	.1.3.6.1.2.1.31.1.1.1.10	$\text{TX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
RX	FCIP	SNMP	.1.3.6.1.2.1.31.1.1.1.16	$\text{RX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$

TABLE 69 Performance statistic counters

Counter name	Type	Protocol	Source OID value	Formula
Uncompressed Tx/Rx MB/sec	FCIP	SNMP	.1.3.6.1.4.1.1588.4.1.1.6	$(\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
TX	EE Monitors	HTTP	PortRX (variable from the return html file)	$\text{TX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
RX	EE Monitors	HTTP	PortTX (variable from the return html file)	$\text{RX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
TX	HBA, CNA	HCM API	N/A	$\text{TX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
RX	HBA, CNA	HCM API	N/A	$\text{RX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
TX	TE	SNMP	.1.3.6.1.2.1.31.1.1.1.6	$\text{TX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
RX	TE	SNMP	.1.3.6.1.2.1.31.1.1.1.10	$\text{RX} = (\text{delta value}^1 / (1000 * 1000)) / (\text{polling interval}^2)$
TX% / RX%	FC, GE, HBA, CNA	N/A	N/A	$\text{TX\% or RX\%} = ((\text{TX or RX}) / ((105000000 * \text{port speed}) * (\text{polling interval}^2))) * 100$ If utilization is less than 1, the value is 0.0.
TX% / RX%	FCIP	N/A	N/A	$\text{TX\% or RX\%} = ((\text{bytes transferred}) / (\text{maximum bytes transmitted})) * 100$ where maximum bytes transmitted = tunnel speed * 134217728 (maximum bytes transmitted 1 Gbps). If utilization is less than 1, the value is 0.0.
TX% / RX% (Pre-Fabric OS 6.4.1 release)	TE	N/A	N/A	$\text{TX\% or RX\%} = ((\text{TX or RX}) / ((105000000 * 10) * (\text{polling interval}^2))) * 100$ If utilization is less than 1, the value is 0.0.
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.
Receive EOF	TE		.1.3.6.1.2.1.16.1.1.1.5	Receive EOF = $\text{delta value}^1 / (1000 * 1000)$
Other ³				Other counters = delta value^1

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 69](#).

[Table 70](#) lists the additional counters for which you can obtain performance statistics.

TABLE 70 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
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
Encode Error Out	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
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.

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 [Chapter 28, “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 [Chapter 28, “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 [Chapter 28, “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 407](#) on page 958.

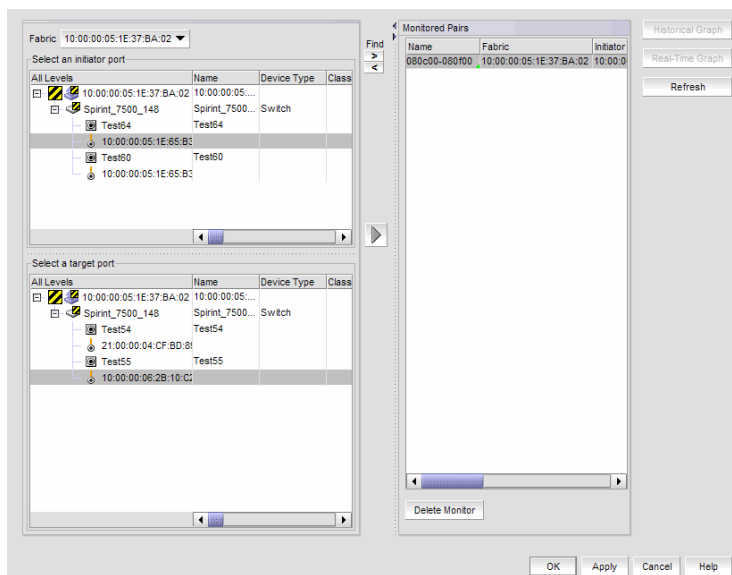


FIGURE 407 Set End-to-End Monitors dialog box

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 959 or “[Displaying end-to-end monitor pairs in a historical graph](#)” on page 960.

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 [Chapter 28, “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 [Chapter 28, “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 [Chapter 28, “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 [Chapter 28, “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.

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 [Chapter 28, “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 6.2 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 [Chapter 28, “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 408](#) on page 963.

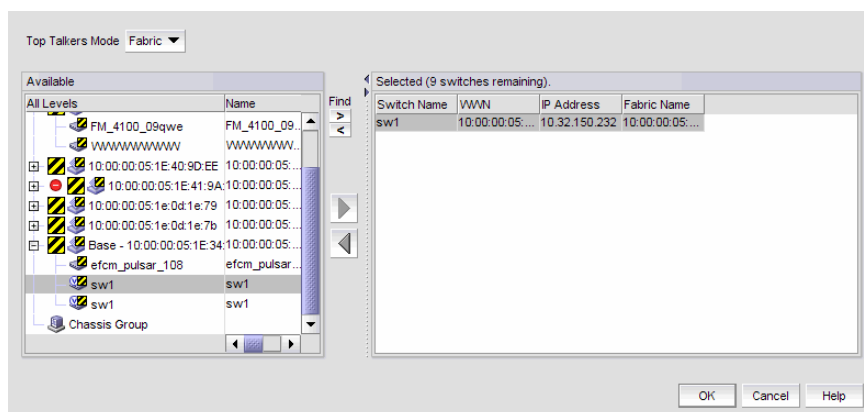


FIGURE 408 Top Talker Selector dialog box

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

8. Click **Destination** to launch the **Port Properties** dialog box for the Destination port.
9. 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 [Chapter 28, “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 [Chapter 28, “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 962 or [“Configuring an F_Port mode Top Talker monitor”](#) on page 964 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 [Chapter 28, “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 962 or [“Configuring an F_Port mode Top Talker monitor”](#) on page 964 to display this dialog box.

2. Click **Pause** at the top of the dialog box.

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 [Chapter 28, “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 962 or [“Configuring an F_Port mode Top Talker monitor”](#) on page 964 to display this dialog box.

2. Click **Continue**.

Bottleneck detection

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 6.4.0 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 409](#). In the figure, port15 and port22 are bottlenecked.

All Levels	Name	Product
Bottleneck	Bottleneck	
Switch Group		
dcm-5100-203	dcm-5100-203	Switch
dcm-5100-204	dcm-5100-204	Switch
10:00:00:03:12:09:13:01		
10:00:00:03:12:09:14:01		
20:0E:00:05:1E:85:9B:40	port14	
20:0F:00:05:1E:85:9B:40	port15	
20:16:00:05:1E:85:9B:40	port22	

FIGURE 409 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 971 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 410](#) on page 970.

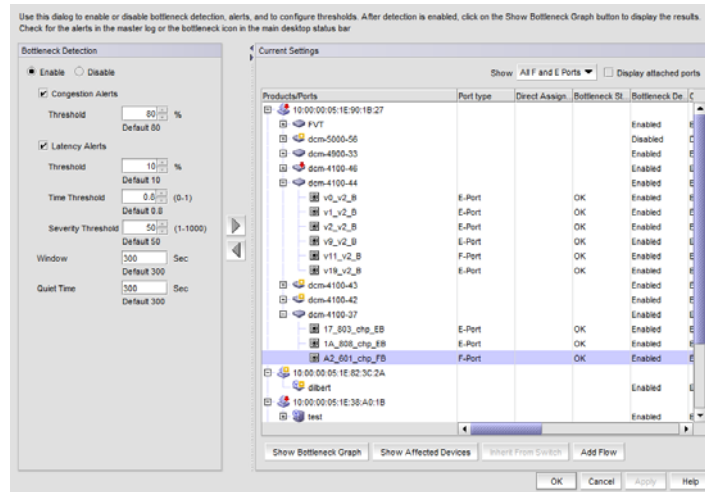


FIGURE 410 Bottlenecks dialog box

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 configure 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 972.
 - **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 [“Flow provisioning and monitoring”](#) on page 994 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 as 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

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 1064.

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 411](#) on page 974.




FIGURE 411 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 412](#) on page 975.

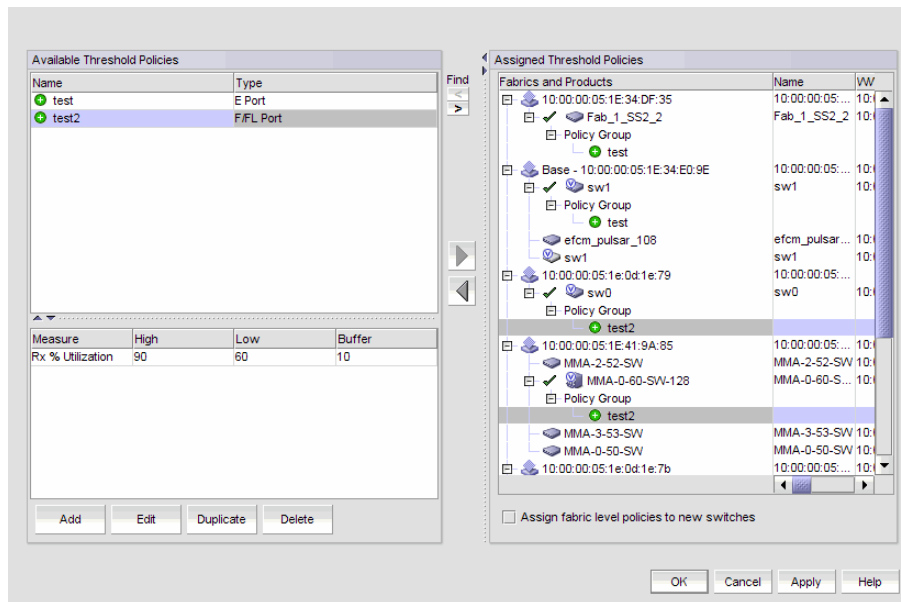


FIGURE 412 Set Threshold Policies dialog box

NOTE

Policies set for switches enabled for Monitoring and Alerting Policy Suite (MAPS) also display in this dialog box.

2. To edit a current policy, select a policy from the **Available Threshold Policies** list and click **Edit**. The **Edit Threshold Policy** dialog box displays, as shown in [Figure 413](#) on page 975.

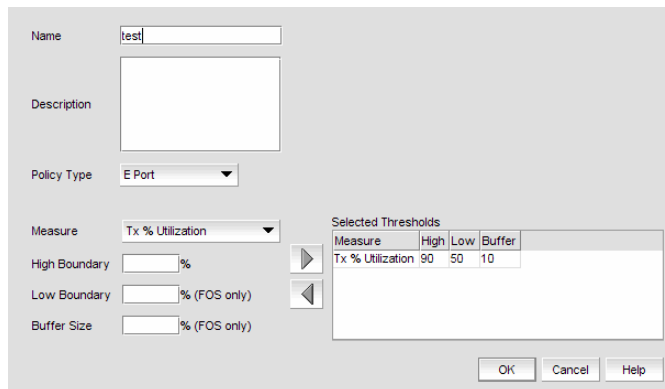


FIGURE 413 Edit Threshold Policy dialog box

3. To add a new policy, perform the following steps.
 - a. Click **Add**.

The **New Threshold Policy** dialog box displays as shown in [Figure 414](#) on page 976.

The screenshot shows a 'New Threshold Policy' dialog box. It contains the following elements:

- Name:** A text input field.
- Description:** A larger text input area.
- Policy Type:** A dropdown menu currently set to 'E Port'.
- Measure:** A dropdown menu currently set to 'Tx % Utilization'.
- High Boundary:** A text input field followed by a '%' sign.
- Low Boundary:** A text input field followed by a '%' sign and '(Fabric OS only)'.
- Buffer Size:** A text input field followed by a '%' sign and '(Fabric OS only)'.
- Selected Thresholds:** A table with columns 'Measure', 'High', 'Low', and 'Buffer'. The table is currently empty.
- Navigation:** Two arrow buttons (right and left) are positioned between the boundary fields and the 'Selected Thresholds' table.
- Buttons:** 'OK', 'Cancel', and 'Help' buttons are located at the bottom right.

FIGURE 414 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 policies 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 415](#) on page 977.




FIGURE 415 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 974. To assign a threshold policy to a fabric or device, refer to [“Assigning a threshold policy”](#) on page 978.

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 416](#) on page 980.




FIGURE 416 Utilization Legend


The colors and their meanings are outlined in [Table 71](#) on page 980.

TABLE 71 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.

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 946.

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 417](#) on page 981.

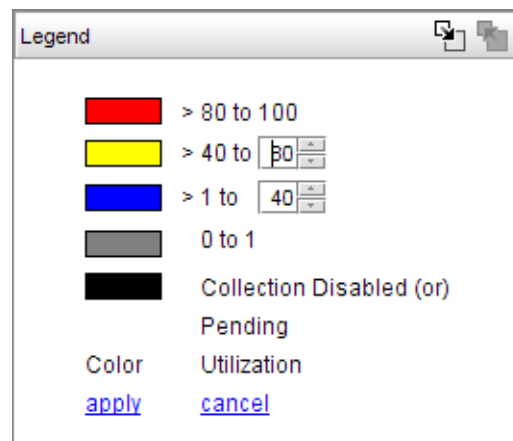


FIGURE 417 Utilization Legend in edit mode

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**.

Configuring the performance graph display

Use the 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 or **Performance Dashboard**.

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 7](#) 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 **Deleted 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 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 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 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 983 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 983 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.

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 or **Performance Dashboard**.

1. Click **Options** on the graph.

The **Graph Options** dialog box displays, as shown in [Figure 418](#) on page 984.

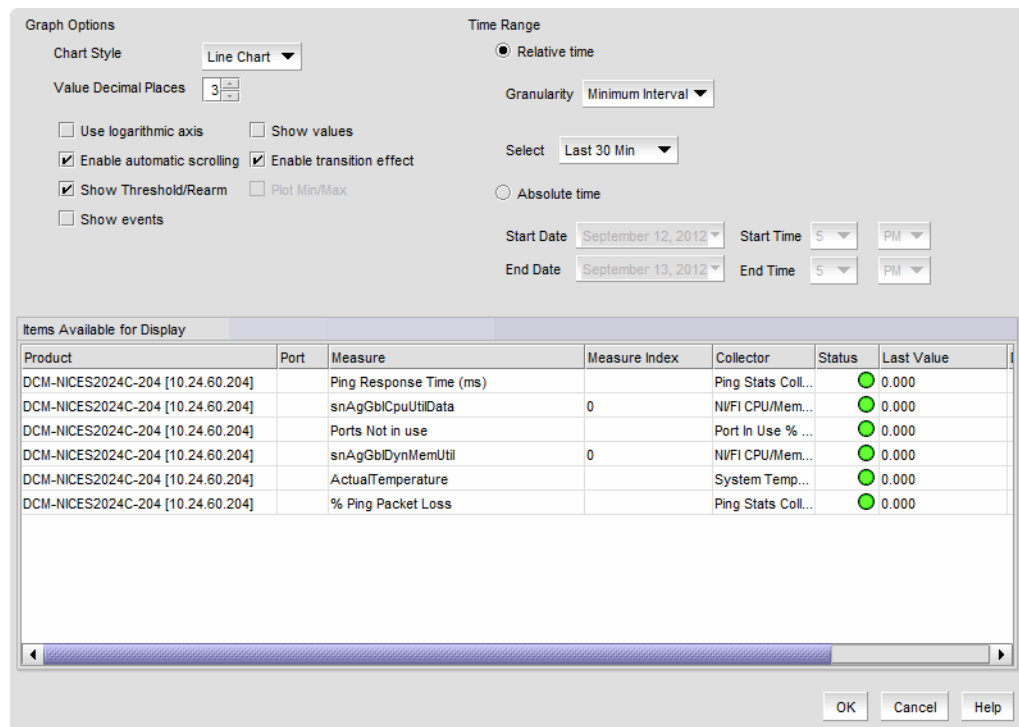


FIGURE 418 Graph Options dialog box (Historical Graphs/Tables dialog box)

NOTE

Figure 418 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** check box — 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** check box — Annotates data point values in the graph.
 - **Enable automatic scrolling** check box — If new data is collected while the chart is in view, the chart will automatically jump to display the new data.
 - **Enable transition effect** check box — 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 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.

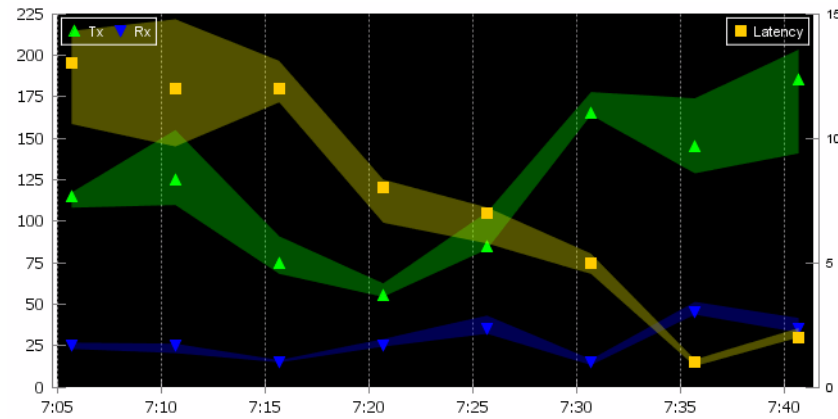


FIGURE 419 Data points graph

- **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 icons a bottom of following graph). Hovering the cursor 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 420](#) on page 985




FIGURE 420 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 **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.
 - b. Select the start time (1 through 12) from the first **Start Time** list.
 - c. Select **AM** or **PM** from the second **Start Time** list.
 - d. Select the end date from the **End Date** list.
 - e. Select the end time (1 through 12) from the first **End Time** list.
 - f. Select **AM** or **PM** from the second **End Time** list.
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 (integer between -2147483648 and 2147483647) in the **Scaling Factor** column for each item in the **Items Available for Display** table.
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 421](#) and [Figure 422](#) for examples.

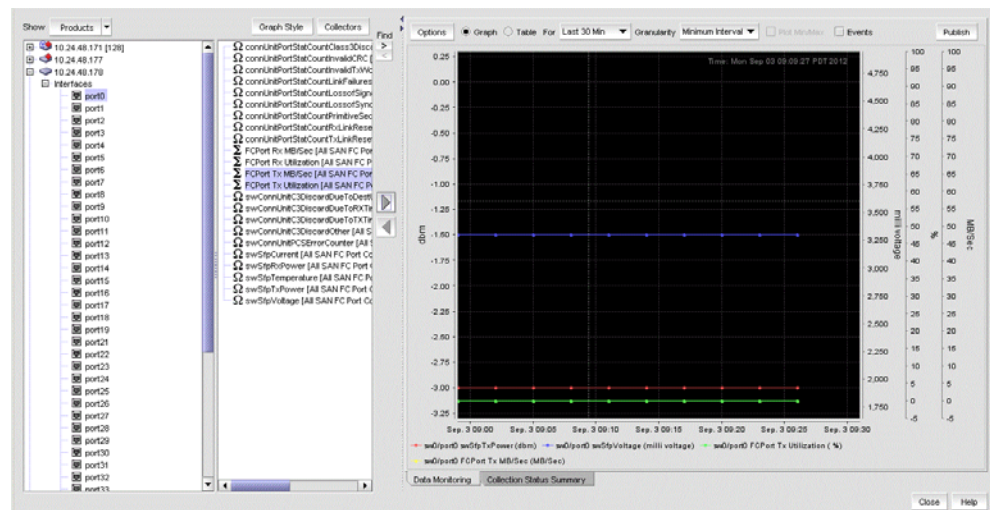


FIGURE 421 SAN Fibre Channel port display

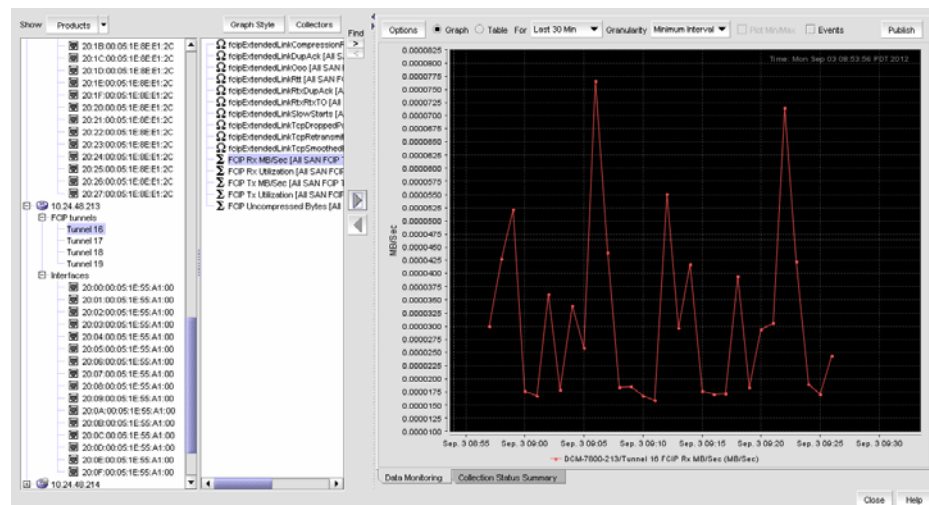


FIGURE 422 SAN FCIP tunnel display

- 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. You can select these collectibles to create applicable historical graphs and tables.

4. Select **Collectors** to open the **Historical Data Collectors** dialog box. Use this dialog box to display, enable or disable, add or delete, and duplicate historical data collectors.
5. To configure the look and feel of the performance graph from the **Historical Graphs/Tables** dialog box, refer to [“Configuring the performance graph display”](#) on page 982.
6. 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. 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.

27 SAN connection utilization

Flow Vision

In this chapter

- [Overview](#) 991
- [Provisioning flows](#) 996
- [Monitoring Flows](#) 1005
- [Dashboard flow performance monitor](#) 1015
- [Flow Vision features](#) 1016
- [Context-based flow definitions](#) 1023
- [Flow parameter and configuration rules and limitations](#) 1024
- [Accessing Flow Vision from other management application features](#) . . . 1027

Overview

Flow Vision is a function of Fabric Vision, which also includes the Monitoring and Alerting Policy Suite (MAPS) feature. For details on MAPS, refer to [Chapter 32, “Monitoring and Alerting Policy Suite”](#) for more information on that feature.

Flow Vision is a network diagnostic tool that provides a unified platform to manage traffic-related applications on Fabric OS devices. Storage administrator can use this platform to simulate, monitor and capture the network’s traffic patterns and to make capacity-planning decisions based on the collected statistical data.

Flow Vision is supported on platforms that support Fabric OS 7.2 and later. It provides the following features:

- Infrastructure for rapid isolation of network issues, service modeling, and accurate measurement.
- Greater visibility into data centers, allowing you to maximize storage network performance and reliability.
- Interfaces to Fabric OS applications to manage traffic and connectivity analysis applications through Flow Vision.

Flow Vision allows you to work in the following areas without using external test equipment, such as a SAN tester or Finisar analyzer.

- Accelerating deployment
- Application performance optimization
- Problem avoidance
- Rapid resolution and recovery

Why Flow Vision exists

As storage networks become larger and more complicated, storage administrators need methods to analyze flows (and dynamic nature of the network) so that they can obtain benefits such as the following:

- Instant identification of the source of performance or other problems. For example, is the problem in an application, SAN, or in storage.
- Tools for pre-deployment verification
 - Hardware component verification - verifying there are no link-level or SFP transceiver issues.
 - Connectivity validation - verifying ASIC blocks such as routing, zoning, and traffic isolation (TI) zoning.
 - Performance verification - verifying link bandwidth and other features.
- Real-time performance and frame monitoring in the SAN.
- Network dynamics (for example, ISL failure).
- Application performance optimization.

Fabric Vision components

[Table 72](#) describes the Fabric Vision components.:

TABLE 72 Fabric Vision components

Component	Description
Flow Vision	Provides: <ul style="list-style-type: none"> • Flow-related data for traffic analysis to various applications • Interfaces to manage flows with flexible parameters for traffic-related applications • Infrastructure to host applications related to traffic functionality • Flow configuration management
Monitoring and Alerting Policy Suite (MAPS)	Responsible for monitoring critical components including port error statistics, traffic analysis, critical resource usage, switch health monitoring, FRU states, and so on. Flow Vision provides data related to flows for traffic analysis in MAPS. For more information on using MAPs, refer to Chapter 32, “Monitoring and Alerting Policy Suite” .
Fabric OS applications	Provides interfaces to Fabric OS applications, specifically diagnostic tools (for example, super ping), allowing you to manage hosted traffic and connectivity applications in Flow Vision for detailed analysis.
Platform services	These populate and update the data in the system database hierarchy.
Switch services	These receive system-related events and interact with hardware.

Flow Vision licensing

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. Refer to [Chapter 2, “Licenses”](#) for more information on using licenses.

Flow Vision support

Flow Vision is supported on platforms using 8 Gbps and 16 Gbps-capable Fibre Channel platforms; there are no platform exclusions. All ports other than those listed under “[Unsupported ports](#)” on page 993 are supported. This includes ICL ports.

For details on Flow Vision feature and parameter support on switch platforms, Access Gateway switches, and virtual fabrics, refer to “[Flow parameter support](#)” on page 1022.

Supported ports

The following port types are supported for Flow Vision. For more information on support for a specific feature, such as Flow Monitor, Flow Mirror, or Flow Generator, refer to appropriate sections on those features.

- E_Ports
- F_Ports
- EX_Ports
- E_Port trunk
- F_Port trunk
- Mirror Ports (M_Ports)

Unsupported ports

The following port types are not supported for Flow Vision. Although flows can be created on unsupported ports, activation will fail. For more information on support for a specific feature, such as Flow Monitor, Flow Mirror, or Flow Generator, refer to appropriate sections on those features.

- Virtual FC Ports (VE_Ports and VEX_Ports)
- Mirror Ports (M_Ports)
- 1 Gigabit Ethernet ports
- FCoE ports
- LISL ports
- N_Ports

Flows

A flow is a set of Fibre Channel (FC) frames or packets that share similar traits, such as a source port identifier or any other data that can be used to uniquely differentiate one set of related frames or packets from another. A flow is provisioned with the following basic parameters:

- **Port Parameters.** These consist of the switch ingress or egress port. Only one port can be specified. The specified port is also called the “point of interest,” or the point where you wish to examine the data flow.
- **Frame Parameters.** The basic frame content provides identification of the source device, destination device, LUN IDs, or Frame Type. You must specify at least one parameter.

- **Feature Parameters.** Specify one or more features, either Flow Monitor, Flow Generator, or Flow Mirror.
- **Direction.** The direction is implicitly defined source device to destination device. All flows except learning flows, can also be bidirectional. You cannot define bidirectional flows with the Generator feature enabled as flows are always from source to destination.

Flow provisioning and monitoring

Flow Vision has two components:

- Flow provisioning
- Flow monitoring

Before monitoring a flow, you must define criteria that uniquely identify the flow. A flow definition includes basic criteria such as a flow name, source identifier (SID), destination identifier (DID), ingress port, egress port and flow direction. You can create new flow definitions and create definitions from existing definitions using **Add Flow Definition** dialog box.

You can also enable the following traffic-related features in the flow definition:

- Flow Generator

This feature simulates and generates traffic for the specified flows. Flow Generator can create and activate multiple custom flows in the fabric, and then identify potential traffic problems using standard traffic monitoring tools while the simulated traffic is flowing. Sub features include default frame generation and custom frame generation (with different frame sizes and patterns).

The Flow Generator allows you to produce data to verify the following:

- Hardware - ASIC, backplane, cable (ISL), SFP
- Connectivity - Routing, ACL
- Performance - Throughput, TI, QoS
- Configuration - With VWWN, you can verify the exact device connectivity and configuration.

For more detailed information on this application, refer to [“Flow Generator”](#) on page 1018.

- Flow Monitor

This feature monitors the network’s traffic pattern and provides statistics for the defined flow(s). Sub features include frame and SCSI statistics.

For more detailed information on this application, refer to [“Flow Monitor”](#) on page 1017.

- Flow Mirror

This feature allows you to select a traffic pattern and mirror this traffic to the CPU, enabling debugging that does not disturb the existing connections. You can also use this to listen or snoop on traffic passing through a port. Flow Mirror supports the following functions:

- Sending mirrored frames to the CPU.
- Mirroring frames in Layer 2 fabric.
- Mirroring frames on F_Ports.

For more detailed information on this application, refer to [“Flow Mirror”](#) on page 1016.

For more information on the Flow Vision features, including limitations and prerequisites, refer to [“Flow Vision features”](#) on page 1016.

After you define flows for a fabric you can monitor these flows using the **Flow Vision** dialog box. Data display for flows includes MAPS violations, values for criteria defined in your flow definitions, SCSI measures, and frame measures. You can adjust the monitor to display data for multiple flows, time durations, hide and display SCSI and frame-related measures, launch graphs of flow data, and change fabrics where flows are monitored. You can also select Flow Monitor, Mirror, and Generator to display measures for sub-flows where these features are active.

To monitor, launch the **Flow Vision** dialog box by right-clicking one of the following objects in Connectivity Map or Product List on the **SAN** tab and selecting **Fabric Vision > Flow > Monitor**.

- Switch port
- Initiator port
- Target port
- Switch that supports Flow Vision
- Fabric with one or more switches that support Flow Vision

You can also launch the **Flow Vision** dialog box by selecting **Monitor > Fabric Vision > Flow > Monitor** from the main menu bar.

Configuring flows using master and slave trunk ports

- A flow using the trunk master port monitors the data traffic for the entire trunk.
- For F_Port and E_Port trunks, “Flow Monitor” monitors only the master port of the trunk. If the master port changes, Flow Monitor automatically starts monitoring the new master port for the specified flow. If a monitor is installed on a port that later becomes a slave port when a trunk comes up, Flow Monitor automatically moves to the master port of the trunk.
- When a trunk port change state change notification (EPORT_CHANGE or FPORT_CHANGE) occurs, the flows on the old master port are deactivated and flows attached to the new master port are activated.
- If you create an active flow on a trunk slave port, the flow will be automatically activated when the slave port becomes master trunk port.
- If you create an inactive flow on a trunk slave port, you will have to manually activate the flow if the slave port becomes trunk master port.
- You should create same flow on all trunk member ports. During configuration replay, flows will be created on both the master and slave ports but only the flows associated with master port would be activated.

Provisioning flows

To provision or define a flow and to configure Flow Vision to monitor that flow, provide a flow name and specify the flow parameters in the **Add Flow Definition** dialog box. These parameters identify the sets of related frames and can either be explicitly defined or Flow Vision can learn them through observation.

Flow Vision allows you to learn a particular set of parameters for a frame. For example, you can define a flow to identify all possible active flows between a defined set of source and destination device pairs that pass through a specific port.

Alternatively, you can define a flow using asterisks (*) for both the source and destination devices, and have Flow Monitor learn all the source device and destination device pairs passing through the switch using a particular ingress port or egress port. This is called learning mode.

Learning source device and destination device values are only supported for 16 Gbps-capable Fibre Channel F_Ports. Only one learning flow is supported per ASIC.

Specifying the source and destination WWN implicitly defines a direction. For example, source WWN = x, destination WWN = y means traffic flowing from x to y. This is true for any flow definition, whether for source to destination device WWN. The bidirectional option causes the flow definition to be monitored in both directions.

A minimum of one frame parameter (source device, destination device, frame type, LUN IDs) and only one port parameter (ingress port, egress port) is required. Entering an asterisk for a value causes all values to be used for that identifier. For example: If the source device is specified with a specific WWN, and the destination device is specified as an asterisk, flows are created from the destination device WWN to all WWNs in the same zone as this source device. Note that you cannot use an asterisk for port parameters.

NOTE

In this guide an asterisk (*) shown in fields with ellipses buttons denotes any selected port or device in PID or WWN format.

You can access the **Add Flow Definition** dialog box on the SAN tab through **Monitor > Fabric Vision > Flow > Add** on the main menu bar and through the **Fabric Vision > Flow > Add** menu when you right-click objects in the products list or connectivity map. The following procedure references these menus. In addition you can access this dialog box through other management application features such as Frame Viewer, Bottleneck Detection, Trace Route and Ping, and Port Connectivity. Refer to [“Accessing Flow Vision from other management application features”](#) on page 1027 for more information.

Use the following steps to define flows for flow monitoring:

1. Select the SAN tab.
2. Select one of the following from the Connectivity Map or Product List, and then select **Monitor > Fabric Vision > Flow > Add**:
 - Switch
 - Switch port
 - Initiator port
 - Target port
 - Fabric with one or more switches that support Flow Vision.

NOTE

You can also right click on either of these objects in the products list or connectivity map and select **Fabric Vision > Flow > Add** from the menu. Selected switches, switch ports, initiator, ports, and target ports must be able to support Flow Vision.

The **Add Flow Definition** dialog box displays. Note that you must select **Advanced Options** to display the Frame Type and LUN IDs fields at the bottom of the dialog box.

FIGURE 423 Add Flow Definition dialog box

Dialog box fields will be populated with criteria and flow identifiers (such as SIDs and DIDs) for the appropriate flow definition based on the object selected to launch dialog box. For example, launching the dialog box from a target port will populate the dialog box with the source device PID, ingress switch port, monitor, and destination device as *. An asterisk (*) denotes any selected device in the same zone as the source device. Refer to [“Context-based flow definitions”](#) on page 1023 for details on criteria populated depending on the context from which the dialog box was launched.

3. Enter a name for the flow using 20 alphanumeric or underscore characters.

NOTE

Use unique names for flows defined in a physical switch. Names do not have to be unique for flows defined in logical switches.

4. Select at least one feature - **Monitor**, **Mirror**, or **Generator**
5. Select **Activate all selected features** to activate features that you have selected.

NOTE

Specific parameters for flow definitions may not be supported by some features. Refer to [“Flow parameter and configuration rules and limitations”](#) on page 1024 for details.

6. Select **Source to Destination** or **Bidirectional** flow.
7. Select **Persist over switch reboots** to persist flow definitions over reboots.

Configuring Basic Options

Use the following steps to configure **Basic Options** on the **Add Flow Definition** dialog box.

1. Select either **Port Address** (Port ID) or **WWN** format for entering the **End Device** identification.
2. Use one of the following methods to enter **End Device Source** and **Destination** identification:
 - Type a port ID or WWN in the **Source** and **Destination** fields.
 - Select the ellipses buttons to the right of **Source** and **Destination** fields. Selecting button displays the **Select Device Ports** dialog box (Figure 424).

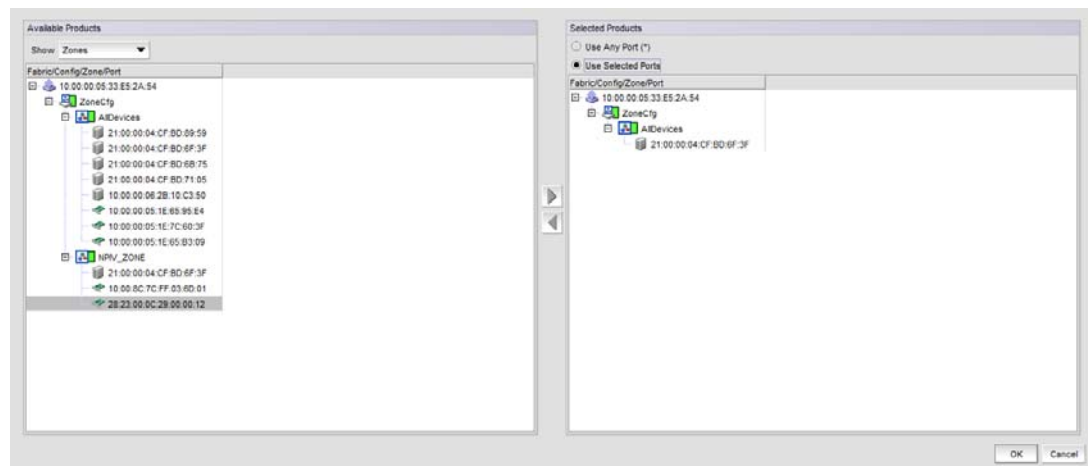


FIGURE 424 Select Device Ports dialog box. (Zones selected)

Use the **Show** list to display either ports in specific zones or available fabrics, devices, and ports for selection.

- Selecting **Zones** from the **Show** list displays available ports under specific zones as shown in Figure 424.
- Selecting **Products and Ports** from the **Show** list displays available fabrics, devices, and ports as in Figure 425.

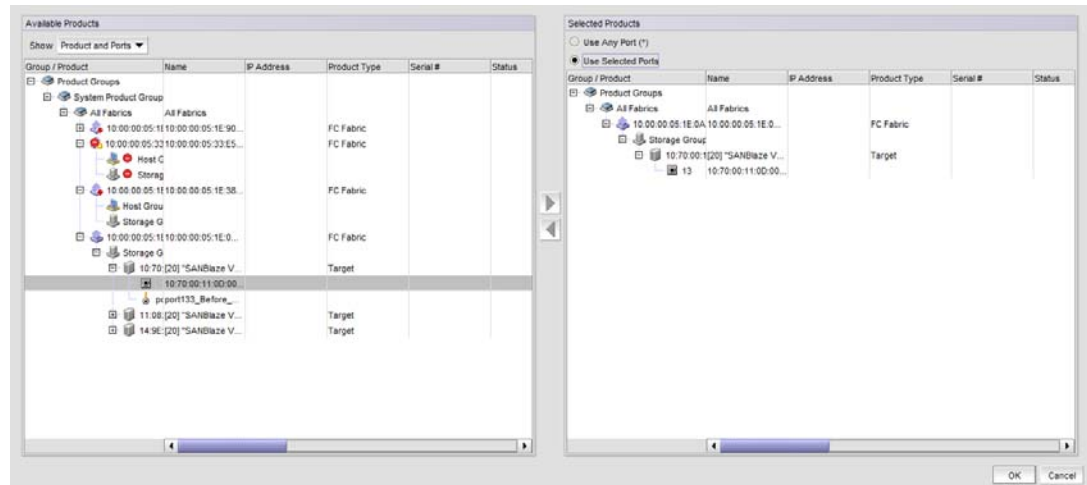


FIGURE 425 Select Device Ports dialog box. (Products and Ports selected)

Select products and ports under the **Available Products** panel and move to the **Selected Products** panel using the right arrow. In the **Selected Products** panel you can select to **Use Any Port (*)** in these products or **Use Selected Ports**. If you choose to use selected ports, select a specific port under the **Selected Products** panel.

Move products and ports from the **Selected Products** panel back to the **Available Products** panel by selecting them and using the left arrow.

Selecting **OK** on the **Select Device Ports** dialog box enters selected product WWNs or port addresses into the **Source** and **Destination** fields on the **Add Flow Definition** dialog box.

3. On the **Add Flow Definition** dialog box, select **<swap>** if you want to swap source and destination device identifications.
4. Select **Port** (slot/port) or **D,I** format.(domain ID, port number) for entering **Switch Ingress** and **Egress** identification.
5. Use the following options to enter identification in the **Switch Ingress** and **Egress** fields.
 - Type identifiers in the fields (entering an asterisk denotes any port) into the fields.

NOTE

You must enter the slot number and 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.

- Click the ellipsis button to the right of the **Ingress** and **Egress** fields to display the **Select Switch Ports** dialog box (Figure 426). Use this dialog box to select switch ports from discovered fabrics.

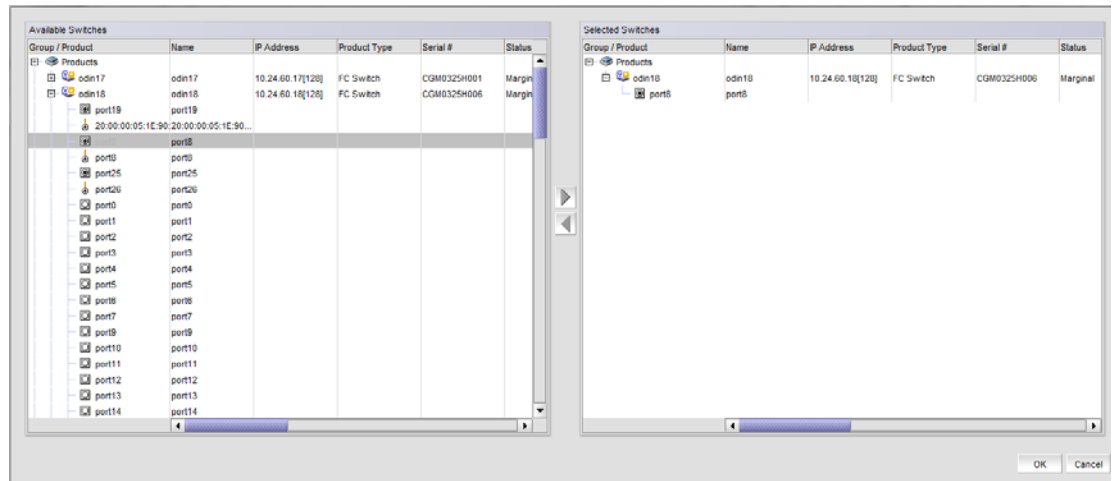


FIGURE 426 Select Switch Ports dialog box.

On the **Select Switch Ports** dialog box, select a port in the **Available Products** panel and move it under the **Selected Products** panel with the right arrow.

Select **OK** to close the **Select Switch Ports** dialog box and return to the **Add Flow Definition** dialog box.

- Select **<swap>** to swap the **Ingress** switch identification and **Egress** switch identification.
6. If a target switch is not entered for the flow definition, select the ellipses button to the right of the **Target Switch** field. The **Select Targets** dialog box displays.

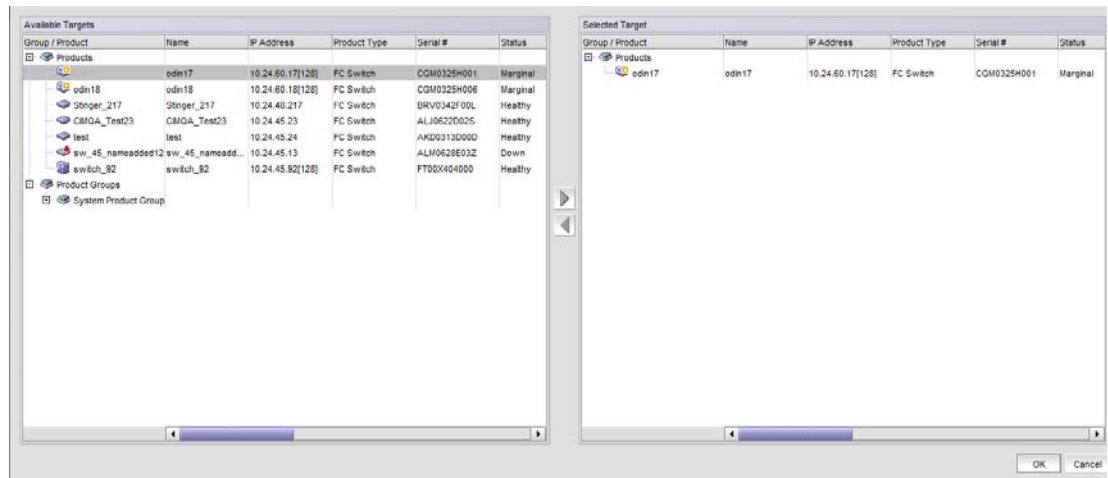


FIGURE 427 Select Targets dialog box

7. Select a switch in the **Available Targets** panel and move to the **Selected Target** panel using the right arrow. Move switches back to the **Available Targets** panel using the left arrow.
8. Select **OK** to close the **Select Targets** dialog box and add the target name to the **Add Flow Definition** dialog box.

Configuring Advanced Options

Use the following steps to configure **Advanced Options** on the **Add Flow Definition** dialog box as needed to define your flow definition.

1. Select the arrows on the **Advanced Options** tab to display configuration options.

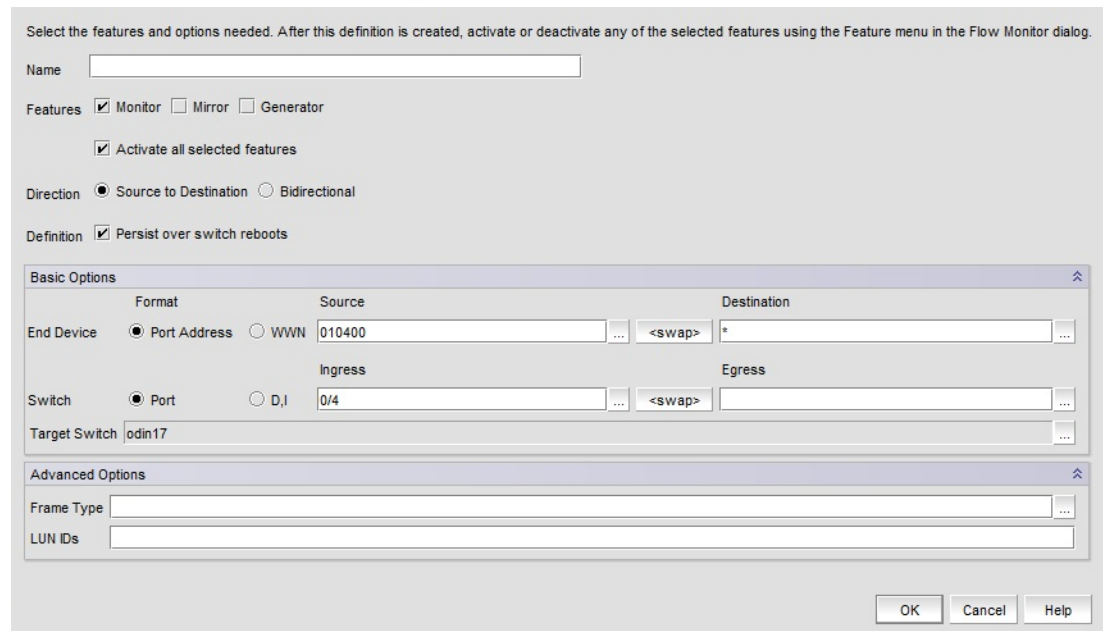


FIGURE 428 Add Flow Definition dialog box

2. Enter an Frame Type or select the ellipses button on the right of the **Frame Type** field to display the **Frame Type Picker** dialog box. Use this dialog box to select available frame command types.

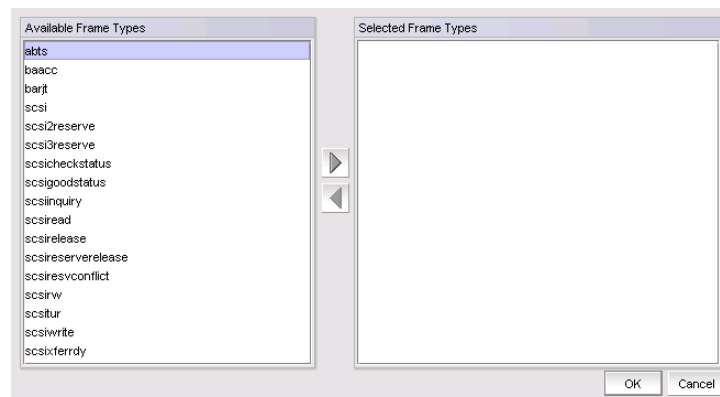


FIGURE 429 Frame Type Picker

Move your selections to the “selected” box using the right arrow and remove them using the left arrow.

The following Frame Type commands are supported:

- SCSI
 - SCSIRead
 - SCSIWrite
 - SCSIRW
 - SCSI2Reserve
 - SCSI3Reserve
 - SCSI2Release
 - SCSI3Release
 - SCSI2ReserveRelease
 - SCSI3ReserveRelease
 - SCSTur
 - SCSIStatus
 - SCSIInquiry
 - SCSIxferry
 - ABTS
 - BAACC
 - SCISGoodStatus
 - SCSICheckStatus
 - SCSIResvConflict
 - BARJT
3. Enter LUN IDs between 0 and 65535 that you want to include in the flow definition.
 4. Select **OK** to save the definition and launch the add flow progress dialog box.

When the flow definition activates, the **Flow Vision** dialog box displays with the created flow selected in the **Flow Definitions** area. Refer to [“Monitoring Flows”](#) on page 1005.

If there is an error in the definition, a message box displays providing the reason for the error.

Flow definition examples

Table 73 on page 1003 provides examples of flow data that you can monitor and criteria that you configure in the **Add** or **Edit Flow Definition** dialog box to define the flow.

NOTE

In this guide an asterisk (*) in fields denotes any selected device or port in PID or WWN format.

TABLE 73 Flow definition examples

Desired flow data	Context for launching Flow Definition dialog box and options	Flow definition
Top talking source for a destination	Select the destination initiator or target from Connectivity Map or Product List, and then select Monitor > Fabric Vision > Flow > Add or right-click the initiator or target, and then select Flow Vision > Flow > Add .	Source Device = * Destination Device = <i>Device port ID</i> Direction = Source to Destination Feature = Monitor
XISL usage as percentage of logical fabric	Select the base port for the logical fabrics from Connectivity Map or Product List, and select Monitor > Fabric Vision > Flow > Add or right click the port and select Flow Vision > Flow > Add .	Source Device = * Destination Device = * Switch Port = <i>port</i> Direction = Source to Destination Feature = Monitor
ISL Monitoring	Select the port participating in the ISL from the Connectivity Map or Product List and select Monitor > Fabric Vision > Flow > Add or right click the port and select Flow Vision > Flow > Add .	Source Device = * Destination Device = * Switch Port = <i>port</i> Direction = Source to Destination Feature = Monitor
End to end (E-E) monitoring	Select a switch to be monitored from the product list and select Monitor > Fabric Vision > Flow > Add . Select the source switch port, device SID, and device DID to be monitored in the port source and destination fields of the Add Flow Definition dialog box. As an option, you can monitor flows from the Flow Vision dialog box. Select a source ID (SID) and destination ID (DID) pair to be monitored, and then select the Name Flow button to create a flow definition for EE monitoring.	Source Device = <i>SID</i> Destination Device = <i>DID</i> Ingress Switch Port = <i>port</i> Direction = Bidirectional Feature = Monitor
E-E monitor (LUN level)	Select a device with a LUN to be monitored from the product list and select Monitor > Fabric Vision > Flow > Add .	Source Device = * Destination Device = * LUN = <i>LUN IDs</i> Direction = Source to Destination Feature = Monitor
LUN level statistics	Select a device with a LUN to be monitored from the product list and select Monitor > Fabric Vision > Flow > Add .	Source Device = <i>PID</i> Destination Device = <i>DID</i> Switch Port = <i>port</i> LUN = <i>LUN IDs</i> Direction = Bidirectional Feature = Monitor

TABLE 73 Flow definition examples

Desired flow data	Context for launching Flow Definition dialog box and options	Flow definition
Top talkers	Select a switch or port on which to monitor the top talkers from the product list and select Monitor > Fabric Vision > Flow > Add or right click the switch or port and select Flow Vision > Flow > Add .	Source Device = * Destination Device = * Switch Port = <i>port</i> Feature = Monitor
Ingress top talkers	Select a switch or port on which to monitor the top talkers from the product list and select Monitor > Fabric Vision > Flow > Add or right click the switch or port and select Flow Vision > Flow > Add .	Source Device = * Destination Device = * Egress Port = <i>port</i> Feature = Monitor
Egress top talkers	Select a switch or port on which to monitor the top talkers from the Connectivity Map or Product List, and then select Monitor > Flow Vision > Flow > Add .	Source Device = * Destination Device = * Egress Port = <i>port</i> Feature = Monitor
Frame monitor	Select a switch port on which you want to monitor frames from the Connectivity Map or Product List, and then select Monitor > Fabric Vision > Flow > Add or right click the switch port and select Flow Vision > Flow > Add .	Source Device = * Destination Device = * Switch Port = <i>port</i> Direction = Source to Destination Frame Type = Aborts Feature = Monitor
Mirroring traffic between a SID and DID	Select a switch from the Connectivity Map or Product List, and then select Monitor > Fabric Vision > Flow > Add . Enter the IDs of the devices to be mirrored in the End Device Source and Destination fields of the Flow Vision > Flow > Add dialog box.	Source Device = <i>SID</i> Destination Device = <i>DID</i> Switch Port = <i>port</i> Direction = Bidirectional Feature = Mirror
Generate traffic from a source port	Select a switch port that will be the source port for generating traffic and select Monitor > Fabric Vision > Flow > Add .	Source Device = <i>SID</i> Destination Device = <i>DID</i> Source Port = <i>SIM port</i> Feature = Generate
End device deployment	Select an end point (initiator or target port) from the Connectivity Map or Product List to which the device is going to be generating traffic, and then select Monitor > Fabric Vision > Flow > Add or right click the end point and select Flow Vision > Flow > Add . This will help you detect if the new device causes bottlenecks in the path	Source Device = <i>SID</i> (could be offline) Destination Device = <i>DID</i> (could be offline) Ingress Port = <i>port</i> Feature = Generate

Monitoring Flows

After you define flows for a fabric you can monitor them using the flow **Flow Vision** dialog box. The dialog box supports all the flows defined on the switches through the management application or CLI commands. Data display for flows includes all measures supported by the flow monitoring features and AMS violations on monitored flows. You can adjust the dialog box to display data for multiple flows, time durations, hide and display SCSI and frame-related measures, launch graphs of flow data, and change fabrics where flows are monitored, You can also select features (Flow Monitor, Mirror, and Generator) to display measures for sub-flows where the feature is active.

To monitor flows in the **Flow Vision** dialog box, use the following steps:

1. Launch the **Flow Vision** dialog box, use one of the following options:
 - Select **Monitor > Fabric Vision > Flow > Monitor** from the SAN tab menu bar.
 - Right-click one of the following objects from the Connectivity Map or Product List, and then select **Flow Vision > Flow > Monitor** from the displayed menu. Note that a port need not be in an active flow definition to launch the dialog box.
 - Switch port
 - Initiator port
 - Target port
 - Switch that supports Flow Vision
 - Fabric with one or more switches that support Flow Vision

The **Flow Vision** dialog box displays (Figure 430).

The screenshot shows the Flow Vision dialog box with two main panels. The left panel, titled 'Flow Definitions', contains a table with columns: Violation, Target Switch, Name, Monitor, Mirror, and Gen. The right panel, titled 'Flows', contains a table with columns: Sub Flow Id, Flow Name, Source, and Source Info.

Violation	Target Switch	Name	Monitor	Mirror	Gen
0	Odin20	testmirror_1		Inactive	
0	Odin20	testmirror_2		Inactive	
0	Odin20	testmirror_3		Active	
0	Odin20	monitor	Active		
0	Odin20	wwn	Active		
0	Odin20	gen			Inac
0	dcm.5100.176	default	Active		
0	dcm.5100.176	stats12_copy	Active		
0	dcm.5100.176	dffgfgf	Active		
0	dcm.5100.176	mirror	Active		
0	dcm.5100.176	test123	Active		
0	dcm.5100.176	test	Active		
0	dcm.5100.176	stats_mirror	Active		
0	Pluto184	mutlu_san_fw	Inactive		
0	Pluto184	bottleneck	Inactive		

Sub Flow Id	Flow Name	Source	Source Info
22	testmirror_3	010200	20.02.00.05:33:E5:2F:D4
20	monitor	011200	[37] "Brocade-825 3.2.0.0 localhost
17	w.wn	011200	[37] "Brocade-825 3.2.0.0 localhost
21	default		
19	stats12_copy	011201	[37] "Brocade-825 3.2.0.0 localhost
16	dffgfgf		
41	mirror		
45	mirror		
9	test123		
1	stats_mirror	3110dc	[28] "SEAGATE ST318304FC FA55"
18	testflow		
15	testflow123_copy		
13	test_copy	020600	[52] "Emulex LPe12002-M8 FV1 10A5 D
11	testflow123	011201	[37] "Brocade-825 3.2.0.0 localhost
49	test		
46	test	020600	[52] "Emulex LPe12002-M8 FV1 10A5 D
50	test		

FIGURE 430 Flow Vision dialog box

The **Flow Definitions** panel displays all current flow definitions for the fabric displayed in the fabric list above the panel. The number of AMS violations detected during the set time duration also display for each flow under the **Violation** column. The **Flows** panel displays statistics and flow data for active and inactive flows being monitored.

If you launch the **Flow Vision** dialog box by right-clicking a switch port on the Connectivity Map or Product List, and then select **Flow Vision > Flow > Monitor**, the management application verifies if there is a flow definition in the fabric with the selected port as the ingress or egress port. If you select an initiator or target, the management application verifies if there is a flow definition in the fabric with the port as a source or destination device. If matching definitions are found, the **Flow Vision** dialog launches with the flows selected. If definitions are not found, the following message displays:

No existing flow was found on the switch. Do you want to define a flow now?

You can select **Yes** to open the **Add Flow Definition** dialog box. You can then define a flow (refer to [“Provisioning flows”](#) on page 996), and it will display under the **Flow Definitions** panel of the **Flow Vision** dialog box.

2. To monitor flows, move flow definitions to the **Flows** panel by selecting a row under **Flow Definitions** and the right arrow key.
3. Move flows being monitored back to the **Flow Definitions** panel by selecting a row under **Flows** and the left arrow key.
4. For details on using options on the dialog box for controlling displayed data, adding flow definitions, and launching additional tools, such as graphs of the displayed flows, refer to [“Using Flow Vision dialog box options”](#).

Using Flow Vision dialog box options

Use the following controls on the **Flow Vision** dialog box to control displayed data, add flow definitions, and launch tools, such as a **Historical Graphs/Tables** dialog box, for selected flows.

- **Monitoring flows** - To monitor flows, move flow definitions to the **Flows** panel by selecting a row under **Flow Definitions** and then the right arrow key. Move flows being monitored back to the **Flow Definitions** side by selecting a row under **Flows** and then the left arrow key.
- **Fabric list** - This is located at the top left of the dialog box. Select a fabric for displaying flow definitions in the **Flow Definitions** panel. Note that the fabric must contain one or more switches that support Flow Vision.
- **Flow** menu - Select one of the following options to administer displayed flow definitions:
 - **Add** - Launches the **Add Flow Definition** dialog box for provisioning flows. If you select a flow definition, and then select **Add**, the **Add Flow Definition** dialog box includes flow criteria already configured. When you select **OK** to close the **Add Flow Definition** dialog box, the left side of the Flow Definitions panel of the **Flow Vision** dialog box refreshes to include any new criteria you have added and the flow displays in the **Flows** panel for monitoring.
 - **Reset** - Resets the flow data for all features on selected flow definitions.
 - **Delete** - Deletes the selected flow definition.
 - **MAPS** - Selecting **Violations** from the submenu displays filtered list of **MAPS Violations** for the selected flows. This list will display only when you select a single flow definition with the Flow Monitoring feature activated. Selecting **Configure** from the submenu displays the **MAPS** dialog box for the selected flow so you can create monitoring and alerting policy suite (MAPS) policies for specific port health, VM violations, and bottleneck events. Refer to [Chapter 32, “Monitoring and Alerting Policy Suite”](#) for more information on using these dialog boxes.

- **Feature** - Select options from the **Monitor**, **Mirror**, and **Generator** submenus to activate, deactivate, or reset the features for selected flow definitions, if these features are included in the flow definitions. Each feature contains a submenu with the following options:
 - **Configure** - For the Flow Generator only, selecting this option launches the **Configure Generate Flow** dialog box.

Using the **Configure Generate Flow** dialog box, you can set the payload size and pattern for traffic generated for the flow. Selecting **Bytes** for **Payload Size**, allows you to set a minimum of 64 and maximum of 2,048 bytes. Selecting **Random** generates a random number of bytes. In the **Pattern** field, you can enter a maximum of 32 hexadecimal characters as the pattern for the generated traffic.
 - **Activate** - Activates the feature if configured for a flow definition and deactivated. The switch updates the statistic counters from the last values maintained.
 - **Deactivate** - Deactivates feature if configured for a flow definition and activated. The switch stops updating the measures for a flow, however the older measures counters are retained.
 - **Reset** - Resets the feature data for the switch if the feature is included for the flow definition.
- **Time duration** - Select a time interval for monitoring flows. Possible values are 30 minutes, 1 hour, 6 hours, 12 hours, 1 day, 3 days, 1 week, and 1 month.
- **Feature** - Select **Monitor**, **Mirror**, or **Generator**. Selecting one of these features will display different measures for the sub-flows under **Flows** where the feature has been defined and activated.
- **SCSI** - Select to add a check mark to display SCSI-related measures. Select to uncheck the box and hide SCSI-related measures. Measures include SCSI read count, write count, read rate, write rate, read data, write data, and read and write frame data.
- **Frame** - Select to add a check mark to display frame-related measures. Select to uncheck the box and hide frame-related measures. 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.
- **Name Flow** - Select a sub-flow row under **Flows** and select **Name Flow** to launch the **Add Flow Definition** dialog box. Dialog box fields will be filled out with information from the selected flow. You can provide a name in the **Name** field for the selected sub-flow.

NOTE

Name Flow is only enabled when you select one sub-flow. The **Active all selected features** option on the displayed **Add Flow Definition** dialog box will be disabled by default.

- **Graph** - Select a row in the **Flows** panel to launch the **Historical Graphs/Tables** dialog box (Performance Graph). Refer to [“Using the Performance Graph”](#) on page 1013 for details on using this graph. Note that the **Graph** option is enabled when one or more sub-flows display in the **Flows** tab.

Flows panel right-click menus

Right-click a sub-flow in the **Flows** panel of the **Flow Vision** dialog box to select the following options:

- **Locate** - displays the following sub options. Select an option to highlight the port or device in the topology map:
 - Ingress Port - Highlights the ingress port.
 - Egress Port - Highlights the egress port.
 - Source Device - Highlights the source device.
 - Destination Device - Highlights the destination device.
- **Table** - displays options to control table functions in the panel, such as copy port, row, and table; export row, export table, print, and select columns.

Flow Vision dialog box overview

This section describes details on information displayed in the **Flow Definitions** and **Flows** panels.

Flow Definitions panel

The **Flow Definitions** panel displays details of all options currently configured for active and inactive flow definitions in the selected fabric.

Violation	Target Switch	Name	Monitor	Mirror	Generate	Source	Source Info	Destination	Destination Info	Source
0	Pluto184	muthu_san_flow	Active			0da200	1401			7/34
0	Pluto184	scsi_cal	Active			0da200	1401		[281] LSI	13,162
0	dcm.5100.176	stats12_copy	Active			311300	1371			0/19
0	dcm.5100.176	stats12	Active			311300	1371		[281]	0/19
0	dcm.5100.176	dst	Active			311300	1281			0/19
0	dcm.5100.176	stats_mirror	Active			311300	1281			0/19
0	sw0	muttstscsi	Active			020600	1521 Emulex			0/6
0	sw0	oracle_flow	Active			020600				0/6
0	Odin20	aldefined	Active	Inactive	Inactive	011200	1371		[371]	0/18
0	Odin20	bna_test_flow	Active	Inactive	Active	010000	20:00:00:05:01	010200	20:02:00:05:01	0/0
0	Odin20	muthu_test_di_f...	Active			010100	1371			1,1
0	Odin20	test4_copy	Active			010600	1391	010900	[281]	0/6
0	Odin20	stats	Active			011200	1371		[281]	0/18
0	Odin20	tstlunmuthu	Active			011200	1371			0/18
0	Odin20	monitor_scsci	Active			011200				0/18

FIGURE 431 Flow Vision dialog box (Flow Definitions panel)

Table 74 describes information displayed in the **Flow Definitions** panel.

TABLE 74 Flows Definitions panel information

Column	information Displayed
Violation	AMS violation for the flow over the time duration selected.
Target Switch	The switch where the flow definition is created.
Name	Name provided for flow when created.
Monitor	Displays one of the following: <ul style="list-style-type: none"> • Active - The flow is defined and the feature is active, • Blank grey cell - Flow is not defined. • Inactive - Flow is defined, but the feature is not active.
Mirror	
Generator	

TABLE 74 Flows Definitions panel information

Column	information Displayed
Source	Source IDs as defined for flow.
Source Info	This field is either empty or displays the icon of inferred destination device (either a VM, host or storage) based on the source ID. The name of the VM, host, or storage displays with a hyper link to the device's property sheet.
Ingress Port	Ingress port as defined for the flow.
Egress Port	Egress port as defined for the flow.
LUN IDs	LUN values port as defined for the flow.
Bi-Directional	Either yes if bidirectional defined for the flow or no .
Flow Definition Persistence	Either yes if "Persist over switch reboots" is defined in the flow or no .
Frame Type	Commands as defined for the flow.
Size	Payload size as defined for the flow.
Pattern	Payload pattern as defined for the flow.

Flows panel

The **Flows** panel displays statistics and data for all monitored flows and sub-flows.

Following are general characteristics of the **Flows** panel:

- Inactive sub-flows will be highlighted in yellow to indicate that statistics have not updated for over 15 minutes.
- A single flow definition might yield data in multiple rows in the **Flows** panel. For example, if a flow definition has 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 will be displayed in the active flows dialog if it was reported in the selected time duration. Reported values may be 0 if the last data point did not report that flow.
- Displayed measures are based on the flow definition. Therefore, not all columns will be populated for the flows being monitored.
- Data will dynamically update every 5 minutes with the data from the switch.
- In the **Flows** panel, you can right-click a sub-flow and select to locate the ingress port, egress port, source device, and destination device defined for the flow in the topology view. Refer to ["Flows panel right-click menus"](#) on page 1008.
- Inactive flows are highlighted in yellow when statistics have not updated for 15 minutes.

Information displayed when Monitor enabled

[Table 75](#) describes information on sub-flows displayed in the **Flows** panel when you select **Monitor** from the Feature list above the **Flows** panel.

TABLE 75 Flows panel information (Monitor feature selected)

Column	Information Displayed
Sub Flow ID	Sub-flow database ID.
Target Switch	Name of the target switch for the flow.
Flow Name	Name of flow as defined.
Source	SIDs as defined or learned.
Source Info	This field is either empty or displays the inferred destination device (either a VM, host or storage) based on the source ID. The name of the VM, host, or storage displays with a hyper link to the device's property sheet.
Destination	DIDs as defined or learned.
Destination Info	This field is either empty or displays the inferred destination device (either a VM, host or storage) based on the source ID. The name of the VM, host, or storage displays with a hyper link to the device's property sheet.
SCSI Read Frame Count (frames)	SCSI read command count as reported in the last data point for the flow.
SCSI Write Frame Count (frames)	SCSI write command count as reported in the last data point for the flow.
SCSI Read Data Rate (Mbps)	SCSI read frame rate as reported in the last data point for the flow.
SCSI Write Data Rate (Mbps)	SCSI write frame rate as reported in the last data point for the flow.
SCSI Read Data (bytes)	SCSI read data as reported in the last data point for the flow.
SCSI Write Data (bytes)	SCSI write data as reported in the last data point for the flow.
SCSI Read Frame Rate (f/s)	SCSI read frame rate in frames per second as reported in the last data point for the flow.
SCSI Write Frame Rate (f/s)	SCSI write frame rate in frames per second as reported in the last data point for the flow.
Transmit Frame Count (frames)	Transmit frame count as reported in the last data point for the flow.
Receive Frame Count (frames)	Receive frame count as reported in the last data point for the flow.
Transmit Frame Rate (f/s)	Transmit frame rate as reported in the last data point for the flow.
Receive Frame Rate (f/s)	Receive frame rate as reported in the last data point for the flow.
Transmit Word Count (bytes)	Transmit word count as reported in the last data point for the flow.
Receive Word Count (bytes)	Receive word count as reported in the last data point for the flow.
Transmit Throughput (Mbps)	Transmit throughput as reported in the last data point for the flow.
Receive Throughput (Mbps)	Receive throughput as reported in the last data point for the flow.
Ingress Port	Source port as defined for the flow.
Egress Port	Destination port as defined for the flow.

TABLE 75 Flows panel information (Monitor feature selected)

Column	Information Displayed
LUN	LUN values as defined for the flow.
Bi-Direction	Yes or No as defined for the flow.
Flow Definition Persistence	Either Yes if “Persist over switch reboots” is defined in the flow or No .
Frame Types	Frame types as defined for the flow.
Size	NA
Pattern	NA
Last Updated Time	The time when the sub-flows were last updated.

Information displayed when Generator enabled

[Table 76](#) describes information on sub-flows displayed in the **Flows** panel when you select **Generator** from the Feature list above the **Flows** panel.

TABLE 76 Flows Panel information (Generator feature selected)

Column	Information Displayed
Sub Flow ID	Sub-flow database ID.
Target Switch	Name of the target switch for the flow.
Flow Name	Name of flow as defined.
Source	SIDs as defined or learned.
Source Info	This field is either empty or displays the inferred destination device (either a VM, host or storage) based on the source ID. The name of the VM, host, or storage displays with a hyper link to the device’s property sheet.
Destination	DIDs as defined or learned.
Destination Info	This field is either empty or displays the inferred destination device (either a VM, host or storage) based on the source ID. The name of the VM, host, or storage displays with a hyper link to the device’s property sheet.
Generator Transmit Frame Count (frames)	Transmit frame count as reported in the last data point for the flow.
Generator Receive Frame Count (frames)	Receive frame count as reported in the last data point for the flow.
Ingress Port	Ingress port as defined for the flow.
Egress Port	Egress port as defined for the flow.
LUN	LUN values as defined for the flow.
Bi-Direction	Yes or No as defined for the flow.
Flow Definition Persistence	Either Yes if “Persist over switch reboots” is defined in the flow or No .
Frame Type	Frame types as defined for the flow.
Size	Payload size as defined for the flow.
Pattern	Payload pattern as defined for the flow.
Last Updated Time	The time when sub-flows were last updated.

Information displayed when Mirror enabled

[Table 77](#) describes information on sub-flows displayed in the **Flows** panel when you select **Mirror** from the Feature list above the **Flows** panel.

TABLE 77 Flows panel information (Mirror feature selected)

Column	Information Displayed
Sub Flow Id	Sub-flow database ID.
Target Switch	Name of the target switch for the flow.
Flow Name	Name of flow as defined.
Source	SIDs as defined or learned.
Source Info	This field is either empty or displays the inferred destination device (either a VM, host or storage) based on the source ID. The name of the VM, host, or storage displays with a hyper link to the device's property sheet.
Destination	DIDs as defined or learned.
Destination Info	This field is either empty or displays the inferred destination device (either a VM, host or storage) based on the source ID. The name of the VM, host, or storage displays with a hyper link to the device's property sheet.
Mirrored Frames Count (frames)	Mirrored frames count as reported in the last data point received for the flow
Mirrored Transmit Frames (frames)	The number of transmitted mirrored frame as reported in the last data point received for the flow.
Mirrored Receive Frames (frames)	The number of received mirrored frame as reported in the last data point received for the flow
Ingress Port	Ingress port as defined for the flow
Egress Port	Egress port as defined for the flow.
LUN	LUN values as defined for the flow.
Bi-Direction	Yes or No as defined for the flow.
Flow Definition Persistence	Either Yes if "Persist over switch reboots" is defined in the flow or No .
Frame Type	Frame types as defined for the flow.
Size	Size of the frame payload.
Pattern	Pattern of the frame payload.
Last Updated Time	The time when the sub-flows were last updated.

NOTE

For 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.

Using the Performance Graph

You can access the **Historical Graphs Tables** dialog box (Figure 432), also called the Performance Graph, on the SAN tab using one of the following options:

- Move flow definitions that you want to graph to the **Flows** panel in the **Flow Vision** dialog box, and then select the **Graph** button.
- Select a device or port from the management application Connectivity Map or Product List, and then select **Monitor > Fabric Vision > Flow > Performance Graph** from the top menu bar.
- Right-click a device or port in the Connectivity Map or Product List, and then select **Fabric Vision > Flow > Performance Graph**.
- Right-click on the Dashboard Flow Performance Monitor, and then select **Graph** from the menu.

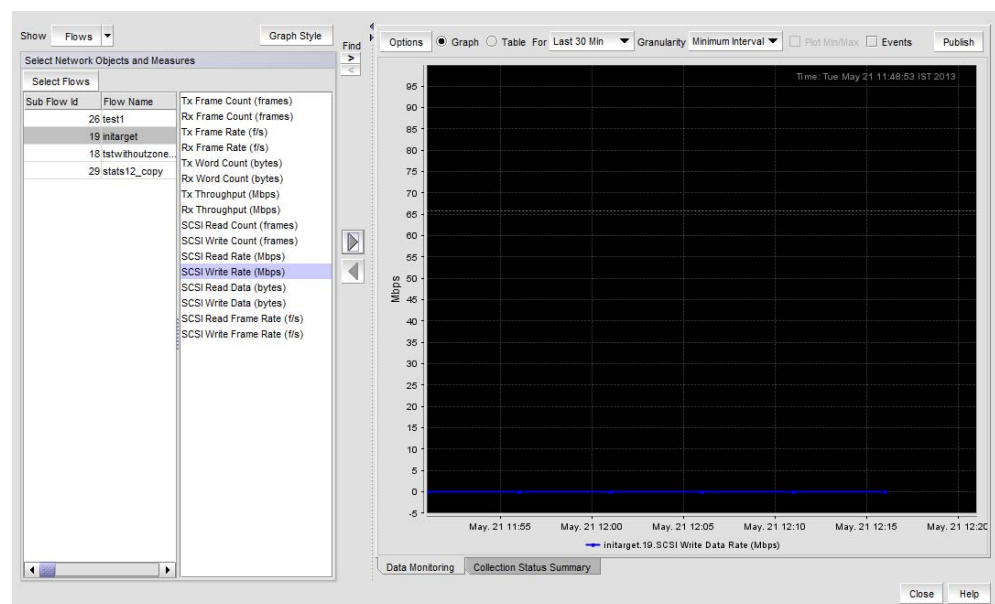


FIGURE 432 Performance Graph (Flows selected)

Figure 432 shows the performance graph when **Flows** is selected from the **Show** list on the **Historical Graphs/Tables** dialog box (default). You can select a sub-flow and measures on the left side of the dialog box and select the right arrow to plot on the graph. The measures list contains all statistics collected for the feature selected in the **Feature** list at the top of the **Flows** panel.

Figure 433 shows the historical graph when **Flow Measures** is selected from the **Show** list. You can select a sub-flow and measures on the left side of the dialog box and select the right arrow to plot on the graph. The measures list contains all statistics collected for the Monitoring feature.

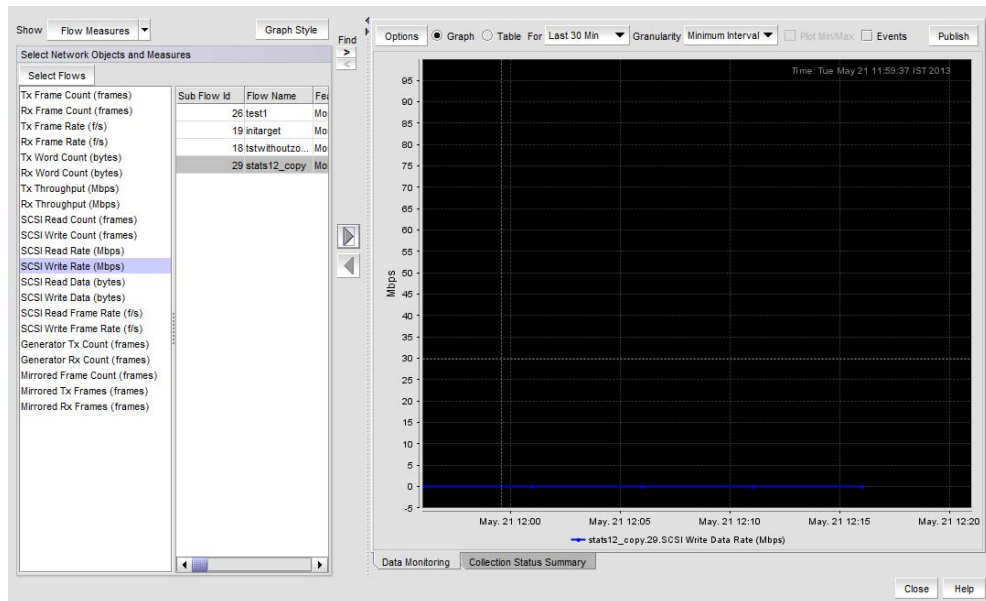


FIGURE 433 Historical Graph (Flow Measures selected)

Control functions on the **Historical Graphs/Tables** dialog box, such as plotting new sub-flows and creating a flow performance monitor on the dashboard, as follows:

- Plot different sub-flows and measures in the graph area by selecting them in the columns under **Select Network Objects and Measures** and moving them to the graph with the right arrow. Move sub-flows and measures out of the graph by selecting them on the graph, and then selecting the left arrow.
- To launch the **Flow Vision** dialog box for sub-flows, select **Select Flows**.
- To create a flow performance monitor on the dashboard, select a flow and one or more measures from the **Select Network Objects and Measures** columns, and then select **Publish** on the tool bar at the top of the dialog box.
- To toggle columns on the left side of the dialog box between flow and measure display, select **Flows** or **Flow Measures** from the **Show** menu. Refer to “[Dashboard flow performance monitor](#)” on page 1015 for more information.
- To use other general control functions on the tool bar at the top of this dialog box, refer to for “[Configuring the performance graph display](#)” on page 982.

Dashboard flow performance monitor

Figure 434 shows a flow performance monitor that you can create using the **Publish** button on the Flow Vision performance graph (**Historical Graphs/Tables** dialog box).




FIGURE 434 Flow performance monitor

Use the dashboard flow performance monitor tool bar and right-click menus to control monitor display, and launch the **Historical Graph/Table and Flow Monitor** dialog box for the flow. You can also modify performance monitors by adding target flows from different fabrics through the **Add Performance Dashboard Monitor** dialog box. For more information on configuring and managing performance monitors, refer to [“Configuring a user-defined traffic flow performance monitor”](#) on page 242.

To display a performance monitor for flows represented by graphs in the **Historical Graphs/Tables** dialog box, perform the following steps:

1. Select flows that you want to graph in the **Historical Graphs/Tables** dialog box.
2. Select **Publish** to create a monitor of the graph data for the dashboard.
3. Refer to [“Configuring a traffic flows monitor from a performance graph”](#) on page 242 for steps to add the performance monitor to the dashboard.

Refer to [“Using Flow Vision dialog box options”](#) on page 1006 for more information on creating a performance graph for flows displayed in the **Flow Vision** dialog box.

Flow Vision features

This section provides detailed information on the Flow Generator, Flow Mirror, and Flow Monitor features.

Flow Mirror

You can define flows with the feature enabled to select a traffic pattern and mirror this traffic to the CPU. You can then monitor sub-flows resulting from the definition to listen or snoop on traffic passing through a port.

Flow Mirror supports the following functions:

- Sending mirrored frames to the CPU.
- Mirroring frames in Layer 2 fabric.
- Mirroring frames only on F_Ports.
- In-depth analysis of flows of interest such as SCSI reservation frames, ABTS frames, flows to bottlenecked devices, and frames during link bring-up.

Limitations and prerequisites

The following limitations and prerequisites apply specifically to the Flow Mirror feature.

- Statistic counts displayed in the Flow Vision dialog box should not be inferred as accurate at this time. Counts greater than zero do imply that the mirrored flow is functioning.
- Mirrored frames are not displayed as with Fabric OS commands.
- This feature is supported only on 16 Gbps-capable Fibre Channel platforms
- You cannot enable Flow Mirror on more than one active flow definition per switch.
- This feature requires a Fabric Vision license on the switch where you are enabling the feature.
- This feature is not supported across Fibre Channel Routing (FCR).
- N_Ports cannot be ingress or egress ports.
- L_Ports cannot be ingress or egress ports.
- Flow Mirroring is supported in VF mode.
- Flow Mirror is not supported on the following ports:
 - DISL ports
 - E_Ports
 - EX_Ports
 - E_Port Trunks and F_Port trunks
 - XISL ports
 - Shared ports
- Flow Mirroring is not supported in Access Gateway mode or across a Fibre Channel Router backbone fabric.
- If a 16-Gbps ingress port is specified when creating a flow, then the rate of outbound traffic from that port may be reduced by as much as half. To prevent this, it is recommended that you specify the ingress port and egress port values only when absolutely needed.

- A maximum of 5 seconds worth of data is stored for any platform.
- This feature is only supported on ports operating at 8 Gbps or less.

Refer to [“Flow parameter support”](#) on page 1022 for more information on Flow Mirror feature limitations for flow definition parameters.

Zoning considerations

Zone checking is not required for the source device or destination device elements of a flow to be mirrored. However regular zoning is enforced by the zoning application on the frames flowing between the devices.

Flow Monitor

Flow Monitor allows you to monitor the network’s traffic pattern and provides statistics to make capacity planning decisions based on the collected data.

Uses for Flow Monitor include:

- Monitoring SCSI commands and write data in flows from a host to a target or LUN.
 - Monitor unidirectional flows from a host to target at an F_Port (either at host or target port).
 - Monitor flows from a host to a target at an individual E_Port or trunk.
- Monitoring unidirectional flows from a target to a host (read data).
- Gathering SCSI statistics such as the following:
 - IOPS - the number of SCSI read, write I/O operations per second.
 - Read, write data in bytes.
 - SCSI frame or data rate in Mbps.
- Create individual flows for each NPIV virtual machine (VM) to a LUN.

Flow Vision provides a single interface to manage flows and unifies different monitor functions, such as end-to-end monitors and frame monitors. The Flow Vision collector polls the switch every 5-minutes. Flow Monitor provides the following monitors:

- End-to-end monitors: Measures the traffic in terms of word count between a pair of ports (host and target) ports using a 5-minute sampling rate.
- Frame based monitoring: Measures the transmitted frame count through a port with specific values in the first 64 bytes of a frame using a 5-minute sampling rate.
- ISL monitoring: Measures the traffic transmitted via an ISL to different destination domains using a 5-minute sampling rate.
- Top Talker monitoring: Measures the flows that are major consumer of bandwidth on a port using a 5-minute sampling rate.

Flow Monitor provides a mechanism to define and monitor any flow parameter. Instead of only the static combinations described above, you can use Flow Monitor to define your own flows using source and destination devices, LUN IDs, and Frame Type as parameters to create a flow definition for a specific use.

For information on replicating these monitors, refer to [“Flow definition examples”](#) on page 1003.

Limitations and prerequisites

The following limitations and prerequisites apply specifically to the Flow Monitor feature.

- Bidirectional flows are supported for F_Ports only. For E_Ports, you must create a separate flow definition for each direction.
- You cannot specify an asterisk in the **Add Flow Definition** dialog box to learn all hosts sending traffic to a LUN.
- You cannot specify an asterisk in the **Add Flow Definition** dialog box to discover all LUNs being accessed from a host.
- A single flow definition to monitor all NPIV flows at a port going to a LUN is not supported.
- N_Ports cannot be ingress or egress ports.
- If switch is in Access Gateway mode, only F_Ports can be used as ingress or egress ports.

Refer to [“Flow parameter support”](#) on page 1022 for more information on Flow Monitor feature limitations for flow definition parameters.

Notes on Flow Monitor and trunk ports

Note the following when using Flow Monitor for trunk ports:

- Flow Monitor will gather statistics for all the flow(s) defined on the slave port on respective master port,
- Flow Monitor gathers statistics on the new master port for all the flow(s) defined on any port in the trunk.

Flow Generator

The Flow Generator application is used for traffic flow diagnostic and analysis purposes. It functions by creating artificial traffic in a SAN fabric so that administrators can stress-test the configuration and the switches in the fabric under extreme traffic load without requiring external equipment.

Use Flow Generator to create custom flows in the fabric and then identify potential traffic problems using standard traffic-monitoring tools while the simulated traffic is flowing. You can produce data to verify several switch metrics such as performance, bandwidth, QoS, routing, zoning and traffic Isolation. Use Flow Generator with Flow Performance Monitor and MAPS to monitor link utilization, CRC errors, link timeouts, link resets, Class 3 discarded frames, and RASlog messages generated during the test.

About test flows

Traffic is generated between a source simulator port (SIM) port enabled on a local switch and destination SIM port enabled on the local switch or a remote switch. Traffic at the port is confined within the switch but appears to enter the port from an external device, such as a host. Traffic appears to egress the remote switch at the destination SIM port to an external device, but is also confined to the switch. Traffic can be simulated between a simulated host and multiple simulated targets.

You can customize test flows by specifying the source, destination, a specific or random payload size, and payload pattern. Flow Generator source and destination ports emulate device entries in the Name Server database, where they are treated as real devices and so can be used to evaluate various switch and fabric operations, such as zoning, QoS, and traffic isolation. Simulated devices will display as virtual end devices. View these in the connectivity view by selecting **View > Connected end devices > include virtual devices**.

Flow Generator can display the traffic measurements of specific flows or all the generated flows to allow you to identify bottlenecks and congestions in the fabric. These measurements include the transmitted frame count from the egress port and the received frame count from the ingress port. A port that generates the simulated traffic can also be the destination of another flow transmitted from another port. In this case the received frames will be discarded and not forwarded. Flow Generator traffic is marked with a unique FC type (0xF6), allowing you to easily distinguish it from the real traffic.

Limitations and prerequisites

The following limitations and prerequisites apply specifically to the Flow Generator feature.

- Flow Generator functionality is based on the 16 Gbps platform Frame Shooter feature which allows 16 Gbps capable Fibre Channel ports to transmit traffic at a line rate. The Frame Shooter feature is fully supported on these ports.
- Flow Generator support for 8 Gbps-capable Fibre Channel ports is limited, as these ports cannot transmit the simulated traffic.
- The port that transmits the simulated traffic (SIM port) must be a 16Gbps-capable Fibre Channel port. The port that receives the simulated traffic can be an 8 Gbps or 16 Gbps-capable Fibre Channel port.
- A SIM port can only be configured on a vacant F_Port (U_Port).
- Access Gateway switches are not supported as Flow Generator sources or destinations.
- Flow Generator works with zoning in Fabric OS Layer 2 fabrics.
- Zoning should be enforced for a learning flow.
- Flow Generator is not supported for Fibre Channel routers as the remote destination device port cannot be checked to confirm that it is a simulator (SIM) port.
- Frame redirection is not supported for SIM ports.
- Source device, ingress port and egress port must be local.
- Source device and ingress port must refer to the same source.
- If destination device and egress port are defined, both must refer to the same destination and that the destination must be local.
- If more than one flow has the same destination, the egress port received frame count will increment based on both flows. In this case, if you want to match the transmitted and received frame counts as one flow, you must add all the transmitted frame counts together.
- If more than one flow has the same source, the ingress port transmitted frame count will not match the egress port received frame count. In this case, if you want to match the transmitted and received frame counts as one flow, you must add all the transmitted frame counts together.
- N_Ports cannot be ingress or egress ports.

- L_Ports cannot be ingress or egress ports.
- A maximum of four flow generator flows is supported per port.

Refer to “[Flow parameter support](#)” on page 1022 for more information on Flow Vision feature limitations for flow definition parameters.

Port characteristics

Ingress Port characteristics

The source port must meet the following criteria. If it does not, the flow will be rejected.

- The port must be a 16 Gbps-capable Fibre Channel port.
- The port cannot be in the base switch.
- The port can be set to long distance or F_Port buffer.
- The port should be a simulator (SIM) port.

Egress Port characteristics

The egress port must meet the following criteria. If it does not, the flow will be rejected.

- The port must be a 8 Gbps or 16 Gbps-capable Fibre Channel port.
- The port cannot be in the base switch.
- The port can be set to long distance or F_Port buffer.
- The port should be a simulator (SIM) port.
- The port cannot be any of the following:
 - Compression port
 - D_Port
 - Disconnected port
 - Disabled port
 - E_Port loopback
 - Encryption port
 - EX_Port
 - Fast write port
 - F_Port trunked
 - L_Port

- M_Port (Mirror port)

SIM ports

To ensure that test flows are not unintentionally transmitted to real devices, Flow Vision requires that you enable the source device and destination device ports in simulated (SIM port) mode before activating the test flows, and the application checks for this setting before activating the test flows. Setting the port to “SIM Port” sets an internal loopback on the port and creates a filter which discards all incoming Flow Generator frames.

On a switch that has live traffic passing through E_Ports or long distance ports, Flow Generator allows you to enable and disable SIM ports that are on the same ASIC as those active ports. Flow Generator can also activate and deactivate the artificial traffic, allowing you to verify the impact of the testing flow on existing traffic.

Prior to creating and activating flows, use the steps under [“Enabling and disabling SIM ports”](#) on page 1022 to set the source device and destination device ports as SIM ports.

Attributes of a SIM-port are as follows:

- Simulates an F_Port on the switch. A SIM port simulates a fabric device using the port WWN or virtual WWN. It is added into the name server database and can be part of a zoning database (needed for learning mode).
- 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.
- Supported on ASICs that support either 8- or 16-Gbps-capable Fibre Channel ports.
- Cannot be on the base switch or an Access Gateway switch.
- A SIM port can only be configured on a vacant F_Port (U_Port). The port has to either be in a “port disabled” status or not connected at the time it is identified as a SIM port.
- A SIM port cannot be configured as any of the following port types; these restrictions also apply at the time a SIM port is enabled.
 - A port running Encryption or Compression
 - D_Port
 - EX_Port
 - F_Port Trunked
 - Fastwrite port
 - FCoE port
 - ICL port
 - L_Port
 - Mirror Port (M_Port)
 - VE_Port
 - VEX_Port
 - GigE port
- The following features of a SIM port are persistent across a reboot:
 - Each SIM port is assigned a PID.
 - Each SIM port is assigned the Switch PWWN as its PWWN.
 - Each SIM port registers itself into Name Server database.

Enabling and disabling SIM ports

Prior to creating and activating flows using the Flow Generator feature, 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 1021.

To enable SIM port mode, use 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 SIM port on the local switch for the destination device in the Product List, and then select **Monitor > Fabric Vision > Flow > SIM Mode > Enable**.

To disable SIM port mode, use the following steps:

1. Select an enabled SIM port on the local switch for the source device in the Product List, and then select **Monitor > Fabric Vision > Flow > SIM Mode > Disable**.
2. Select an Enabled SIM port on the local switch for the destination device in the Product List, and then select **Monitor > Fabric Vision > Flow > SIM Mode > Disable**.

Zoning considerations

For learning flows, Flow Generator requires several levels of zoning support when simulating traffic. By default, the source and destination devices must both be in a defined zone (but not necessarily the same zone) for the simulated traffic to reach the destination device.

As with real devices, the source and destination devices can be zoned using “domain, index” identification, WWN identification, or session-based identification (mixed zoning). By default, the emulated device is assigned the port WWN of the port. Alternatively, you can bind the port to a virtual WWN by using Dynamic Fabric Provisioning. The advantage of using a WWN binding is that you can then deploy and test a zoning database which can be used for both the real and simulated devices without requiring changes.

Flow parameter support

Parameters used to configure flow definitions have different support based on the switch platform, enabled Flow Vision feature, and switch configuration mode.

Supported parameters by platform and switch configuration mode

Table 78 lists the supported configurations for each Flow Vision feature for the basic flow identification parameters, such as ingress port, source device, egress port, and destination device.

TABLE 78 Configurations Supported

Feature	Platforms		Switch Configuration Mode	
	16 Gbps-capable Fibre Channel	8 Gbps-capable Fibre Channel	Access Gateway	Virtual Fabric
Flow Generator	Supported	Supported	Not Supported	Supported
Flow Mirror	Supported	Not Supported	Not Supported	Supported
Flow Monitor	Supported	Supported	Supported	Supported

Supported flow parameters by Flow Vision feature

[Table 79](#) lists the supported **basic** flow configuration parameters for the Flow Vision Generator, Monitor, and Mirror features.

TABLE 79 Supported basic flow parameters for Flow Vision features

Parameter	Flow Generator	Flow Monitor	Flow Mirror
Ingress port	Supported	Supported	Supported
Egress port	Supported	Supported	Supported
Source device	Supported	Supported	Supported
Destination device	Supported	Supported	Supported
Bidirectional	Not applicable	Supported	Supported

[Table 80](#) lists the supported **advanced** flow configuration parameters for the Flow Vision Generator, Monitor, and Mirror applications.

TABLE 80 Supported advanced flow parameters for Flow Vision features

Parameter	Flow Generator	Flow Monitor	Flow Mirror
LUN IDs	Not applicable	Supported	Supported
Frame Type	Not applicable	Supported	Supported

Context-based flow definitions

Add Flow Definition dialog box fields and options are populated automatically to create the most appropriate flow according to the context from which you launched the dialog box. You can launch the dialog box from a switch port, initiator port, target port, or switch context using these options:

- Select the switch port, initiator port, target port, or switch on the products list or connectivity map, then select **Monitor > Flows > Add** from the main menu bar.
- Right-click the switch port, initiator port, target port, or switch on the products list or connectivity map, then select **Flows > Add** from the menu.

[Table 81](#) describes the flow created and definition created in the Add Flow Definition dialog box when launching the dialog box from a specific context.

NOTE

In general, * will be populated in source and destination device fields when the target switch supports learning. Otherwise, these fields will be blank.

TABLE 81 Flow definition based on context where Add Flow Definition dialog box launched

Context	Flow Created	Definition parameters in dialog box
Switch port	Flow definition created to monitor the traffic on the selected port without inferring with connected device	Ingress switch port = <port> Source device = * Destination device = * Features = Monitor
Switch port	Flow definition created to monitor the traffic on the selected port while inferring connected device.	Ingress switch port = <port> Source device = <device PID> Destination device = * Features = Monitor
Initiator port	Flow definition created on the switch attached to the initiator port.	Ingress switch port = <port> Source Device = <device PID> Destination Device = * Features = Monitor
Target port	Flow definition created on the switch attached to the target port.	Ingress switch port = <port> Source Device = <device PID> Destination Device = * Features = Monitor
Switch	Flow definition created on the selected switch.	Source device = * Destination device = * Features = Monitor

Flow parameter and configuration rules and limitations

Parameters and values for configuring flows may only be supported on specific switch platforms and when enabling specific Flow Vision applications, such as Flow Generator, Flow Mirror, and Flow Monitor.

General flow parameter rules

Following are general rules for provisioning flows in the **Add Flow Definition** dialog box. Note that noncompliance will cause error messages.

- Use unique names for flows defined in a physical switch. Names do not have to be unique for flows defined in logical switches.
- Changing between Port Address (PID) and WWN format for source and destination end devices. This may cause an error if WWN to PID cannot be resolved.
- Source and destination device identifiers must be in the same format - either WWN or Port Address.
- Ingress and egress port identifiers cannot be provided in the same flow.
- Flow definitions are only supported on Flow Vision capable switches. Refer to [“Supported parameters by platform and switch configuration mode”](#) on page 1022 for information on supported platforms.
- Use unique flow definitions. Duplicate flow definitions are not allowed.

- An asterisk can only be specified for an end device source and destination for 16 Gbps switch platforms.
- If you are using at least one advanced parameter (LUN IDs or Frame Type), then feature-specific rules apply. Refer to the following links for specific details.
 - [Flow Generator supported flow identification parameter combinations](#)
 - [Flow Mirror supported flow identification parameter combinations](#)
 - [Flow Monitor supported flow parameter combinations](#)

Refer to “[Flow parameter support](#)” on page 1022 for more information on switch platform, switch configuration mode (such as Access Gateway), and Flow Vision feature limitations for flow definition parameters.

Supported basic flow parameter combinations

[Table 82](#) lists the supported flow identification parameter combinations for **Add Flow Definition** dialog box fields.

TABLE 82 Basic flow identification rules

Parameter	Rules
Ingress Port	<ul style="list-style-type: none"> • Both cannot be specified.
Egress Port	<ul style="list-style-type: none"> • Values can only be fixed value.
Source Device	<ul style="list-style-type: none"> • Either one or both can be specified.
Destination Device	<ul style="list-style-type: none"> • Values can be fixed or “*” (“*” indicates all matching flows)

Flow Generator supported flow identification parameter combinations

To enable Flow Generator, you must define one of the following parameter combinations:

- Source device, destination device, egress port

NOTE

You cannot specify both the ingress and egress port in a flow definition.

- Source device, destination device, egress port
- Source device, egress port
- Destination device, ingress port
- Source device, ingress port
- Destination device, egress port

Values for these parameters define the following:

- The source device and destination device values mark the content of the frame (the frame Source ID and Destination ID as contained in the frame header).
- The ingress port and the egress port values represent the port IDs for the traffic flow endpoints.
- The source device and ingress port values indicate the origination point of the test traffic. For Flow Generator, these should be the same port.

- The destination device and egress port indicate the destination of the test traffic. For Flow Generator, these should be the same port.
- The source is on the local switch, but the destination can be local or remote.

Bidirectional, and the advanced flow identification parameters (LUN IDs or Frame Type) are not supported for Flow Generator.

Entering an asterisk for a flow identifier value causes Flow Generator to use all values for that identifier. For example, if a specific WWN is used for the source device, and an asterisk is specified for the destination device, Flow Generator will create flows from the source device WWN to all destination device WWNs that are in the same zone as the source device.

Flow Mirror supported flow identification parameter combinations

Table 83 lists the supported flow parameter combinations for Flow Mirror.

TABLE 83 Advanced Flow ID parameter combinations for Flow Mirror

Source Device	Destination Device	Ingress Port	Egress Port	Bidirectional	LUN IDs	Frame Type	Description/ Notes
Fixed value, Not specified, or Any (*)	Fixed value, Not specified, or Any (*)	Not specified	Fixed value	Not specified	Fixed value or Not specified	Fixed value	Frame Type and Bidirectional cannot be supported together
Fixed value, Not specified, or Any (*)	Fixed value, Not specified, or Any (*)	Not specified	Fixed value	Fixed value or Not specified	Fixed value or Not specified	Not specified	At least one frame parameter should be specified in the flow.
Fixed value, Not specified, or Any (*)	Fixed value, Not specified, or Any (*)	Fixed value	Not specified	Fixed value or Not specified	Fixed value or Not specified	Not specified	Ingress port and Frame Type cannot be supported together

Flow Monitor supported flow parameter combinations

Table 84 lists the supported flow parameter combinations for Flow Monitor.

NOTE

In the following combinations, an asterisk (*) is also supported as source and destination device values for learning flows only on 16 Gbps-capable Fibre Channel platforms.

TABLE 84 Flow parameter combinations for Flow Monitor

Source Device	Destination Device	Ingress Port	Egress Port	LUN IDs	Frame Type	Bidirectional	Description/ Notes
Fixed value or Not specified	Fixed value or Not specified	Fixed value	Not specified	Not specified	Not specified	Fixed value	Bidirectional is not supported with LUN IDs or Frame Type values
Fixed value or Not specified	Fixed value or Not specified	Not specified	Fixed value	Not specified	Not specified	Fixed value	

TABLE 84 Flow parameter combinations for Flow Monitor

Source Device	Destination Device	Ingress Port	Egress Port	LUN IDs	Frame Type	Bidirectional	Description/ Notes
Fixed value or Not specified	Fixed value or Not specified	Fixed value	Not specified	Fixed value or Not specified	Fixed value	Not specified	Ingress port is only available on 16 Gbps-capable Fibre Channel platforms. Frame Type is supported without source device or destination device values.
Fixed value or Not specified	Fixed value or Not specified	Not specified	Fixed value	Fixed value or Not specified	Fixed value	Not specified	Frame Type is supported without using source device or destination device
Fixed value	Fixed value	Fixed value	Not specified	Fixed value	Not specified	Not specified	LUN monitor without Frame Type needs both source device and destination device values (ASIC limitation)
Fixed value	Fixed value	Not specified	Fixed value	Fixed value	Not specified	Not specified	LUN monitor without Frame Type needs both source device and destination device values (ASIC limitation)

Accessing Flow Vision from other management application features

Access Flow Vision features through the following management application features:

- Frame Viewer
- Bottleneck detection
- Fibre Channel trace route, ping, and port connectivity troubleshooting features.

Frame Viewer

Select a frame from the **Discarded Frames** dialog box for a device or port and select **Add Flow** to launch the **Add Flow Definition** dialog box. The **Add Flow Definition** dialog box will be populated with the following information when available:

- Source Device - Source ID
- Destination Device - Destination ID
- Ingress port - Transmit port
- Egress Port - Receive port
- Monitor - Bidirectional

NOTE

If the source or destination ID are not available, the source and destination device fields will contain asterisks if the target switch supports learning.

Refer to the Frame Viewer section of SAN Device Configuration chapter for more information on accessing and using Frame Viewer and viewing discarded frames. Refer to [“Frame viewer”](#) on page 386.

MAPS

Through the Traffic Monitoring category in the Monitoring and Alerting Policy Suite (MAPS) feature, you can view a filtered list of MAPS violations for imported flows or sub-flows. You can also create MAPS policies for events related to flows and sub-flows. For more information on accessing and using the MAPS features, refer to [Chapter 32, “Monitoring and Alerting Policy Suite”](#).

Bottleneck Detection

Select a port from the **Bottlenecks** dialog box and select **Add Flow** to launch the **Add Flow Definition** dialog box. The **Add Flow Definition** dialog box will be populated with the following information when available:

- Source Device - *
- Destination Device - *
- Ingress port - Transmit port
- Monitor - Bidirectional

[Table 85](#) describes how options on the **Add Flow Definition** dialog box are populated according to the port selected on the **Bottlenecks** dialog box.

TABLE 85 Add Flow Definition dialog box options populated per Bottlenecks selection

Port Selected	Options populated
E_Port	Source Device = * if port is on a 16-Gbps-capable switch, otherwise empty. Destination Device = * if port is on a 16-Gbps switch, otherwise empty. Ingress Port = selected Port Number Target Switch = Selected switch
Initiator	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. Target Switch = Selected port's connected switch
Target	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. Target Switch = Selected port's connected switch
If port other than E_Port, initiator, or target port selected	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. Target Switch = Selected switch

NOTE

Entering * uses all values available devices from the ellipses button on the right of the **Source Device** and **Destination Device** field.

For more information on using Bottleneck Detection, refer to [“Bottleneck detection”](#) on page 967.

Trace route and ping

Select an row from the **Forward Route**, **Reverse Route**, and **FC Ping** tabs on the **Trace Route Summary** dialog box and select **Add Flow** to launch the **Add Flow Definition** dialog box. Use the **Add Flow Definition** dialog box to configure and monitor flows established by your selections on the **Trace Route Summary** dialog box. For more information on accessing and using trace route and ping on the **Trace Route Summary** dialog box, refer to [“Tracing FC routes”](#) on page 922.

If rows are not selected on the **Trace Route Summary** dialog box, the **Add Flow Definition** dialog box is populated with the following information when available:

- Source Device - Source ID for ingress port
- Destination Device - Destination ID for egress port
- Target Switch - Switch for ingress port
- Direction - Bidirectional

If you select a row on the **FC Ping** tab the **Add Flow Definition** dialog box is populated with the following information when available:

- Source Device - Source ID
- Destination Device - Destination ID
- Target Switch - Switch for source WWN
- Direction - Bidirectional

[Table 86](#) describes how options are populated on the **Add Flow Definition** dialog box are populated according to the row selected on the **Forward Route** tab.

TABLE 86 Add Flow Definition dialog box options populated per Forward Route row selection

Row Selected	Options populated
No rows selected	Source Device = Source ID from source device port Destination Device = Destination ID from destination device port. Target Switch = Switch in selected row. Direction = Bidirectional
If first row is selected (where In Port Address is source device port's connected switch port address)	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. Target Switch = Switch of selected row. Direction = Bidirectional
If last row is selected (where Out Port Address is destination device port's connected switch port address)	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. Target Switch = Switch of selected row. Direction = Bidirectional

[Table 87](#) describes how options are populated on the **Add Flow Definition** dialog box are populated according to the row selected on the **Reverse Route** tab,

TABLE 87 Add Flow Definition dialog box options populated per Reverse Route selection

Port Selected	Options populated
No rows selected	Source Device = Source ID from source device port Destination Device = Destination ID from destination device port. Target Switch = Switch in selected row. Direction = Bidirectional
If first row is selected (where In Port Address is destination device port's connected switch port address)	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. Target Switch = Switch of selected row. Direction = Bidirectional
If last row is selected (where Out Port Address is source device port's connected switch port address)	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. Target Switch = Switch of selected row. Direction = Bidirectional
If selected row is neither the first or last 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. Target Switch = Switch of selected row. Direction = Bidirectional

Port connectivity

Select a port from the **Port Connectivity View** dialog box and select **Add Flow** to launch the **Add Flow Definition** dialog box. Use this dialog box to define a flow between the source and destination devices.

[Table 87](#) describes how options are populated on the **Add Flow Definition** dialog box are populated according to the row selected on the **Reverse Route** tab,

TABLE 88 Add Flow Definition dialog box options populated per Reverse Route selection

Row Selected	Options populated
E_Port	Source Device = * if port is on a 16-Gbps switch, otherwise empty. Destination Device = * if port is on a 16-Gbps switch, otherwise empty. Target Switch = Selected switch. Direction = Bidirectional
Initiator	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. Target Switch = Selected port's connected switch. Direction = Bidirectional

TABLE 88 Add Flow Definition dialog box options populated per Reverse Route selection

Row Selected	Options populated
Target row	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. Target Switch = Selected port's connected switch. Direction = Bidirectional
If selected row is neither the first or last 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 = Selected port number. Target Switch = Selected switch. Direction = Bidirectional

For more information on accessing and using the **Port Connectivity View** dialog box, refer to [“Viewing port connectivity” on page 389](#).

Top Talkers

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 [“Provisioning flows” on page 996](#) 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.

End-to-End Monitors

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 [“Provisioning flows” on page 996](#) 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.

28 Accessing Flow Vision from other management application features

Frame Monitor

In this chapter

- [Frame Monitor](#) 1033
- [Creating a custom frame monitor](#) 1035
- [Editing a frame monitor](#) 1037
- [Assigning a frame monitor to a port](#) 1037
- [Finding frame monitor assignments](#) 1038
- [Removing a frame monitor from a port](#) 1038
- [Removing a frame monitor from a switch](#) 1039

Frame Monitor

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 89](#).

TABLE 89 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 *Fabric OS 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 435).

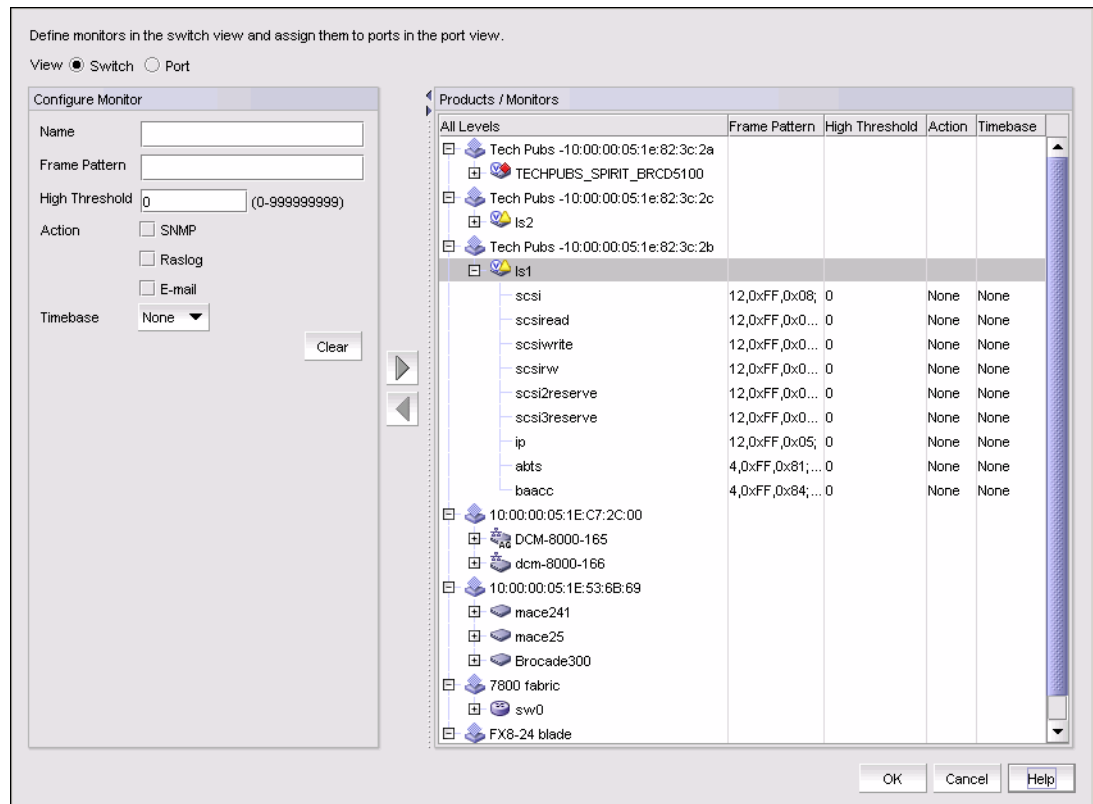


FIGURE 435 Frame Monitor dialog box

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 436).

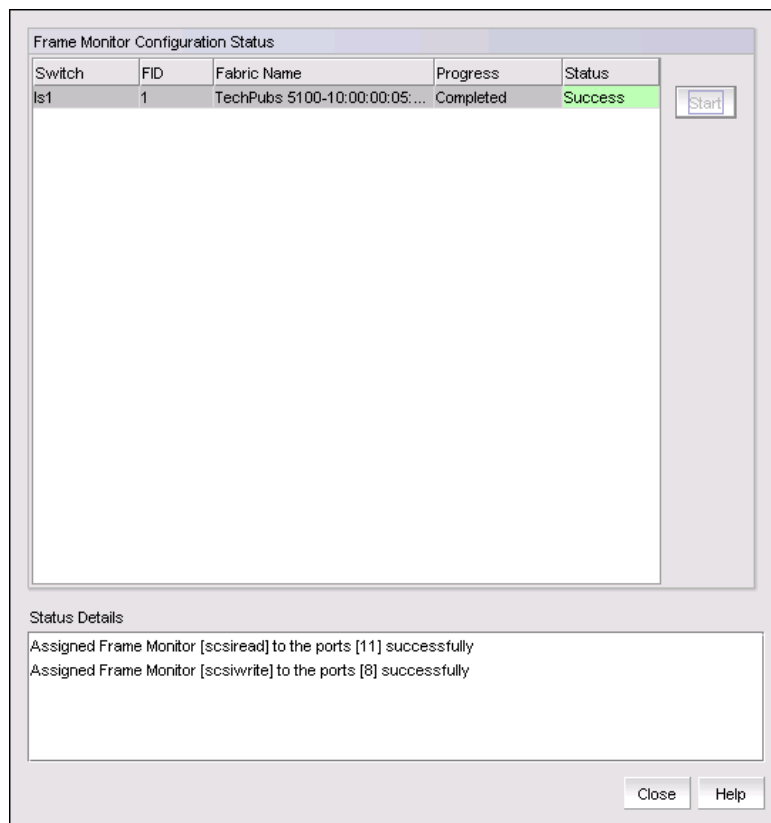


FIGURE 436 Frame Monitor Configuration Status dialog box

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).

29 Removing a frame monitor from a switch

Policy Monitor

In this chapter

- [Policy monitor overview](#) 1041
- [Preconfigured policy monitors](#) 1047
- [Viewing policy monitor status](#) 1048
- [Viewing existing policy monitors](#) 1048
- [Adding a policy monitor](#) 1049
- [Policy monitor scheduling](#) 1055
- [Editing a policy monitor](#) 1056
- [Deleting a policy monitor](#) 1057
- [Running a policy monitor](#) 1057
- [Viewing a policy monitor report](#) 1058
- [Viewing historical reports for all policy monitors](#) 1061
- [Viewing historical reports for a policy monitor](#) 1062

Policy monitor overview

Use the Policy Monitor 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 monitors 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 policy monitors

Fabric policy monitors enable you to set the following policy monitors on SAN (refer to [“Adding a policy monitor”](#) on page 1049):

- **Check zoning status** — This fabric policy monitor 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 policy monitor report shows a violation, the Administrator can use the **Zoning** dialog box (**Configure > Zoning > Fabric**) to fix the violation. Refer to [“Zoning”](#) on page 755.

For example, if you use the policy monitor 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 policy monitor 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 policy monitor report shows a violation, the Administrator can use the **Zoning** dialog box (**Configure > Zoning > Fabric**) to fix the violation. Refer to [“Zoning”](#) on page 755.

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 policy monitor enables 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 policy monitor 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 policy monitor 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 policy monitor 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 755).

For example, if you use the policy monitor 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 policy monitors

Switch and router policy monitors enable you to set the following policy monitors on switches and routers.

- **Check connections: redundant connections to neighboring switches (SAN only)** — This switch and router policy monitor 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 policy monitor 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 policy monitor 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 policy monitor 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 policy monitor 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 policy monitor 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 1063.

If the policy monitor 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 1063.

- **Check if the product is configured to send Upload Failure Data Capture to an FTP server (SAN only)** — This switch and router policy monitor 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 1195.

- **Check for SSH (secure Telnet) configuration** — This switch and router policy monitor 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 policy monitor 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 policy monitor 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 policy monitor report shows a violation, configure SNMPv3 on the device. Remove SNMPv1 and SNMPv2 settings on the device, if configured.

Host policy monitors

Host policy monitors enable you to set the following checks on host devices.

- **Check for multiple fabrics connections** — This host policy monitor 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 [“Discovering fabrics”](#) on page 39)
 - Make sure there are Brocade HBAs on the host.
 - Make sure to configure the host port mapping. (refer to [“Host port mapping overview”](#) on page 423)
- Host adaptor discovery with 2.1 or later driver for host to unmanaged fabric connections (refer to [“Host discovery”](#) on page 58)
 - Make sure there are Brocade HBAs on the host.

- Fabric plus Host adapter discovery with 2.1 or earlier driver (refer to [“Host discovery”](#) on page 58)
Make sure there are Brocade HBAs on the host.
- Fabric plus VM Manager for hosts discovered through vCenter (refer to [“VM Manager discovery”](#) on page 68)
Make sure there are Brocade HBAs on the host.
Make sure you discover the associated fabrics.
- VM Manager plus Host adapter discovery (refer to [“VM Manager discovery”](#) on page 68)
Make sure there are Brocade HBAs on the host.
Make sure you discover the associated fabrics.

Rule Violation Fix — If the policy monitor 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 policy monitor 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 58)
Make sure there are Brocade HBAs (with a 2.1 or later driver) on the host.
- Fabric plus Host discovery (refer to [“Discovering fabrics”](#) on page 39)
Make sure there are Brocade HBAs on the host connected to the fabric.
Make sure to configure the host port mapping (refer to [“Host port mapping overview”](#) on page 423).
- Fabric plus VM Manager discovery (refer to [“Discovering fabrics”](#) on page 39)
Make sure there are Brocade HBAs (with a 2.1 or later driver) on the host connected to the fabric.
- VM Manager plus Host discovery (refer to [“VM Manager discovery”](#) on page 68)
Make sure there are Brocade HBAs (with a 2.1 or later driver) on the host.
Make sure you discover the associated fabrics.

Rule Violation Fix — If the policy monitor 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 policy monitor

The management policy monitor enables you to set a policy monitor on the Management application.

Check to see if the server backup is enabled and working – This management policy monitor 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 policy monitor 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 77.

Preconfigured policy monitors

The Management application provides preconfigured policy monitors. The preconfigured policy monitors 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** – Weekly
- **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 PDT 2012.
- **Last Run** – Empty
- **Result** – Empty
- **Rule** – The default SAN policy is configured with the following rules:
 - Zoning status
 - 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 policy monitor status

You can view policy monitor status from the main Management application window or from the **Policy Monitor** dialog box.

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

TABLE 90 Policy Monitor Icons

Icon	Description
	Passed — Displays when all policy monitors, excluding un-alerted and acknowledged monitors, pass. Pause on icon to display flyover detail: Policy monitor is OK.
	Failed — Displays when at least one policy monitor failed. Pause on icon to display flyover detail: The last run of <i>number</i> policy monitor(s) has one or more failures.

To view more detail regarding policy monitor status, click the **Policy Monitor** icon. The **Policy Monitor** dialog box displays. For more information, refer to “[Viewing existing policy monitors](#)” on page 1048

Viewing existing policy monitors

To view existing policy monitors, complete the following steps.

1. Select **Monitor > Policy Monitor** (Figure 437).

The **Policy Monitor** dialog box displays.

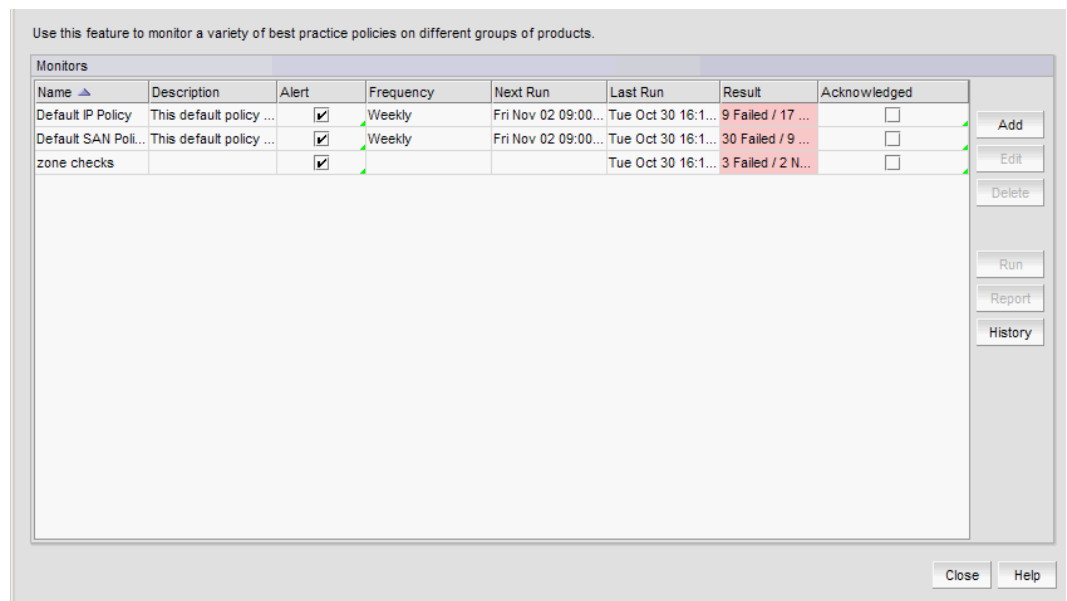


FIGURE 437 Policy Monitor dialog box

2. Review the policy monitor 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 policy monitor 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 last policy monitor 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 **Alert** check box is cleared.
3. To add a policy monitor, click **Add** (refer to [“Adding a policy monitor”](#) on page 1049).
4. To edit the selected policy monitor, click **Edit** (refer to [“Editing a policy monitor”](#) on page 1056).
5. To delete the selected policy monitor, click **Delete** (refer to [“Deleting a policy monitor”](#) on page 1057).
6. To run the selected policy and view the report, click **Run** (refer to [“Running a policy monitor”](#) on page 1057).
7. To open the last executed report for a selected policy monitor, select a policy monitor and click **Report** (refer to [“Viewing a policy monitor report”](#) on page 1058).
8. To view the report history for all policy monitors, click **History** (refer to [“Viewing historical reports for a policy monitor”](#) on page 1062).
9. To view the report history for a selected policy monitor, select a policy monitor and click **History** (refer to [“Viewing historical reports for a policy monitor”](#) on page 1062).
10. Click **Close** on the **Policy Monitor** dialog box.

Adding a policy monitor

To view existing policy monitors, complete the following steps.

1. Select **Monitor > Policy Monitor**.
The **Policy Monitor** dialog box displays.
2. Click **Add**.
The **Add Monitor** dialog box displays ([Figure 438](#)).

The screenshot shows the 'Add Policy Monitor' dialog box with the 'Fabric Checks' tab selected. At the top, there are input fields for 'Name' and 'Description'. Below these is a 'Schedule' section with a 'Use' checkbox and a dropdown menu currently showing '<Not configured>'. The main area contains four tabs: 'Fabric Checks', 'Switch / Router Checks', 'Host Checks', and 'Management Checks'. Under the 'Fabric Checks' tab, there are 'Zoning Checks' with three checkboxes: 'Check zoning status' (radio buttons for 'Enabled' and 'Disabled'), 'Check that all zones belong to at least one zone config', and 'Check the number of initiator ports zoned to each storage port'. The last checkbox is checked, and an 'Initiator Port Limit' spinner is set to '20' with '(1-2147483647)' next to it. Below the checkboxes are two tables: 'Available Fabrics' and 'Selected Fabrics'. Both tables have columns for 'Name', 'Seed Switch', 'Status', and 'Last Discovery'. The 'Available Fabrics' table contains one entry, 'All Fabrics'. The 'Selected Fabrics' table is empty. At the bottom right, there are 'OK', 'Cancel', and 'Help' buttons.

FIGURE 438 Add Policy Monitor 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 policy monitor:
 - To configure deployment to run only once, refer to [“Configuring a one-time policy monitor schedule”](#) on page 1055.
 - To configure hourly deployment, refer to [“Configuring an hourly policy monitor schedule”](#) on page 1055.
 - To configure daily deployment, refer to [“Configuring a daily policy monitor schedule”](#) on page 1055.

- To configure weekly deployment, refer to [“Configuring a weekly policy monitor schedule”](#) on page 1056.
 - To configure monthly deployment, refer to [“Configuring a monthly policy monitor schedule”](#) on page 1056.
6. To set policy monitors 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 policy monitors”](#) on page 1042.
 - 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 policy monitors”](#) on page 1042.
 - 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 policy monitors”](#) on page 1042.
 - 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 policy monitors”](#) on page 1042.
 - e. Enter the initiator port limit in the **Initiator Port Limit** field.

The default recommended threshold ratio is 20:1 (20 initiator ports to 1 target port). Therefore, if the ratio for the storage port is equal to or higher than 20:1, the policy monitor 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** table 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 policy monitors for switches, select the **Switch/Router Checks** tab (Figure 439) and complete the following steps.

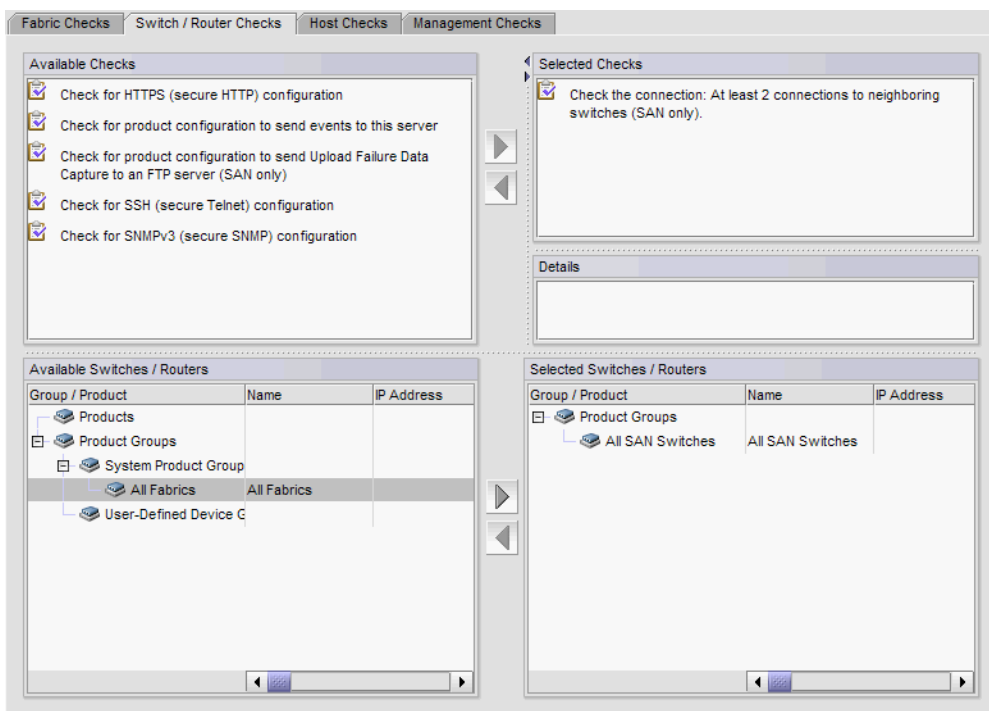


FIGURE 439 Add Policy Monitor dialog box, Switch/Router Checks tab

- a. Select one or more of the following checks in the **Available Checks** list to include them in the policy monitor:

For more information about these checks and fixes for rule violations, refer to [“Switch and router policy monitors”](#) on page 1043.

- Select the **Check if the product is configured to send events to this server** check 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 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 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 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 to check each target to see if SSH is enabled for device data transmission.
- Select the **Check for SNMPv3 (secure SNMP) configuration** check to check each target to see if SNMPv3 is active for device data transmission and SNMPv1 and SNMPv2 are not configured.

- 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, enter the minimum number of connections allowed between a switch pair in the **Minimum Connections** field.

The default recommended is 2. Valid values are from 2 through 512.

- d. Select the switches or routers to which you want to apply this policy in the **Available Switches/Routers** list and click the right arrow button.

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 policy monitors for hosts, select the **Host Checks** tab (Figure 440) and complete the following steps.

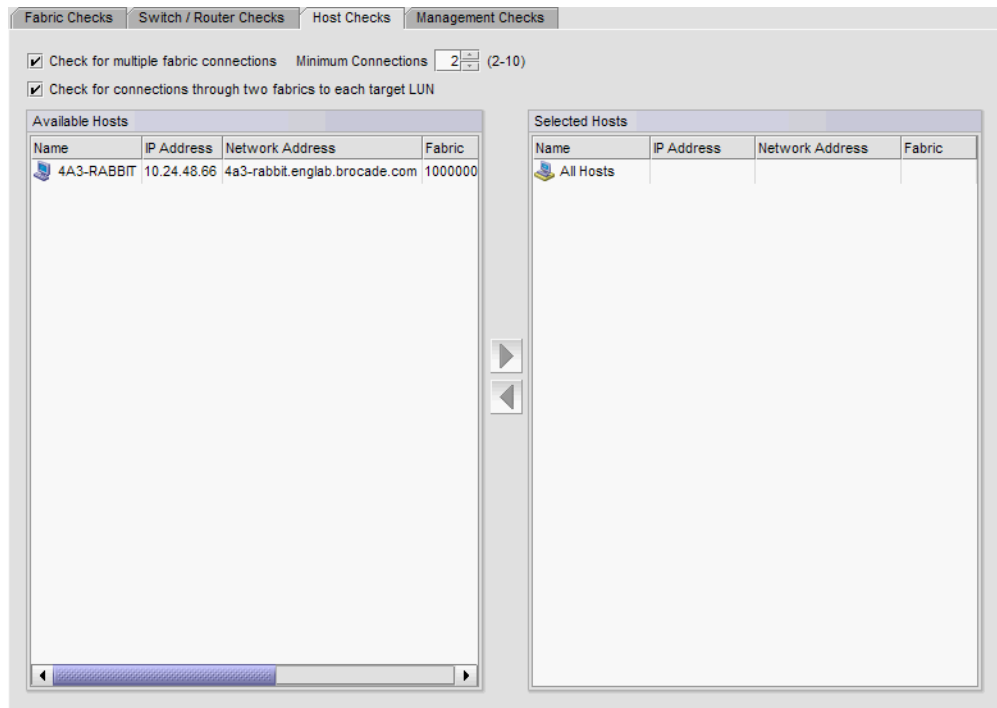


FIGURE 440 Add Policy Monitor dialog box, Hosts Checks tab

- 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 policy monitors”](#) on page 1045.

- 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 policy monitors”](#) on page 1045.

- 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 Host 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 policy monitors for the Management application ([Figure 441](#)), complete the following steps.




FIGURE 441 Add Policy Monitor dialog box, Management Checks tab

- a. Select the **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 policy monitor”](#) on page 1047.

10. Click **OK** on the **Add Monitor** dialog box.

The **Policy Monitor** dialog box displays with the new policy monitor in the **Monitors** list.

11. Click **Close** on the **Policy Monitor** dialog box.

Policy monitor scheduling

You can schedule a policy monitor to run automatically. For step-by-step instructions, refer to the following procedures:

- [“Configuring a one-time policy monitor schedule”](#) on page 1055
- [“Configuring an hourly policy monitor schedule”](#) on page 1055
- [“Configuring a daily policy monitor schedule”](#) on page 1055
- [“Configuring a weekly policy monitor schedule”](#) on page 1056
- [“Configuring a monthly policy monitor schedule”](#) on page 1056

Configuring a one-time policy monitor 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 policy monitor, return to [step 6](#) of [“Adding a policy monitor”](#) on page 1049.

Configuring an hourly policy monitor 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 policy monitor, return to [step 6](#) of [“Adding a policy monitor”](#) on page 1049.

Configuring a daily policy monitor 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 policy monitor, return to [step 6](#) of “Adding a policy monitor” on page 1049.

Configuring a weekly policy monitor 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 policy monitor, return to [step 6](#) of “Adding a policy monitor” on page 1049.

Configuring a monthly policy monitor 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 policy monitor, return to [step 6](#) of “Adding a policy monitor” on page 1049.

Editing a policy monitor

To edit an existing policy monitor, complete the following steps.

1. Select **Monitor > Policy Monitor**.

The **Policy Monitor** dialog box displays.

2. Select the policy you want to edit in the **Monitors** list and click **Edit**.

The **Edit Policy Monitor** dialog box displays. The **Edit Policy Monitor** dialog box has the same fields and components as the **Add Policy Monitor** 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 policy monitor checks, repeat [step 5](#) through [step 9](#) of “Adding a policy monitor” on page 1049.
6. Click **OK** on the **Edit Monitor** dialog box.
The updated policy monitor displays in the **Monitors** list of the **Policy Monitor** dialog box.
7. Click **Close** on the **Policy Monitor** dialog box.

Deleting a policy monitor

To delete an existing policy monitor, complete the following steps.

1. Select **Monitor > Policy Monitor**.
The **Policy Monitor** 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 **Policy Monitor** dialog box.

Running a policy monitor

Before you run a policy monitor, make sure your policy monitors are valid. Valid policy monitors must have at least one policy selected with one or more targets. Management checks do not require a target.

To run an existing policy monitor, complete the following steps.

1. Select **Monitor > Policy Monitor**.
The **Policy Monitor** dialog box displays.
2. Select the policy you want to run in the **Monitors** list.
3. Click **Run**.

When the policy monitor check is complete, the *Policy_Name - Policy Monitor Report* displays ([Figure 442](#)) in a web browser.

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.

FIGURE 442 *Policy_Name* - Policy Monitor Report

4. Review the report details (refer to “[Viewing a policy monitor report](#)” on page 1058).
To export a report, refer to “[Exporting a policy monitor report](#)” on page 1061.
To e-mail a report, refer to “[Exporting reports to e-mail recipients](#)” on page 1206.
5. Click the close button (X) on the *Policy_Name* - **Policy Monitor Report** browser window.
6. Click **Close** on the **Policy Monitor** dialog box.

Viewing a policy monitor report

NOTE

You must run the policy monitor at least once before you can view a report.

To view an existing policy monitor report, complete the following steps.

1. Select **Monitor > Policy Monitor**.
The **Policy Monitor** 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* - **Policy Monitor Report** displays (Figure 442) in a web browser.

4. Review the report details:

- **Name** — Name of the policy monitor report.
- **Date** — Date and time the report was finished.
- **Export** button — To export a report, refer to “[Exporting a policy monitor report](#)” on page 1061.
- **E-Mail** button — To e-mail a report, refer to “[Exporting reports to e-mail recipients](#)” on page 1206.
- **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 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.
- 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 — 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.

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* - **Policy Monitor Report** browser window.
- 6. Click **Close** on the **Policy Monitor** dialog box.

Exporting a policy monitor 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* - **Policy Monitor Report** browser window.

Viewing historical reports for all policy monitors

1. Select **Monitor > Policy Monitor**.
The **Policy Monitor** 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 policy monitor.
 - **Name** — Name of the policy monitor.
 - **Date** — Date and time the report was finished.
 - **Result** — Result of the policy monitor 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* - **Policy Monitor Report** displays in a web browser. For detailed information about reports, refer to “[Viewing a policy monitor report](#)” on page 1058.
4. Click the close button (X) on the *Policy_Name* - **Policy Monitor Report** browser window.
5. Click **Close** on the **Report History** dialog box.

Viewing historical reports for a policy monitor

1. Select **Monitor > Policy Monitor**.

The **Policy Monitor** 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 policy monitor.

- **Name** — Name of the policy monitor.
- **Date** — Date and time the report was finished.
- **Result** — Result of the policy monitor 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* - **Policy Monitor Report** displays in a web browser. For detailed information about reports, refer to [“Viewing a policy monitor report”](#) on page 1058.

4. Click the close button (X) on the *Policy_Name* - **Policy Monitor Report** browser window.

5. Click **Close** on the **Report History** dialog box.

Fault Management

In this chapter

• Fault management overview.....	1063
• Event notification.....	1064
• Defining filters.....	1066
• SNMP traps.....	1069
• SNMP informs.....	1082
• Syslogs.....	1083
• Event action definitions.....	1088
• Pseudo events.....	1103
• Event custom reports.....	1114
• Event custom report schedules.....	1123
• Event logs.....	1126

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 283.

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 443](#)) displays.

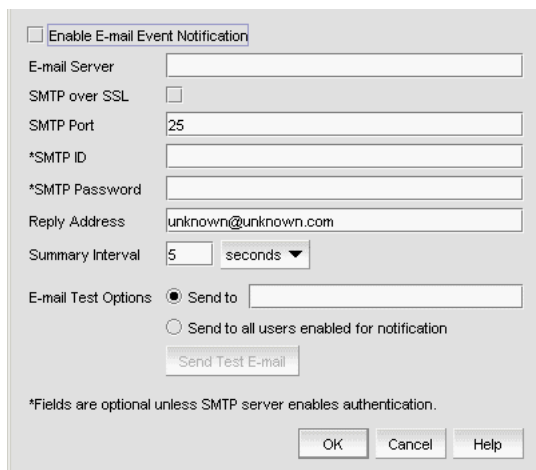


FIGURE 443 E-mail Event Notification Setup dialog box

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 444](#), 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 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, shown in [Figure 444](#), displays.

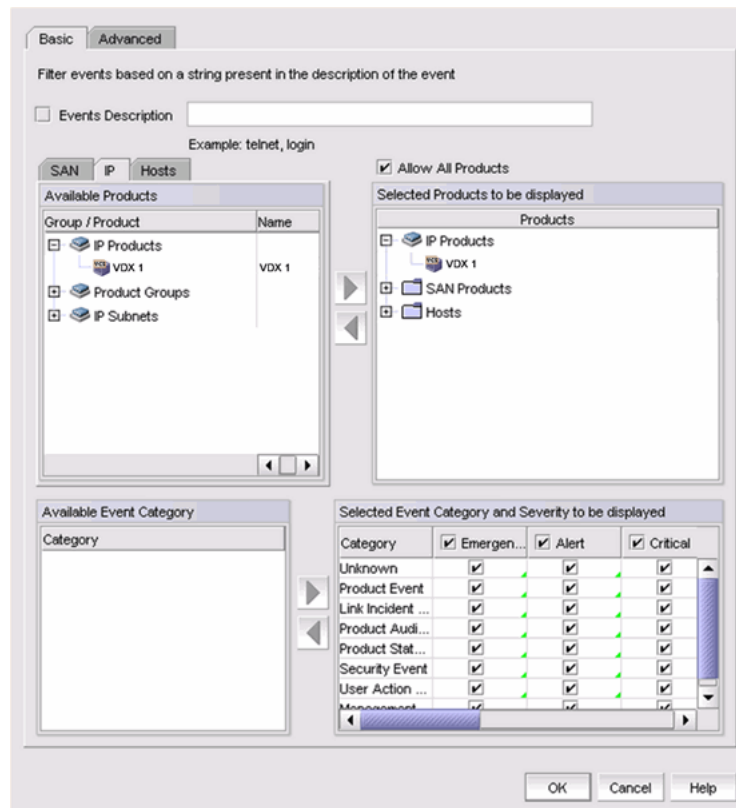


FIGURE 444 Define Filter dialog box

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 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.
7. Select one or more event categories from the **Available Event Category** list and click the right arrow button to move it 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 445](#), displays.

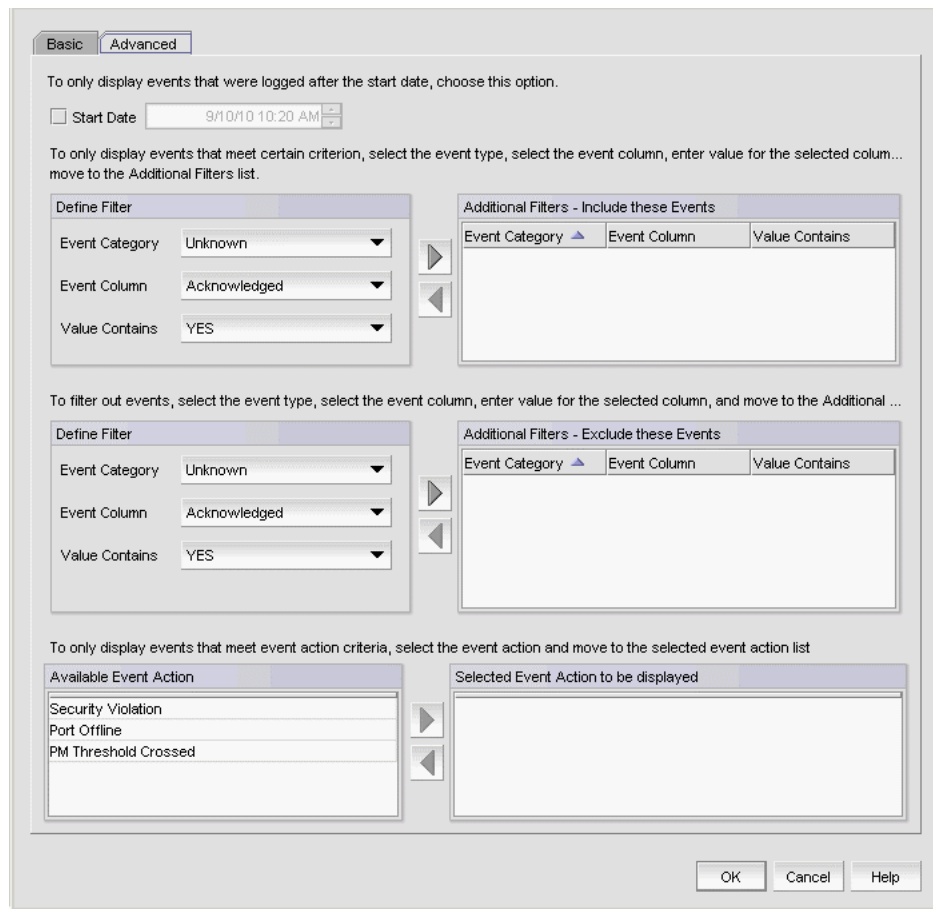


FIGURE 445 Define Filter dialog box - Advanced tab

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](#).

- To exclude events from the event filter, complete the following steps.

NOTE

You can configure a maximum of ten filters to be included.

- Select the event type you want to remove from the **Event Category** list.
All event types are listed in alphabetical order.
 - Select the event column for the event from the **Event Column** list.
All event columns are listed in alphabetical order.
 - Enter all or part of the event type value in the **Value Contains** field.
 - Click the right arrow button to move the event type to the **Additional Filters - Exclude these Events** list.
 - To remove additional filters, repeat [step a](#) through [step d](#).
- 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.
 - 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.

- Right-click a switch from the device tree or connectivity map.
- 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”](#) on page 1070
- [“Removing a trap recipient from one or more switches”](#) on page 1071
- [“SNMP trap forwarding”](#) on page 1071
- [“Adding a trap destination”](#) on page 1072
- [“Adding a new trap filter”](#) on page 1073
- [“Event reception”](#) on page 1075
- [“Adding an SNMP v3 credential”](#) on page 1077

- “Adding an SNMP v1 or v2c community string” on page 1078
- “Importing a new MIB into the Management application” on page 1079
- “Trap customization” on page 1080
- “Unregistering a registered trap” on page 1081
- “Customizing a registered trap definition” on page 1082
- “Reverting the customization of a registered trap to default” on page 1082

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 446](#), displays.

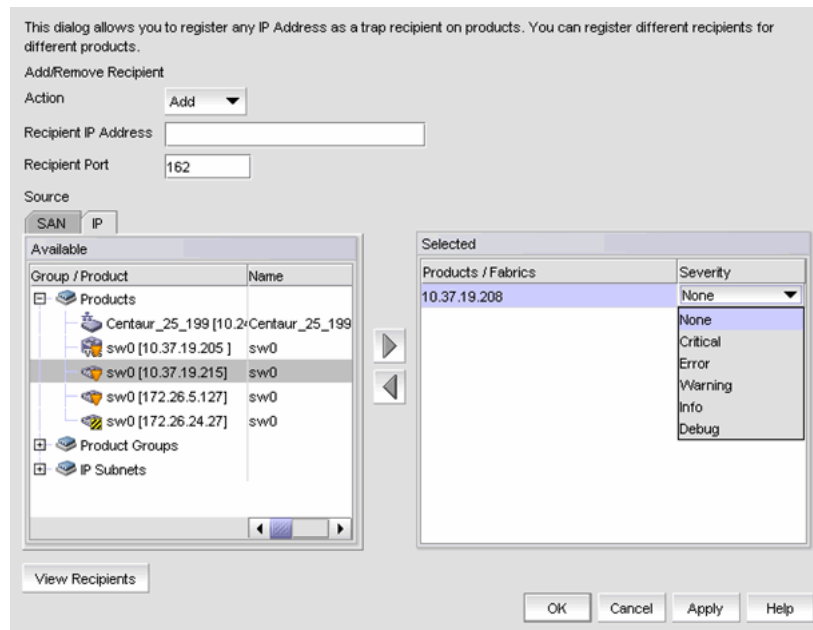


FIGURE 446 SNMP Trap Recipients dialog box

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 446](#), 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.

1. Select **Monitor > SNMP Setup > Trap Forwarding**.

The **SNMP Trap Forwarding** dialog box, shown in [Figure 447](#), displays.

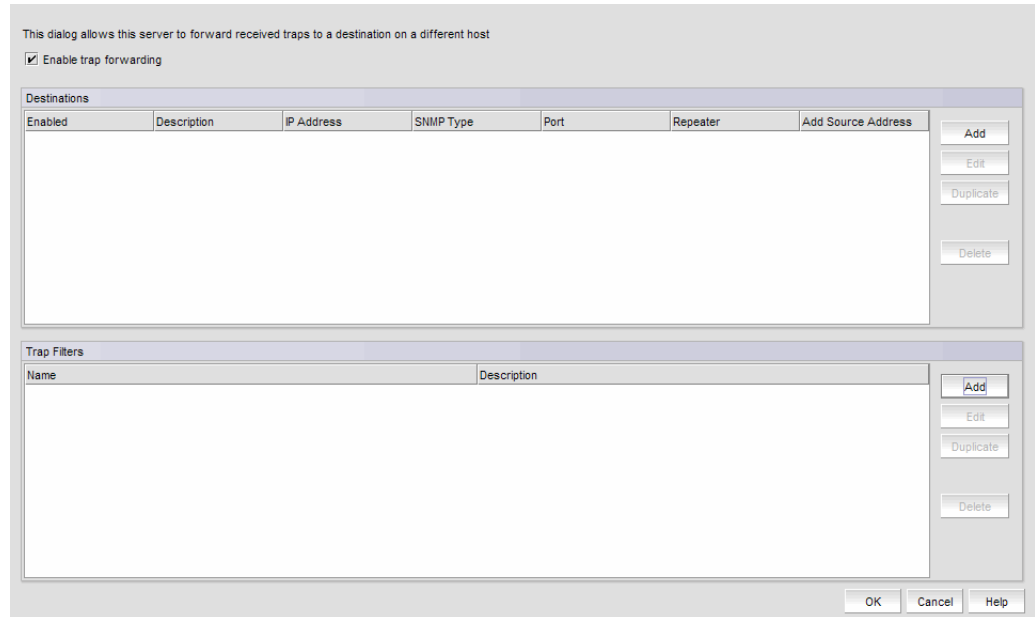


FIGURE 447 SNMP Trap Forwarding dialog box

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 447](#), 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 448](#), displays.

FIGURE 448 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, the Open View Source Name 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.
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 449](#), displays.

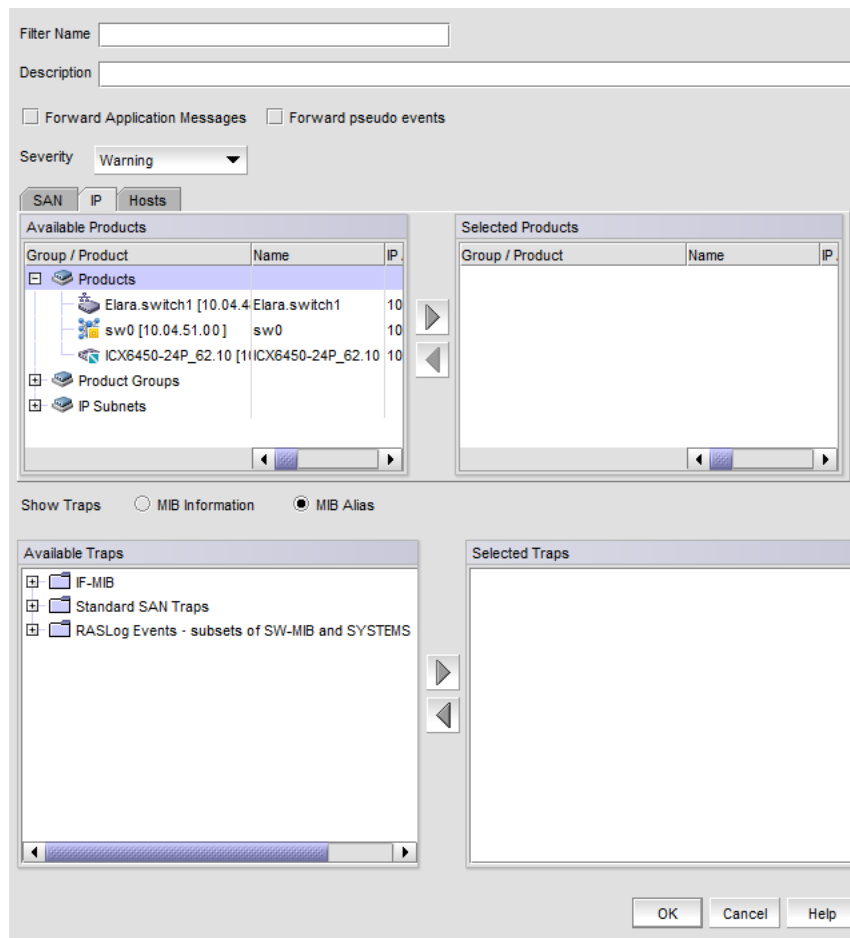


FIGURE 449 Add Trap Filter dialog box

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.
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 450](#), displays.

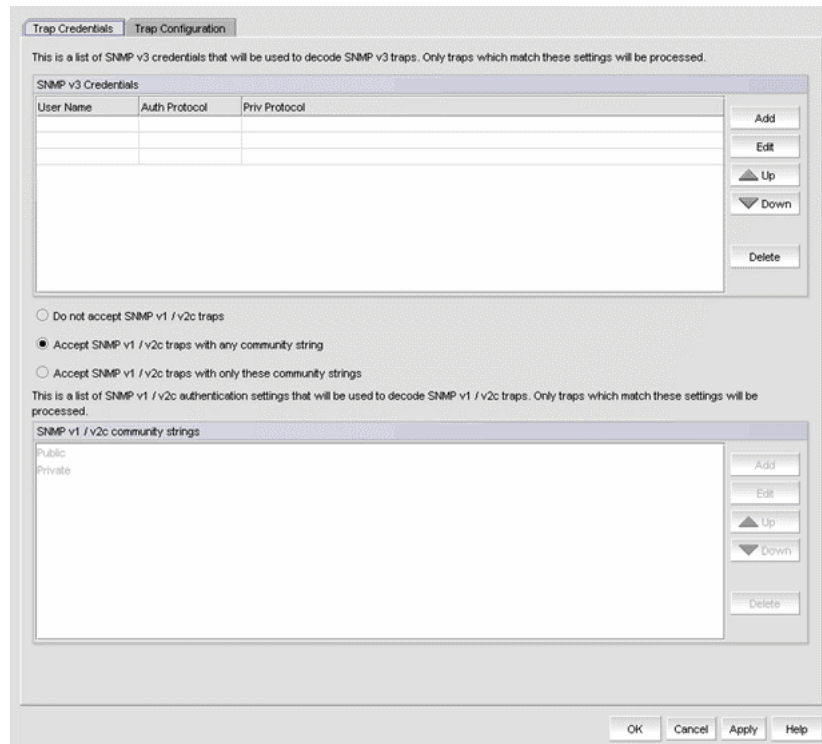


FIGURE 450 Event Reception dialog box - Trap Credentials tab

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.

[Table 91](#) explains the combinations of security and authentication, which will help you when you make your SNMP credentials configuration decisions.

TABLE 91 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 1077 or [“Adding an SNMP v1 or v2c community string”](#) on page 1078.

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 451](#), displays.

FIGURE 451 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 452](#), displays.

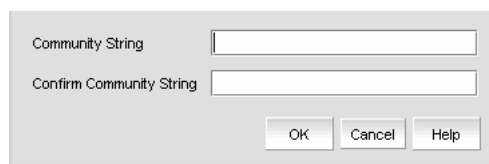


FIGURE 452 SNMP v1/v2 Community String dialog box

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 453](#), 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.

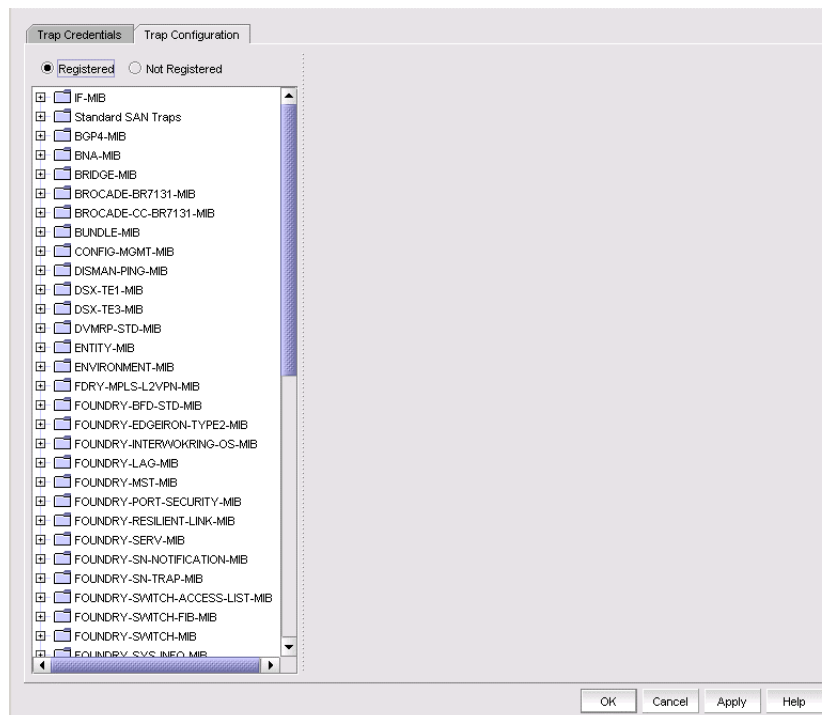


FIGURE 453 Trap Configuration tab of the Event Reception dialog box

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 acknowledgement 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

IronWare OS 6910 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 454](#), displays.

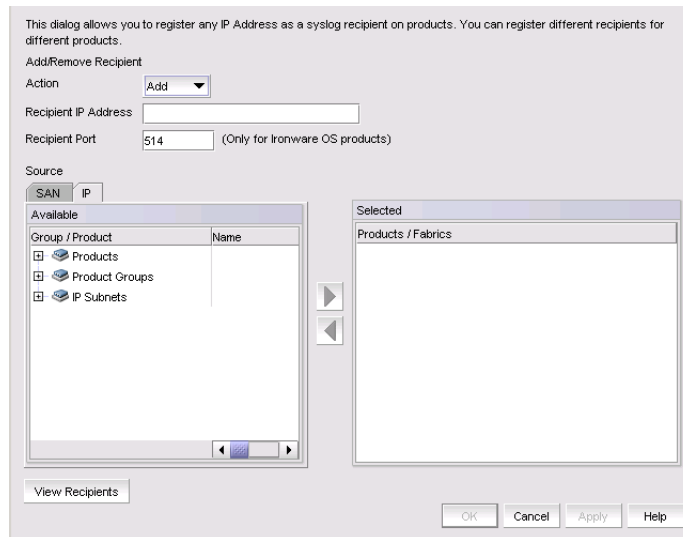


FIGURE 454 Syslog Recipients dialog box

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 455](#), displays.

This dialog allows this server to forward received syslog events to a destination on a different host

Enable syslog forwarding

Enable	Description	IP Address	Port	Repeater
Yes	Forwarding to 220	192.1.1.220	514	Yes

Name	Description
Filter 1	Sample filter 1

FIGURE 455 Syslog Forwarding dialog box

2. Select the **Enable syslog forwarding** check box.
3. Click **Add**.

The **Add Syslog Destination** dialog box, shown in [Figure 456](#), displays. The **Enable** and **Syslog Repeater** check boxes are selected by default.

FIGURE 456 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 457](#), displays.

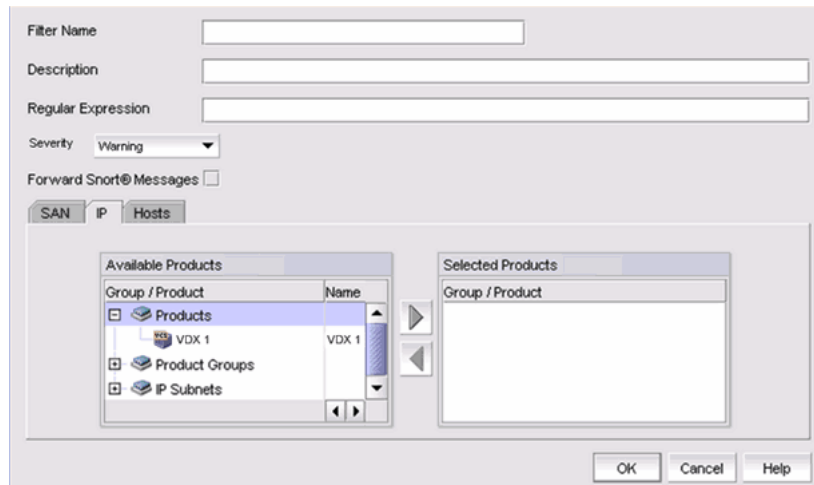


FIGURE 457 Add Syslog Filter dialog box

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 1088 for more information.
9. 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.

10. Select the product from the **Available Products** list and click the right arrow button to move it to the **Selected Products** list.
11. 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 1086).

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 458](#), 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 458](#), displays.

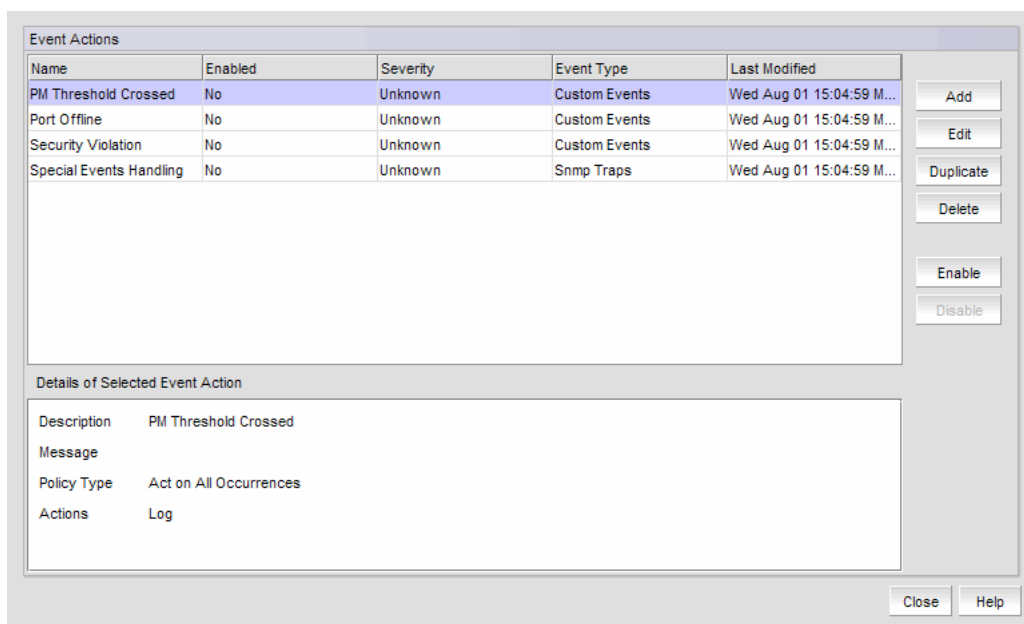


FIGURE 458 Event Actions dialog box

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 **Next** to advance to the **Events** pane.

By default, the **Events** pane of the **Add Event Action** dialog box displays, shown in [Figure 459](#).

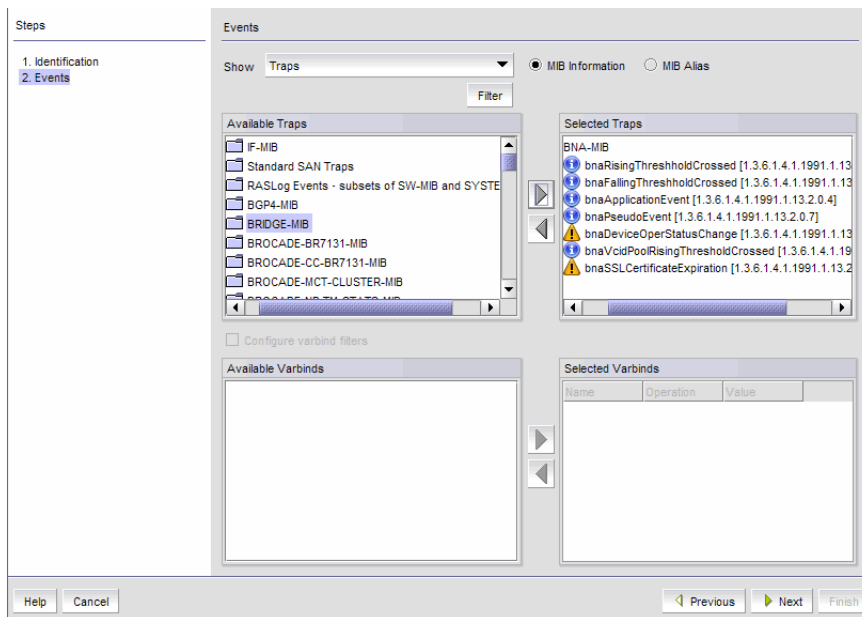


FIGURE 459 Add Event Action dialog box - Events pane

3. Select one of the following event types from the **Show** list:

- Traps (default)
- Application Events
- Pseudo Events
- Custom Events
- Snort® Message

Depending on what event type you select, a box listing the available events or pseudo events displays.

4. 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 doing any of the following:

- Click one of the following buttons:
 - **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.

5. Click the **Filter** button to launch the **Trap Filters** dialog box, which allows you to find the trap you want.

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.

7. If you selected **Application Events** in [step 3](#), select the application events in the left table and use the arrow button to move them to the right.

8. If you selected **Pseudo Events** in [step 3](#), 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 3](#), 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 3](#), 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 1091 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 3](#) of [“Creating an event action definition”](#) on page 1088.

The varbinds for the selected trap are listed in the **Available Varbinds** list, shown in [Figure 460](#).

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.

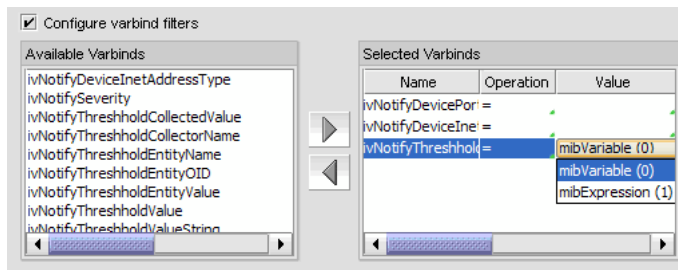


FIGURE 460 Available Varbinds and Selected Varbinds lists

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 460](#).
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 461](#), 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.

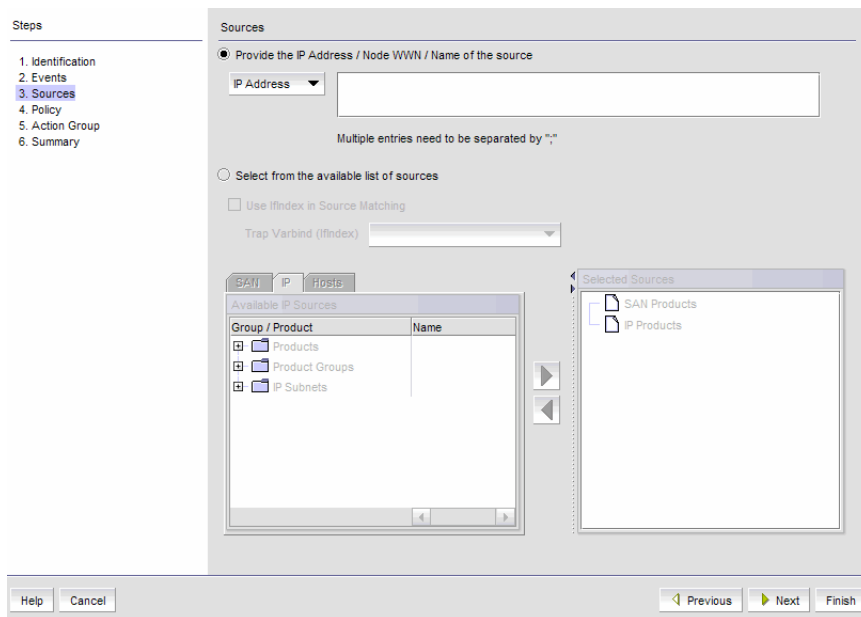


FIGURE 461 Sources pane of the Add Event Action dialog box

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.
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. Click the **Select from the available list of sources** button as an alternative to manually entering the IP address, WWN, or name of the source. You can select source address products or ports to use as event senders from the available list of sources.
5. Select the **Use If index in source matching** check box if you want to use if Index to filter traps on a specific port of a product; otherwise, the filter is applied globally on a product.
6. If the **Use Ifindex in source matching** check box is selected, select the varbind to be used from the **Trap Varbind (Ifindex)** list.
7. 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.

8. 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 462](#), 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 462 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** - 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 or hours. 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”](#).

Editing event actions

The **Action Group - Actions** pane of the **Edit Event Action** dialog box, shown in [Figure 463](#), defines what action the Management application takes when the criteria are met.

The screenshot shows the 'Action Group - Actions' pane of the 'Edit Event Action' dialog box. The interface is organized into several sections:

- Steps:** A sidebar on the left lists the configuration steps: 1. Identification, 2. Events, 3. Sources, 4. Policy, 5. Action Group, and 5.1. Actions (the current step).
- Actions:**
 - Apply as Logging Policy**
 - Log** **Drop**
 - Auto Acknowledge**
 - Enable Troubleshooting (IP only)**
 - In case of maintenance, events can be suppressed up to 168 hours (7 days).
 - Time: (0-168 Hours) (0-59 Minutes)
 - Alert by E-mail**
 - Run Policy Monitor** Target:
 - Launch a Script**
 - Send Event parameters as arguments (Level, Source Name, Source Address, Type and Description)
 - Broadcast to Client**
 - Mark as Special Events**
- Tech Support:**
 - Collect support save (only for event sender)**
- Deployment:**
 - Deploy CLI Configuration** Selected Configuration:
 - Has Parameters:
 - Parameters table:

Parameter	Source	Transformation
 - Deploy Product Configuration** Selected Configurations:
 - Target Source:
 - Targets: Source: Transformation:

At the bottom of the dialog box, there are buttons for **Help**, **Cancel**, **Previous**, **Next**, and **Finish**.

FIGURE 463 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 this also helps to avoid cluttering Master Log with unwanted messages without modifying filters.

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.

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.

See [step 7](#) of “[Editing event actions](#)” on page 1095 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 following status bar icon displays:



FIGURE 464 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 465](#).
The Special Events dialog box, shown in [Figure 465](#), lists the most recent 1000 events that have been identified as special events.

31 Event action definitions




FIGURE 465 Special Events dialog box

2. Select the **Acknowledged** check box that corresponds to the special event you want to acknowledge.

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.

3. 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 466](#), allows you to select e-mail recipients from a list, add new e-mail recipients, and compose e-mail messages.

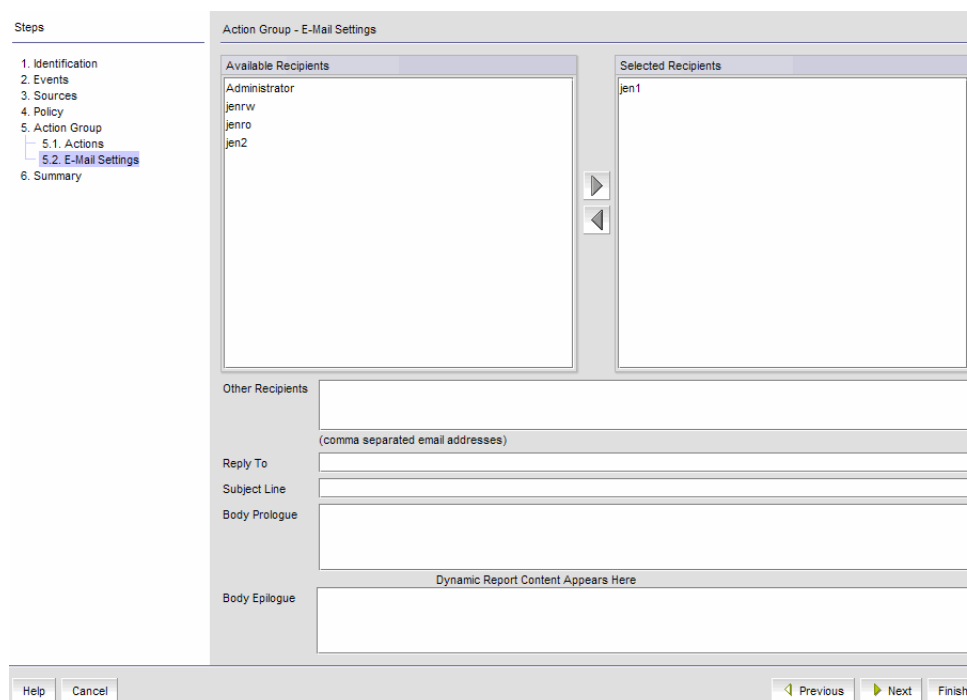


FIGURE 466 Action Group - E-Mail Settings pane of the Add Event Action dialog box

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 semicolon. 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, enter the text in the **Subject Line** 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 this 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 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.

6. Click **Finish**.

The **Summary** pane of the **Edit Event Action** dialog box displays an overview of the e-mail configuration you are creating.

7. 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

**CAUTION**

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. See [“Creating an event action definition”](#) on page 1088 for complete instructions on event actions.

The **Events** pane of the **Add Event Action** dialog box displays, shown in [Figure 467](#). Snort® Message is the default in the **Show** list.

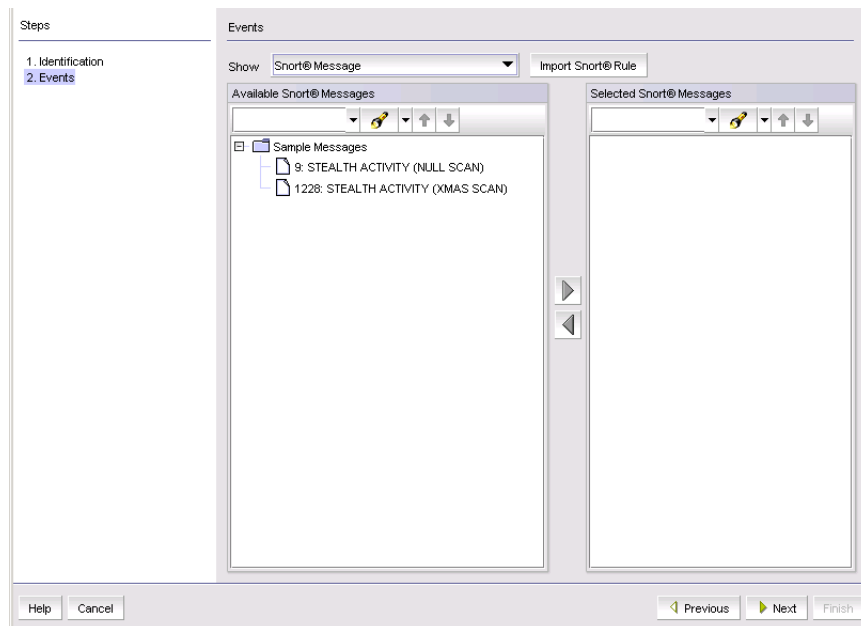


FIGURE 467 Events pane of the Add Event Action dialog box

2. Click the **Import Snort® Rule** button.

The **Import Snort® Rule File** dialog box displays, shown in [Figure 468](#).




FIGURE 468 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 469](#), displays.

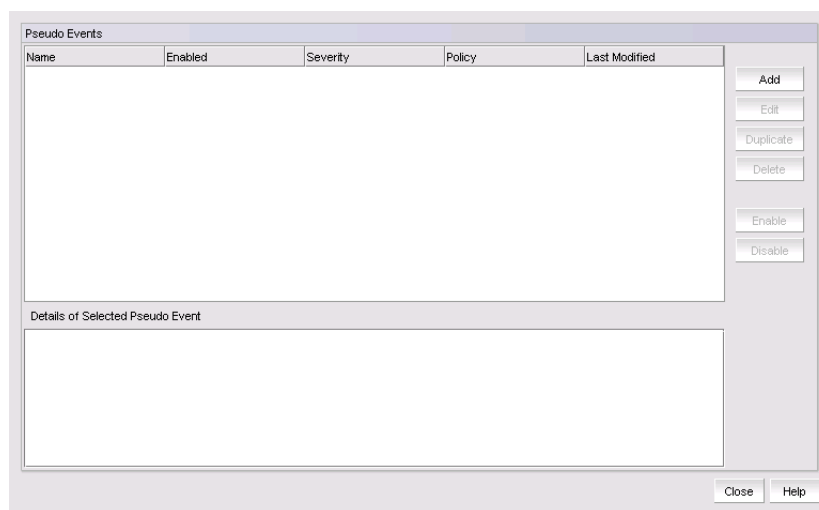


FIGURE 469 Pseudo Events dialog box

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 469](#), 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 470](#), displays.

Setting pseudo event policies

The **Policy** pane of the **Add Pseudo Event** dialog box, shown in [Figure 470](#), 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 470 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 1108 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 1111 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 1112 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 **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 471](#), displays.

Refer to the following topics for specific procedures using this dialog box.

- [“Creating pseudo event definitions”](#) on page 1103
- [“Editing a pseudo event definition”](#) on page 1107

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 469](#), displays.

2. Click **Add**.

The **Events** pane of the **Add Pseudo Event** dialog box, shown in [Figure 471](#), displays.

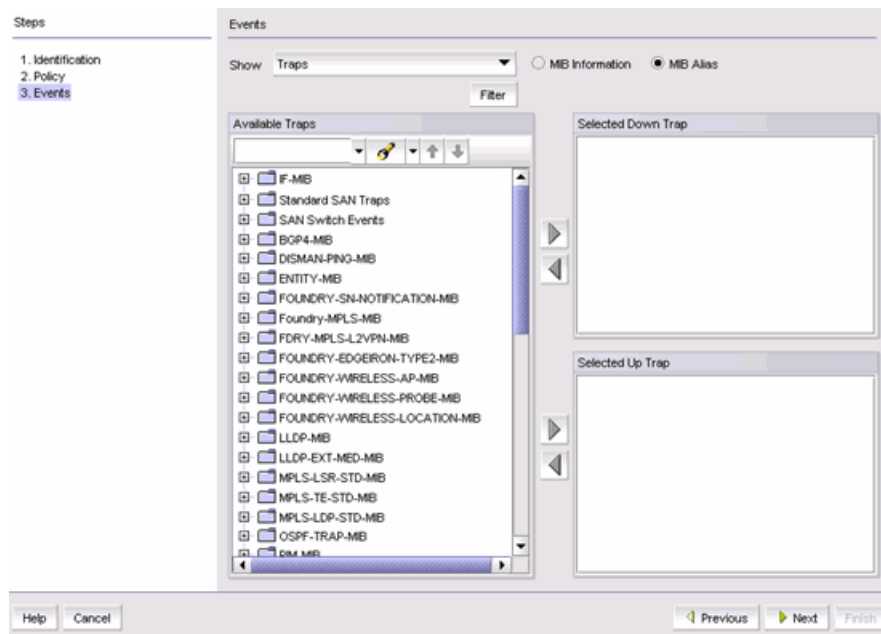


FIGURE 471 Events pane of the Add Pseudo Event dialog box

3. From the **Available Traps** list, select the trap for the down state of a product or interface.
By default, all traps known to the Management application 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:
 - Click one of the following buttons:
 - **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.
6. After limiting the list of available traps, expand the MIB folder to which the trap you want belongs under the **Available Traps** list, or right-click to select that trap. Click the right arrow button to move it to the **Selected Traps** list.
7. Select a trap for the up state of the condition.

NOTE

You must select a down and an up trap. You cannot select the same trap for the up and down conditions.

8. Click **Next** to advance to the **Summary** pane.
9. 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

You can create 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 469](#), 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 1103 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 469](#), 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 1103 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 469](#), 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.

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 469](#), 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 1088
- [“Creating a new event action definition by copying an existing definition”](#) on page 1100
- [“Creating an event action with a pseudo event on the escalation policy”](#) on page 1109
- [“Creating an event action with a pseudo event on the resolving policy”](#) on page 1111
- [“Creating an event action with a pseudo event on the flapping policy”](#) on page 1112

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 1088.

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 1088.

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 1088.

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 472](#), displays.

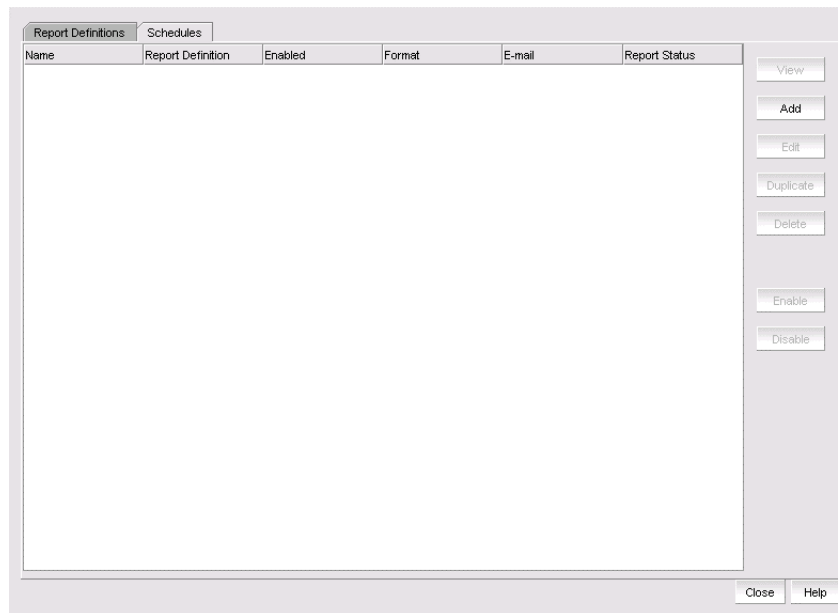


FIGURE 472 Event Custom Reports dialog box - Report Definitions tab

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 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 473](#), displays.
4. Click the **Result Settings** tab.

The **Add/Edit Report Definition** dialog box - **Result Settings** tab displays.

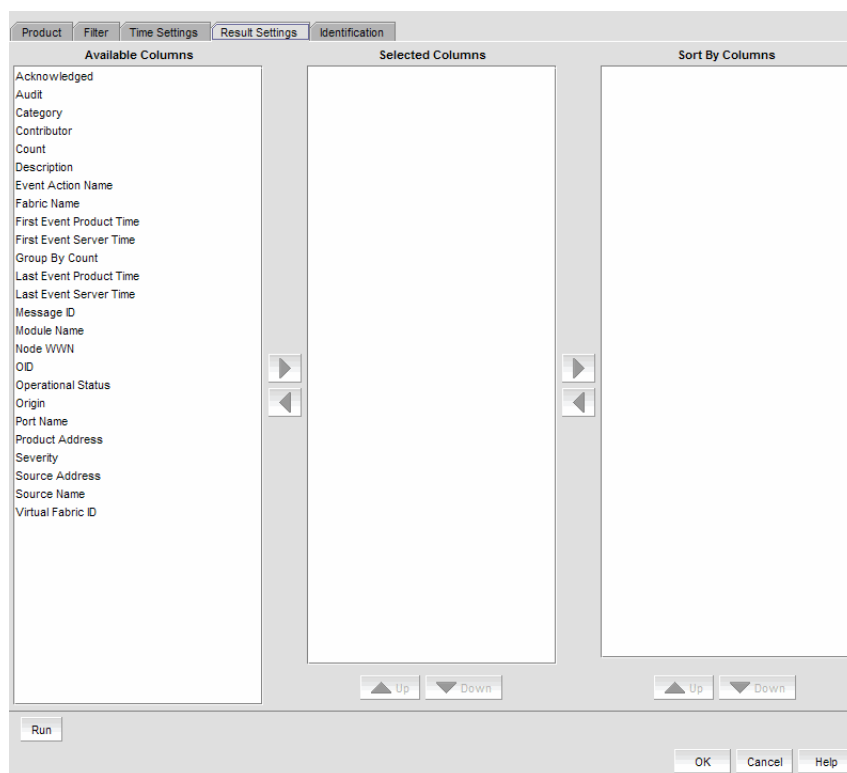


FIGURE 473 Add/Edit Report Definition dialog box - Result Settings tab

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 474](#), displays.

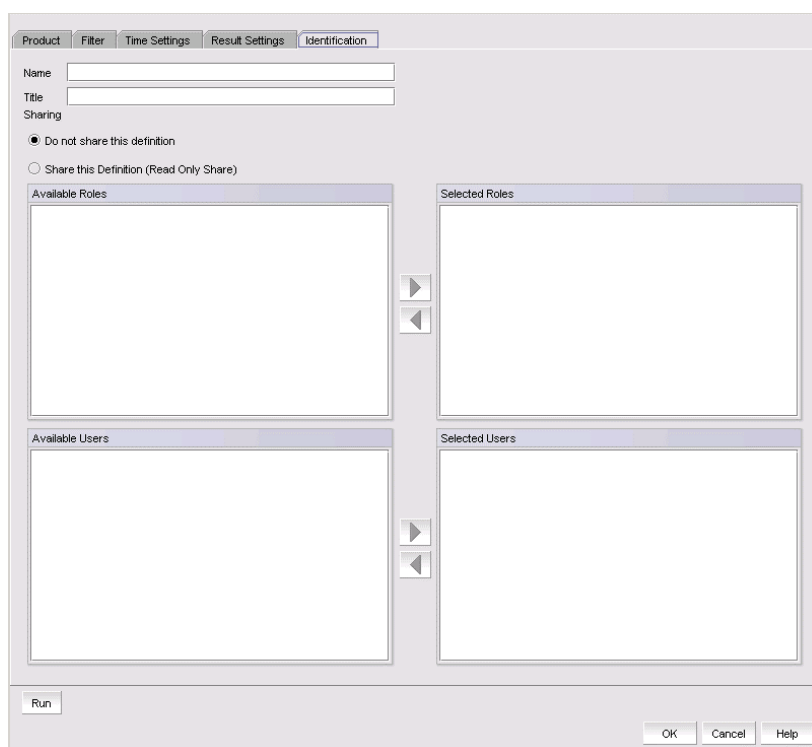


FIGURE 474 Add/Edit Report Definition dialog box - Identification tab

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 1237.

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 475](#), displays.

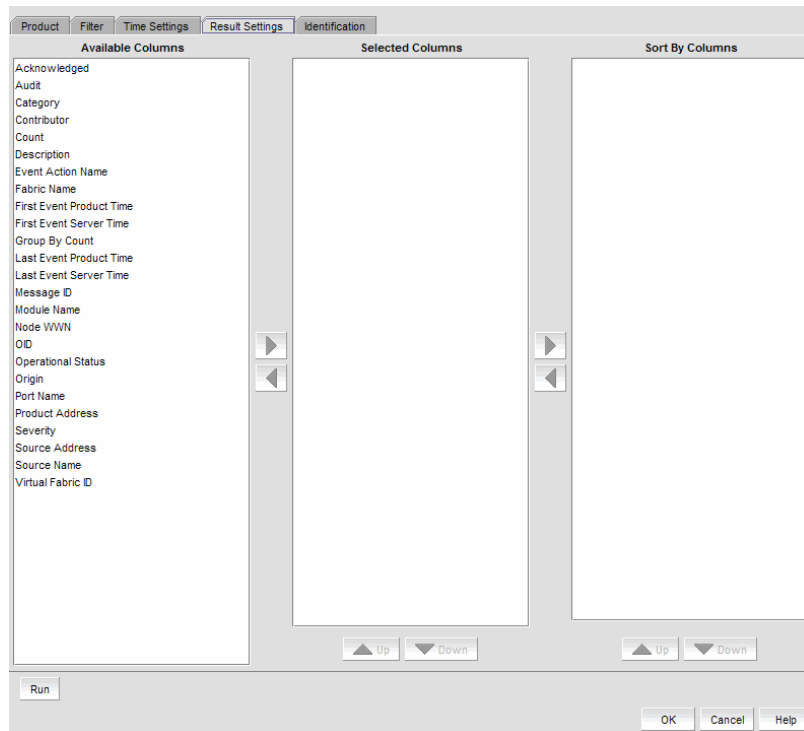


FIGURE 475 Add/Edit Report Definition dialog box - Product tab

4. Click the **Filter** tab.

The **Add/Edit Report Definition** dialog box - **Filter** tab, shown in [Figure 476](#), displays.

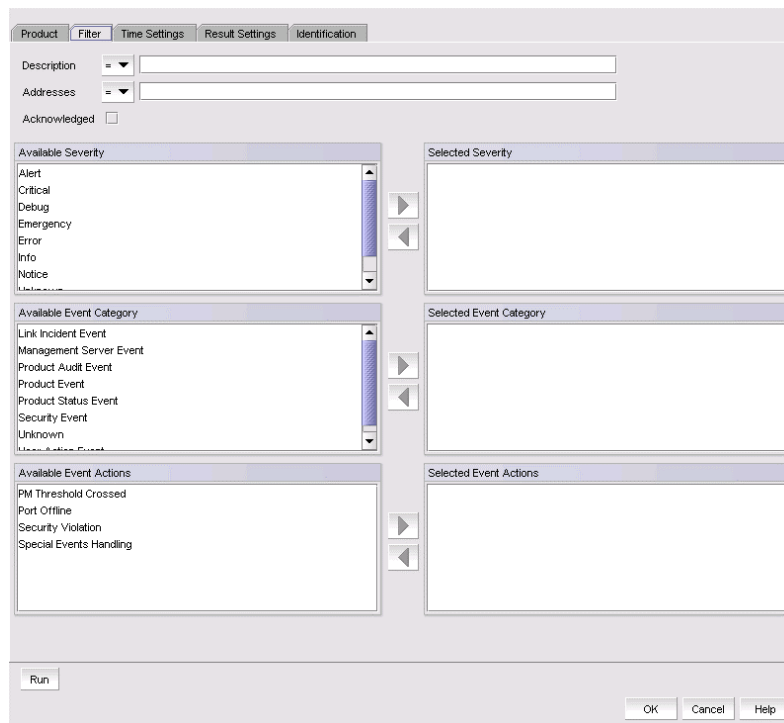


FIGURE 476 Add/Edit Report Definition dialog box - Filter tab

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 1237.

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 477](#), displays.

FIGURE 477 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 1118.)
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 478](#), to display its contents. The **Schedules** list shows the definitions that have been scheduled to automatically run at a specified date and time.




FIGURE 478 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 479](#), 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 479 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 1064.

For information about the Master Log interface, fields, and icons, refer to [“Master Log”](#) on page 255.

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**.

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 1064.

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 1064.

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, shown in [Table 92](#), displays.

3. Review the information.

TABLE 92 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 1237: <ul style="list-style-type: none"> • 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.

TABLE 92 Event Properties (Continued)

Event Field	Description
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.

4. Click **Close** to close the **Event Properties** dialog box.

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 255.

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 480 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.

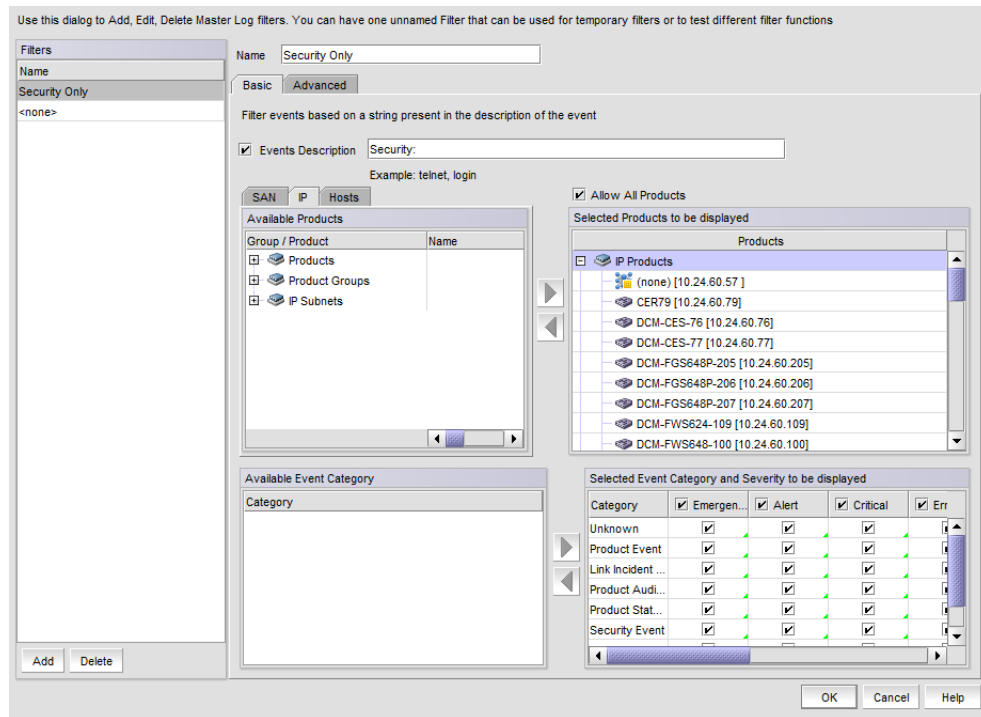


FIGURE 481 Define Filter dialog box - Basic tab, IP tab selected

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.
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 in 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.

Monitoring and Alerting Policy Suite

In this chapter

- [Monitoring and Alerting Policy Suite overview](#) 1135
- [MAPS interoperability with other features](#) 1138
- [MAPS category, object, and measure hierarchy](#) 1143
- [MAPS monitoring categories](#) 1146
- [MAPS policies](#) 1152
- [MAPS rules](#) 1154
- [MAPS conditions](#) 1154
- [MAPS actions](#) 1155
- [MAPS groups](#) 1171
- [MAPS violations](#) 1179
- [MAPS events](#) 1181
- [MAPS integration with other features](#) 1184

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 SAN 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.

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.

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 or 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 7.1 or earlier with the Fabric Watch and Performance Monitor license.

MAPS is supported on Fabric OS devices running 7.2 or later with the Fabric Vision license. MAPS must be enabled on the device (refer to [“Enabling MAPS on a device”](#) on page 1137).

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 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.

In addition to the read actions, MAPS Management read and write permissions enable 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

You can enable MAPS on one or more devices at the same time. Enabling MAPS on a device converts existing Fabric Watch thresholds to MAPS policies and the active thresholds currently monitored by Fabric Watch will continue to be monitored through MAPS.

1. Select **Monitor > Fabric Vision > MAPS > Enable**.

The **Enable MAPS** dialog box displays (Figure 482).

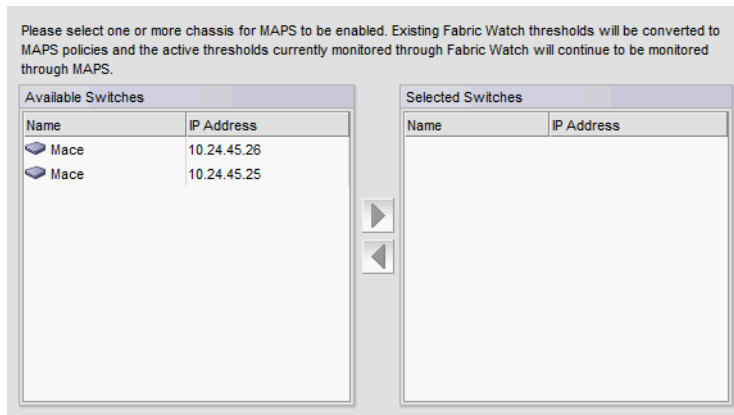


FIGURE 482 Enable MAPS dialog box

Select one or more devices on which you want to enable MAPS in the **Available Switches** list.

2. 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.

3. Click **OK** on the **Enable MAPS** dialog box.

A confirmation message displays.

4. 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.

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 higher 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 optic 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.
- 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 191)
- **Port Health Violations** widget (refer to "[Port Health Violations widget](#)" on page 193)
- **Violations** dialog box (refer to "[Viewing MAPS violations](#)" on page 1179)

[Table 95](#) details the supported RAS log event identifiers that display in the MAPS dashboard widgets and the **Violations** dialog box.

TABLE 93 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	

TABLE 93 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
FRU Health	PS_STATE – Power supply state changes	N/A ²	1440
			1441
			1442
			1443
			1444
	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	
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	

TABLE 93 Fabric Watch supported RAS event IDs (Continued)

Category	Measure	Unit label	RAS event ID		
Switch Resources	TEMP – Temperature sensor	N/A	1002		
			1003		
			1004		
	FLASH_USAGE – Flash usage	%	1402		
	CPU – CPU usage	%	1404		
FCIP	CIR_STATE – FCIP circuit state changes	Count	3020		
			CIR_UTIL – FCIP circuit utilization	%	3012
			CIR_PKTLOSS – FCIP packet loss	%	3016
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		
	CRED_ZERO – Time Tx credit zero	Million	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			

1. Only valid for 16 Gbps, 10 Gbps, and QSFP transceivers.
2. Unit label does not display. You can set the value option from a list (such as, IN_RANGE and OUT_OF_RANGE for temperature measures or FRU states for FRU measures).

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.

An example of a very simple category, object, measure hierarchy follows.

Traffic (category)

FC Port (object)

Receive Bandwidth usage % (measure)

Transmit Bandwidth usage % (measure)

Trunk Utilization (measure)

For a list of all categories, objects, and measures, refer to [“MAPS categories, measures, and actions”](#) on page 1144.

MAPS structural elements

MAPS has the following structural elements: categories, groups, rules, and policies. [Table 94](#) provides a brief description of each structural element.

TABLE 94 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 1155.
Category	A grouping of similar elements that can be monitored (for example “Security Violations”.) For more information, refer to “MAPS monitoring categories” on page 1146.
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 1154.
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 the section “MAPS monitoring categories” on page 1146.
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 1144.
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 1171.
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 1154.
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 1152.

MAPS categories, measures, and actions

Table 95 details the object types for each category, the threshold measures for each object type, and the action you can configure when a threshold is crossed.

TABLE 95 Monitors and actions by category

Category	Objects	Measures	Possible actions
Port Health	FC Port	<ul style="list-style-type: none"> • 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 	<ul style="list-style-type: none"> • RAS Log Event • Fence • SNMP Trap • E-mail
	SFP transceiver	<ul style="list-style-type: none"> • CURRENT — SFP transceiver current • RXP — SFP transceiver receive power • TXP — SFP transceiver transmit power • VOLTAGE — SFP transceiver voltage • SFP_TEMP — SFP transceiver temperature • PWR_HRS¹ — SFP transceiver power on hours 	<ul style="list-style-type: none"> • RAS Log Event • SNMP Trap • E-mail • SFP Marginal
Switch Policy Status	Chassis	<ul style="list-style-type: none"> • BAD_PWR — Absent or faulty power supply • BAD_TEMP — Temperature sensors outside range • BAD_FAN — Absent or faulty fans • FLASH_USAGE² — Flash usage • WWN_DOWN — WWN faulty or down • DOWN_Core — Core blade monitoring • FAULTY_BLADE — Faulty blades • HA_SYNC — HA monitoring 	<ul style="list-style-type: none"> • Status Critical • Status Marginal
	Local Switch	<ul style="list-style-type: none"> • 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 	
Fabric State Changes	Local Switch	<ul style="list-style-type: none"> • 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 	<ul style="list-style-type: none"> • RAS Log Event • SNMP Trap • E-mail
FRU Health	Power Supply	<ul style="list-style-type: none"> • PS_STATE — Power supply state 	<ul style="list-style-type: none"> • RAS Log Event • SNMP Trap • E-mail
	Fan	<ul style="list-style-type: none"> • FAN_STATE — Fan state 	
	Blade	<ul style="list-style-type: none"> • BLADE_STATE — Blade state 	
	SFP transceiver	<ul style="list-style-type: none"> • SFP_STATE — SFP transceiver state 	
	WWN	<ul style="list-style-type: none"> • WWN — World Wide Name state 	

TABLE 95 Monitors and actions by category (Continued)

Category	Objects	Measures	Possible actions
Security Health	Local Switch	<ul style="list-style-type: none"> • 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 	<ul style="list-style-type: none"> • RAS Log Event • SNMP Trap • E-mail
Switch Resources	Temperature sensor	<ul style="list-style-type: none"> • TEMP – Temperature 	<ul style="list-style-type: none"> • RAS Log Event • SNMP Trap • E-mail
	Local Chassis	<ul style="list-style-type: none"> • FLASH_USAGE² – Flash usage • CPU – CPU usage • MEMORY_USAGE – Memory usage 	<ul style="list-style-type: none"> • RAS Log Event • SNMP Trap • E-mail
FCIP	Circuit	<ul style="list-style-type: none"> • CIR_STATE – FCIP circuit state changes • CIR_UTIL – FCIP circuit utilization • CIR_PKTLOSS – FCIP packet loss 	<ul style="list-style-type: none"> • RAS Log Event • Fence (CIR_STATE) • SNMP Trap • E-mail
Traffic / Flows Performance	FC Port	<ul style="list-style-type: none"> • RX – Receive bandwidth usage percentage • TX – Transmit bandwidth usage percentage • UTIL – Trunk utilization 	<ul style="list-style-type: none"> • RAS Log Event • SNMP Trap • E-mail
	Flow	<ul style="list-style-type: none"> • TX_FCNT – Tx Frame Count • RX_FCNT – Rx Frame Count • TX_THPUT – Tx Throughput • CRED_ZERO – Time Tx credit zero • RX_THPUT – Rx Throughput • IO_RD – IO Read Command Count • IO_WR – IO Write Command Count • IO_RD_BYTES – IO Read Data • IO_WR_BY_TES – IO Write Data 	

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.

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.

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.

[Table 96](#) lists measures in the Port category and describes each measure.

TABLE 96 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.
Class 3 timeouts (C3TXTO)	The number of Class 3 discards frames because of timeouts.
State changes (STATE_CHG)	The state of the port has changed for one of the following reasons: <ul style="list-style-type: none"> • The port has gone offline. • The port has come online. • The port is faulty.

TABLE 96 Port measures (Continued)

Measure	Description
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 (μW). 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 (μW). 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.
SFP Power on hours (PWR_HRS)	The number of hours the 16 Gbps SFP transceiver is powered on.

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.

[Table 97](#) 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 97 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).
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.
Faulty Ports ¹ (FAULTY_PORTS)	Hardware-related port faults.

TABLE 97 Switch status measures

Measure	Description
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 98](#) lists measures in the Fabric category and describes each measure.

TABLE 98 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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.

FRU monitoring category

The FRU category enables you to define rules for field replaceable units (FRU), including ports, power supplies, and flash memory.

[Table 99](#) 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 99 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.
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 100](#) lists measures in the Security category and describes what each measure indicates.

TABLE 100 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.

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 101](#) lists measures in the Resource category and describes what each measure indicates.

TABLE 101 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.

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 102](#) lists measures in the FCIP category and describes what each measure indicates.

TABLE 102 FCIP measures

Measure	Description
FCIP circuit state changes (CIR_STATE)	The state of the circuit has changed for one of the following reasons: <ul style="list-style-type: none"> • 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.

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 103](#) lists measures in the Traffic/Flows category and describes each measure.

TABLE 103 Traffic measures

Measure	Description
Receive bandwidth usage percentage (RX)	The percentage of received word frames that exceeds the configured thresholds.
Transmit bandwidth usage percentage (TX)	The percentage of transmitted word frames that exceeds the configured thresholds.
Trunk utilization (UTIL)	The percentage of utilization for the trunk at the time of the last poll.
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.
Time Tx credit zero (CRED_ZERO)	The number of times that the port was unable to transmit frames because the transmit BB credit was zero.
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.

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 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 thresholds values between the aggressive and conservative policies.

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 1164.

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 1161.

Fabric Watch legacy policies

You cannot return 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 1171 for additional information.
- The measure to be monitored.
Refer to “[MAPS monitoring categories](#)” on page 1146 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 1154 for additional information.
- The actions to take if the condition is evaluated to be true.
Refer to “[MAPS actions](#)” on page 1155 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 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 1157. 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.

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 1158. 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** command to send the trap to the management station. You can configure SNMP notifications using the Management application (refer to “[Event notification](#)” on page 1064).

For information on configuring the SNMP agent using the **snmpConfig** command, see 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 104](#).

TABLE 104 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

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. 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. Options include:

- **RAS Log Event** — Use to log a RAS event.
- **SNMP Trap** — Use to send an SNMP trap event.
- **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.

For a complete list of categories and the associated measures and actions, refer to “[MAPS categories, measures, and actions](#)” on page 1144.

4. Click **OK** on the **MAPS Policy Actions** dialog box.
5. Click **Close** on the **MAPS Configuration** 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. 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 **OK** on the **MAPS E-Mail Setup** dialog box.
7. Click **Close** on the **MAPS Configuration** dialog box.

Viewing MAPS policy data

You can view the MAPS-capable devices and the associated MAPS policies and actions.

1. Right-click a device in the Product List or Connectivity Map and select **Fabric Vision > MAPS > Configure**.

The **MAPS Configuration** dialog box displays (Figure 483).

Sort the contents by clicking the column header. Click the same column header again to reverse the sort order.

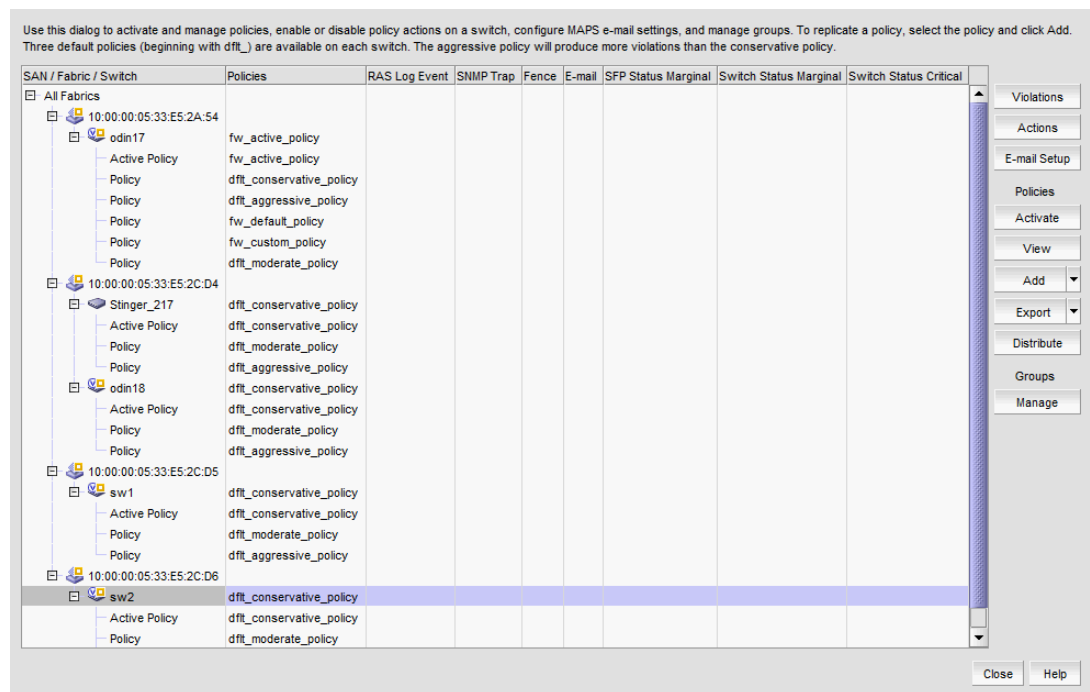


FIGURE 483 MAPS Configuration dialog box

2. Review the MAPS data:
 - MAPS policies list — Lists the MAPS-capable devices and associated policies and actions.
 - **SAN/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.
 - **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.
 - **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.
 - **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 1179.
 - **Actions** button — Select an object (switch, fabric, or all fabrics) and click to enable or disable actions for all policy rules on an object. For more information, refer to [“Enabling or disabling policy actions for all policies”](#) on page 1157.
 - **E-mail Setup** button — Select to configure e-mail notification. For more information, refer to [“Configuring e-mail notification”](#) on page 1158.
 - **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 1166.
 - **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 1168.
 - **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 1161 and [“Cloning a MAPS policy”](#) on page 1164.
 - **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 1164.
 - **Delete** button — Select one or more policies and click to delete. For more information, refer to [“Deleting a MAPS policy”](#) on page 1168.
 - **Export** button — Click to export a policy definition to an XML file. For more information, refer to [“Exporting a MAPS policy”](#) on page 1167.
 - **Import** option (on the **Export** button) — Click to import a policy definition. For more information, refer to [“Importing a MAPS policy”](#) on page 1168.
 - **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 1167.
 - **Manage** button in the **Groups** area — Select the 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 [“Editing multiple groups”](#) on page 1177.
3. Click **Close**.

Configuring a MAPS policy

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 (Figure 484).

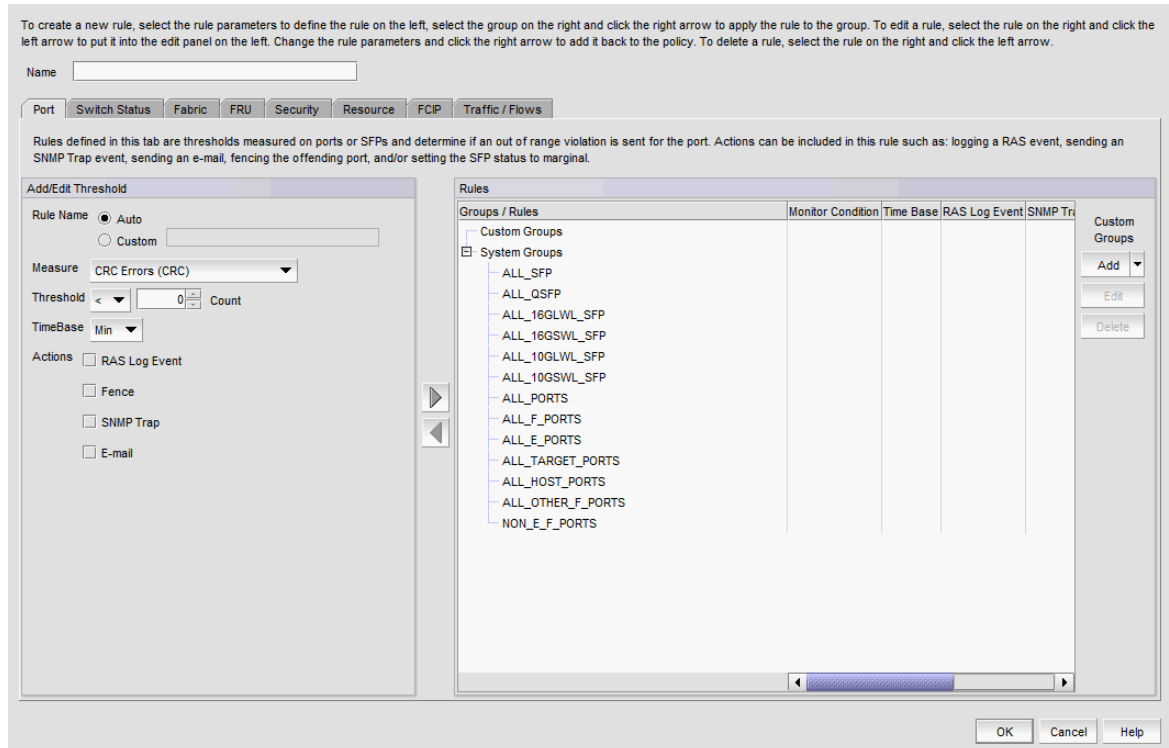


FIGURE 484 Add Policy dialog box

3. Enter a name for the policy in the **Name** field.

The policy name can be up to 32 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 1144. 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.

5. 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 , and hour, and D represents day.

Actions are represented in a bitwise 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), and SFP marginal (S). 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.

- 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 40 characters and can only contain of alphanumeric and underscore characters.

6. 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 1144.

7. 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).

8. 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.
9. 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**.
10. 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**, **Fence**, **SNMP Trap**, **E-mail**, and **SFP Marginal**. For a complete list of categories and the associated measures and actions, refer to [“MAPS categories, measures, and actions”](#) on page 1144.
11. Add the rule to a group 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** and **FCIP** tabs. For more information, refer to [“Configuring a group”](#) on page 1172.

Rules display below the appropriate group node based on rule targets.
12. (**Port** and **FCIP** tabs 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 1172.
13. Click **OK** to add the policy to the **MAPS Configuration** dialog box.
14. Click **Close** on the **MAPS Configuration** dialog box.

Editing a MAPS policy

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 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 1161 and complete [step 5](#) through [step 11](#)).
4. To edit groups, refer to [“Editing a group”](#) on page 1174.
5. Click **OK** on the **Edit Policy** dialog box to update the policy and return to the **MAPS Configuration** dialog box.
6. Click **Close** on the **MAPS Configuration** dialog box.

Cloning a MAPS policy

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 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 32 characters and can only contain of alphanumeric and underscore characters.

4. To create a new policy, refer to [“Configuring a MAPS policy”](#) on page 1161 and complete [step 4](#) through [step 12](#).
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. 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. 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.

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.
- 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.
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 1161 and complete [step 5](#) through [step 11](#)).
 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. 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. 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.

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. 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 **Distribution Options** dialog box displays.
3. Set the destination by choosing one of the following options:
 - **All fabric distribution** – Select to replicate the policy on all MAPS-capable devices in the SAN.
 - **Within fabric distribution** – Select to replicate the policy on all MAPS-capable devices in the selected Fabric.
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 **Distribution Options** 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 the target devices and the source device, if necessary.
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 the 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. 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 select the **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.

Viewing MAPS policy rules

You can open more than one **View Policy** dialog box at the same time.

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 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 485](#)).

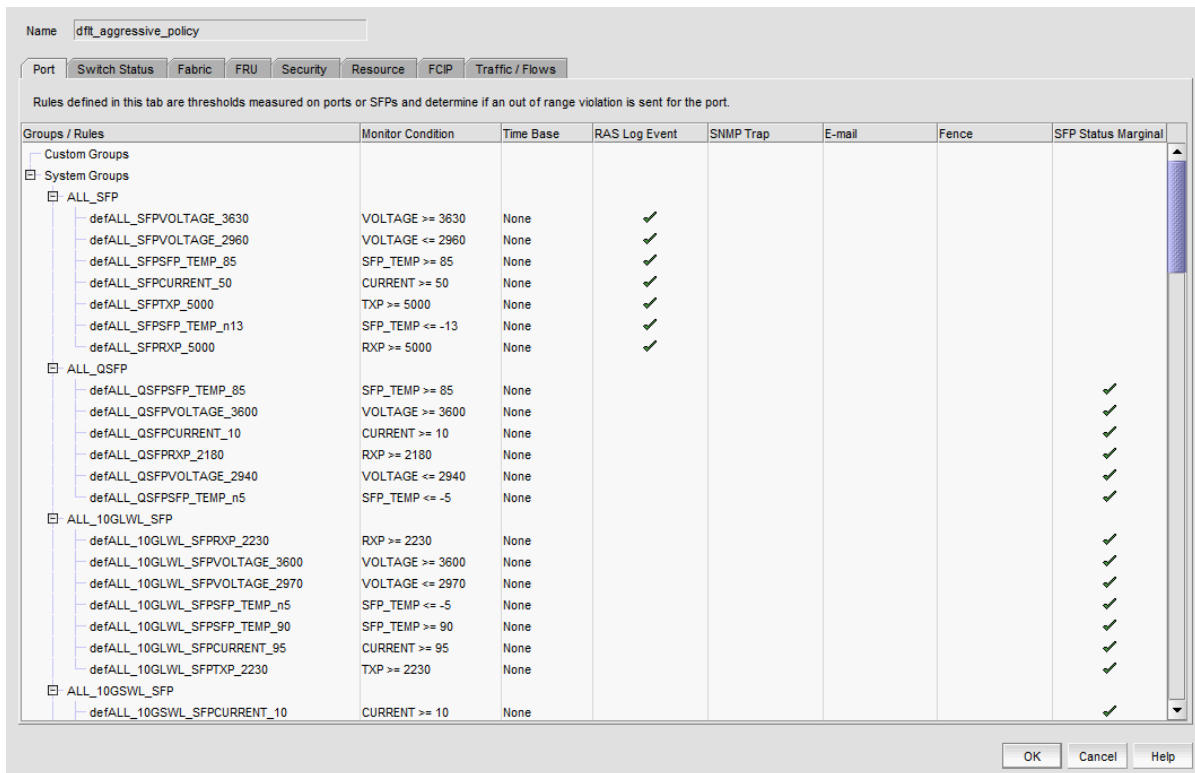


FIGURE 485 View Policy dialog box

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 1144. 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.

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 1172. Rules display below the appropriate group node based on rule targets.
 - **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
 - Fence (Port Fencing)
 - SNMP Trap
 - E-mail
 - **SFP Marginal** (Port SFP status)
4. Click **OK** on the **View Policy** dialog box.
 5. Click **Close** on the **MAPS Configuration** dialog box.

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 1161.

[Table 105](#) lists the preconfigured groups and their descriptions.

TABLE 105 Pre-configured groups

Pre-configured 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_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_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_SLOTS	Slot	All slots present in the chassis.
ALL_SW_BLADES	Blade	All port and application blades in the chassis.
ALL_CORE_BLADES	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	Circuit	All FCIP circuits present in the logical switch.
SWITCH	Switch	Default group used to define rules on global parameters for the entire switch; for example, security violations or fabric health.

TABLE 105 Pre-configured groups

Pre-configured group name	Element type	Description
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.

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, SFPs, or circuits 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 1172.

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:

- **(Port tab)** Create a port group by clicking **Add > Port** in the **Custom Groups** area.
- **(Port tab)** Create a SFP group by clicking **Add > SFP** in the **Custom Groups** area.
- **(FCIP tab)** Create a FCIP circuit group by clicking **Add** in the **Custom Groups** area.

The **Add Group** dialog box displays ([Figure 486](#)).

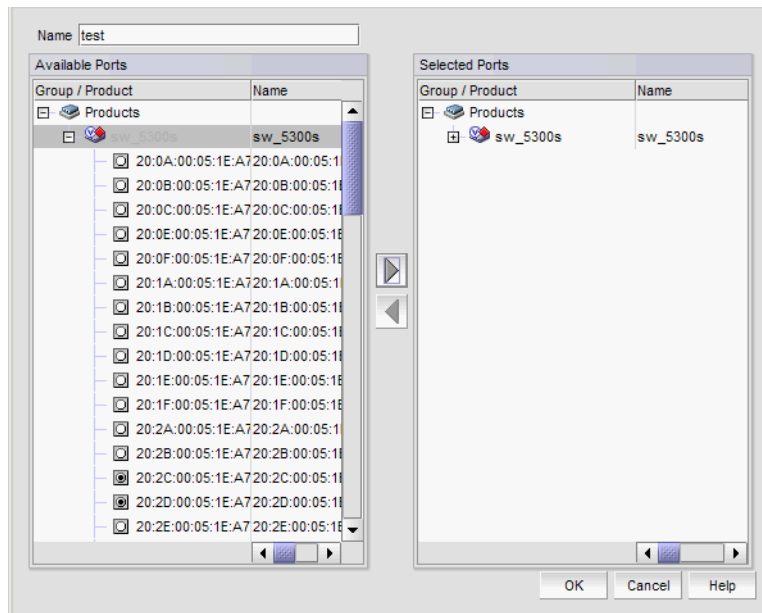


FIGURE 486 Add Group dialog box

4. Enter a unique name for the group in the **Name** field.
The name can be up to 32 characters and can only contain of alphanumeric and underscore characters.
5. Add objects to the group by selecting the object (port 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.
6. Remove objects from the group by selecting the object (port 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.
7. Click **OK** on the **Add Group** dialog box.
The new group displays in the **Custom Groups** folder of the **Rules** area.
8. Configure policies and rules for the group. For more information, refer to [“Configuring a MAPS policy”](#) on page 1161.
9. Click **OK** on the **Add Policy** dialog box.
10. 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. 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** 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 487).

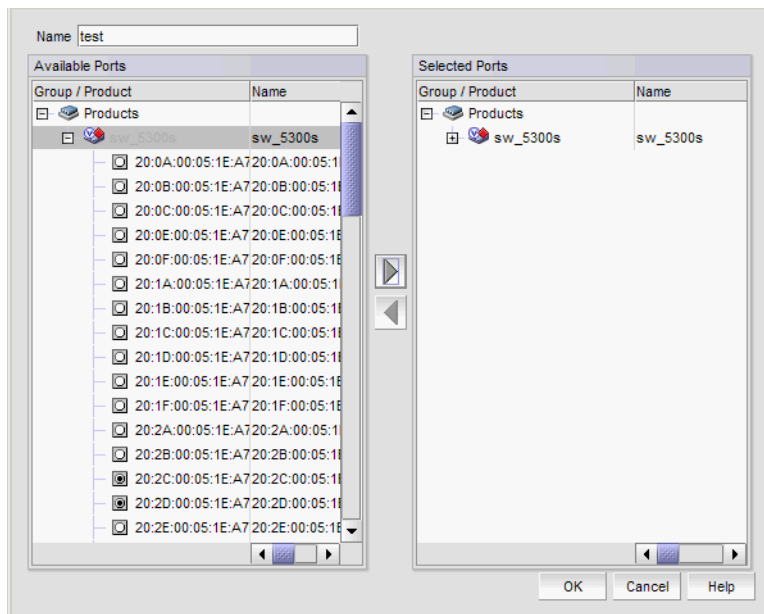


FIGURE 487 Edit Group dialog box

4. Add objects to the group by selecting the object (port 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 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 1161.
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. 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** or **FCIP** 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.

Viewing all groups on a 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 **SAN/Fabric/Switch** list and click **Manage**.
The *Fabric/Device _Name - Manage MAPS Groups* dialog box displays ([Figure 488](#)) with a list of all configured Port, SFP, or FCIP groups on the selected fabric or device in the **Groups** area.

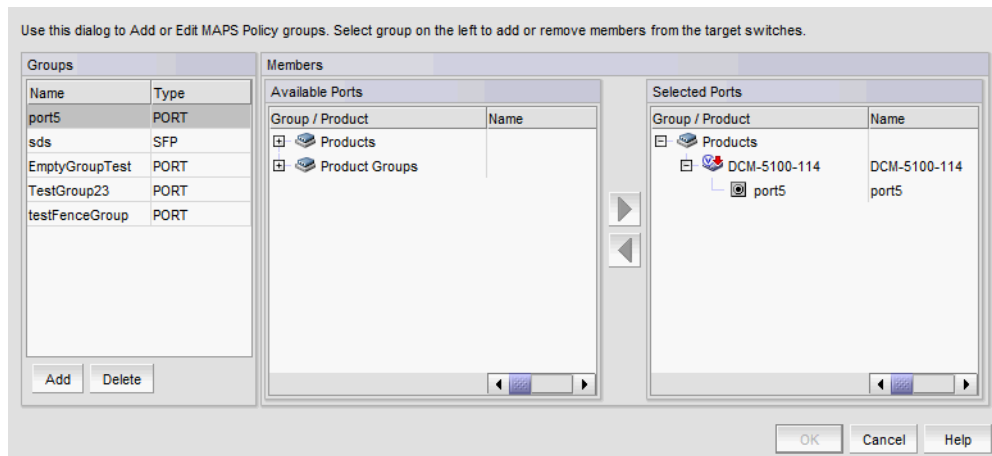


FIGURE 488 *Fabric/Device_Name*- Manage MAPS Groups dialog box

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.

- **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 1177.
- **Delete** button — Click to delete the selected group from the **Groups** list ([“Deleting a group”](#) on page 1178).

4. Click **OK** on the *Fabric/Device_Name* - **Manage MAPS Groups** dialog box.
5. 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 **SAN/Fabric/Switch** list and click **Manage**.

The *Fabric/Device _Name - Manage MAPS Groups* dialog box displays with a list of all configured Port, SFP, or Circuit groups on the selected fabric or device in the **Groups** area.

3. Click **Add**.

The **Add Group** dialog box displays.

4. 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, SFP, and Circuit.

6. Add objects to the group by selecting the object (port, 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 **SAN/Fabric/Switch** list and click **Manage**.

The *Fabric/Device _Name - Manage MAPS Groups* dialog box displays with a list of all configured Port, SFP, or FCIP groups on the selected fabric or device in the **Groups** list.

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 Ports/Circuits** list.

4. Add objects to the group by selecting the object (port, 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 Ports/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 **SAN/Fabric/Switch** list and click **Manage**.

The *Fabric/Device _Name - Manage MAPS Groups* dialog box displays with a list of all configured Port, 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 489).

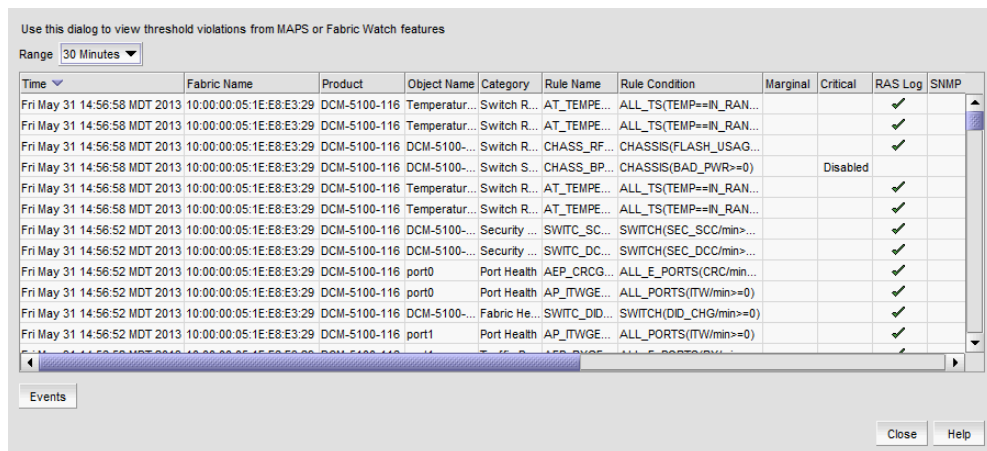


FIGURE 489 Violations dialog box

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. Tooltip displays as “Configured and triggered”.
 - Disabled – The action is configured, but disabled, on the device. Tooltip displays as “Configured in rule, but disabled at switch”.
 - Empty – The action is not configured on the device. Tooltip 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**
 - **Fence** (Port Fencing)
 - **E-mail**
 - **SFP Marginal** (Port SFP status)
 - **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 1181).
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 1063.

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 > Violations**.

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 490).

Severity	Acknowledged	Last Event Server Time	Description
Warning	No	Fri May 31 2013 15:03:52 MDT	Monitoring and Alerting System notification - Rule AP_ITWGE0_M_Rxxxxxx viol
Warning	No	Fri May 31 2013 15:03:52 MDT	Monitoring and Alerting System notification - Rule SWITC_DCCGE0_M_Rxxxxxx
Warning	No	Fri May 31 2013 15:03:52 MDT	Dashboard Category=Port Health, RuleName=AEP_CRCGE0_M_Rxxxxxx, Condi
Warning	No	Fri May 31 2013 15:03:52 MDT	Monitoring and Alerting System notification - Rule SWITC_SCCGE0_M_Rxxxxxx
Warning	No	Fri May 31 2013 15:03:52 MDT	Monitoring and Alerting System notification - Rule SWITC_DIDGE0_M_Rxxxxxx v
Warning	No	Fri May 31 2013 15:03:52 MDT	Monitoring and Alerting System notification - Rule AEP_CRCGE0_M_Rxxxxxx vik
Warning	No	Fri May 31 2013 15:03:52 MDT	Monitoring and Alerting System notification - Rule AEP_CRCGE0_M_Rxxxxxx vik
Warning	No	Fri May 31 2013 15:03:52 MDT	Dashboard Category=Port Health, RuleName=AEP_CRCGE0_M_Rxxxxxx, Condi
Warning	No	Fri May 31 2013 15:03:52 MDT	Monitoring and Alerting System notification - Rule AEP_RXGE0_0_M_Rxxxxxx v
Warning	No	Fri May 31 2013 15:03:52 MDT	Dashboard Category=Traffic Performance, RuleName=AEP_RXGE0_0_M_Rxxxx
Warning	No	Fri May 31 2013 14:56:58 MDT	AT_TEMPEQIN_RAN_N_Rxxxxxx - ALL_TS(TEMP==IN_RANGE)
Warning	No	Fri May 31 2013 14:56:58 MDT	CHASS_BPWRGE0_N_xxxxCxx - CHASSIS(BAD_PWR>=0)
Warning	No	Fri May 31 2013 14:56:52 MDT	SWITC_SCCGE0_M_Rxxxxxx - SWITCH(SEC_SCC/min>=0)
Warning	No	Fri May 31 2013 14:56:52 MDT	AP_ITWGE0_M_Rxxxxxx - ALL_PORTS(TW/min>=0)
Warning	No	Fri May 31 2013 14:56:52 MDT	AEP_CRCGE0_M_Rxxxxxx - ALL_E_PORTS(CRC/min>=0)
Warning	No	Fri May 31 2013 14:55:58 MDT	CHASS_RFUSGGE0_N_Rxxxxxx - CHASSIS(FLASH_USAGE>=0)
Warning	No	Fri May 31 2013 14:55:58 MDT	AT_TEMPEQIN_RAN_N_Rxxxxxx - ALL_TS(TEMP==IN_RANGE)
Warning	No	Fri May 31 2013 14:55:52 MDT	AEP_CRCGE0_M_Rxxxxxx - ALL_E_PORTS(CRC/min>=0)
Warning	No	Fri May 31 2013 14:55:52 MDT	AEP_CRCGE0_M_Rxxxxxx - ALL_E_PORTS(CRC/min>=0)
Warning	No	Fri May 31 2013 14:54:58 MDT	Dashboard Category=Switch Resource, RuleName=CHASS_RFUSGGE0_N_Ro

FIGURE 490 MAPS Violation Master Log Events dialog box

The events display for the selected time range (50% of the events before the selected violation(s) 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 violation(s).

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 violation(s).

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.

Event field	Description
Severity	The MAPS event severity is Warning.
Acknowledged	N/A
Last Event Server Time	The time range selected in the MAPS Violations dialog box.
Description	The rule name and condition.
Origin	N/A
Source Name	The source name of the event.
Source Address	N/A
Category	The event category is MAPS.
Count	The number of times the violation occurred on the device.
Module Name	N/A
Message ID	N/A
Product Address	N/A
Contributor	N/A
Node WWN	N/A
Fabric Name	The fabric name.
Port Name	The port name on which the violation occurred.
Operational Status	N/A
First Event Product Time	N/A
Last Event Product Time	N/A
First Event Server Time	N/A
Audit	N/A
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 190). 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 191).
- Port Health Violations widget — Table view of out of range port health violations (refer to [“Port Health Violations widget”](#) on page 193). 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 1129.

Performance graphs and monitors

You can enable events on historical graphs and monitors. For instructions, refer to [“Configuring the performance graph display”](#) on page 982.

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

In this chapter

- [Server and client support save](#) 1185
- [Device technical support](#) 1189
- [Upload failure data capture](#) 1195

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 343.

Scheduling technical support information collection

You can capture technical support and event information for up to 50 devices. 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 5.3.X 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 126.

NOTE

Fabric OS switches must be running Fabric OS 5.2.X or later to collect technical support data.

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 1243.

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 Support Save data for Fabric OS devices is saved to the following directory:
Install_Home\data\ftproot\technicalsupport

Technical Support Save 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 5.3.X or later, 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 126.

NOTE

Fabric OS switches must be running Fabric OS 5.2.X 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 1243.

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 may 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 and 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.

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 106](#) details the default locations for the technical support save files.

TABLE 106 Technical support save defaults

Type	Default location	Default naming convention
Client SupportSave	<i>User_Home/ServerIP/Managed Product Name/support</i>	DCM-Client-SS- <i>Time_Stamp</i>
Server SupportSave	<i>Install_Home\support</i>	DCM-SS- <i>Time_Stamp</i>
Host (discovered from the SAN tab)	<i>Install_Home\data\ftproot\technicalsupport\host</i>	Supportinfo-HostName-Day-mm-dd-yyyy-hh-mm-ss
SAN Product	<i>Install_Home\data\ftproot\technicalsupport\</i>	Supportinfo-HostName-Day-mm-dd-yyyy-hh-mm-ss
Auto Trace Dump	<i>Install_Home\data\ftproot\tracedump\</i>	

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 1194.
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 1194.
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 1193.
Delete button	Click to delete the support data file. For the procedure, refer to “Deleting technical support files from the repository” on page 1195.

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 1064.

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.

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.

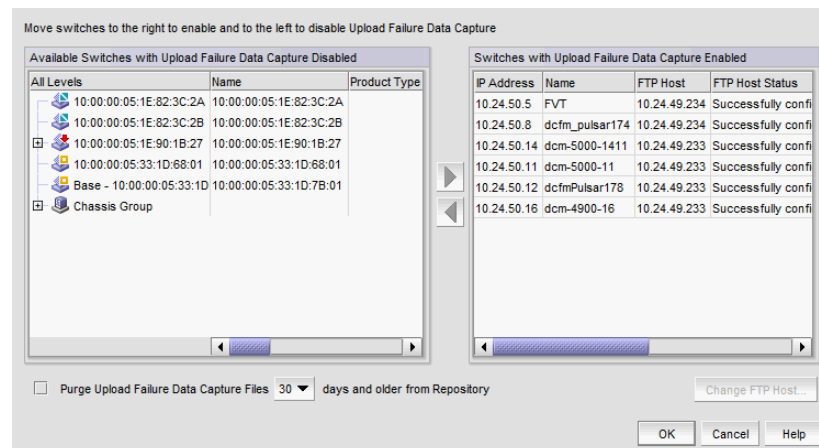


FIGURE 491 Upload Failure Data Capture dialog box

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 1199
- SAN report types 1200
- Generating SAN reports 1200
- Viewing SAN reports 1201
- Exporting SAN reports 1202
- Printing SAN reports 1202
- Deleting SAN reports 1203
- Generating SAN performance reports 1203
- Generating SAN zoning reports 1205

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 1243.

Browser requirements

SAN reports can be printed from a web browser. Reports are supported in the following browsers:

- Windows Internet Explorer 9, 10, or later on Windows
- Firefox 19 or later on Windows or Linux
- Google Chrome

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.

The following device-specific reports are available through the **Monitor (Monitor > Performance > Historical Report)** or **Reports** menu 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 Ports
 - Fabric Summary
3. Select the fabrics 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.









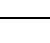
The **View Reports** dialog box displays.

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 1200.

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.

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 1200.
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 1200.

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.
 - i. 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 14 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 1202.

To export the selected report, refer to [“Exporting SAN reports”](#) on page 1202.

To delete the selected report, refer to [“Deleting SAN reports”](#) on page 1203.

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 935.

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 1202.

To export the selected report, refer to [“Exporting SAN reports”](#) on page 1202.

To delete the selected report, refer to [“Deleting SAN reports”](#) on page 1203.

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 755.

Exporting reports to e-mail recipients

You can e-mail a report in a CSV or HTML file format. To export reports to an e-mail recipient, you must configure e-mail event notification (refer to “[Configuring e-mail notification](#)” on page 1064).

If you want to export the report to an e-mail recipient, complete the following steps.

1. Select one of the following from the **E-mail** list:
 - Select **E-mail as HTML**.
 - Select **E-mail as CSV**.

The **Report via E-mail** dialog box displays ([Figure 492](#)).

FIGURE 492 Report via E-Mail dialog box

2. Click the ellipsis button next to the **E-mail Recipients** field.
The **Users** dialog box displays.
3. Select the preconfigured e-mail user account from the list and click **OK**.
4. Enter additional e-mail addresses in the **Other Recipients** field.
5. Enter text in the **Subject** field to change the subject of the e-mail.
6. Enter text in the **Body** field to send a message with the report.
7. Click **Send** to send the report.

NOTE

Mozilla Firefox Browser does not support the window close script. Click the browser close button (X) to cancel.

Application menus

In this appendix

- [Dashboard main menus](#) 1207
- [SAN main menus](#) 1208
- [SAN shortcut menus](#) 1218

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.	

A SAN main menus

Menu	Command	Command Options
	License	Select to view or change your License information.
	About	Select to view the application information, such as the company information 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.

Menu	Command	Command Options
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.
		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.
		<i>MyCustomList</i> — Lists all custom views.

A SAN main menus

Menu	Command	Command Options
	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.	
		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.	

Menu	Command	Command Options
	Host Adapters	— Select to discover hosts.
	VM Managers	— Select to discover VM managers.
	VCEM Managers	— Select to discover Virtual Connect Enterprise 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.
	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.
	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.

A SAN main menus

Menu	Command	Command Options
		Configuration Repository — (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.	
	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.

Menu	Command	Command Options
		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.
	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.

A SAN main menus

Menu	Command	Command Options
Monitor Menu		
	Fabric Vision — Select to configure MAPS or Flow Vision.	
		<p>Flow Vision — Select to define or monitor network traffic by choosing one of the following options:</p> <ul style="list-style-type: none"> • Monitor — Select to monitor network traffic and provides 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.
		<p>MAPS — Select to configure or monitor Monitoring and Alerting Policy Suite policies by choosing one of the following options:</p> <ul style="list-style-type: none"> • 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.	
		Dashboard — Select to launch the Performance Dashboard.
		View Utilization — (Trial and Licensed version Only) Select to display connection utilization.
		View Bottlenecks — (Trial and Licensed version Only) Select to display bottlenecks.
		<p>Historical Data Collection — (Trial and Licensed version Only) Select how to monitor historical data by choosing one of the following options:</p> <ul style="list-style-type: none"> • Enable SAN Wide • Enable Selected • Disable All
		End-to-End Monitors — (Trial and Licensed version Only) Select to monitor -end connections.
		Bottlenecks — Select to monitor bottlenecks.
		Clear Counters — Select to clear all port statistics counters.
		Favorites — Select a custom favorite.

Menu	Command	Command Options
		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.
		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 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 all accept changes 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/logout).
		Fabric — Select to display the events related to the selected fabric.

A SAN main menus

Menu	Command	Command Options
		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 port auto disable flag on individual FC_ports or all ports on a selected device, as well as unblock currently blocked ports.	
	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.	

Menu	Command	Command Options
		<p>FC — Select how to troubleshoot FC by choosing one of the following options:</p> <ul style="list-style-type: none"> • 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.
		<p>FCIP — Select how to troubleshoot FCIP by choosing one of the following options:</p> <ul style="list-style-type: none"> • 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		
	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 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 subsequent 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
	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	

Component	Menu/Submenu Commands	Comments
	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 > Occupied Product Ports UnOccupied Product Ports Attached Ports Switch to Switch Connections	Only available from Product List.
	Table > Copy 'Fabric_Name' 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 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 > Occupied Product Ports UnOccupied Product Ports Attached Ports Switch to Switch Connections	Only available from Product List.

A SAN shortcut menus

Component	Menu/Submenu Commands	Comments
	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 or Expand	Only available from Connectivity Map
Fabric OS Switch/Chassis/Access Gateway		
	Element Manager > Hardware Ports Admin Router Admin Name Server	
	Enable / Disable > Enable 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	

Component	Menu/Submenu Commands	Comments
	Swap Blades	
	Virtual Fabric > Disable Logical Switches Locate Logical Switches > <i>List_of_Logical_Switches</i> (Fabric OS only) (Virtual Fabric-capable switches only)	
	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

A SAN shortcut menus

Component	Menu/Submenu Commands	Comments
	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	
	Show Connections	
	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	
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 > Enable Disable Logical Switches Locate Chassis	Only available from Product List.
	Events	

Component	Menu/Submenu Commands	Comments
	Technical Support > (Fabric OS only) Product/Host SupportSave Upload Failure Data Capture View Repository	
	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	
DCB		
	Element Manager > Hardware Ports Admin Router Admin Name Server	Launches Web Tools.
	Configuration > Save Save Running to Startup Restore Configuration Repository Schedule Backup Replicate > Configuration Security	
	Enable / Disable > Enable Disable	
	Firmware Management	
	Swap Blades	Only available from chassis.
	Zoning	
	DCB	
	FCoE	
	VLAN	

A SAN shortcut menus

Component	Menu/Submenu Commands	Comments
	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.
	<Other Ports > <Fabric Name 1> <Fabric Name 2>	Visible only for AGs that are attached to multiple fabrics.
	Show Ports	
	Accept Changes	
	Show Connections	
	Port Display > Occupied Product Ports UnOccupied Product Ports Attached Ports Switch to Switch Connections	Only available from Product List.
	Properties	

Component	Menu/Submenu Commands	Comments
	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.
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 > Real Time Graphs	Disabled when all ports are offline. 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
	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.

A SAN shortcut menus

Component	Menu/Submenu Commands	Comments
	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.
	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	
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	

Component	Menu/Submenu Commands	Comments
	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	
Switch Port FC		
	Enable / Disable > Enable 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.

A SAN shortcut menus

Component	Menu/Submenu Commands	Comments
	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.
	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.

Component	Menu/Submenu Commands	Comments
	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		
	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	

A SAN shortcut menus

Component	Menu/Submenu Commands	Comments
Giga-Bit Ethernet Port		
	Performance > Real-Time Graph	
	Modify	Launches Element Manager.
	IP Troubleshooting > Ping Trace Route Performance (Trial and Licensed version Only)	
	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	
Connection		
	Properties	
FCIP Tunnel		
	Properties	
Trunk		
	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	
VE Port (physical)		
	Enable / Disable >	
	Enable	
	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	

A SAN shortcut menus

Component	Menu/Submenu Commands	Comments
White Area of the Connectivity Map		
	Accept All Changes	
	Zoom	
	Zoom In	
	Zoom Out	
	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 ' <i>Component</i> '	
	Copy Row	
	Copy Table	
	Export Row	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	
Product List	Table >	
	Copy ' <i>Component</i> '	Some form of this shortcut menu is available for all tables in the Management interface.
	Copy Table	
	Export Table	
	Search	
	Select All	
	Size All Columns To Fit	
	Expand All	
	Collapse All	
	Customize	

Call Home Event Tables

In this appendix

This appendix provides information about the specific events that display when using Call Home. This information is shown in the following Event Tables.

- [# CONSRV Events](#) 1233
- [# Thermal Events](#) 1234
- [Fabric OS Events](#) 1234

TABLE 107 # 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

B Call Home Event Tables

TABLE 108 # 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
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 109 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
1402	FW-1402	Flash usage is out of range (Fabric OS version 6.0 or earlier).	3

TABLE 109 Fabric OS Events (Continued)

Event reason code	FRU code/Event type	Description	Severity
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
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

B Call Home Event Tables

Event Categories

In this appendix

This section provides information about the events that display in each of the following categories:

- [Link incident events](#) 1237
- [Product status events](#) 1237
- [Product audit events](#) 1238
- [Security events](#) 1239
- [User action events](#) 1240
- [Management server events](#) 1240
- [Product events](#) 1241

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]

C Product audit events

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.

Security events for IP devices

For IOS devices, if the event OID starts with any of the following OIDs, then the event is categorized as a security event.

- 1.3.6.1.2.1.14.16.2.6 [ospfIfAuthFailure]
- 1.3.6.1.2.1.14.16.2.7 [ospfVirtIfAuthFailure]
- 1.3.6.1.4.1.1991.0.9 [snOspfIfAuthFailure]
- 1.3.6.1.4.1.1991.0.10 [snOspfVirtIfAuthFailure]
- 1.3.6.1.4.1.1991.0.75 [snTrapUserLogin]
- 1.3.6.1.4.1.1991.0.76 [snTrapUserLogout]
- 1.3.6.1.4.1.1991.0.77 [snTrapPortSecurityViolation]
- 1.3.6.1.4.1.1991.0.78 [snTrapPortSecurityShutdown]
- 1.3.6.1.4.1.1991.0.85 [snTrapMacAuthEnable]
- 1.3.6.1.4.1.1991.0.86 [snTrapMacAuthDisable]
- 1.3.6.1.4.1.1991.0.87 [snTrapMacAuthMACAccepted]
- 1.3.6.1.4.1.1991.0.88 [snTrapMacAuthMACRejected]
- 1.3.6.1.4.1.1991.0.89 [snTrapMacAuthPortDisabled]
- 1.3.6.1.4.1.1991.0.110 [snTrapClientLoginReject]
- 1.3.6.1.4.1.1991.0.131 [snTrapDot1xSecurityViolation]
- 1.3.6.1.4.1.1991.0.143 [snTrapMacAuthRadiusTimeout]
- 1.3.6.1.4.1.1991.0.144 [snTrapDot1xRadiusTimeout]
- 1.3.6.1.4.1.1991.1.5.1.1.2.1.0.36 [swPortSecurityTrap]
- 1.3.6.1.4.1.1991.1.6.1.1.5.2.3 [sysRadiusServerChanged]
- 1.3.6.1.4.1.1991.1.6.1.1.6.2.3 [sysRadiusServerChanged]
- 1.3.6.1.4.1.1991.1.6.1.7.4.2.3 [dot11StationAuthentication]
- 1.3.6.1.4.1.1991.1.6.1.7.4.2.7 [dot1xMacAddrAuthSuccess]
- 1.3.6.1.4.1.1991.1.6.1.7.4.2.8 [dot1xMacAddrAuthFail]
- 1.3.6.1.4.1.1991.1.6.1.7.4.2.9 [dot1xAuthNotInitiated]
- 1.3.6.1.4.1.1991.1.6.1.7.4.2.10 [dot1xAuthSuccess]
- 1.3.6.1.4.1.1991.1.6.1.7.4.2.11 [dot1xAuthFail]
- 1.3.6.1.4.1.1991.1.6.1.7.4.2.12 [localMacAddrAuthSuccess]

C User action events

- 1.3.6.1.4.1.1991.1.6.1.7.4.2.13 [localMacAddrAuthFail]
- 1.3.6.1.4.1.1991.1.6.1.7.4.2.14 [pppLogonFail]
- 1.3.6.1.4.1.1991.1.6.1.7.4.2.18 [dot1xSupplicantAuthenticated]
- 1.3.6.1.4.1.1991.1.7.2.2.2.9 [apAuthFailureTooMany]
- 1.3.6.1.4.1.1991.1.8.2.1.4.0.2 [userLoginNotification]
- 1.3.6.1.4.1.1991.1.8.2.1.4.0.3 [userLogOffNotification]
- 1.3.6.1.4.1.1991.1.8.2.1.4.0.4 [userLoginFailNotification]
- 1.3.6.1.4.1.1991.1.11.1.1.2.2.2.32 [mwlAuthFailure]
- 1.3.6.1.4.1.1991.1.11.1.1.2.2.2.33 [mwlRadiusServerSwitchover]
- 1.3.6.1.4.1.1991.1.11.1.1.2.2.2.34 [mwlRadiusServerSwitchoverFailure]
- 1.3.6.1.4.1.1991.1.11.1.1.2.2.2.35 [mwlRadiusServerRestored]
- 1.3.6.1.4.1.1991.1.11.1.1.2.2.2.36 [mwlAcctRadiusServerSwitchover]
- 1.3.6.1.4.1.1991.1.11.1.1.2.2.2.37 [mwlAcctRadiusServerSwitchoverFailure]
- 1.3.6.1.4.1.1991.1.12.1.1.5.100.0.4 [portSecurityViolation]
- 1.3.6.1.4.1.1991.1.12.1.1.5.109.0.1 [portSECViolation]
- 1.3.6.1.4.1.1991.1.12.1.1.111.1.0.3 [unauthorizedAccessViaCLI]
- 1.3.6.1.6.3.1.1.5.5 [authenticationFailure]

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 110](#), occur when users select the option to forward events to the vCenter during VM Manager discovery.

TABLE 110 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.

C IP Performance monitoring events

User Privileges

In this appendix

- [About user privileges](#) 1243
- [About Roles and Access Levels](#) 1260

About user privileges

The Management application provides the User Administrator with a high level of control over what functions individual users can see and use. This section 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 (AOR). 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.

- [Application privileges and behavior](#) 1244
- [SAN privileges and application behavior](#) 1255

D About user privileges

TABLE 111 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.
Configuration Management	Allows you to access the Configuration Management dialog box and perform configuration upload and replication.	Disables Save , Restore , Configuration Repository , and Schedule Backup under Configure > Switch and the Configuration command under Configure > Switch > Replicate .	Enables Configuration Repository under Configure > Switch . Only viewing of saved configuration is supported. Configuration upload and replication are disabled.	Enables all commands under Configure > Switch . Allows you to perform configuration upload, download and restore.
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 111 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Element Manager - Product Administration	An Element Manager privilege that enables most functionality.	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.

D About user privileges

TABLE 111 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • configure max events • configure event purging policy • apply changes Enables the following functions from the Master Log right-click menu: <ul style="list-style-type: none"> • 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 111 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.

D About user privileges

TABLE 111 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Policy Monitor	Allows you to configure policy monitors.	Disables Policy Monitor command on the Monitor menu.	Enables Policy Monitor command on the Monitor menu. Allows you to open the Policy Monitor 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 Policy Monitor command on the Monitor menu. Allows you to open the Policy Monitor dialog box and enables all controls.
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: <ul style="list-style-type: none"> • 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. <<OK>>	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.

TABLE 111 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
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.
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.

D About user privileges

TABLE 111 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
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.
Virtual Network Management	Allows you to perform VMM based host discovery and management.	Disables the VM Manager Discover menu.	Enables the VM Manager Discover menu. Disables all functions on the dialog box.	Enables the VM Manager 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.

TABLE 111 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Web Services	Allows you to use Web Services API.			
Zoning Activation (Fabric and offline zone database)	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.
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.				

D About user privileges

TABLE 111 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Zoning Online	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.
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.				

TABLE 111 Application privileges and behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
Zoning Offline	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.
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.				

D About user privileges

TABLE 111 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 <i>LSAN_<FabricName></i> as an entry.	Enables the Zoning > LSAN Zoning (Device Sharing) command on the Configure menu. In Zoning dialog box, the Zoning Scope list includes <i>LSAN_<FabricName></i> as an entry, if discovered. If <i>LSAN_<FabricName></i> 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 112 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 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.
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.

D About user privileges

TABLE 112 SAN privileges and application behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
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 - 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.
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.

TABLE 112 SAN privileges and application behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
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.
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.

TABLE 112 SAN privileges and application behavior (Continued)

Privilege	Description	No Privilege	Read-Only	Read/Write
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: <ul style="list-style-type: none"> 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.
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: <ul style="list-style-type: none"> 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 112 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: <ul style="list-style-type: none"> • creating a new encryption group • adding 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 112 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.

About 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 146 for instructions.)

- [Application Features and Role Access Levels](#) 1260
- [SAN Features and Role Access Levels](#) 1262

TABLE 113 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 Administrator, Network Administrator, Host Administrator, Security Administrator	Operator
Configuration Management	SAN System Administrator, Network 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,	

TABLE 113 Application Features and Role Access Levels (Continued)

Feature	Roles with Read/Write Access	Roles with Read-Only Access
Event Management	SAN System Administrator, Network Administrator	Operator
Fabric Watch	SAN System Administrator,	
Fault Management	SAN System AdministratorNetwork Administrator	Operator
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
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 114 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- 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

In this appendix

- [SAN device properties](#) 1264
- [Viewing VC module properties](#) 1276
- [Host properties](#) 1278
- [Properties customization](#) 1280

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 1281.

Fields containing a green triangle (▲) in the lower right corner are editable.

TABLE 115 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 1281.
Edit button	Click to edit a user-defined property. For more information, refer to “Editing a property field” on page 1281.
Delete button	Click to delete a user-defined property. For more information, refer to “Deleting a property field” on page 1282.

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 1281.

Fields containing a green triangle (▲) in the lower right corner are editable.

Depending on the device type, some of the properties listed in the following table may not be available for all products.

TABLE 116 Device properties

Field/Component	Description
Addressing Mode	The addressing mode of the switch.
Back to Edge Routing Supported	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.
Firmware	The firmware version.
GigE Port	The GigE port of the FCIP tunnel.
Host Name	The Host name.

TABLE 116 Device properties (Continued)

Field/Component	Description
IKE Policy #	The IKE policy number. Also includes the following information: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 port of the trunk.
Sequence number	The sequence number of the switch.
Serial #	The hardware serial number.

TABLE 116 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 gigabytes 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 1281.
Edit button	Click to edit a user-defined property. For more information, refer to "Editing a property field" on page 1281.
Delete button	Click to delete a user-defined property. For more information, refer to "Deleting a property field" on page 1282.

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 1272.

4. Click **OK** on the **Properties** dialog box to close.

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.

Field	Description
(Status)	Lists two kinds of data: the LUN’s health and the state of the LUN’s disks. The colored icon in the lower-left corner indicates the LUN’s 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 icons: Normal. All disks are operating normally and online. Transitioning. One or more disks are in a transitioning state. For example, rebuilding or binding. RAID type is 1 in this case. Faulted/Offline. One or more disks is offline or faulted. RAID type is 10 in this case. 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.

Field	Description
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.
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/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: <ul style="list-style-type: none"> • 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.

Field	Description
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 1281.
Edit button	Click to edit a user-defined property. For more information, refer to “Editing a property field” on page 1281.
Delete button	Click to delete a user-defined property. For more information, refer to “Deleting a property field” on page 1282.

4. Click **OK** on the **Properties** dialog box to close.

Viewing iSCSI Properties dialog box

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.

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).

Field	Description
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/Down) of the product.
Name (Product)	The name of the product.
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.

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 on a port and select **Properties**, or double-click the port.

The port **Properties** dialog box displays (Figure 493).

To add user-defined property labels, refer to “Adding a property field” on page 1281.

Fields containing a green triangle (▲) in the lower right corner are editable.

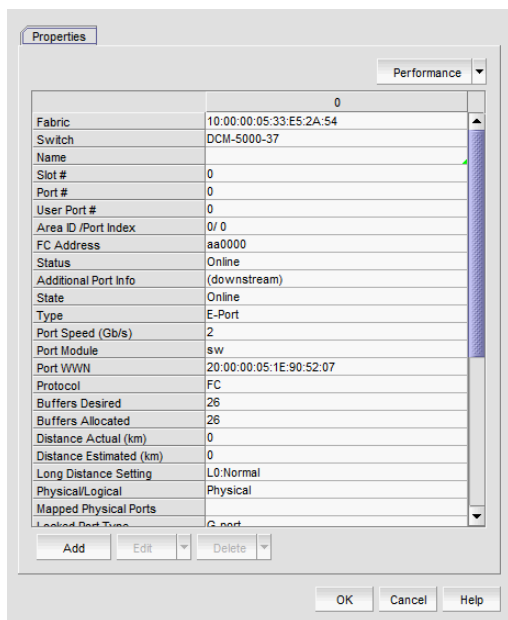


FIGURE 493 Port Properties dialog box

NOTE
Depending on the port type, some of the following properties may not be available for all products.

TABLE 117 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.

TABLE 117 Port properties (Continued)

Field	Description
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.
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.
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.

TABLE 117 Port properties (Continued)

Field	Description
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: <ul style="list-style-type: none"> • 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.
IPSec Policy #	The IPSec policy number. Also includes the following information: <ul style="list-style-type: none"> • 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.

TABLE 117 Port properties (Continued)

Field	Description
Performance list	Select to launch the dialog box of one of the following performance options: <ul style="list-style-type: none"> • 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	Number of NPIV ports.
Port Speed (Gb/s)	The port speed, in Gbits 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.
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 port of the trunk.
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 Gbits 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.

E Viewing VC module properties

TABLE 117 Port properties (Continued)

Field	Description
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: <ul style="list-style-type: none">• 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.
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 1281.
Edit button	Click to edit a user-defined property. For more information, refer to “Editing a property field” on page 1281.
Delete button	Click to delete a user-defined property. For more information, refer to “Deleting a property field” on page 1282.

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 118 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 119 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 120 NPIV WWNs tab

Field	Description
NPIV Port WWN	The world wide name of the NPIF port.
NPIV Node WWN	The world wide name of the NPIF 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.

TABLE 120 NPIV WWNs tab

Field	Description
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 1281).

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 31](#) details the properties of the selected port.

TABLE 31 Adapter port properties

Field	Description
<i>Port Attributes</i>	
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.

TABLE 31 Adapter port properties (Continued)

Field	Description
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.
<i>Configuration</i>	
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.
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.
IOC ID	The IO Controller identifier.
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.

TABLE 31 Adapter port properties (Continued)

Field	Description
Operating State	Displays details about the state of the following operating parameters: <ul style="list-style-type: none"> • Beacon State • Link Beacon State • MPIO Mode State • Path Time Out • Logging Level • Target Rate Limit • Default Rate Limit
<i>FC-SP</i>	
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.
<i>QoS</i>	
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).
Add button	Click to add a user-defined property. For more information, refer to “Adding a property field” on page 1281.
Edit button	Click to edit a user-defined property. For more information, refer to “Editing a property field” on page 1281.
Delete button	Click to delete a user-defined property. For more information, refer to “Deleting a property field” on page 1282.

3. Click **OK** to close.

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

In this appendix

This appendix presents a summary of Unicode regular expression constructs that you can use in the Management application.

- [Characters](#) 1283
- [Character classes](#) 1284
- [Predefined character classes](#) 1284
- [POSIX character classes \(US-ASCII only\)](#) 1284
- [java.lang.Character classes \(simple java character type\)](#) 1285
- [Classes for Unicode blocks and categories](#) 1285
- [Boundary matches](#) 1285
- [Greedy quantifiers](#) 1286
- [Reluctant quantifiers](#) 1286
- [Possessive quantifiers](#) 1286
- [Logical operators](#) 1286
- [Back references](#) 1287
- [Special constructs \(non-capturing\)](#) 1287

TABLE 1 Characters

Construct	Matches
x	The character x
\\	The backslash character
\\On	The character with octal value On (0 <= n <= 7)
\\Onn	The character with octal value Onn (0 <= n <= 7)
\\Omnn	The character with octal value Omnn (0 <= m <= 3, 0 <= n <= 7)
\\xhh	The character with hexadecimal value Oxhh
\\uhhhh	The character with hexadecimal value Oxhhhh
\\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')

TABLE 1 Characters

Construct	Matches
\e	The escape character ('\u001B')
\cx	The control character corresponding to x

TABLE 2 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)
[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 3 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 4 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]

TABLE 4 POSIX character classes (US-ASCII only)

Construct	Matches
<code>\p{Blank}</code>	A space or a tab: [\t]
<code>\p{Cntrl}</code>	A control character: [\x00-\x1F\x7F]
<code>\p{XDigit}</code>	A hexadecimal digit: [0-9a-fA-F]
<code>\p{Space}</code>	A whitespace character: [\t\n\x0B\f\r]

TABLE 5 java.lang.Character classes (simple java character type)

Construct	Matches
<code>\p{javaLowerCase}</code>	Equivalent to java.lang.Character.isLowerCase()
<code>\p{javaUpperCase}</code>	Equivalent to java.lang.Character.isUpperCase()
<code>\p{javaWhitespace}</code>	Equivalent to java.lang.Character.isWhitespace()
<code>\p{javaMirrored}</code>	Equivalent to java.lang.Character.isMirrored()

TABLE 6 Classes for Unicode blocks and categories

Construct	Matches
<code>\p{InGreek}</code>	A character in the Greek block (simple block)
<code>\p{Lu}</code>	An uppercase letter (simple category)
<code>\p{Sc}</code>	A currency symbol
<code>\P{InGreek}</code>	Any character except one in the Greek block (negation)
<code>[\p{L}&&[^\p{Lu}]]</code>	Any letter except an uppercase letter (subtraction)

TABLE 7 Boundary matches

Construct	Matches
<code>^</code>	The beginning of a line
<code>\$</code>	The end of a line
<code>\b</code>	A word boundary
<code>\B</code>	A non-word boundary
<code>\A</code>	The beginning of the input
<code>\G</code>	The end of the previous match
<code>\Z</code>	The end of the input but for the final terminator, if any
<code>\z</code>	The end of the input

TABLE 8 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 9 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 10 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 11 Logical operators

Construct	Matches
XY	X followed by Y
X Y	Either X or Y
(X)	X, as a capturing group

TABLE 12 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 13 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

F Regular Expressions

Troubleshooting

In this chapter

- Application Configuration Wizard troubleshooting 1290
- Browser troubleshooting 1290
- Client browser troubleshooting 1291
- Fabric tracking troubleshooting 1291
- FICON troubleshooting 1292
- Firmware download troubleshooting 1292
- Launch Client troubleshooting 1294
- Names troubleshooting 1296
- Patch troubleshooting 1296
- Performance troubleshooting 1297
- Port Fencing troubleshooting 1301
- Professional edition login troubleshooting 1301
- Server troubleshooting 1301
- Server Management Console troubleshooting 1302
- Supportsave troubleshooting 1303
- Technical support data collection troubleshooting 1304
- View All list troubleshooting 1304
- Zoning troubleshooting 1305

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	<p>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:</p> <p>The following are the various ways we can disable UAC in Vista:</p> <p>Disable using msconfig by completing the following steps.</p> <ol style="list-style-type: none"> 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. <p>A command window displays and runs the disable UAC command. When the command is complete, close the window.</p> <ol style="list-style-type: none"> 6 Close the System Configuration Utility. 7 Restart the computer to apply changes. <p>NOTE: You can re-enable UAC using the above procedure and selecting the Enable UAC tool name in step 4.</p> <p>Disable using regedit by completing the following steps.</p> <p>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.</p> <ol style="list-style-type: none"> 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. <p>NOTE: You can re-enable UAC using the above procedure and changing the Value data field to 1 in step 5.</p>

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.	<p>Mozilla Firefox Browser does not support window close script.</p> <p>Click the browser Close button to cancel.</p> <p>NOTE: The Cancel button still displays on all Report via E-mail dialog boxes.</p>

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	<p>If the JNLP file does not launch automatically, use one of the following options:</p> <ul style="list-style-type: none"> • Complete the following steps. <ol style="list-style-type: none"> 1 Save the JNLP file to the local host. 2 Launch the JNLP file manually. • In Internet Explorer 7, complete the following steps. <ol style="list-style-type: none"> 1 Select Tools > Internet Options. 2 Click the Advanced tab. 3 Clear the Do not save encrypted pages to disk check box. <p>If the browser warns you about the security certificate, use the fully qualified hostname to launch the web page.</p>

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.	<p>Choose from one of the following options:</p> <ul style="list-style-type: none"> • 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.	<p>FICON Unsupported Configurations:</p> <ul style="list-style-type: none"> • 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.	<p>Choose from one of the following options:</p> <ul style="list-style-type: none"> • 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.

Problem	Resolution
<p>Firmware download using SCP/SFTP does not work because of one of the following issues:</p> <ul style="list-style-type: none"> • For internal SCP/SFTP server, the application was uninstalled and reinstalled without migration • For external SCP/SFTP server, the SSH handshake keypair is changed <ul style="list-style-type: none"> - 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 	<p>Clear the SSH (SCP/SFTP) server IP address or hostname from the known_hosts table of the device.</p> <ul style="list-style-type: none"> • For Fabric OS devices, use the following command: <pre>sw0:FID128:admin> sshutil delknownhost</pre> IP Address/Hostname to be deleted: <i>SSH_server_IP_address</i> where <i>SSH_server_IP_address</i> 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: <pre>sw0# clear ssh-key SSH_server_IP_address</pre> where <i>SSH_server_IP_address</i> 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: <pre>sw0# execute-script sshdeleteknownhost</pre> IP Address/Hostname to be deleted: <i>SSH_server_IP_address</i> where <i>SSH_server_IP_address</i> is the IP address of the SSH server you want to delete.
<p>Firmware download using an external SCP server does not work after reinstalling the application or the external SCP server on the Host machine.</p>	<p>Clear the SCP server IP address or hostname from the known_hosts table of the device using the following command: <pre>sw0# FID10:root> ssh-keygen -R Host_Name</pre> where <i>Host_Name</i> is the IP address or host name of the external SCP server.</p>

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.	<p>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.</p> <ol style="list-style-type: none"> 1 Clear the previous version from the Java cache. <ol style="list-style-type: none"> 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. <ol style="list-style-type: none"> 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, <i>IP_Address:Port_Number</i>. 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 <i>Domain\User_Name</i> 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
<p>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).</p>	<p>Use one the following procedures to configure the IP address in the host file.</p> <p>Windows operating systems</p> <ol style="list-style-type: none"> 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: <i>IP_Address Host_Name</i>. For example, 127.0.0.1 localhost 6 Save and exit the file. <p>Unix operating systems</p> <ol style="list-style-type: none"> 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: <i>IP_Address Host_Name</i>. For example, 127.0.0.1 localhost 4 Save and exit the file.
<p>Unable to launch the remote client (the SSL setting, web server port number, or server starting point number changed during the server upgrade).</p>	<p>To remove the old link and launch the correct remote client version, complete the following steps.</p> <ol style="list-style-type: none"> 1 Clear the previous version from the Java cache,. <ol style="list-style-type: none"> 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. <ol style="list-style-type: none"> 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, <i>IP_Address:Web_Server_Port_Number</i>. 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 <i>Domain\User_Name</i> 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.	<p>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:</p> <ul style="list-style-type: none"> • “Searching for a device by name” on page 99 in the Configure Names dialog box • “Searching for a device by WWN” on page 100 in the Configure Names dialog box • “Searching for a device” on page 269

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	<p>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:</p> <p>The following are the various ways we can disable UAC in Vista:</p> <p>Disable using msconfig by completing the following steps.</p> <ol style="list-style-type: none"> 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. <p>A command window displays and runs the disable UAC command. When the command is complete, close the window.</p> <ol style="list-style-type: none"> 6 Close the System Configuration Utility. 7 Restart the computer to apply changes. <p>NOTE: You can re-enable UAC using the above procedure and selecting the Enable UAC tool name in step 4.</p> <p>Disable using regedit by completing the following steps.</p> <p>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.</p> <ol style="list-style-type: none"> 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. <p>NOTE: You can re-enable UAC using the above procedure and changing the Value data field to 1 in step 5.</p>

Performance troubleshooting

The following section states a possible issue and the recommended solution for Performance errors.

Problem	Resolution
<p>An error message with the following text displays: Real Time statistics collection has failed. Please see master log for details.</p>	<p>Make sure that the following prerequisites for Performance Monitoring Data collection are met.</p> <ol style="list-style-type: none"> <li data-bbox="618 516 1466 1386"> <p>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.</p> <p>For example, data collection occurs in the following cases.</p> <p>Case 1: Default access control list is empty</p> <pre>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</pre> <p>Case 2: Management Server IP included in access control list</p> <pre>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</pre> <p>Verification and Troubleshooting.</p> <p>To add the server IP address to the access control list, use the following command from the switch CLI:</p> <pre>FCRRouter:admin> snmpconfig --set accesscontrol</pre> <p>To set the default access control, use the following command from the switch CLI:</p> <pre>FCRRouter:admin> snmpconfig --default accesscontrol</pre>

Problem	Resolution
<p>An error message with the following text displays: Real Time statistics collection has failed. Please see master log for details.</p>	<p>2 To collect data, the SNMP credentials in the Management application and switch must match.</p> <p>SNMP v1 or v3: The community strings entered in the Address Properties dialog box - SNMP tab must match the one entered in the switch.</p> <p>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.</p> <p>To view the switch SNMP value, use one of the following commands from the switch CLI:</p> <pre>HCLSwitch:admin> snmpconfig --show snmpv1 HCLSwitch:admin> snmpconfig --show snmpv3</pre> <p>To set the switch SNMP value, use one of the following commands from the switch CLI:</p> <pre>HCLSwitch:admin> snmpconfig --set snmpv1 HCLSwitch:admin> snmpconfig --set snmpv3</pre> <p>Example</p> <pre>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]</pre>

Problem	Resolution
<p>An error message with the following text displays: Real Time statistics collection has failed. Please see master log for details.</p>	<p>3 To collect GigE port and FCIP statistics, you must enable the FCIP-MIB capability.</p> <p>Verification and Troubleshooting</p> <p>To verify that FCIP-MIB capability is enabled, use the following command from the switch CLI:</p> <pre>FCRRouter:admin> snmpconfig --show mibcapability FCIP-MIB: YES</pre> <p>To enabling FCIP-MIB capability, use the following command from the switch CLI:</p> <pre>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]</pre> <p>4 To collect FCIP or GE statistics, you must configure SNMPv3 credentials in the Address Properties dialog box - SNMP tab.</p> <p>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:</p> <pre>sw1:FID128:admin> snmpconfig --show snmpv3</pre> <p>SNMPv3 USM configuration:</p> <pre>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</pre>

Problem	Resolution
<p>An error message with the following text displays: Real Time statistics collection has failed. Please see master log for details.</p>	<p>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.</p> <p>To verify the Fabric OS user (verify Role-LF List), use the following command from the switch CLI:</p> <pre>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</pre> <p>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</p> <p>FCIP tunnels: command: portshow fcipunnel <Ge port number> <tunnel no> -perf Example Sprint-65:root> portshow fcipunnel ge0 1 -perf</p>
<p>An error message with the following text displays: Real Time statistics collection has failed. Please see master log for details.</p>	<p>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:</p> <pre>snmpconfig --show secLevel</pre> <p>Example</p> <pre>snmpconfig --show secLevel GET security level = 0, SET level = 0 SNMP GET Security Level: No security SNMP SET Security Level: No security</pre> <p>To set the security level respectively, use the following command from the switch CLI:</p> <pre>snmpconfig --set secLevel</pre> <p>Example</p> <pre>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]</pre>

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 14 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.

Problem	Resolution
Management server exits unexpectedly on Red hat Linux 6.1	<p>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.</p> <p>Therefore, if M = Amount of RAM in GB and S = Amount of swap in GB, then</p> <p>If $M < 2$</p> $S = M * 2$ <p>Else</p> $S = M + 2$

Server Management Console troubleshooting

The following section states a possible issue and the recommended solution for server management console errors.

Problem	Resolution
Unable to launch the SMC on a Windows Vista, Windows 7, or Windows 2008 R2 system	<p>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:</p> <p>The following are the various ways we can disable UAC in Vista:</p> <p>Disable using msconfig by completing the following steps.</p> <ol style="list-style-type: none"> 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. <p>A command window displays and runs the disable UAC command. When the command is complete, close the window.</p> <ol style="list-style-type: none"> 6 Close the System Configuration Utility. 7 Restart the computer to apply changes. <p>NOTE: You can re-enable UAC using the above procedure and selecting the Enable UAC tool name in step 4.</p> <p>Disable using regedit by completing the following steps.</p> <p>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.</p> <ol style="list-style-type: none"> 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. <p>NOTE: You can re-enable UAC using the above procedure and changing the Value data field to 1 in step 5.</p>

Problem	Resolution
Unable to launch the SMC on a Windows Vista or Windows 7 system continued	<p>Disable using the Group Policy by completing the following steps.</p> <p>You can perform this procedure on your 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.</p> <p>To disable using the Local Group Policy editor, complete the following steps.</p> <ol style="list-style-type: none"> 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. <p>To disable using the Active Directory-based GPO editor, complete the following steps.</p> <ol style="list-style-type: none"> 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.	<p>Capture support show by running the batch file from the <i>Install_Home/bin/supportshow.bat</i> from Windows and UNIX systems.</p> <ol style="list-style-type: none"> 1 <code>Open Install_Home\bin\supportsave.bat.</code> 2 <code>Edit file supportsave dbuser dbpasswd [tareget-dir] [pause-option].</code> <p>NOTE: Unreachable switches increase the time needed to collect supportSave data.</p>

Technical support data collection troubleshooting

The following section states a possible issue and the recommended solution for technical support data collection errors.

Problem	Resolution
<p>Technical support data collection using SCP/SFTP does not work because of one of the following issues:</p> <ul style="list-style-type: none"> For internal SCP/SFTP server, the application was uninstalled and reinstalled without migration For external SCP/SFTP server, the SSH handshake keypair is changed <ul style="list-style-type: none"> 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 	<p>Clear the SSH (SCP/SFTP) server IP address or hostname from the known_hosts table of the device.</p> <ul style="list-style-type: none"> For Fabric OS devices, use the following command: <code>sw0:FID128:admin> sshutil delknownhost</code> IP Address/Hostname to be deleted: <i>SSH_server_IP_address</i> where <i>SSH_server_IP_address</i> 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: <code>sw0# clear ssh-key <i>SSH_server_IP_address</i></code> where <i>SSH_server_IP_address</i> 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: <code>sw0# execute-script sshdeleteknownhost</code> IP Address/Hostname to be deleted: <i>SSH_server_IP_address</i> where <i>SSH_server_IP_address</i> is the IP address of the SSH server you want to delete.
<p>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.</p>	<p>Clear the SCP server IP address or hostname from the known_hosts table of the device using the following command: <code>sw0# FID10:root> ssh-keygen -R <i>Host_Name</i></code> where <i>Host_Name</i> is the IP address or host name of the external SCP server.</p>

View All list troubleshooting

The following section states a possible issue and the recommended solution for **View All** list errors.

Problem	Resolution
<p>View All list does not display.</p>	<p>The View All list does not display until you discover a fabric. To discover a fabric, refer to "Discovering fabrics" on page 39.</p>
<p>View All list does not display and there are discovered fabrics.</p> <p>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.</p>	<p>To select another view, select View > Manage View > Display View > <i>View_Name</i>.</p>

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 118.

G Zoning troubleshooting

Database Fields

In this appendix

- [Database tables and fields](#) 1307
- [Views](#) 1490

Database tables and fields

NOTE

The primary keys are marked by an asterisk (*)

TABLE 15 **ACH_CALL_CENTER**

Field	Definition	Format	Size
ID *	Unique generated database identifier.	int	
NAME	Name of the Call Center.	varchar	256

TABLE 16 **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 17 **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 17 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 18 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 19 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 20 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 21 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_NETW ORK_ADDRESS	The LDAP Server Network Address from which the Active directory group is fetched	varchar	255

TABLE 22 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 23 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 24 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 25 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 26 AOR_FABRIC_MAP

Field	Definition	Format	Size
AOR_ID	ID of AOR	int	
FABRIC_ID	FABRIC ID which is in the AOR	int	

TABLE 27 AOR_HOST_MAP

Field	Definition	Format	Size
AOR_ID	ID of AOR	int	
HOST_ID	HOST ID which is in the AOR	int	

TABLE 28 AOR_INM_PORT_GROUP_MAP

Field	Definition	Format	Size
AOR_ID	ID of AOR	int	
PORT_GROUP_ID	IP of port group	int	

TABLE 29 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 30 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 31 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. <ul style="list-style-type: none"> • 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

TABLE 31 BIRTREPORT_PARAMETER (Continued)

Field	Definition	Format	Size
DATA_TYPE	Data type of the parameter. Possible values are: <ul style="list-style-type: none"> • 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 32 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	
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 33 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. <ul style="list-style-type: none"> • 0 indicates output will be in HTML • 1 indicates PDF • 2 indicates CSV 	int	

TABLE 34 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	
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 35 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 36 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 37 CAPABILITY_

Field	Definition	Format	Size
NAME *	Name of the capability.	varchar	256
DESCRIPTION	Optional detailed description about the capability.	varchar	512

TABLE 38 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

TABLE 38 CARD (Continued)

Field	Definition	Format	Size
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 management 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. <ul style="list-style-type: none"> • 1 = true. • 0 = false. • Default value is 0. 	smallint	
FCIP_LICENSED	FCIP Advanced Extension Licensing is available. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • -1 = not supported. • Default value is -1. 	int	
MAX_FCIP_CIRCUITS	Describes the maximum number of circuits that can be created in this slot. <ul style="list-style-type: none"> • -1 = not supported. • Default value is -1. 	int	
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	

TABLE 38 CARD (Continued)

Field	Definition	Format	Size
HEADER_VERSION	The OEM or vendor-assigned version number.	int	
GIGE_MODE	Determines the port operating mode for GE ports. <ul style="list-style-type: none"> • 0 - 1G • 1 - 10G • 2 - Dual mode • 3 - Failover mode Default value -1 means it is not applicable.	smallint	

TABLE 39 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 40 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 41 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 42 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 43 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	

TABLE 43 CEE_PORT (Continued)

Field	Definition	Format	Size
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	This field 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 44 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 45 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

TABLE 45 CLIENT_VIEW_COLUMN (Continued)

Field	Definition	Format	Size
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 46 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 47 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 48 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 49 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 50 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	
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 Index	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 51 CNA_PRODUCT_CONNECTIVITY

Field	Definition	Format	Size
CNA_PORT_ID	CNA Port identifier.	int	
INTERFACE_ID	Interface Identifier.	int	

TABLE 52 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

TABLE 52 CNA_ETH_PORT_CONFIG (Continued)

Field	Definition	Format	Size
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 53 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

TABLE 54 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 55 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 56 COLLECTOR_SNMP_EXPRESSION_ENTRY

Field	Definition	Format	Size
COLLECTOR_SNMP_EXPRESSION_ENTRY_ID	Primary key autogenerated ID.	int	
COLLECTOR_ID	ID of the PERF_COLLECTOR.	int	
EXPRESSION_ID	Id of the SNMP_EXPRESSION.	int	

TABLE 57 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 deviceld, and for port level it will use interfaceld.	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 58 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	

TABLE 58 CORE_SWITCH (Continued)

Field	Definition	Format	Size
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
CREATION_TIME	Core switch record creation time. This tells us when the initial 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	Determines whether the switch is registered for sending syslog traps. <ul style="list-style-type: none"> • 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 . <ul style="list-style-type: none"> • 1 is registered • 0 is not registered. 	smallint	
USER_IP_ADDRESS	Only for 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.	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 recipients.	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. <ul style="list-style-type: none"> • 1 is enabled • 0 is disabled 	smallint	
VF_SUPPORTED	Determines whether virtual fabric is supported on the switch. <ul style="list-style-type: none"> • 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 58 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 59 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 60 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 61 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 62 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

TABLE 62 CORE_SWITCH_DETAILS (Continued)

Field	Definition	Format	Size
FC_MASK	FC IP Address ethernet mask.	char	64
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. <ul style="list-style-type: none"> • 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 62 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 63 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 64 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. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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 64 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 64 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_PROVISIONING_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 65 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	

TABLE 65 CRYPTO_SWITCH (Continued)

Field	Definition	Format	Size
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	
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 66 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

TABLE 66 CRYPTO_TARGET_CONTAINER (Continued)

Field	Definition	Format	Size
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	
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 67 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 68 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	

TABLE 68 DASHBOARD (Continued)

Field	Definition	Format	Size
CREATION_TIME	Time when dashboard was created.	timestamp	
LAST_OPENED_TIME	Time when dashboard was last opened.	timestamp	

TABLE 69 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 70 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 71 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	

TABLE 71 DASHBOARD_PROVIDER

Field	Definition	Format	Size
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 72 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	

TABLE 72 DASHBOARD_WIDGET (Continued)

Field	Definition	Format	Size
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	

TABLE 73 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	

TABLE 74 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 74 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 75 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	

TABLE 76 DEPLOYMENT_CONFIGURATION

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
NAME	Name of the configuration	varchar	255
CONFIGURATION_TYPE	Identifies the save configuration type. <ul style="list-style-type: none"> • 1 - Not applicable • 1 - Running • 2 - Startup • 3 - Running & Startup 	smallint	
DEPLOY_OPTION	Identifies the deployment options. <ul style="list-style-type: none"> • 1-Deploy Now • 2-Save & Deploy • 3-Save deployment only • 4-Scheduled 	smallint	
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	

TABLE 76 DEPLOYMENT_CONFIGURATION (Continued)

Field	Definition	Format	Size
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 <ul style="list-style-type: none"> • 1-Pre snapshot • 2-Post snapshot • 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

TABLE 77 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
PRIVILEGE_ID	Comma separated privilege IDs	varchar	64

TABLE 78 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	

TABLE 78 DEPLOYMENT_PRODUCT_STATUS (Continued)

Field	Definition	Format	Size
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-Successful 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 79 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	
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 80 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

TABLE 80 DEPLOYMENT_TARGET_MAP (Continued)

Field	Definition	Format	Size
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 81 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
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)

TABLE 81 DEVICE (Continued)

Field	Definition	Format	Size
IS_SLB	Flag to identify whether the device supports server load balancing or not.	num	(1,0)
FIRST_SEEN_TIME		vchar	64
LAST_SEEN_TIME	Time when the device is getting discovered by recent collection.	vchar	64
LAST_PROBE_TIME		vchar	64
LAST_PROBE_STATUS		vchar	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.	vchar	1
SNMPV3_RO_AUTH_USERNAME	SNMP V3 read only authentication user name.	vchar	512
SNMPV3_RO_AUTH_PASSWORD	SNMP V3 read only authentication password.	vchar	512
SNMPV3_RO_PRIV_PROTOCOL	SNMP V3 read only privacy protocol.	vchar	1
SNMPV3_RO_PRIV_PASSWORD	SNMP V3 read only privacy password.	vchar	512
SNMPV3_RW_AUTH_TYPE	SNMP V3 read write authentication type.	vchar	1
SNMPV3_RW_AUTH_USERNAME	SNMP V3 read write authentication user name.	vchar	512
SNMPV3_RW_AUTH_PASSWORD	SNMP V3 read write authentication password.	vchar	512
SNMPV3_RW_PRIV_PROTOCOL	SNMP V3 read write privacy protocol.	vchar	1
SNMPV3_RW_PRIV_PASSWORD	SNMP V3 read write privacy password.	vchar	512
LOCAL_USERNAME_PORT_CFG	Agent user name configured in device used for port configuration.	vchar	512
LOCAL_PASSWORD_PORT_CFG	Agent password configured in device used for port configuration.	vchar	512
LOCAL_USERNAME_READ_ONLY	Local user name for read only access.	vchar	512
LOCAL_PASSWORD_READ_ONLY	Local password for read only access.	vchar	512
RADIUS_USERNAME_PORT_CFG	RADIUS user name configured in device used for port configuration.	vchar	512
RADIUS_PASSWORD_PORT_CFG	RADIUS password configured in device used for port configuration.	vchar	512
RADIUS_USERNAME_READ_ONLY	RADIUS user name configured in device used for read only access.	vchar	512
RADIUS_PASSWORD_READ_ONLY	RADIUS password configured in device used for read only access.	vchar	512
TAC_USERNAME_PORT_CFG	TACACS username for port configuration.	vchar	512
TAC_PASSWORD_PORT_CFG	TACACS password for port configuration.	vchar	512
TAC_USERNAME_READ_ONLY	TACACS username for read only access.	vchar	512
TAC_PASSWORD_READ_ONLY	TACACS password for read only access.	vchar	512
TACPLUS_USERNAME_PORT_CFG	TACACS+ username for port configuration.	vchar	512
TACPLUS_PASSWORD_PORT_CFG	TACACS+ password for port configuration.	vchar	512

TABLE 81 DEVICE (Continued)

Field	Definition	Format	Size
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
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. <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> 0 indicates not registered. 1 indicates registered. 	num	1

TABLE 81 DEVICE (Continued)

Field	Definition	Format	Size
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 <ul style="list-style-type: none"> • 0 is for unknown • 1 is for fixed configuration device • 2 is for chassis device • 3 is for stack device (logical) 	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: <ul style="list-style-type: none"> • 0=Netconf not supported by this device • 1=SSH • 2=HTTPS • 3=HTTP • 4=WING_HTTPS • 5=WING_HTTP 	smallint	

TABLE 81 DEVICE (Continued)

Field	Definition	Format	Size
POE_CAPABLE	The POE capability of device. Possible values are: <ul style="list-style-type: none"> 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	
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 Management Cluster. 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 are the values and their enums UNKNOWN("vcs-unknown-cluster"), STAND_ALONE("vcs-stand-alone"), FABRIC_CLUSTER("vcs-fabric-cluster"), MANAGEMENT_CLUSTER("vcs-management-cluster")	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.	smallint	

TABLE 81 DEVICE (Continued)

Field	Definition	Format	Size
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	
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 Management Cluster. States can be Online, Offline, Rejoining etc.. For all other devices this column will be empty.	varchar	64

TABLE 82 DEVICE_ENCLOSURE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of the Device enclosure.	varchar	256
TYPE	Type of Device enclosure - Storage Array/Server.	varchar	32

TABLE 82 DEVICE_ENCLOSURE (Continued)

Field	Definition	Format	Size
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_VALUE1	User-defined custom value.	varchar	256
USER_DEFINED_VALUE2	User-defined custom value.	varchar	256
USER_DEFINED_VALUE3	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. 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_MODULE	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 Enclosure.	varchar	256
SYSLOG_REGISTERED	SysLog flag that indicates if syslog has been enabled or not.	smallint	

TABLE 82 DEVICE_ENCLOSURE (Continued)

Field	Definition	Format	Size
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	

TABLE 83 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 84 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
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

TABLE 84 DEVICE_FDMI_DETAILS (Continued)

Field	Definition	Format	Size
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 85 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: <ul style="list-style-type: none"> • 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 86 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 87 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

TABLE 87 DEVICE_NODE (Continued)

Field	Definition	Format	Size
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
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_STAT E	Default value is 0.	smallint	

TABLE 88 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	

TABLE 88 **DEVICE_PORT**

Field	Definition	Format	Size
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
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 nullable 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 89 **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	

TABLE 89 DEVICE_PORT_GIGE_PORT_LINK (Continued)

Field	Definition	Format	Size
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 90 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 91 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	
STATE	Administrative state for this engine. 0 = disabled, 1 = enabled. The default value is 0.	smallint	

TABLE 91 ENCRYPTION_ENGINE (Continued)

Field	Definition	Format	Size
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 definition 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 92 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	
LEADER_SWITCH_WWN	The Node WWN of the current group leader switch. Each encryption group has one group leader switch.	char	23

TABLE 92 ENCRYPTION_GROUP (Continued)

Field	Definition	Format	Size
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	
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	

TABLE 92 ENCRYPTION_GROUP (Continued)

Field	Definition	Format	Size
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 93 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 94 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
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 95 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 96 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 97 ETHERNET_INTERFACE

Field	Definition	Format	Size
INTERFACE_ID		int	

TABLE 98 ETHERNET_ISL

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SOURCE_PORT_ID	The unique id of the source port.	int	
DEST_PORT_ID	The unique id of the destination port.	int	

TABLE 98 ETHERNET_ISL (Continued)

Field	Definition	Format	Size
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 99 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 99 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	

TABLE 100 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 101 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 102 EVENT_DESCRIPTION

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
DESCRIPTION	Holds the description of the Event.	varchar	1024

TABLE 103 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 104 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

TABLE 104 EVENT_INSTANCE (Continued)

Field	Definition	Format	Size
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	
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 105 EVENT_POLICY_SOURCE_ENTRY

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
EVENT_POLICY_ID	Foreign Key to Event_Policy Table	int	
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 <ul style="list-style-type: none"> • 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

TABLE 106 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 107 EVENT_RULE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Name of the Event Rule.	varchar	255
TYPE	Event Rule Type: <ul style="list-style-type: none"> • 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
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: <ul style="list-style-type: none"> • 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

TABLE 107 EVENT_RULE (Continued)

Field	Definition	Format	Size
LAST_MODIFIED_TIME	Rules last edited time.	timestamp	
SELECTED_TIME_UNIT	Timestamp unit of the selected rule: <ul style="list-style-type: none"> • 0 = second • 1 = Minutes • 2 = Hours Default value is 1.	smallint	

TABLE 108 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	

TABLE 108 FABRIC (Continued)

Field	Definition	Format	Size
LAST_UPDATE_TIME	Time when fabric was last updated. Default value is 'now()'.	timestamp	
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	Denotes 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 Network Advisor 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	

TABLE 109 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 110 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 111 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 112 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 113 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 114 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 115 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
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	

TABLE 115 FAVORITES (Continued)

Field	Definition	Format	Size
CUSTOM_SELECTION_OBJECT_TYPE	Represents the selected filter type. <ul style="list-style-type: none"> • 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 116 FCIP_CIRCUIT_PORT_MAP

Field	Definition	Format	Size
CIRCUIT_ID		int	
SWITCH_PORT_ID	SWITCH_PORT_ID of one end of the circuit	int	

TABLE 117 FCIP_PORT_TUNNEL_MAP

Field	Definition	Format	Size
SWITCHPORT_ID*	Switch Port ID.	int	
TUNNEL_ID*	FCIP Tunnel ID.	int	

TABLE 118 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	

TABLE 118 FCIP_TUNNEL (Continued)

Field	Definition	Format	Size
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	
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

TABLE 118 FCIP_TUNNEL (Continued)

Field	Definition	Format	Size
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	
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

TABLE 118 FCIP_TUNNEL (Continued)

Field	Definition	Format	Size
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	

TABLE 119 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 TapeAcceleration 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	
PATH_MTU_DISCOVERY	MTU Discovery Path. The default value is 0.	smallint	

TABLE 119 FCIP_TUNNEL_CIRCUIT (Continued)

Field	Definition	Format	Size
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	
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

TABLE 119 FCIP_TUNNEL_CIRCUIT (Continued)

Field	Definition	Format	Size
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	

TABLE 120 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	

TABLE 120 FCIP_TUNNEL_PERFORMANCE (Continued)

Field	Definition	Format	Size
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 121 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 122 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 123 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 124 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 125 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 126 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
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 127 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 the Management application repository.	varchar	1024
FIRMWARE_REPOSITORY_TYPE	Repository type to identify the FTP server: 0 = internal FTP. 1 = external FTP.	smallint	

TABLE 128 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 129 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	

TABLE 129 FOUNDRY_DEVICE (Continued)

Field	Definition	Format	Size
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 130 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)
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 131 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 132 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	

TABLE 132 FOUNDRY_PHYSICAL_PORT (Continued)

Field	Definition	Format	Size
MEDIA_TYPE	The media type for the port. Refer snSwPortInfoMediaType of foundry.mib for more details and possible values.	smallint	
GIG_TYPE		smallint	

TABLE 133 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 134 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 135 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

TABLE 135 FRU (Continued)

Field	Definition	Format	Size
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_DATE	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
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 136 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 137 GBIT_ETHERNET_INTERFACE

Field	Definition	Format	Size
INTERFACE_ID		int	

TABLE 138 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

TABLE 138 GIGE_PORT (Continued)

Field	Definition	Format	Size
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 139 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 140 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 141 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 142 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 143 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 144 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

TABLE 144 HBA (Continued)

Field	Definition	Format	Size
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
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

TABLE 144 HBA (Continued)

Field	Definition	Format	Size
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
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: <ul style="list-style-type: none"> • HBA/CNA/AnyIO/Mezzanine • HBA/Mezzanine CNA/Mezzanine AnyIO 	varchar	32
VENDOR	Adapter vendor name.	varchar	128

TABLE 145 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 146 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
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 port 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	

TABLE 146 HBA_PORT (Continued)

Field	Definition	Format	Size
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
MAX_PCIF	Maximum number of Pci functions.	smallint	
SYNTHETIC_FC	Synthetic FC is applicable for Windows only: <ul style="list-style-type: none"> • 0 - Unknow • 1 - Yes • 2 - No. 	int	

TABLE 147 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	

TABLE 147 HBA_PORT_DETAIL (Continued)

Field	Definition	Format	Size
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	
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	
RECEIVE_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

TABLE 147 HBA_PORT_DETAIL (Continued)

Field	Definition	Format	Size
QOS_CONFIGURED_STATE	Indicates whether QoS is enabled or disabled. The default value is 0.	smallint	
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. If 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 148 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 149 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

TABLE 149 HBA_PORT_FCOE_DETAILS (Continued)

Field	Definition	Format	Size
SUPPORTED_CLASS	The classes supported on the FCoE port. For example, Class2 and Class3.	varchar	256
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 150 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 regarding 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

TABLE 150 HBA_REMOTE_PORT (Continued)

Field	Definition	Format	Size
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
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_REENTRY_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_COMPLETION_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 151 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	

TABLE 151 HBA_REMOTE_PORT_LUN (Continued)

Field	Definition	Format	size
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
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 device	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 152 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 if 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 153 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	

TABLE 153 HOST_DISCOVERY_OPTION (Continued)

Field	Definition	Format	Size
CIM_IMPL	CIM implementation 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	
<i>Application_Name_USER_NAME</i>	Management application User Name of the user who generated the last operation on the request	varchar	255
<i>Application_Name_SERVER_ADDRESS</i>	Management application Server address which generated the last operation on this request	varchar	50

TABLE 154 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 155 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	

H Database tables and fields

TABLE 156 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 157 INM_IP_INTERFACE

Field	Definition	Format	Size
IP_INTERFACE_ID		int	
IP_ROUTING_SERVICE_ID		int	
INTERFACE_ID		int	
DEVICE_ID		int	
IP_SUBNET_ID		int	
IP_ADDRESS		varchar	40
SUBNET_MASK		varchar	40
PRIMARY_IP	Indicates if the IP address is the primary IP address of the Interface. 1 - Primary 0 - Secondary.	smallint	

TABLE 158 INTERFACE

Field	Definition	Format	Size
INTERFACE_ID		int	
SWITCH_SERVICE_ID		int	
DEVICE_ID		int	
NAME		varchar	255
IDENTIFIER		varchar	255
TABLE_SUBTYPE		varchar	255
TAG_MODE		smallint	
VLAN_TAG_TYPE		int	
UNTAGGED_VLAN_ID		smallint	
IF_NAME		varchar	64

TABLE 158 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	The type of the device that is connected to the edge switch port. -1 : NA, 0 : connected to device with unknown type, 1 : connected to managed Brocade branded AP, 2 : connected to standalone Brocade branded AP.	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

TABLE 159 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 160 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 160 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	

TABLE 161 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 162 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_TY PE	

TABLE 163 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 164 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 165 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. <ul style="list-style-type: none"> • 0 denotes untrusted • 1 denotes trusted. 	smallint	
CREATION_TIME	Creation time of the ISL record in the Management application database.	timestamp	
MISSING	Denotes whether ISL link is missing or not. <ul style="list-style-type: none"> • 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 166 ISL_CONNECTION

Field	Definition	Format	Size
ID	The primary key of the table.	int	
FABRIC_ID	This is the fabric ID	int	

TABLE 166 ISL_CONNECTION

Field	Definition	Format	Size
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	
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	
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	

TABLE 167 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	

TABLE 167 ISL_TRUNK_GROUP (Continued)

Field	Definition	Format	Size
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_STAT US	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 168 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 169 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 170 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 171 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	
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 172 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_SUBTY PE		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 173 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 174 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 175 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	
CONFIG_XPATH	The xpath string.	varchar	1024
LAST_UPDATE_TIME	Timestamp returned by the device for this particular xpath.	bigint	

TABLE 176 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. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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 177 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 is applicable, 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: <ul style="list-style-type: none"> • 0 - Trial, • 1 - Permanent. The default value is 0.	smallint	
SUB_TYPE	Sub Type of license: <ul style="list-style-type: none"> • 0 - Base, • 1 - Addon. The default value is 0.	smallint	
VALID	Is this license still considered: <ul style="list-style-type: none"> • 0 - No, • 1 - Yes. The default value is 1.	smallint	

TABLE 178 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. <ul style="list-style-type: none"> • 1 - currently active downgrade • 0 - inactive or older downgrade. 	smallint	

TABLE 179 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	

TABLE 180 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 181 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 182 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 183 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 184 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	
DEVICE_PORT_WWN	Device port WWN of physical device.	char	23
PHYSICAL_PID	PID of physical device.	char	6

TABLE 185 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 186 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 187 LSAN_ZONE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
BB_FABRIC_ID	Backbone fabric DB ID.	int	
EDGE_FABRIC_ID	FID assigned to edge fabric.	int	
NAME	LSAN zone name.	varchar	128
BACKBONE	0= is not a backbone lsan zone, 1= is a backbone lsan zone. Default value is 0.	smallint	

TABLE 188 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	

TABLE 189 LSAN_ZONE_MEMBER

Field	Definition	Format	Size
LSAN_ZONE_ID*	LSAN_ZONE record reference.	int	
MEMBER_PORT_WWN*	Zone member WWN.	char	23

TABLE 190 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: <ul style="list-style-type: none"> • Deployed(0) • Undeployed(1) 	smallint	
VCN_MEMBER_ID	Virtual Cluster Node member Cluster id foreign key.	int	

TABLE 192 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 193 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 194 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 195 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	

TABLE 196 MAPS_EVENT_DETAILS

Field	Definition	Format	Size
ID	The primary key of the table.	int	
MAPS_EVENT_ID	The corresponding maps_event.	int	
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

TABLE 196 MAPS_EVENT_DETAILS (Continued)

Field	Definition	Format	Size
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	

TABLE 197 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 198 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 199 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 200 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	

TABLE 200 MESSAGE_RECIPIENT (Continued)

Field	Definition	Format	Size
RECIPIENT_TYPE_ID	Recipient Type (Syslog or SNMP).	int	
ENABLED	If forwarding to destination is enabled.	smallint	
SOURCE_ADDRESS_ADD ED	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 201 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 202 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 203 MODULE_TYPE

Field	Definition	Format	Size
MODULE_TYPE_ID		int	
MODULE_TYPE		num	(4,0)
NAME		varchar	32
DESCRIPTION		varchar	128
NUM_PORTS		num	(4,0)

TABLE 204 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 205 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 206 NETWORK_VLAN

Field	Definition	Format	Size
VLAN_DB_ID	Database ID of the VLAN instance which is associated with the Network.	int	

TABLE 207 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 208 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 209 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.	int	
DFID	Destination fabric ID	int	
SRCDOMAIN	Source domain ID	int	
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

TABLE 209 NP_FLOW_DEFINITION (Continued)

Field	Definition	Format	Size
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	

TABLE 210 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	int	
DFID	Destination fabric ID	int	
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

TABLE 211 OUI_GUESSED_DEVICE_MAP

Field	Definition	Format	Size
OUI*	Vendor OUI.	char	6
TYPE	Guessed device type for this vendor.	varchar	32

TABLE 212 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
VENDOR_CATEGORY	Default is 'none'.	varchar	32

TABLE 213 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 214 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 214 PHYSICAL_DEVICE (Continued)

Field	Definition	Format	Size
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)
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)

TABLE 215 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)
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	

TABLE 215 PHYSICAL_INTERFACE (Continued)

Field	Definition	Format	Size
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 216 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 217 PM_COLLECTOR_MEASURE_SETTING

Field	Definition	Format	Size
COLLECTOR_ID	ID of the data_collector.	int	
MEASURE_ID	ID of the measure.	int	

TABLE 218 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 219 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 220 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	
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 modified 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	

TABLE 220 PM_DASHBOARD_WIDGET (Continued)

Field	Definition	Format	Size
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 bands 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	
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

TABLE 220 PM_DASHBOARD_WIDGET (Continued)

Field	Definition	Format	Size
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 221 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 222 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 223 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 224 PM_WIDGET_MEASURE_TYPE

Field	Definition	Format	Size
Type	Primary key column.	int	
NAME	Storing the NAME of the measure.	varchar	64

TABLE 225 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 226 PM_WIDGET_MONITOR_TYPE

Field	Definition	Format	Size
Type	Primary key column.	int	
NAME	Storing the NAME of the monitor type.	varchar	64

TABLE 227 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 228 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 229 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 230 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 231 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. <ul style="list-style-type: none"> • 0 = product level • 1 = port level 	smallint	

TABLE 231 POE_THRESHOLD (Continued)

Field	Definition	Format	Size
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. <ul style="list-style-type: none"> 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 <ul style="list-style-type: none"> 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 232 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 233 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	

TABLE 233 PORT_BOTTLENECK_CONFIG (Continued)

Field	Definition	Format	Size
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 234 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 235 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	

TABLE 235 PORT_COMMISSION_CIMOM_SERVER (Continued)

Field	Definition	Format	Size
LAST_CONTACTED_TIME	Last time CIMOM server contacted. This will be updated when we test the reachability 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 236 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	<ul style="list-style-type: none"> • 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	
DEFAULT_POLICY	<ul style="list-style-type: none"> • 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 237 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	<ul style="list-style-type: none"> • 0 = All Fabric • 1 = Fabric • 2 = Core Switch Group • 3 = Switch • 4 = Port Type • 5 = Port List 	smallint	
SUB_LEVEL	<ul style="list-style-type: none"> • 1 = E_Port • 2 = F_Port • 3 = FL_Port, Fabric WWN, Switch WWN 	char	23

TABLE 237 PORT_FENCING_POLICY_MAP (Continued)

Field	Definition	Format	Size
NODE	WWN of Node which policy assigned.	char	23
INHERITANCE	Directly assigned or inherited from root level. <ul style="list-style-type: none"> • 0 = Directly assigned • 1 = Indirectly assigned 	smallint	

TABLE 238 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 239 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 240 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 241 QRTZ_BLOB_TRIGGERS

Field	Definition	Format	Size
TRIGGER_NAME*	Name of the trigger.	varchar	80
TRIGGER_GROUP*	Name of the trigger group.	varchar	80
BLOB_DATA	The Scheduler info.	bytea	

TABLE 242 QRTZ_CALENDARS

Field	Definition	Format	Size
CALENDAR_NAME*	Name of the Calendar.	varchar	80
CALENDAR	Calendar object.	bytea	

TABLE 243 QRTZ_CRON_TRIGGERS

Field	Definition	Format	Size
TRIGGER_NAME*	Name of the trigger.	varchar	80
TRIGGER_GROUP*	Name of the trigger group.	varchar	80
CRON_EXPRESSION	The CRON trigger Expression (ex:"0 0 12 * * ?" - meaning:Fire at 12pm (noon) every day).	varchar	80
TIME_ZONE_ID	Given "cron" expression resolved with respect to the TimeZone.	varchar	80

TABLE 244 QRTZ_FIRED_TRIGGERS

Field	Definition	Format	size
ENTRY_ID*	Fired instance ID.	varchar	95
TRIGGER_NAME	Name of the trigger.	varchar	80
TRIGGER_GROUP	Name of the trigger group.	varchar	80
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	80
FIRED_TIME	The trigger fired time.	num	(13,0)
STATE	The fired trigger job state.	varchar	16
JOB_NAME	Name of the job.	varchar	80
JOB_GROUP	Name of the job group.	varchar	80
IS_STATEFUL	Whether the job implements the interface StatefulJob.	boolean	
REQUESTS_RECOVERY	True or false.	boolean	

TABLE 245 QRTZ_JOB_DETAILS

Field	Definition	Format	Size
JOB_NAME*	Name of the job.	varchar	80
JOB_GROUP*	Name of the job group.	varchar	80
DESCRIPTION	Description of the job (optional).	varchar	120

TABLE 245 QRTZ_JOB_DETAILS (Continued)

Field	Definition	Format	Size
JOB_CLASS_NAME	The instance of the job that will be executed.	varchar	128
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	

TABLE 246 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 247 QRTZ_LOCKS

Field	Definition	Format	Size
LOCK_NAME*	Resource identification name assigned by user.	varchar	40

TABLE 248 QRTZ_PAUSED_TRIGGER_GRPS

Field	Definition	Format	Size
TRIGGER_GROUP*	Name of the trigger group.	varchar	80

TABLE 249 QRTZ_SCHEDULER_STATE

Field	Definition	Format	Size
INSTANCE_NAME*	Instance of the scheduler.	varchar	80
LAST_CHECKIN_TIME	Last fired time in milliseconds.	num	(13,0)
CHECKIN_INTERVAL	Repeat interval.	num	(13,0)
RECOVERER	Misfire instruction.	varchar	80

TABLE 250 QRTZ_SIMPLE_TRIGGERS

Field	Definition	Format	size
TRIGGER_NAME*	Name of the trigger	varchar	80
TRIGGER_GROUP*	name of the trigger group	varchar	80
REPEAT_COUNT	number of times to repeat	num	(13,0)

TABLE 250 QRTZ_SIMPLE_TRIGGERS (Continued)

Field	Definition	Format	size
REPEAT_INTERVAL	interval for first and second job	num	(13,0)
TIMES_TRIGGERED	Number of times the corresponding trigger fired	num	(13,0)

TABLE 251 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 252 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

TABLE 253 RECIPIENT_TYPE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
TYPE	Type of the recipient (Syslog or SNMP).	varchar	20

TABLE 254 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 255 REPORT_TYPE

Field	Definition	Format	Size
ID*	Meta Data for available reports.	int	

TABLE 255 REPORT_TYPE (Continued)

Field	Definition	Format	Size
NAME	Report name.	varchar	128
DESCRIPTION	Report type description.	varchar	256

TABLE 256 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	

TABLE 257 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 258 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 259 RESOURCE_GROUP

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	

TABLE 259 RESOURCE_GROUP (Continued)

Field	Definition	Format	Size
NAME	Resource group name.	varchar	128
DESCRIPTION	Resource group description.	varchar	512

TABLE 260 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 261 ROLE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
NAME	Role name.	varchar	128
DESCRIPTION	Role description.	varchar	512
HIDDEN	Field to identify whether the role is Hidden from users or not. Values: <ul style="list-style-type: none"> • 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 262 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 263 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 263 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 264 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 265 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 265 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 266 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 267 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 268 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 applciation and values filled in by Switch Asset Collector	timestamp	
LAST_UPDATE_TIME	provides the record last updated time, standard columns for Management applciation 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 268 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 269 SLOT

Field	Definition	Format	Size
SLOT_ID		int	
PHYSICAL_DEVICE_ID		int	
CORE_SWITCH_ID		int	
SLOT_NUM		num	(4,0)

TABLE 270 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	
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

TABLE 270 SMART_CARD (Continued)

Field	Definition	Format	Size
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 271 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 272 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: <ul style="list-style-type: none"> • 1-Pre snapshot • 2-Post snapshot 	int	
SNAPSHOT_TIME	Time when this pre/post snapshot occurred.		
MESSAGE	Detailed message for snapshot report.	text	

TABLE 273 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	

TABLE 273 SNMP_CREDENTIALS (Continued)

Field	Definition	Format	Size
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	64
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	64
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
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 274 SNMP_DATA

Field	Definition	Format	Size
ID	Primary key column.	serial	
MIB_OBJECT_ID	MIB Object ID.	int	

TABLE 274 SNMP_DATA (Continued)

Field	Definition	Format	Size
TARGET_TYPE	Target type of the SNMP collector data. The target type for, <ul style="list-style-type: none"> • 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 deviceld, and for port level it will use interfaceld.	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	Corresponding collector table ID.	int	
MIB_INDEX	Index value for a MIB variable. For scalar value it will be empty.	varchar	256

TABLE 275 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, <ul style="list-style-type: none"> • 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 276 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, <ul style="list-style-type: none"> • 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 277 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, <ul style="list-style-type: none"> • 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 occurred 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 278 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. The target type for, <ul style="list-style-type: none"> • 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	Corresponding collector table ID.	int	

TABLE 279 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, <ul style="list-style-type: none"> • 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	

TABLE 279 SNMP_EXPR_DATA_1DAY (Continued)

Field	Definition	Format	Size
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 280 SNMP_EXPR_DATA_2HOUR

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
EXPRESSION_ID		int	
TARGET_TYPE		smallint	
TARGET_ID		int	
VALUE		double precision	
TIME_IN_SECONDS		int	
COLLECTOR_ID		int	

TABLE 281 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 occurred in seconds	int	
COLLECTOR_ID	DB Id of the collector object used for collection	int	

TABLE 282 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 283 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	64
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	64
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 283 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 284 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 285 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 286 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 287 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	

TABLE 287 STATS_AGING (Continued)

Field	Definition	Format	Size
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 policy. <ul style="list-style-type: none"> 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 288 SUB_INTERFACE

Field	Definition	Format	Size
SUB_INTERFACE_ID		int	
INTERFACE_ID		int	
SUB_INTERFACE_VC_ID		int	

TABLE 289 SUB_PORT_VLAN

Field	Definition	Format	Size
VLAN_DB_ID		int	
PORT_VLAN_DB_ID		int	
IS_DYNAMIC		num	(1,0)

TABLE 290 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	

TABLE 290 SWITCH_BOTTLENECK_CONFIG

Field	Definition	Format	Size
LAST_UPDATE_TIME	Last update time of the record	timestamp	
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 291 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 292 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 293 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	
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

TABLE 293 SWITCH_MODEL (Continued)

Field	Definition	Format	Size
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 capable 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 294 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
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	

TABLE 294 SWITCH_PORT (Continued)

Field	Definition	Format	Size
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	
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	

TABLE 294 SWITCH_PORT (Continued)

Field	Definition	Format	Size
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	
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

TABLE 294 SWITCH_PORT (Continued)

Field	Definition	Format	Size
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_OF_FSET	Maximum offset supported in fame monitor for switch port.	int	
v	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
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	

TABLE 294 SWITCH_PORT (Continued)

Field	Definition	Format	Size
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	

TABLE 295 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 296 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 297 SYSTEM_PROPERTY

Field	Definition	Format	Size
NAME*	The name of the property.	char	64
VALUE	The value for the property.	varchar	2048

TABLE 298 TARGET_TYPE

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
TYPE	Type of the target device. Some possible values are <ul style="list-style-type: none"> • Switch • Device • Port • Host • Port Group • Product Group • VLAN • Fabric 	varchar	64

TABLE 299 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 300 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 301 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

TABLE 301 TIME_SERIES_DATA (Continued)

Field	Definition	Format	Size
ME_ID	ME_ID of the target.	int	
VALUE	30 mins aggregated data.	double precision	

TABLE 302 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 interfaced.	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 303 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 interfaced.	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 304 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	

TABLE 304 TIME_SERIES_DATA_30MIN

Field	Definition	Format	Size
TARGET_ID	Target ID of the PM collector data. For device level collector it will use deviceId, for port level it will use interfacedId.	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 305 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 interfacedId/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 306 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	
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 307 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	

TABLE 307 TIME_SERIES_DATA_1_2HOUR (Continued)

Field	Definition	Format	Size
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 308 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	
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 309 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 310 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	

TABLE 310 TIME_SERIES_DATA_2_1DAY (Continued)

Field	Definition	Format	Size
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 311 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	
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 312 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/fcipcTunnelId/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 313 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/fcipcTunnelId/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 313 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 314 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 315 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	

TABLE 315 TIME_SERIES_DATA_3_2HOUR (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 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 316 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 317 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	
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	

TABLE 318 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	

TABLE 318 TIME_SERIES_DATA_4_1DAY (Continued)

Field	Definition	Format	Size
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 319 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 320 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	

TABLE 320 TIME_SERIES_DATA_4_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	
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 321 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 322 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 323 TRUNK_GROUP_INTERFACE

Field	Definition	Format	Size
INTERFACE_ID		int	
VLAG	Specifies whether the lag is a vlag or not	smallint	

TABLE 324 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	

TABLE 325 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: <ul style="list-style-type: none"> • 0=Disabled • 1=Enabled Default value is 1.	smallint	
SOURCE_OF_CREATION	To identify the source for creating the user account. <ul style="list-style-type: none"> • 0= User created through Management applciation 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_PASS WORD	User CLI credential login password.	varchar	768

TABLE 325 USER_ (Continued)

Field	Definition	Format	Size
IP_PRODUCT_ENABLE_USE R_NAME	User CLI credential enable user name.	varchar	256
IP_PRODUCT_ENABLE_PAS SWORD	User CLI credential enable password.	varchar	768

TABLE 326 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 327 USERDEFINED_NETWORK_SCOPE

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
NAME	Name of the Scope	varchar	128
USER_ID	Foreign Key USER_.ID. ID of the user who created the Custom Dashboard.	int	

TABLE 328 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	

TABLE 328 USERDEFINED_NETWORK_SCOPE_MEMBERSHIP

Field	Definition	Format	Size
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	

TABLE 329 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 330 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 331 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 332 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 333 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	vchar	32
FCP_INITIATOR	The role of the virtual port; for example, FCP Initiator	vchar	256
SWITCH_IP	IP of the switch, the V port is connected to	vchar	128
VF_ID	VF ID for the V port	smallint	

TABLE 334 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	vchar	64
PORT_SPEED	Will be 10G.	vchar	32
PORT_TYPE	Will be Virtual-FCoE-Port	vchar	16
ENABLED	Enabled/disabled	smallint	
STATUS	Status	vchar	64
TRUNK_INDEX	Trunk index	smallint	
PORT_NUMBER	Port number	smallint	
NAME	Name	vchar	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	vchar	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	vchar	

TABLE 335 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.	vchar	64

TABLE 336 VIRTUAL_FCOE_PORT_STAT

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SWITCH_ID		int	

TABLE 336 VIRTUAL_FCOE_PORT_STAT (Continued)

Field	Definition	Format	Size
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	
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 337 VIRTUAL_FCOE_PORT_STAT_2HR

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SWITCH_ID		int	
PORT_ID		int	
TX		double precision	
RX		double precision	
TX_UTILIZATION		double precision	

TABLE 337 VIRTUAL_FCOE_PORT_STAT_2HR (Continued)

Field	Definition	Format	Size
RX_UTILIZATION		double precision	
CREATION_TIME		double precision	
ACTIVE_STATE		timestamp	
LINK_FAILURES		double precision	
TX_LINK_RESETS		double precision	
RX_LINK_RESETS		double precision	
SYNC_LOSSES		double precision	
SIGNAL_LOSSES		double precision	
SEQUENCE_ERRORS		double precision	
INVALID_TX		double precision	
CRC_ERRORS		double precision	
DATA_GAPS_5MIN		smallint	
DATA_GAPS_30MIN	Data gap in 30 minutes table	smallint	

TABLE 338 VIRTUAL_FCOE_PORT_STAT_30M

Field	Definition	Format	Size
ID	Unique generated database identifier.	int	
SWITCH_ID		int	
PORT_ID		int	
TX		double precision	
RX		double precision	
TX_UTILIZATION		double precision	
RX_UTILIZATION		double precision	
CREATION_TIME		smallint	

TABLE 338 VIRTUAL_FCOE_PORT_STAT_30M (Continued)

Field	Definition	Format	Size
ACTIVE_STATE		double precision	
LINK_FAILURES		double precision	
TX_LINK_RESETS		double precision	
RX_LINK_RESETS		double precision	
SYNC_LOSSES		double precision	
SIGNAL_LOSSES		double precision	
SEQUENCE_ERRORS		double precision	
INVALID_TX		double precision	
CRC_ERRORS		double precision	
DATA_GAPS_5MIN	Data gap in 5 minutes table	smallint	

TABLE 339 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. <ul style="list-style-type: none"> • 1-Enabled • 0-disabled 	smallint	
USER_VPWWN	User created VPWWN.	char	23
AUTO_VPWWN	VPWWN created by Switch.	char	23

TABLE 339 VIRTUAL_PORT_WWN_DETAILS (Continued)

Field	Definition	Format	Size
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 340 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. <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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 . <ul style="list-style-type: none"> • 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 340 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. <ul style="list-style-type: none"> • 0 is native • 2 is McData • 3 is open fabric. 	smallint	
CRYPTO_CAPABLE	Stores the switch capability for crypto support . <ul style="list-style-type: none"> • 1 is capable • 0 is not capable. 	smallint	
FCR_CAPABLE	Stores the switch capability for FCR support . <ul style="list-style-type: none"> • 1 is capable • 0 is not capable. 	smallint	
FCIP_CAPABLE	Stores the switch capability for FCIP support . <ul style="list-style-type: none"> • 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, <ul style="list-style-type: none"> • 1 is capable • 0 is not capable. 	smallint	
PORT_BASED_ROUTING	Indicates whether the port based routing is present. <ul style="list-style-type: none"> • 1 is present • 0 is absent. 	smallint	
IN_ORDER_DELIVERY	Indicates whether in order delivery is enabled or disabled. <ul style="list-style-type: none"> • 1 is enabled • 0 is disabled. 	smallint	

TABLE 340 VIRTUAL_SWITCH (Continued)

Field	Definition	Format	Size
INSISTENT_DID_MODE	Indicates whether persistent domain ID is enabled on the switch. <ul style="list-style-type: none"> • 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 340 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 Management Cluster. 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 are the values and their enums UNKNOWN("vcs-unknown-cluster"), STAND_ALONE("vcs-stand-alone"), FABRIC_CLUSTER("vcs-fabric-cluster"), MANAGEMENT_CLUSTER("vcs-management-cluster").	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	

TABLE 341 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 342 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 343 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 344 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	32
TABLE_SUBTYPE	Table subtype possible value is VLAN.	varchar	32

TABLE 345 VLAN_DYNAMIC_INTERFACE_MEMBER

Field	Definition	Format	Size
VLAN_INTERFACE_RELATION_ID	Database ID of the VLAN_INTERFACE_RELATION instance which is associated with the dynamic interface member.	int	

TABLE 346 VLAN_EXCLUDED_INTERFACE

Field	Definition	Format	Size
VLAN_INTERFACE_RELATION_ID	Database ID of the VLAN_INTERFACE_RELATION instance which is associated with the excluded interface member.	int	

TABLE 347 VLAN_INTERFACE_MEMBER

Field	Definition	Format	Size
VLAN_INTERFACE_RELATION_ID	Database ID of the VLAN_INTERFACE_RELATION instance which is associated with the interface member.	int	

TABLE 348 VLAN_INTERFACE_RELATION

Field	Definition	Format	Size
VLAN_INTERFACE_RELATION_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.		

TABLE 349 VLAN_STATIC_INTERFACE_MEMBER

Field	Definition	Format	Size
VLAN_INTERFACE_RELATION_ID	Database ID of the VLAN_INTERFACE_RELATION instance which is associated with the static interface member.	int	

TABLE 350 VLL_DEVICE_RELATION

Field	Definition	Format	Size
MPLS_SERVICE_DEVICE_RELATION_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 Tagged-2.	int	

TABLE 351 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 352 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

TABLE 352 VMOTION_EVENT (Continued)

Field	Definition	Format	Size
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
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 Datacenter 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 353 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

TABLE 353 VMOTION_PNIC_DETAILS (Continued)

Field	Definition	Format	Size
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 354 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 355 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	

TABLE 355 VM_DATASTORE_DETAILS

Field	Definition	Format	Size
STORAGE_IORM_SUPPORTED	Indicates whether the datastore supports Storage I/O Resource Management. 0 = no 1 = yes.	smallint	
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 356 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_STATUS	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	

TABLE 356 VM_DV_PORT (Continued)

Field	Definition	Format	Size
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
KEY	The key for the port	varchar	64
MOR_ID	The managed object reference number assigned by the hypervisor	int	

TABLE 357 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 358 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

TABLE 358 VM_DV_SWITCH (Continued)

Field	Definition	Format	Size
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	vvarchar	256
PRODUCT_NAME	Short form of the product name	vvarchar	256
VENDOR_NAME	Name of the vendor of this product	vvarchar	256
VERSION	Dot-separated version string. For example, "1.2"	vvarchar	256
FORWARDING_CLASS	Forwarding class of the distributed virtual switch	vvarchar	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	vvarchar	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.	vvarchar	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 359 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 360 VM_FC_HBA

Field	Definition	Format	Size
ID	Unique generated database identifier.	serial	
NODE_WWN	The world wide node name for the adapter	vchar	23
PORT_WWN	The world wide port name for the adapter	vchar	23
PORT_TYPE	The type of the fiber channel port. One of : <ul style="list-style-type: none"> • Fabric • Loop • Point to point • Unknown 	smallint	
SPEED	The current operating speed of the adapter in bits per second.	vchar	64
BUS	The host bus number	int	
DEVICE_NAME	The device name of host bus adapter	vchar	256
DRIVER	The name of the driver	vchar	256
MODEL	The model name of the host bus adapter	vchar	256
PCI	The Peripheral Connect Interface (PCI) ID of the device representing the host bus adapter	vchar	256
STATUS	The operational status of the adapter. Valid values include : <ul style="list-style-type: none"> 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 361 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 362 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
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_STATU S	<ul style="list-style-type: none"> • 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 363 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 364 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	vchar	256
DVS_UUID	The uuid of the DistributedVirtualSwitch that the HostProxySwitch is a part of	vchar	256
KEY_	The proxy switch key	vchar	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	vchar	256

TABLE 365 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	vchar	256
UPLINK_PORT_GROUP_KEY	The key of the portgroup to be connected to the physical NIC	vchar	256
UPLINK_PORT_KEY	The key of the port to be connected to the physical NICs	vchar	256
UPLINK_PORT_NAME	The name of the port to be connected to the physical NICs	vchar	256

TABLE 366 VM_HOST_VIRTUAL_NIC

Field	Definition	Format	Size
ID	Unique Auto Generated DB ID.	serial	
DEVICE_NAME	Device Name for the virtual NIC.	vchar	256
MAC	The media access control (MAC) address of the virtual network adapter	vchar	64

TABLE 366 VM_HOST_VIRTUAL_NIC (Continued)

Field	Definition	Format	Size
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
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 367 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 368 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	
POLICY	Network adapter teaming policy includes failover and load balancing. It can be one of the following: <ul style="list-style-type: none"> 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 369 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	

TABLE 369 VM_PATH (Continued)

Field	Definition	Format	Size
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	
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 370 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
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 371 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	

TABLE 371 VM_SECURITY_POLICY (Continued)

Field	Definition	Format	Size
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_SETTING_ID	ID of network settings table.	int	

TABLE 372 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 373 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: <ul style="list-style-type: none"> • 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 374 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_VIRTUA L__SWITCH_ID	Foreign Key to the vm_standard_virtual_switch table. The standard virtual switch on which this port group exists.	int	
MOR_ID	The managed object reference number assigned by the hypervisor.	int	

TABLE 375 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	
DASTORE_URL	The unique locator for the datastore.	varchar	256
DASTORE_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 376 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 377 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_SETTING_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 378 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 Network Advisor 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	

TABLE 378 VM_VCENTER (Continued)

Field	Definition	Format	Size
MANAGED_ELEMENT_ID	A foreign key referencing MANAGED_ELEMENT(ID).	int	
FAULT_MONITORING_ST ATE	Flag to indicate whether fault monitoring is registered or not for a VM Host. Possible values are: 1.Not registered 2.Registered (Default)	smallint	
NAME	The name of the VCenter.	varchar	64
UUID	Unique identifier for vCenter server instance.	varchar	64

TABLE 379 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 380 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	

TABLE 380 VM_VIRTUAL_ETHERNET_ADAPTER (Continued)

Field	Definition	Format	Size
ADDRESS_TYPE	MAC address type. Valid values for address type are: <ul style="list-style-type: none"> • 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
WAKE_ON_LAN_ENABLE D	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: <ul style="list-style-type: none"> • E1000 • Vmxnet • Pcnnet32 	smallint	
VM_STD_VSWITCH_POR T_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 381 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	vvarchar	80
DESCRIPTION	Optional user-entered notes describing the VM. (Annotation in VMware terminology.)	vvarchar	256
OS	Operating system name and version.	vvarchar	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.	vvarchar	64
MEM_RESOURCES	Summary of memory resource configuration. Format may depend on VM vendor.	vvarchar	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.	vvarchar	32
HOSTNAME	The primary hostname assigned to this VM.	vvarchar	128
BOOT_TIME	The date and time the VM was last started.	timestamp	
DATSTORE_NAME	The user-assigned name for the VMs datastore. The datastore holds the VMs virtual disks, swap file, and configuration data.	vvarchar	80
DATSTORE_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.	vvarchar	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		vvarchar	64
BINARY_IP	IP address in binary format.	bytea	
CONNECTION_STATE	The connectivity state of a virtual machine. <ul style="list-style-type: none"> • 0 = not available • 1 = connected • 2 = disconnected • 3 = inaccessible • 4 = invalid • 5 = orphaned 	smallint	
COMMITTED_STORAGE	Used storage by a particular virtual machine.	vvarchar	64

TABLE 381 VM_VIRTUAL_MACHINE (Continued)

Field	Definition	Format	Size
UNCOMMITTED_STORAGE	Additional Provisioned storage for a particular virtual machine.	varchar	64
UNSHARED_STORAGE	Exclusive storage for a particular virtual machine.	varchar	64

TABLE 382 VM_VIRTUAL_MACHINE_DATASTORE_MAP

Field	Definition	Format	Size
VM_DATASTORE_DETAIL_S_ID	A foreign key referencing VM_DATASTORE_DETAILS(ID).	int	
VIRTUAL_MACHINE_ID	A foreign key referencing VM_VIRTUAL_MACHINE(ID).	int	
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 383 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 384 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 385 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
IP_ADDRESS		varchar	128
STATUS		varchar	256
FIRMWARE_VERSION		varchar	128
CREATION_TIME		timestamp	
LAST_UPDATE_TIME		timestamp	

TABLE 386 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 387 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 388 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	
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 389 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 390 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	
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 391 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 392 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 393 ZONE

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
ZONE_DB_ID	PK the owning ZONE_DB.	int	
NAME	The zone name.	varchar	64

TABLE 393 ZONE (Continued)

Field	Definition	Format	Size
TYPE	The zone type.	int	
SUB_TYPE	The zone subtype.	int	
ACTIVATE	For TI zones only, zone is activated. Default value is 0.	smallint	
CONFIGURED_FAILOVER	Configured Failover state of the TI Zone.	smallint	
CONFIGURED_ACTIVATE	Configured active state of the TI Zone.	smallint	
ENABLED_FAILOVER	Enabled Failover state of the TI Zone.	smallint	
ENABLED_ACTIVATE	Enabled Active state of the TI Zone.	smallint	

TABLE 394 ZONE_ALIAS

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
ZONE_DB_ID	PK of the owning ZONE_DB.	int	
NAME	The zone alias name.	varchar	64

TABLE 395 ZONE_ALIAS_IN_ZONE

Field	Definition	Format	Size
ZONE_ALIAS_ID*	PK of the zone alias.	int	
ZONE_ID*	PK of the zone.	int	23

TABLE 396 ZONE_ALIAS_MEMBER

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
TYPE	Zone alias member type: 2 = WWN 4 = D,P	smallint	
VALUE	Member value (D,P or WWN).	varchar	256
ZONE_ALIAS_ID	PK of the owning zone alias.	int	

TABLE 397 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	

TABLE 397 ZONE_DB (Continued)

Field	Definition	Format	Size
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 398 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	
ACTIVE_CONTENT	Active zone raw config string	text	
TI_ZONE_CONTENT	TI zone raw config string	text	

TABLE 399 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 400 ZONE_IN-ZONE_SET

Field	Definition	Format	Size
ZONE_SET_ID*	PK of the owning zone set.	INT	
ZONE_ID*	PK of the owning zone.	INT	

TABLE 401 ZONE_MEMBER

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
TYPE	Member type: 2 = WWN 4 = D,P	smallint	

TABLE 401 ZONE_MEMBER (Continued)

Field	Definition	Format	Size
VALUE	Member value (D,P or WWN).	varchar	256
ZONE_ID	PK of owning zone.	int	

TABLE 402 ZONE_SET

Field	Definition	Format	Size
ID*	Unique generated database identifier.	int	
ZONE_DB_ID	PK of owning zone DB.	int	
NAME	Zone set name.	varchar	64
ACTIVE	1 = active zone set 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';

```

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.DESCRPTION,
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 VIEW dashboard_preferences_info AS
select
    DASHBOARD.NAME as DASHBOARD_NAME,
    DASHBOARD.DESCRPTION as DASHBOARD_DESC,
    DASHBOARD.CREATED_BY,
    DASHBOARD_CANVAS.NAME as CANVAS_NAME,
    DASHBOARD_CANVAS.DESCRPTION 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
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 as ID,
    DEPLOYMENT_CONFIGURATION.NAME as NAME,
    DEPLOYMENT_CONFIGURATION.DESCRPTION as DESCRIPTION,
    DEPLOYMENT_HANDLER.MODULE as MODULE,
    DEPLOYMENT_HANDLER.SUB_MODULE as SUB_MODULE,
    DEPLOYMENT_STATUS.DEPLOYMENT_TIME as DEPLOYMENT_TIME,
    DEPLOYMENT_CONFIGURATION.DEPLOY_OPTION as DEPLOYMENT_OPTION,
    DEPLOYMENT_STATUS.STATUS as STATUS,
    DEPLOYMENT_STATUS.DEPLOYED_BY as DEPLOYED_BY,
    DEPLOYMENT_CONFIGURATION.CREATED_BY as CREATOR,
    DEPLOYMENT_CONFIGURATION.SCHEDULE_ENABLED as SCHEDULE_ENABLED,
    DEPLOYMENT_CONFIGURATION.SNAPSHOT_ENABLED as 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
from
    DEPLOYMENT_CONFIGURATION
    join DEPLOYMENT_HANDLER on DEPLOYMENT_CONFIGURATION.DEPLOYMENT_HANDLER_ID
= DEPLOYMENT_HANDLER.ID
    left outer join DEPLOYMENT_STATUS on
        (DEPLOYMENT_STATUS.DEPLOYMENT_TIME =
            (select
                max(DEPLOYMENT_STATUS.DEPLOYMENT_TIME)
            from
                DEPLOYMENT_STATUS
            where
                DEPLOYMENT_STATUS.DEPLOYMENT_CONFIGURATION_ID =
DEPLOYMENT_CONFIGURATION.ID))
    left outer join SCHEDULE_ENTRY on
        SCHEDULE_ENTRY.IDENTITY = cast(DEPLOYMENT_CONFIGURATION.ID as
varchar(16))
        and SCHEDULE_ENTRY.TABLE_NAME = 'DEPLOYMENT_CONFIGURATION'
    left outer 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.DESCRPTION,

```

```

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 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) 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 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;

```

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,
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
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;

```

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_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_PORT_FULL_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_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.DESCRPTION,
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.DESCRPTION,
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;

```

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,
    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
from
    ISL_CONNECTION,
    SWITCH_PORT SOURCE_SWITCH_PORT,
    SWITCH_PORT DEST_SWITCH_PORT
where
    ISL_CONNECTION.SOURCE_SWITCH_PORT_ID = SOURCE_SWITCH_PORT.ID
and ISL_CONNECTION.TARGET_SWITCH_PORT_ID = DEST_SWITCH_PORT.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, 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) 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_ORIGIN.ID as ORIGIN,
EVENT_CATEGORY.ID as EVENT_CATEGORY,
EVENT_DESCRIPTION.DESCRPTION 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
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_ORIGIN.ID as ORIGIN,
EVENT_CATEGORY.ID as EVENT_CATEGORY,
EVENT_DESCRIPTION.DESCRPTION 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.DESCRPTION,
FABRIC.CREATION_TIME,
FABRIC.LAST_SCAN_TIME,
FABRIC.LAST_UPDATE_TIME,
FABRIC.TRACK_CHANGES,
FABRIC.TYPE,
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, FABRIC_MEMBER
where
    FABRIC.SEED_SWITCH_WWN = VIRTUAL_SWITCH.WWN and
    VIRTUAL_SWITCH.CORE_SWITCH_ID = CORE_SWITCH.ID and
    FABRIC_MEMBER.FABRIC_ID = FABRIC.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_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.DESCRPTION,
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.COMPRESSION_MODE,
FCIP_TUNNEL.TURBO_WRITE_ENABLED,
FCIP_TUNNEL.TAPE_ACCELERATION_ENABLED,
FCIP_TUNNEL.IPSEC_ENABLED,
FCIP_TUNNEL.PRESHARED_KEY,
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,
PORT.WWN as VIRTUAL_PORT_WWN,
PORT.REMOTE_PORT_WWN as REMOTE_PORT_WWN,
PORT.REMOTE_NODE_WWN as REMOTE_NODE_WWN,
PORT.ID as SWITCH_PORT_ID,
PORT.PORT_NUMBER as SWITCH_PORT_NUMBER,
PORT.USER_PORT_NUMBER as USER_PORT_NUMBER,
PORT.PORT_INDEX,
PORT.STATUS_MESSAGE
from
  FCIP_TUNNEL
  left outer join
    FCIP_PORT_TUNNEL_MAP on
      FCIP_PORT_TUNNEL_MAP.TUNNEL_ID = FCIP_TUNNEL.ID
  left outer 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,
    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);

```

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_REENTRY_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_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.DESCRPTION,
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;

```

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;

```

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,
    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;

```

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
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 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.FCIP_CIRCUIT_ID,
  MAPS_EVENT.FRU_NAME,
  MAPS_EVENT.VM_ID,
  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,
FCIP_TUNNEL_CIRCUIT.SLOT_NUMBER as FCIP_SLOT_NUMBER,
FCIP_TUNNEL_CIRCUIT.VE_PORT_NUMBER as FCIP_PORT_NUMBER,
MAPS_EVENT_CAUSE_ACTION.CAUSE,
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;

```

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.DESCRPTION as MODULE_TYPE_TXT
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.DESCRPTION
  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;

```

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;

```

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.DESCRPTION,
    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.DESCRPTION 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;

```

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.DESCRPTION,
  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_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_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,
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.DESCRPTION,
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_PORT_INFO

```

CREATE VIEW switch_port_info AS
  select
    SWITCH_PORT.ID,
    SWITCH_PORT.VIRTUAL_SWITCH_ID,

```

H Views

```
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.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_ENABLED,
SWITCH_PORT.FEATURES_ACTIVE,
SWITCH_PORT.DISABLED_REASON,
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.MANAGED_ELEMENT_ID,
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 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.DESCRPTION,
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;

```

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;

```

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_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_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,

```

H Views

```
VM_HOST.MEM_RESOURCES      as HOST_MEM_RESOURCES,
VM_HOST.LICENSE_SERVER,
VM_HOST.BOOT_TIME          as HOST_BOOT_TIME,
VM_HOST.CLUSTER_NAME as 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.DESCRPTION 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
from
  VM_DATA_CENTER,
  VM_HOST
  left join VM_VIRTUAL_MACHINE
    on VM_HOST.DEVICE_ENCLOSURE_ID = VM_VIRTUAL_MACHINE.HOST_ID
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;

-- Name: ap_usage; Type: VIEW; Schema: dcm; Owner: dcmadmin
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;

-- Name: events; Type: VIEW; Schema: dcm; Owner: dcmadmin
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)));

-- Name: sflow; Type: VIEW; Schema: dcm; Owner: dcmadmin
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);

-- Name: sflow_minute_bgp_view; Type: VIEW; Schema: dcm; Owner: dcmadmin
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;

-- Name: sflow_minute_l3_view; Type: VIEW; Schema: dcm; Owner: dcmadmin
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);

-- Name: sflow_minute_mac_view; Type: VIEW; Schema: dcm; Owner: dcmadmin
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);

-- Name: sflow_minute_vlan_view; Type: VIEW; Schema: dcm; Owner: dcmadmin
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);

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;

create 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;

-- Name: wired_interface; Type: VIEW; Schema: dcm; Owner: dcmadmin
CREATE VIEW wired_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.table_subtype, l2.tag_mode, l2.speed_in_mb, l2.physical_address,
    l2.duplex_mode, l3.ip_id, l3.ip_interface_id, l3.ip_address, l3.subnet_mask 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, i.table_subtype, i.tag_mode, 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));

-- Name: wireless_interface; Type: VIEW; Schema: dcm; Owner: dcmadmin
CREATE 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));

```

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.DESCRPTION,
    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;

```

VM_VNETWORK_INFO

This view provides 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
    PNIC_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;

```

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,
    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  
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;
```

H Views

Index

Symbols

“Viewing MAPS violations” on page 1934, 192

A

- access levels
 - defined, 1260
 - features, 1260–1261, 1262–??
 - roles, 1260
- accessing
 - FTP server folder, 124
- ACK emulation, device level, 841
- activating
 - Allow/Prohibit Matrix configuration, 900
 - LSAN zones, 785
 - zone configuration, 771
- active session management, roles and access levels, 1260
- active sessions, viewing, 9
- adapter software
 - using to manage driver files, 448
- adapters
 - HBA models, 440
 - types of, 440
 - types of converged network adapters, 441
 - types of fabric adapters, 441
 - types of HBAs, 440
- Adaptive Rate Limiting (ARL), 814
- add authentication card dialog box, 568
- Add Flow Definition dialog box
 - Bottlenecks Detection, 1028
 - End-to-End Monitors, 1031
 - Frame Viewer, 1027
 - Port Connectivity View, 1030
 - Top Talkers, 1031
 - Trace Route and Ping, 1029
- add/delete properties, roles and access levels, 1262
- Adding
 - C3 discard frames threshold, 869
 - state change threshold, 875, 882
- adding
 - detached devices to fabric binding, 859
 - invalid CRCs thresholds, 870
 - invalid words thresholds, 872
 - ISL protocol thresholds, 874
 - link reset thresholds, 873
 - members to LSAN zone
 - LSAN zone
 - adding members, 784
 - property labels, 267, 1281
 - storage ports to storage array, 432
 - switches to fabric binding, 859
 - thresholds, 869
 - traffic isolation zone members, 789
 - zone members, 761
 - zones, 770
 - adding a node to a cluster, 629
 - administrator access, defined, 1260
 - advanced filtering
 - setting up, 1067
 - alerts, zone configuration comparison, 796
 - allow/prohibit matrix
 - configuring, 895
 - Allow/Prohibit Matrix configuration
 - activating, 900
 - copying, 898, 899
 - deleting, 901
 - Allow/Prohibit Matrix display
 - changing, 901
 - AnyIO technology, 442
 - asset polling, configuring, 120
 - assigned thresholds
 - finding, 882
 - assigning
 - event filter to a device, 301
 - event filters to call home centers, 301
 - threshold policies, 978
 - thresholds, 877
 - associating HBAs to servers, 426
 - Authentication type
 - PAP, CHAP, 328
 - auto rekey
 - viewing time left, 731

B

- backbone fabrics
 - about, 540
 - connecting to edge fabrics, 542
- backup
 - changing interval, 81
 - configuration repository, 366
 - configuring to writable CD, 78
 - data, 77
 - disabling, 81
 - enabling, 80
 - immediate, 82
 - management server, 77
 - reviewing events, 82
 - roles and access levels, 1261
 - starting, 82
 - status, determining, 257
 - switch configuration, 366
 - viewing status, 81
- base switch, creating, 552
- blade processor
 - links, 577
- blade processors
 - configuring links, 577
- boot image repository
 - and host adapters, 451
 - backing up files, 453
 - deleting image, 453
 - downloading an image to a selected host, 452
 - importing, 451
- boot LUN zones
 - about, 792
 - creating, 792
 - deleting, 793
 - modifying, 793
- bottleneck detection, 967

C

- C3 Discard Frames threshold, 867
- CA certificate
 - loading onto the GL, 586
 - uploading onto the DPM, 584
- call home, 284

- centers
 - assigning a device, 299
 - assigning event filters, 301
 - disabling, 297
 - editing, 289
 - Brocade International, 289
 - e-mail, 290
 - EMC, 294
 - HP LAN, 295
 - IBM, 289
 - enabling, 296
 - enabling support save, 296
 - hiding, 288
 - removing a device, 299
 - removing all devices and filters, 299
 - removing event filters, 302
 - test connection, 297
 - viewing, 288
 - configuring, 284
 - roles and access levels, 1260
 - status, determining, 258
 - system requirements, 285
 - viewing status, 298
- call home event filters table
 - removing event filters, 303
- certificates
 - backing up, 626
 - cascaed FICON fabric
 - configuring, 903
 - cascaed FICON fabrics, 887
 - merging, 907
 - CEE management, roles and access levels, 1260
- certificates
 - storing the public key, 645
- changing
 - Allow/Prohibit Matrix display, 901
 - database passwords, 19, 23
 - port display, 280
 - port label, 280
 - product label, 280
 - view options, 262
 - changing connection utilization, 981
 - CHAP, 328
 - clearing fabric zone database, 797
 - clearing port counters, performance, 945
 - client authentication audit trail, displaying, 341
 - client export port, configuring, 113
 - client registration
 - manual enrollment of identities, 581, 585
 - client/server
 - firewall requirements, 3

- CNA
 - product overview, 441
- collapsing groups, 281
- color, changing, 278, 279
- community strings
 - reverting to default, 45
- comparing
 - zone databases, 794
- concepts, FCIP, 806
- configuration
 - Allow/Prohibit Matrix
 - activating, 900
 - deleting, 901
 - Allow/Prohibit Matrix , copying, 898, 899
 - FICON CUP, 887
 - storage encryption privileges, 565
 - storage port mapping, 431
- configuration file
 - searching, 371
 - viewing, 370
- configuration files, saving, 364, 365
- configuration management
 - roles and access levels, 1260
- configuration repository
 - backup, 366
- configuration repository management, overview, 363
- configuration status results
 - understanding, 669
- Configure menu, 1211
- configuring
 - allow/prohibit matrix, 895
 - asset polling, 120
 - bottleneck alert parameters, 968
 - call home, 284
 - cascaded FICON fabric, 903
 - client export port, 113
 - discovery, 39
 - e-mail notification, 1064
 - encrypted storage in a multi-path environment, 687
 - explicit server IP address, 116
 - external FTP server, 126
 - FCIP advanced settings, 837
 - FCIP tunnels, 830
 - FICON emulation, 841
 - FTP server, 123
 - HA clusters using BNA, 677
 - internal FTP server, 124
 - internal SCP server, 125
 - internal SFTP server, 125
 - IP interfaces, 830
 - IP routes, 830
 - IPSec and IKE policies, 839
 - LDAP server, 331
 - login banner, 102
 - login security, 102
 - LSAN zoning, 782
 - memory allocation, 118
 - Radius server, 328
 - routing domain IDs, 544
 - security authentication using the GUI, 465
 - server name, 101
 - server port, 128
 - smart cards, 566
 - SNMP credentials, 44
 - software, 106
 - support mode settings, 129
 - Switch authentication, 339
 - traffic isolation zoning, 788
 - UNIX authentication, 340
 - Windows authentication, 340
 - zoning, 759
- configuring zoning, 759
- connected ports, showing, 396
- connecting edge fabrics to backbone fabrics, 542
- connection utilization
 - changing, 981
 - disabling, 981
 - enabling, 980
 - overview, 979
 - supported on, 979
- connections
 - status, determining, 257
- connections between a switch and an LKM key vault, 748
- connections, changing display of, 278
- connections, monitoring utilization, 979
- copying
 - Allow/Prohibit Matrix configuration, 898, 899
 - log entries, 1127
 - log entry parts, 1127
 - master log, 1130
 - master log parts, 1130
 - threshold policies, 977
 - zones, 764
- copying views, 275
- creating
 - LSAN zone, 783
 - new members in LSAN zone
 - LSAN zone
 - creating new members, 785
 - storage array, 432
 - threshold policies, 974
 - traffic isolation zone, 788

- views, 271
- zone, 760
- zone alias, 766
- zone configuration, 769
- zone databases, 775
- zone members, 762
- zone sets, 769

CSR

- exporting from properties, 736
- submitting to a CA, 583

CUP, FICON, 887

- customized views, copying, 275
- customized views, deleting, 274
- customized views, editing, 273
- customizing, product list columns, 271, 273

D

- data
 - historical performance, 946
 - real time performance, 942
- data backup, 77
- data backup and restore, 469
- data collection
 - historical performance, 946
 - historical performance graph, 948
 - historical performance graph configuration, 948
- data restore, 83
- database fields
 - Sybase and Derby, 1307
- database, restoring, 342
- deactivating zone configuration, 773
- decommissioned IDs
 - deleting, 723
 - displaying, 723
- default background color, changing, 279
- default community strings, 45
- default desktop color, changing, 279
- default zone (fabrics)
 - disabling, 765
 - enabling, 765
- defining, event filter, 300
- delete
 - switch configuration, 372
- deleting
 - Allow/Prohibit Matrix configuration, 901
 - end-to-end monitoring pairs, 961
 - fabrics, 46, 47, 65
 - FCIP tunnels, 849
 - historical performance graph, 953
 - hosts, 66, 72
 - logical switches, 557
 - offline zone database, 776
 - property labels, 268, 1282
 - reports, 1203
 - storage arrays, 434
 - technical support information, 1195
 - threshold policies, 978
 - VM managers, 71
 - zone alias, 769
 - zone configuration, 774
 - zones, 764
- deleting firmware files from
 - firmware repository, 385
- deleting servers, 425
- deleting views, 274
- deployment configuration
 - deleting, 917
 - deploying, 917
 - duplicating, 916
 - editing, 916
 - snapshot report, 918
- deployment logs
 - viewing, 917
- Deployment Manager, 915
- deployment report, generating, 918
- Derby database fields, 1307
- device
 - adding names, 96
 - assigning event filters, 301
 - removing name, 97
- device icons, 258, 259
- device properties, 1265
- device properties dialog boxes, customizing, 1265
- device shortcut menu
 - adding options, 318
 - changing options, 319
 - removing options, 320
- device tips
 - configuring, 89
- device tips, turning on and off, 92
- device tips, viewing, 92
- diagnostics
 - types of tests, 443
- dialog box
 - iSCSI properties, 1270
 - port properties, 1271
 - storage properties, 1268
- directory structure overview, backing up, 78
- disabling

- bottleneck detection, 973
- call home centers, 297
- default zone for fabrics, 765
- fabric binding, 858
- FCIP tunnels, 848, 849
- historical performance data collection, 948
- login banner, 103
- port connectivity view filter, 394
- ports, 393
- traffic isolation zone, 790
- traffic isolation zone failover, 791
- disabling backup, 81
- disabling connection utilization, 981
- Discover menu, 1210
- discovering a fabric, 37
- discovery, 37
 - configuring, 39
 - description of, 443
 - in-band, enabling, 39
 - out-of-band, enabling, 39
 - setting up, 39
 - SNMP version, 39
 - state, 47
 - troubleshooting, 48, 67, 73
- discovery setup
 - roles and access levels, 1262
- disk devices
 - decommissioning, 722
- disk luns
 - decommissioning, 723
 - rekeying manually, 725
 - setting rekey all, 726
 - viewing rekey details, 727
- display
 - end nodes, 86
- display, FICON, 84
- displaying
 - bottleneck statistics, 972
 - event details, 1128, 1129
 - FCIP performance graphs for Ethernet ports, 850
 - FCIP performance graphs for FC ports, 850
 - firmware repository, 383
 - master log event details, 1128, 1129
- downloading
 - firmware, 381
- dual network cards, configuration, 117
- duplicate names, fixing, 93
- duplicating
 - zone alias, 769
 - zone configuration, 774
 - zones, 764

Dynamic Load Sharing (DLS), 861

E

- edge fabrics
 - about, 540
- Edit menu, 1208
- editing
 - property fields, 1282
 - property labels, 268, 1281
 - storage array properties, 434
 - thresholds, 878
 - views, 273
 - zone alias, 767
- EE state
 - disabling from properties, 737
 - enabling from properties, 737
- Element Manager, launching
 - launching Element Manager, 307
- e-mail event notification setup
 - roles and access levels, 1260
- e-mail filter override, 468
- e-mail notification
 - configuring, 1064
- emailing
 - technical support information, 1194
- enabling
 - bottleneck detection, 968
 - call home centers, 296
 - default zone for fabrics, 765
 - fabric binding, 857
 - FCIP tunnels, 848, 849
 - historical performance data collection, 946
 - port connectivity view filter, 394
 - ports, 393
 - support save for call home centers, 296
 - traffic isolation zone, 790
 - traffic isolation zone failover, 791
- enabling backup, 80
- enabling connection utilization, 980
- encryption
 - certificate generation, 578
 - configuration planning for the management application, 576, 632
 - configure dialog box, 564
 - configuring in a multi-path environment, 687
 - gathering information before using the setup wizard, 576, 632
 - launching the encryption targets dialog box, 720
 - node initialization, 578

- preparation, 632
- selecting mode for LUNs, 701
- viewing and editing group properties, 737
- encryption engine
 - rebalancing, 709
- encryption engines
 - adding to HA clusters, 746
 - effects of zeroizing, 719
 - recovering from zeroizing, 719
 - removing from HA clusters, 746
 - support for tape pools, 750
 - zeroizing, 719
- encryption group
 - adding a switch using the management application, 670
 - confirming configuration status, 662
 - creating using the encryption setup wizard, 633
 - switch connection requirements, 581, 608
- encryption group properties
 - editing, 737
 - using the restore master key, 720
 - viewing, 737
 - viewing encryption group properties, 737
- encryption group properties dialog box
 - authentication cards, 567
 - General tab, 742
 - HA Clusters tab, 678, 746
 - Link Keys tab, 748
 - Members tab, 742
- encryption groups
 - creating, 633
 - replacing an EE, 676
- encryption node
 - setting initialization, 578
- encryption properties
 - viewing properties, 732
- encryption switch or group, removing using the management application, 743
- encryption targets
 - adding, 681
 - adding to virtual targets and virtual initiators within the encryption switch, 680
 - configuring hosts for, 689
 - using the dialog box, 720
 - using the dialog box to add Disk LUNs, 723
- end nodes
 - display, 86
- end-to-end monitoring
 - configuring pair, 958
 - displaying pairs, 959, 960
 - overview, 957
 - refreshing, 960
- end-to-end monitoring pairs
 - deleting, 961
- engine operations tab
 - encryption group properties
 - engine operations tab, 752
- ESXi
 - adding host adapter credentials, 446
 - CIM listener ports, 445
- ESXi hosts
 - updating drivers, 448
- ESXi systems
 - management application support for, 445
- Ethernet events
 - disabling, 87
 - enabling, 87
- event action
 - handling special events, 1096
- event action definition
 - creating, 1088
 - deleting, 1101
- event action definitions
 - configuring e-mail settings, 1099
 - configuring varbind filters, 1091
 - creating a new definition, 1100
 - modifying an existing definition, 1100
- event actions, handling special events, 1097
- event custom reports
 - adding a report schedule, 1124
 - copying an existing definition, 1122
 - defining report settings, 1115
 - defining the report identity, 1116
 - deleting a report definition, 1123
 - editing a report definition, 1122
 - filtering a report definition, 1118
 - filtering events by date and time, 1120
- event details
 - displaying, 1128, 1129
- event filter
 - assigning, 301
 - assigning to a device, 301
 - defining, 300
 - overwriting, 302
 - removing from device, 303
 - searching for, 303
- event filtering, advanced, 1067
- event filters table
 - removing event filters, 303
- event logs, 1126
 - copying entries, 1127
 - copying parts, 1127

- exporting entries, 1128
- viewing, 1126
- event management
 - overview, 1063
 - roles and access levels, 1261
- event notification
 - configuring e-mail notification, 1064
 - overview, 974
- event notification, description, 1064
- event policies
 - viewing events, 1069
- event types, 467
- events
 - Ethernet, 87
 - event types, 467
 - filtering, 468, 1131
 - monitoring methods, 1063
 - storage, 88
 - viewing, 1069
- expanding groups, 281
- explicit server IP address
 - configuring, 116
- export
 - switch configuration, 372
- exporting
 - log entries, 1128
 - master log, 1131
 - real time performance data, 945, 952
 - reports, 1202
 - storage port mapping, 428, 437
 - zone alias, 768
 - zone databases, 780
- Extended Fabrics license, 907
- external FTP server
 - configuring, 126

F

- fabric assigned WWN
 - adding AG ports, 461
 - auto-assigning to a switch or AG port, 459
 - deleting from a switch or AG port, 460
 - disabling on switch or AG port, 459
 - manually assigning to a switch or AG port, 460
 - modifying on a switch or AG port, 460
 - moving across switches, 462
 - on attached AG ports, 461
- fabric binding
 - adding detached devices, 859
 - adding switches, 859

- disabling, 858
- enabling, 857
- removing switches, 860
- roles and access levels, 1262
- Fabric OS
 - seed switch version, 55
- Fabric OS feature listing, 27
- fabric properties, 1264
- fabric properties dialog boxes, customizing, 1264
- fabric tracking
 - roles and access levels, 1262
- Fabric Vision
 - components, 992
- Fabric Watch, launching, 310
- fabrics
 - deleting from discovery, 46, 65
 - deleting permanently, 47
 - discovering, 37
 - IPv6 discovery, 37
 - monitoring, 51
 - status, determining, 257
 - zone database, clearing, 797
- failback
 - invoking, 680
- Fastwrite, 820
 - determining when to enable, 827
- fault management
 - roles and access levels, 1261
- FC Address
 - for inactive iSCSI devices, 400, 1272
- FC routing module, 308
- FC-FC routing
 - about, 540
 - setting up, 541
 - supported switches, 539
- FCIP
 - advanced settings
 - configuring, 837
 - connection properties
 - viewing, 843
 - DSCP, 819
 - Ethernet connection
 - troubleshooting, 854
 - Ethernet port properties
 - viewing, 845
 - Fastwrite, 820
 - FC port properties
 - viewing, 845
 - IPsec implementation, 816
 - L2CoS, 819

- management
 - roles and access levels, 1262
- performance graphs, Ethernet ports
 - displaying, 850
- performance graphs, FC ports
 - displaying, 850
- properties
 - viewing, 844
- QoS implementation, 818
- services
 - licensing, 806
- Tape Pipelining, 820
- tunneling, 806
- tunnels
 - configuring, 830
 - deleting, 849
 - disabling, 848, 849
 - enabling, 848, 849
 - modifying, 832
- VE/VEX port properties
 - viewing, 845
- FCIP configuration
 - advanced settings, 837
 - IP interfaces, 830
 - IP routes, 830
- FCIP configuration, guidelines, 829
- FCIP Design Considerations
 - 7800 switch and FX8-24 blade, 811
- FCIP trunking, 808
- FCoE management, roles and access levels, 1261
- FCR configuration, launching, 308
- feature
 - active session management, 1260
 - add/delete properties, 1262
 - backup, 1261
 - call home, 284, 1260
 - CEE management, 1260
 - configuration management, 1260
 - discovery setup, 1262
 - e-mail event notification setup, 1260
 - event management, 1261
 - fabric binding, 1262
 - fabric tracking, 1262
 - fault management, 1261
 - FCIP management, 1262
 - FCoE management, 1261
 - FICON management, 1262
 - firmware management, 1261
 - high integrity fabric, 1262
 - host management, 1261
 - license update, 1261
 - licensing requirements, 27
 - Logical Switch Configuration, 1262
 - LSAN zoning, 1261
 - map port to storage, 1262
 - performance, 1261
 - properties edit, 1261
 - report, 1261
 - routing configuration, 1262
 - security, 1261
 - setup tools, 1261
 - software configuration properties, 1261
 - storage encryption configuration, 1262
 - storage encryption key operations, 1262
 - storage encryption security, 1262
 - technical support data collection, 1261
 - user management, 1261
 - view management, 1262
 - zoning activation, 1261
 - zoning offline, 1261
 - zoning online, 1261
 - zoning set edit limits, 1261
- feature-to-firmware requirements, 27
- Fibre Channel over IP, 806
- FICON
 - best practices for virtual fabrics, 548
 - cascaded fabrics, 887
 - configurations, 887
 - configuring a switch, 888, 890
 - configuring emulation, 841
 - CUP, 887
 - display
 - resetting, 85
 - setting, 84
 - planning switch configuration, 888
- FICON management
 - roles and access levels, 1262
- filtering
 - events for users, 468
 - master log events, 1131
 - port connectivity view results, 393
 - real time performance data, 944
- finding
 - assigned thresholds, 882
- firmware
 - deleting files from repository, 385
 - downloading, 381
 - management, overview, 380
 - overwriting, 382
- firmware management
 - roles and access levels, 1261
- firmware repository
 - deleting firmware files, 385

- displaying, 383
- importing into, 384
- first-time configurations, 584
- Flow Generator
 - limitations, 1016, 1018, 1019
 - prerequisites, 1016, 1018, 1019
- flow vision performance graph, 1013
- flyovers
 - configuring, 89
 - turning on and off, 92
 - viewing, 92
- FSPF link cost calculation when ARL is used, 814
- FTP
 - overview, 123
 - server
 - accessing the folder, 124
 - configuring, 123
 - testing, 127

G

- general tab
 - encryption group properties
 - general tab, 738
- generating
 - performance graph, 942
 - performance reports, 1203
 - reports, 1200
 - zoning reports, 1205
- graphing
 - end-to-end monitor pairs, historical, 960
 - end-to-end monitor pairs, real time, 959
 - historical performance data collection, 948
- graphs
 - FCIP performance for Ethernet ports, 850
 - FCIP performance for FC ports, 850
- group background color, changing, 278
- grouping
 - overview, 281
- groups
 - collapsing, 281
 - expanding, 281
 - overview, 281
- groups, changing color, 278
- groups, icons, 259
- guidelines
 - FCIP configuration, 829

H

- HA clusters
 - adding engines, 677
 - configuration rules, 677
 - creating using BNA, 677
 - removing engines, 679
 - removing engines from, 679
 - requirements for, 677
 - rules, 677
 - swapping engines in, 679
- HA clusters tab
 - encryption group properties
 - HA clusters tab, 746
- HBAs
 - associating to servers, 426
 - unassociating, 428
- HCM
 - features, 443
 - software overview, 442
 - statistics monitoring, 443
- HCM Agent, launching, 310
- Help menu, 1207, 1217
- high availability
 - deployment, 586
- high integrity fabric
 - roles and access levels, 1262
- high integrity fabric configuration settings, 903
- high integrity fabrics (HIF), requirements, 887
- historical performance data
 - disabling collection, 948
 - enabling collection, 946
 - graphing, 948
 - overview, 946
 - saving graph configuration, 948
- historical performance graph
 - deleting, 953
- host adapter
 - discovery, 444
- host adapters
 - adding a port configuration, 455
 - and boot image repository, 451
 - and connectivity map, 447
 - and driver repository, 449
 - and fault management, 467
 - and performance management, 464
 - and port mapping, 447
 - and role-based access control, 463
 - and security authentication, 465
 - and supportSave, 467
 - and syslog forwarding, 468

- and view management, 447
- bulk port configuration, 454
- configuring FAWWNs, 458
- configuring ports, 454
- deleting a driver file from the repository, 450
- deleting a port configuration, 457
- duplicating a port configuration, 457
- editing a port configuration, 457
- filtering event notifications, 468
- port WWN virtualization, 458
- host connectivity manager
 - about the application, 442
 - features, 443
- host discovery
 - state, 66, 72
- host group, icons, 260
- host management, description of, 439
- host management, remote, 442
- host management, roles and access levels, 1261
- hosts
 - configuring for encryption targets, 689
 - deleting permanently, 66, 72
- http
 - [//www.gemalto.com/readers/index.html](http://www.gemalto.com/readers/index.html), 566

I

- IBM z/OS Global Mirror (z Gm), 821
- icons
 - device, 258, 259
 - products, 258, 259
- IFL. See interfabric links
- IKE policies
 - configuring, 839
- immediate technical support information collection, 1191
- import
 - switch configuration, 372
- importing
 - firmware files and release notes, 384
 - storage port mapping, 435
 - zone databases, 780
- inactive iSCSI devices, identifying, 400, 1272
- in-band discovery, enabling, 39
- insistent domain ID (IDID), 860
- installation
 - Windows, ODBC driver, 21
- installing a patch, 33
- interfabric links (IFLs), 829

- internal FTP server
 - configuring, 124
- internal SCP server
 - configuring, 125
- internal SFTP server
 - configuring, 125
- Invalid CRCs threshold, 868
- Invalid CRCs thresholds
 - editing, 878, 879
- invalid CRCs thresholds
 - adding, 870
- Invalid words threshold, 868
- invalid words thresholds
 - adding, 872
 - editing, 880, 881
- IP frames, 806
- IP interfaces, configuring, 830
- IP routes, configuring, 830
- IPsec
 - FCIP, 816
- IPSec policies
 - configuring, 839
- iSCSI devices, identifying inactive, 400, 1272
- iSCSI properties dialog box, 1270
- ISL protocol threshold, 868
 - adding, 874

K

- KAC
 - importing signed certificate, 584, 625
- KAC certificate
 - setting expiry, 583
 - uploading onto the DPM, 585
- KAC CSR
 - exporting, 582
- keep
 - switch configuration, 373
- key vaults
 - connection from switch, 748
 - entering the IP address or host name for, 638, 644, 649, 654, 659

L

- launch
 - remote client, 4
- launching
 - Server Management Console, 323
 - SMIA Configuration Tool, 345
- launching Fabric Watch, 310
- launching FCR configuration, 308
- launching HCM Agent, 310
- launching Name Server, 309
- launching Telnet, 306
- launching Web Tools, 307
- layout, changing, 277
- layout, overview, 276
- LDAP server
 - configuring, 331
- learning mode, 996
- license update
 - roles and access levels, 1261
- licensing
 - FCIP services, 806
- Lifetime Key Manager (LKM)
 - description of, 587
- link keys tab
 - encryption group properties
 - link keys tab, 748
- link keys, creating, 748
- link reset threshold, 868
- link reset thresholds
 - adding, 873
- listing
 - un-zoned members, 800
 - zone members, 799
- LKM
 - creating link keys, 748
 - support for high availability (HA), 590, 598
- Load leveling and failover, 810
- log entries
 - copying, 1127
 - copying parts, 1127
 - exporting, 1128
- logging in
 - remote client, 4
 - remote SMIA configuration tool, 347
 - server, 3
- logical switch
 - assigning ports, 555
 - changing to base switch, 560
 - creating, 552

- deleting, 557
 - moving, 559
 - removing ports, 556
- Logical Switch Configuration
 - roles and access levels, 1262
- login banner
 - configuring, 102
 - disabling, 103
- login security
 - configuring, 102
- logon conflicts, 771
- logs
 - event, 1126
- LSAN zone
 - creating, 783
- LSAN zones
 - activating, 785
- LSAN zoning
 - configuring, 782
 - overview, 781
 - roles and access levels, 1261
- LUN
 - choosing to be added to an encryption target container, 700

M

- Main window
 - master log, 255
 - minimap, 256
- main window
 - SAN tab, 247
- Management application
 - server and client, 3
- management application
 - main window, 2, 248
 - user interface, 1
- Management application feature listing, 27
- Management application services
 - monitoring and managing, 324
- management information base (MIB), importing into the Management application, 1079
- managing
 - zone configuration comparison alerts, 796
- manual identity enrollment, 585
- manual rekey
 - viewing progress, 728
- map port to storage
 - roles and access levels, 1262

- master key
 - active, 711
 - alternate, 711
 - backup, 711
 - create new master key, 711
 - creating a new, 710
 - description of, 710
 - reasons they are disabled, 710
 - restore master key, 711
 - saving to a file, 712
- master keys
 - actions, 711
 - active, 711
 - alternate, 711
 - creating, 718
 - overview, 710
 - restoring from a file, 715
 - restoring from a key vault, 716
 - restoring from a smart card set, 717
 - saving to a key vault, 713
 - saving to a smart card set, 714
- master log, 255
 - copying, 1130
 - copying parts, 1130
 - displaying, 1128, 1129
 - exporting, 1131
 - filtering events, 1131
- members tab
 - encryption group properties
 - members tab, 742
- membership list, fabric binding
 - adding detached devices, 859
 - adding switches, 859
 - removing switches, 860
- memory allocation
 - configuration, 118
 - configuring asset polling, 120
 - viewing status, 121
- menu bar
 - Configure, 1211
 - Discover, 1210
 - Edit, 1208
 - Help, 1207, 1217
 - Monitor, 1214
 - Server, 1207, 1208
 - Tools, 1217
 - View, 1207, 1209
- M-EOS feature listing, 27
- merging
 - cascaded FICON fabrics, 907
 - zone databases, 777

- merging zones, 758
- metaSAN, definition, 540
- minimap, 256
 - anchoring, 256
 - attaching, 256
 - detaching, 256
 - floating, 256
 - resizing, 256
- modifying
 - FCIP tunnels, 832
- Monitor menu, 1214
- monitoring
 - connection utilization, 979
 - end-to-end, 957
 - end-to-end, configuring, 958
 - end-to-end, displaying, 959, 960
- monitoring fabrics, 51
- monitoring pairs
 - deleting, 961
 - refreshing, 960
- monitoring statistics, 443
- multi-path configuration for encrypted storage using the
 - Management application, 687
- multi-path environments
 - configuring encrypted tape storage, 702

N

- Name Server, launching, 309
- names
 - adding to existing device, 96
 - adding to new device, 97
 - editing, 98
 - exporting, 98
 - fixing duplicates, 93
 - importing, 99
 - removing from device, 97
 - searching by, 99
 - setting as non-unique, 93
 - setting as unique, 93
 - viewing, 95
- names, overview, 92
- naming conventions, 758
- NetApp Lifetime Key Manager (LKM), description of, 587
- NetApp LKM key vaults
 - effects of zeroizing, 720
- network size status, determining, 257
- new device, adding name, 97

O

objects

- removing thresholds, 884

offline ports, display, 400

offline zone database

- deleting, 776

out-of-band discovery

- setting up, 39

overwriting

- firmware, 382

overwriting, event filter, 302

P

PAP, 328

passwords

- database, changing, 19, 23

patch

- install, 33

- uninstall, 34

performance

- clearing port counters, 945

- roles and access levels, 1261

performance data

- real time, 942

performance database, 953

performance graph, 1013

- generating, 942

- saving historical configuration, 948

performance monitoring

- overview, 935

- thresholds, 974

performance reports

- generating, 1203

physical map

- customizing views, 271

- default background color, changing, 279

- displaying connections, 278

- group background color, changing, 278

- layout, changing, 277

- layout, overview, 276

- levels of detail, 263

- port display, changing, 280

- port label, changing, 280

- product label, changing, 280

- showing connected ports, 396

- viewing port types, 396

- viewing ports, 395

- zooming in, 262

- zooming out, 262

port connection properties, viewing, 396

port connectivity view

- disabling filter, 394

- enabling filter, 394

- filtering results, 393

- refreshing, 393

- resetting filter, 394

- viewing details, 395

port connectivity, viewing, 389

port display, changing, 280

port fencing inheritance

- avoiding, 878

port fencing, description, 863

port label, changing, 280

port optics

- refreshing, 402

- viewing, 400

port properties, 395, 1271

port properties dialog box, 1271

port status, determining, 400

port status, viewing, 11

port types, viewing, 396

port-based routing, 860

ports, 389

- determining status, 400

- disabling, 393

- enabling, 393

- showing connected, 396

- view connectivity, 389

- viewing, 395

- viewing connection properties, 396

- viewing types, 396

primary FCS, 37

printing

- reports, 1202

priorities, threshold, 866

privileges

- user, 1243

privileges, user, 565

product label, changing, 280

Product list, 251

- columns, 251

product list

- customizing columns, 271, 273

product overview, 441

products

- icons, 258, 259

- status, determining, 257

- properties
 - FCIP connection, 843
 - FCIP Ethernet port, 845
 - FCIP FC port, 845
 - FCIP VE/VEX port, 845
 - general FCIP, 844
 - storage array
 - editing, 434
 - viewing, 435
 - storage port
 - viewing, 434
- properties dialog box
 - iSCSI tab, 1270
 - Port tab, 1271
 - Storage tab, 1268
- properties edit
 - roles and access levels, 1261
- property fields
 - editing, 1282
- property labels
 - adding, 267, 1281
 - deleting, 268, 1282
 - editing, 268, 1281
- pseudo event definitions, 1103
 - adding an escalation policy, 1108
 - adding on the flapping policy, 1112
 - copying an existing definition, 1107
 - creating, 1103
 - creating an event action on the flapping policy, 1112
 - creating an event action on the resolving policy, 1111
 - deleting, 1107
 - filtering traps, 1105
 - modifying an existing definition, 1107
 - setting policies, 1104
- public key certificate
 - importing from properties, 736

Q

- QoS implementation in FCIP, 818
- QoS priorities per FCIP circuit, 814

R

- Radius server
 - configuring, 328
- RBAC
 - user privileges, 1243
- RDR application considerations, 827

- real time performance, 942
 - exporting data, 945, 952
 - filtering data, 944
 - graph, 942
- real time performance data
 - thresholds, 974
- reassigning
 - storage ports to storage array, 433
- redirection zones, 721
- refreshing
 - end-to-end monitoring pairs, 960
 - port optics view, 402
 - zone databases, 776
- refreshing the port connectivity view, 393
- registering SNMP traps, 1080
- remote client
 - launch, 4
 - logging in, 4
- remote host management, 442
- remote replication
 - metadata requirements, 697
- remote replication luns
 - SRDF, 696
- remote SMIA configuration tool
 - logging in, 347
- removing
 - members from zone, 763
 - objects from zone alias, 768
 - servers, 425
 - switches from fabric binding, 860
 - thresholds, 884
 - thresholds from individual objects, 884
 - thresholds from table, 884
 - zone from zone configuration, 771
 - zones from zone configuration, 771
- removing event filters
 - call home centers, 302
 - call home event filters table, 303
 - devices, 303
- renaming
 - zone alias, 768
 - zone configuration, 773
 - zones, 763
- renaming servers, 425
- replacing
 - zone members, 801
- replicate
 - switch configuration, 373, 377
- report
 - roles and access levels, 1261
- report types, 1200

- reports
 - deleting, 1203
 - exporting, 1202
 - exporting to e-mail recipients, 1206
 - generating, 1200
 - performance, 1203
 - printing, 1202
 - viewing, 1201
 - zoning, 1205
- resetting
 - port connectivity view filter, 394
- restore
 - switch configuration, 368
- restore data, 83
- restore master key wizard, 720
- restoring
 - database, 342
- reviewing
 - backup events, 82
- role based access control. See RBAC.
- roles, 1260
 - access levels, 1260
- rolling back changes
 - zone databases, 781
- routing configuration
 - roles and access levels, 1262
- routing domain IDs, configuring, 544

S

- SAN tab, 247
- saving
 - historical performance graph configuration, 948
 - switch configuration files, 364, 365
 - zone databases to switch, 780
- scheduling
 - technical support information collection, 1189
- search
 - names, 99
 - WWN, 100
- searching
 - configuration file, 371
 - members in zones, 798
 - Potential Members list, 798
 - zones in zone configuration, 799
 - Zones list, 799
- security
 - configuring, 101
 - roles and access levels, 1261

- security authentication
 - configuring using the GUI, 465
- security tab
 - encryption group properties
 - security tab, 744
- security tab on management application
 - using to back up a master key, 745
 - using to create a master key, 745
 - using to restore a master key, 745
- seed switch, 37, 55
 - change requirements, 56
 - changing, 57
 - FCS policy, 38
- sequential devices, 820, 821
- server configuration, 628
- server IP address, explicit, 116
- Server Management Console
 - about, 323
 - launching, 323
- Server menu, 1207, 1208
- server name
 - configuring, 101
- server name, determining, 258
- server port
 - configuring, 128
- server port numbers, changing, 327
- server properties, viewing, 10
- servers
 - associating to HBAs, 426
 - determining name, 258
 - logging in, 3
 - removing, 425
 - renaming, 425
- setting up
 - advanced filtering, 1067
 - discovery, 39
- setup tools, 305
 - adding menu options, 316
 - adding to device shortcut menu, 318
 - changing menu options, 317
 - changing option on device shortcut menu, 319
 - changing server address, 315
 - removing menu options, 317
 - removing option from device shortcut menu, 320
 - roles and access levels, 1261
- showing levels of detail, physical map, 263
- showing ports
 - connected, 396
 - procedure, 395

- SIM ports
 - enabling and disabling, 1022
 - general information, 1021
- smart card set
 - overview, 715
- smart cards
 - configuring, 566
 - removing using the management application, 575
 - saving to a file, 575
- SMIA Configuration Tool
 - launching, 345
- SNMP credentials
 - adding and editing SNMP v3, 1077
- SNMP credentials, configuring, 44
- SNMP trap forwarding
 - adding a trap filter, 1073
- SNMP trap recipients
 - adding to switches, 1070
 - removing from switches, 1071
- SNMP traps
 - description of, 1069
 - importing a new MIB, 1079
 - modifying the definitions of registered traps, 1082
 - registering, 1080
 - reverting a trap to its default, 1082
 - unregistering, 1081
- SNMP v3, adding and editing credentials, 1077
- software configuration, 106
- software configuration properties
 - roles and access levels, 1261
- special events handling, 1096, 1097
- SRDF pairs, 697
- start monitoring, 54
- state change threshold, 868
- status
 - backup, 81
 - discovery, 47
 - host discovery, 66, 72
 - memory allocation, 121
- status bar, 257
- stop monitoring, 52
- storage array
 - adding storage ports to, 432
 - creating, 432
 - deleting, 434
 - reassigning storage ports to, 433
 - unassigning storage ports from, 433
- storage array properties
 - editing, 434
 - viewing, 435
- storage arrays
 - configuring, 696
- storage encryption
 - configuration privileges, 565
 - configuring, 682
 - confirming the configuration status, 687
 - selecting the encryption engine for configuration, 683
 - selecting the hosts, 684
 - specifying a name for the target container, 685
- storage encryption configuration
 - roles and access levels, 1262
- storage encryption key operations
 - roles and access levels, 1262
- storage encryption security
 - privileges for, 565
 - roles and access levels, 1262
- storage events
 - configuring, 88
- storage port mapping
 - exporting, 428, 437
 - importing, 435
- storage port mapping configuration, description, 431
- storage port properties
 - viewing, 434
- storage ports
 - adding to storage array, 432
 - reassigning to storage array, 433
 - unassigning from storage array, 433
- storage properties, 1268
- support mode
 - configuring, 129
- Switch authentication
 - configuring, 339
- switch configuration
 - backup, 366
 - deleting, 372
 - exporting, 372
 - file, search content, 371
 - file, view content, 370
 - importing, 372
 - keeping past age limit, 373
 - replicating, 373, 377
 - restore, 368
- switch connection control (SCC) policy, 861
- switch encryption configuration
 - confirm configuration using the management application, 673
 - designate switch membership using the management application, 671
 - specify public key certificate file name using the management application, 672

- switch encryption properties
 - editing, 732
 - viewing, 732
- switch removal, consequences of, 743
- Sybase database fields, 1307
- syslog forwarding, 1085
 - adding a destination, 1085
 - adding a filter, 1086
 - description, 468
- syslogs
 - adding a recipient, 1083
 - removing a recipient, 1084

T

- tab
 - Authentication (SMC), 329, 331, 334, 339, 340, 341
 - Services (SMC), 342
- tab Ports (SMC), 327
- tab Technical Support Information (SMC), 343
- tab, Services (SMC), 324
- table
 - # Brocade events, 1234
 - # CONSRV event, 1233
 - # thermal event reason codes, 1234
 - call home event, 1240
 - features, user groups access levels, 1260–1261, 1262–??
 - privileges and application behavior, 1255–1260
- tables
 - config database fields, ??–1434
 - GigE port stats database fields, 1373–??
 - Meta SAN database fields, ??–1453
 - UI database fields, ??–1416
 - zoning 1 database fields, 1416–??
- tape lun
 - read ahead, 703
 - write early, 703
- tape lun statistics
 - clearing container, 705
 - clearing for specific tape luns, 706
 - clearing for tape luns in a container, 707
 - viewing container, 705
 - viewing for specific tape luns, 706
 - viewing for tape luns in a container, 707
- Tape Pipelining, 820
- tape pipelining, 821
- tape pools
 - adding, 750
 - description of, 750

- identifying using a name or a number, 751
 - modifying, 749
 - overview, 750
 - removing, 749
- tape pools tab
 - encryption group properties
 - tape pools tab, 749
- tape read and write acceleration, 820
- tape write acceleration, 821
- target disk luns
 - adding, 691
- target tape luns
 - adding, 698
- targets
 - moving, 701
- technical support data collection
 - roles and access levels, 1261
- technical support information
 - copying to an external FTP server, 1194
 - deleting, 1195
 - emailing, 1194
 - immediate, 1191
- technical support information collection
 - scheduling, 1189
- technical support information, capturing, 343
- technical support information, viewing, 1193
- Telnet
 - launching session, 306
- testing
 - FTP server, 127
- thin provisioned luns
 - general, 730
- thin provisioning
 - support, 730
- third-party tools
 - adding, 305
 - adding menu option, 316
 - adding to device shortcut menu, 318
 - changing menu options, 317
 - changing option on device shortcut menu, 319
 - changing server address, 315
 - removing menu options, 317
 - removing option from device shortcut menu, 320
 - starting, 306
- threshold
 - adding, 869
 - adding C3 discard frames, 869
 - adding state change, 875, 882
 - C3 Discard Frames, 867
 - Invalid CRCs, 868
 - Invalid words, 868

- ISL protocol, 868
- link reset, 868
- state change, 868
- threshold policies
 - assigning, 978
 - copying, 977
 - creating, 974
 - deleting, 978
- threshold priorities, 866
- thresholds, 866
 - assigning, 877
 - editing, 878
 - finding specific, 882
 - overview, 974
 - removing, 884
 - viewing, 883
 - viewing on a specific device, 883
- thresholds table
 - removing thresholds, 884
- TIN/TUP emulation, 841
- tips, turning on and off, 92
- tips, viewing, 92
- TLS certificates, 595
- tool tips, turning on and off, 92
- tool tips, viewing, 92
- toolbox, 251
- tools
 - adding, 305
 - adding menu options, 316
 - adding to device shortcut menu, 318
 - changing menu options, 317
 - changing option on device shortcut menu, 319
 - changing server address, 315
 - removing menu options, 317
 - removing option from device shortcut menu, 320
- Tools menu, 1217
- tooltips
 - configuring, 89
- topolgy
 - viewing ports, 395
- topology
 - changing port display, 280
 - changing port label, 280
 - changing product label, 280
 - customizing views, 271
 - displaying connections, 278
 - group background color, changing, 278
 - showing connected ports, 396
 - viewing port types, 396
- topology, changing layout, 277
- topology, overview, 276

- topology, See also physical map
- total user count, 258
- traffic isolation zone
 - adding members, 789
 - creating, 788
 - disabling, 790
 - disabling failover, 791
 - enabling, 790
 - enabling failover, 791
- traffic isolation zoning, 786
 - configuring, 788
- troubleshooting
 - discovery, 48, 67, 73
 - FCIP Ethernet connections, 854
- tunnels, configuring, 830

U

- unassigning
 - storage ports from storage array, 433
- unassociating, HBA to server, 428
- uninstalling a patch, 34
- universal IDs
 - displaying, 725
- UNIX authentication
 - configuring, 340
- unregistering an SNMP trap, 1081
- un-zoned members
 - listing, 800
- user
 - privileges, 1243
- User Administrator, 1243
- user ID, determining, 258
- user interface, description, 1
- user management
 - roles and access levels, 1261
- user privileges
 - defined, 565, 1243
 - RBAC, 1243
 - resource groups, 565
- users
 - access levels, 1260
 - disconnecting, 10
 - filtering events for, 468
 - privileges, 1243
- users, total, 258
- using from encryption group properties dialog, 720

V

VE_Ports, 829

VEX_Port, 829

view management, 271

roles and access levels, 1262

View menu, 1207, 1209

view options, changing, 262

View window

product list, 251

View window, toolbox, 251

viewing

call home status, 298

configuration file, 370

disabling port connectivity filter, 394

enabling port connectivity filter, 394

event logs, 1126

events, 1069

FCIP connection properties, 843

FCIP Ethernet port properties, 845

FCIP FC port properties, 845

FCIP VE/VEX port properties, 845

filtering port connectivity, 393

general FCIP properties, 844

iSCSI properties, 1270

offline ports, 400

port connectivity, 389

port connectivity details, 395

port optics, 400

port properties, 395, 1271

port types, 396

ports, 395

reports, 1201

resting port connectivity filter, 394

storage array properties, 435

storage port properties, 434

storage properties, 1268

technical support information, 1193

thresholds, 883

thresholds on a specific device, 883

zooming in, 262

zooming out, 262

viewing ports

connection properties, 396

views

copying, 275

creating, 271

deleting, 274

editing, 273

Virtual Fabrics, 545

FICON best practices, 548

VM Manager

deleting, 445

discovery, 444

editing, 445

VM managers

deleting from discovery, 71

VMware vSphere Update Manager, using to update drivers

on ESXi hosts, 448

W

Web Tools, launching, 307

Windows authentication

configuring, 340

Windows installation

ODBC driver installation, 21

WWN

searching by, 100

Z

zeroization

setting, 720

zeroizing

effects of using on encryption engine, 719

zone

adding to configuration, 770

alias, 766

creating, 760

creating LSAN, 783

database size, 758

merging, 758

removing, 771

traffic isolation, adding members, 789

traffic isolation, creating, 788

traffic isolation, disabling, 790

traffic isolation, disabling failover, 791

traffic isolation, enabling, 790

traffic isolation, enabling failover, 791

zone alias

creating, 766

deleting, 769

editing, 767

exporting, 768

zone alias, duplicating, 769

zone alias, removing objects, 768

zone alias, renaming, 768

zone configuration

- activating, 771
- adding zones, 770
- creating, 769
- deactivating, 773
- deleting, 774
- duplicating, 774
- finding member in Zones list, 799
- removing a zone, 771
- removing zones, 771
- renaming, 773
- zone configuration comparison alerts
 - managing, 796
- zone configuration member
 - finding in Zones list, 799
- zone database
 - automatic checkout, undoing, 797
- zone databases
 - comparing, 794
 - creating, 775
 - exporting, 780
 - importing, 780
 - merging, 777
 - refreshing, 776
 - rolling back changes, 781
 - saving to switch, 780
- zone members
 - adding to zone, 761
 - creating in zone, 762
 - finding in Potential Members list, 798
 - finding in zones, 798
 - listing, 799
 - removing from zone, 763
 - replacing, 801
- zone set
 - creating, 769
 - naming conventions, 758
- zone set. See zone configuration
- zones
 - deleting, 764
 - duplicating, 764
 - finding in zone configuration, 799
 - removing from zone configuration, 771
 - renaming, 763
- zoning
 - configuration overview, 759
 - configuring, 759
 - invalid names, 758
 - LSAN, 781
 - naming conventions, 758
 - offline, 757
 - online, 757
 - overview, 755
 - traffic isolation, 786
 - traffic isolation, configuring, 788
- zoning activation
 - roles and access levels, 1261
- zoning administration, 794
- zoning configuration
 - overview, 759
- zoning offline
 - roles and access levels, 1261
- zoning online
 - roles and access levels, 1261
- zoning reports
 - generating, 1205
- zoning set edit limits, roles and access levels, 1261
- zooming in, 262
- zooming out, 262



Printed in USA