

BROCADE



EDUCATION SOLUTIONS

BCSM in a Nutshell Study Guide for Exam 143-360

Exam Preparation Materials

Revision September 2009

Corporate Headquarters - San Jose, CA USA

T: (408) 333-8000
info@brocade.com

European Headquarters - Geneva, Switzerland

T: +41 22 799 56 40
emea-info@brocade.com

Asia Pacific Headquarters - Singapore

T: +65-6538-4700
apac-info@brocade.com

© 2009 Brocade Communications Systems, Inc. All Rights Reserved. 09/09

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, IronPoint, IronShield, IronView, IronWare, JetCore, NetIron, SecureIron, ServerIron, StorageX, and Turbolron are registered trademarks, and DCFM, Extraordinary Networks, and SAN Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

Revision: September 2009

BCSM in a Nutshell 8 Gbit/sec Edition



Objective: The BCSM Nutshell guide is designed to help you prepare for the BCSM Certification, exam number 143-360.

Audience: The BCSM Nutshell self-study guide is intended for those who have successfully completed the CSM 270 Brocade SAN Management and Monitoring Training course, and who wish to undertake self-study or review activities before taking the actual BCSM exam. The BCSM guide is not intended as a substitute for classroom training or hands-on time with Brocade products.

How to make the most of the BCSM guide: The BCSM guide summarizes the key topics on the BCSM exam for you in an easy to use format. It is organized closely around the exam objectives. We suggest this guide be used in conjunction with our free online knowledge assessment test - CSM 266 WBT BCSM 8 Gbit/sec Knowledge Assessment. To benefit from the BCSM guide, we strongly recommend you have successfully completed the CSM 270 Brocade SAN Management and Monitoring Training course.

We hope you find this useful in your journey towards BCSM Certification, and we welcome your feedback by sending an email to jcannata@brocade.com.

Helen Lautenschlager
Director of Education Solutions

Joe Cannata
Certification Manager

A handwritten signature in black ink, appearing to read "Helen Lautenschlager".

A handwritten signature in blue ink, appearing to read "Joe Cannata".

Table of Contents

1 - Security	7
SAN Device Security	7
DCC Policies	7
Authentication Policies	7
Access Gateway Policies	8
Switch Access Security	9
Active Directory	9
RADIUS	10
Policy Database Distribution	12
Managing SAN Data Security	13
Devices	13
Usage Limitations	13
Master Key Management	13
CryptoTargets	14
Encryption Group Configuration	14
2 - Management Tools	15
SAN Management Concepts and Tools	15
HBA Management	15
DCFM	15
DCFM Professional	16
DCFM Enterprise	16
SAN Health	17
Advanced Performance Monitor	18
3 - Troubleshooting	19
Troubleshooting and Data Gathering	19
LSANs	19
Ports, Media and ISLs	19
Performance Baselines	19
Zoning and Segmentation Issues Across a Fabric	20
Audit Event Logging	20
Fabric State	21
Problem Escalation	21
4 - SAN Management and Monitoring	22
Monitoring SAN Performance	22
Adaptive Networking	22
Port Fencing	22
SMI-S	22
Fabric Watch	23
TI Zones	24
ISL Trunking	25
Monitoring and Managing Fabrics	26
NPIV	26
Virtual Fabrics	27
Logical Switches	28
ICLs	30
Access Gateway F_Port Trunking	30
5 - SAN Configuration and Maintenance	31
Routine Maintenance	31
QoS Zones	31
Brocade Branded USB Device	32
Buffer-to-Buffer Credits	33
Web Tools - B-Series Element Manager	33
FC-FC Routing	34
FCIP	35
Virtual Ports and FCIP Tunnels	36
Taking the Test	37

List of Tables

DCC Policy States	7
Tunnel and Virtual Port Numbering	36

List of Figures

LDAP Example	9
RADIUS Example	11
Event Classes	20
Fabric Watch Alarms.....	23
Sample TI Zone	24
B-Series ISL Trunking.....	25
NPIV Example	26
Virtual Fabric Example with Logical Switches	27
DCX with ICLs	30
QoS Zones	31
Brocade Branded (supported) USB Device	32
MetaSAN with Edge-to-Edge, Backbone Fabrics and LSAN Zones.....	34
Sample NDA	37
Sample Question.....	38
Sample Exam Score Report.....	39

1 - Security

SAN Device Security

DCC Policies

Multiple DCC policies can be used to restrict which device ports can connect to which switch ports. The devices can be initiators, targets, or intermediate devices such as SCSI routers and loop hubs. By default, all device ports are allowed to connect to all switch ports; no DCC policies exist until they are created. For information regarding DCC policies and F_Port trunking, refer to the *Access Gateway Administrator's Guide*.

Each device port can be bound to one or more switch ports; the same device ports and switch ports may be listed in multiple DCC policies. After a switch port is specified in a DCC policy, it permits connections only from designated device ports. Device ports that are not specified in any DCC policies are allowed to connect only to switch ports that are not specified in any DCC policies.

When a DCC violation occurs, the related port is automatically disabled and must be re-enabled using the `portenable` command.

Policy State	Characteristics
No policy	Any device can connect to any switch port in the fabric.
Policy with no entries	Any device can connect to any switch port in the fabric. An empty policy is the same as no policy.
Policy with entries	<p>If a device WWN is specified in a DCC policy, that device is only allowed access to the switch if connected by a switch port listed in the same policy.</p> <p>If a switch port is specified in a DCC policy, it only permits connections from devices that are listed in the policy.</p> <p>Devices with WWNs that are not specified in a DCC policy are allowed to connect to the switch at any switch ports that are not specified in a DCC policy.</p> <p>Switch ports and device WWNs may exist in multiple DCC policies. Proxy devices are always granted full access and can connect to any switch port in the fabric.</p>

Table 1: DCC Policy States

Virtual Fabric Considerations:

The DCC policies that have entries for the ports that are being moved from one logical switch to another will be considered *stale* and will not be enforced. You can choose to keep *stale* policies in the current logical switch or delete the *stale* policies after the port movements. Use the `secpolicydelete` command to delete stale DCC policies.

Authentication Policies

By default, Fabric OS v6.1.0 and later use DH-CHAP or FCAP protocols for authentication. These protocols use shared secrets and digital certificates, based on switch WWN and public key infrastructure (PKI) technology, to authenticate switches. Authentication automatically defaults to FCAP if both switches are configured to accept FCAP protocol in authentication. To use FCAP on both switches, PKI certificates have to be installed.

Access Gateway Policies

The Brocade policy-based approach lets you restrict or filter traffic on standard Fabric OS switches and switches in Access Gateway mode. You can enable the following policies on a switch in Access Gateway mode:

- Advanced Device Security policy (ADS)
- Automatic Port Configuration policy (APC)
- Port Grouping policy (PG)

The ADS policy is supported on AG F_Ports. Fabric OS v6.2.0 extends the DCC policy to switches in AG mode to provide an additional level of security. It does this by extending the DCC policy to the physical F_Ports and the NPIV logins on F_Ports. As more physical servers become virtual, servers can become vulnerable and security becomes an integral part of server I/O virtualization. This security policy is a mechanism that restricts fabric connectivity to a set of devices that you can specify or allow to log in to the fabric connected through a switch in AG mode. By default, the ADS policy is not enabled. After you set a switch in AG mode, you can enable the ADS policy, and then specify which devices to allow at login on a per F_Port basis.

Security enforcement can also be done in the Enterprise fabric; the DCC policy in the Enterprise fabric takes precedence over the ADS policy. When you enable the ADS policy, it applies to all the ports on the switch. By default, all devices have access to the fabric on all ports.

You can determine which devices are allowed to log in on a per F_Port basis by specifying the device's port WWN (PWWN). Use the `ag --adsset` command to determine which devices are allowed to log in to a specified set of F_Ports. Lists must be enclosed in double quotation marks. List members must be separated by semicolons. Replace the WWN list with an asterisk (*) to indicate all access on the specified F_Port list. Replace the F_Port list with an asterisk (*) to add the specified WWNs to all the F_Ports' allow lists. A blank WWN list ("") indicates no access. The ADS policy must be enabled for this command to succeed. Note the following characteristics of the Allow List:

- The maximum device entries allowed in the Allow List is twice the per port max login count
- Each port can be configured to "not allow any device" or "to allow all the devices" to log in
- If the ADS policy is enabled, by default, every port is configured to allow all devices to log in
- The same Allow List can be specified for more than one F_Port

Switch Access Security

Active Directory

In Fabric OS v6.0 and higher, one AAA option is an external Lightweight Directory Access Protocol (LDAP) server running the Microsoft Active Directory service. Fabric OS v6.0 and higher includes an Active Directory/LDAP client. An Active Directory (AD) structure is a hierarchical framework of objects, which fall into three broad categories: resources (e.g. printers), services (e.g. e-mail) and users (user accounts and groups). The AD provides information on the objects, organizes the objects, controls access and sets security. Active Directory enables centralized, secure management of an entire network through a single user name and password combination that follows the user throughout the network. LDAP is an application protocol for querying and modifying directory services running over TCP/IP. Active Directory supports LDAPv3 and LDAPv2. With LDAP, user credentials are passed between the switch and the Active Directory server via a TLS-encrypted channel.

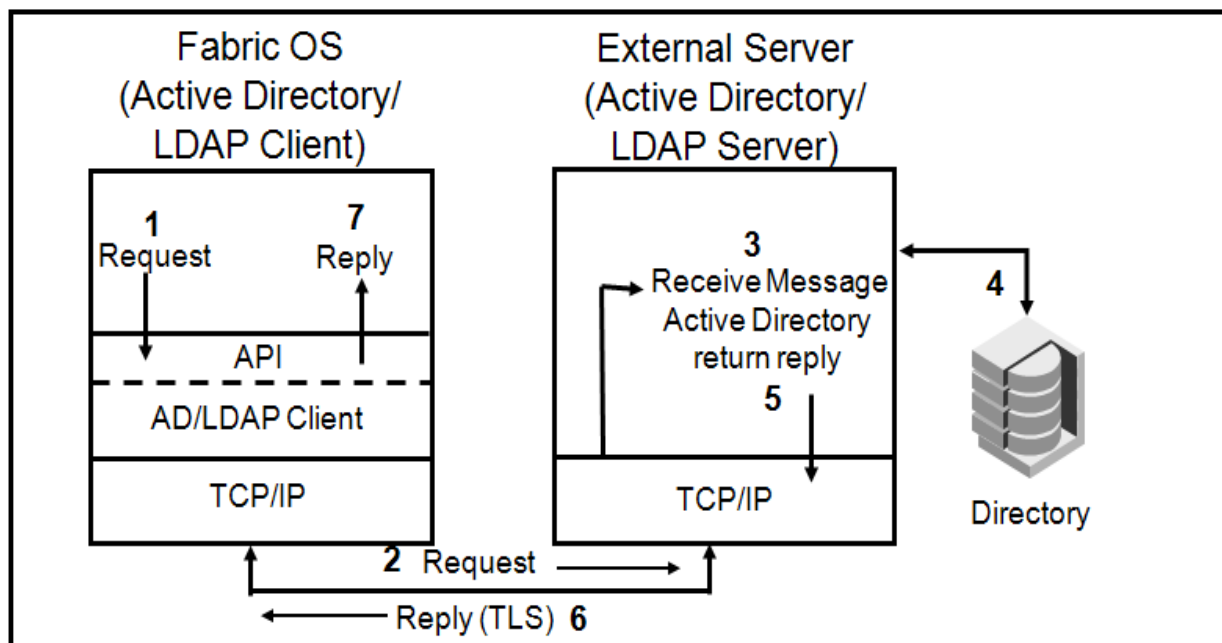


Figure 1: LDAP Example

The Transport Layer Security (TLS) protocol allows applications to communicate across a network in a way designed to prevent eavesdropping, tampering, and message forgery. TLS provides endpoint authentication and communication privacy over the Internet using cryptography. Typically, only the server is authenticated (i.e., its identity is ensured) while the client remains unauthenticated; this means that the end user (whether an individual or an application, such as a Web browser) can be sure with whom they are communicating.

The steps followed in authenticating and authorizing a user login with an AD/LDAP server are:

1. The authentication/authorization request is sent from Fabric OS on the switch
2. The AD/LDAP client in Fabric OS passes the request through the TCP/IP stack, across the IP network, and to the AD/LDAP Server
3. The request passes through the TCP/IP stack to the AD service, which receives the request
4. The AD/LDAP directory is queried, and the appropriate reply is determined
5. The reply is sent from the AD service to the TCP/IP stack
6. The reply passes through a TLS-encrypted channel to the switch
7. The reply passes through the TCP/IP stack on the switch to Fabric OS, which handles the information appropriately

Active Directory is used only for user authentication

- Cannot change switch passwords through Active Directory
- No automatic migration of new users from local switch database to the Active Directory server
- Active Directory authentication is only for a local switch, and not for the entire fabric

AAA services can be provided with any of the following combinations:

- Local switch database only (default)
- RADIUS only
- RADIUS and local
- AD/LDAP only
- AD/LDAP and local
- You cannot combine Active Directory with RADIUS

RADIUS

To centrally control user logins, Fabric OS supports the open-standard RADIUS protocol

- Provides remote user access authentication, authorization, and accounting
- Client/server model: A Brocade switch running Fabric OS acts as a RADIUS client to a RADIUS server
- Network Security: All RADIUS client/server traffic is authenticated via a shared secret key
- Focused on user logins, not FC device logins or switch attachment
- Simplifies the management of user accounts and passwords for larger fabrics

When RADIUS is enabled on a switch:

- All logins are authenticated through a RADIUS server (bypasses local database)
- All switch passwords are managed through the RADIUS server - the switch/Director local password database is bypassed
- Monitor user logins on a RADIUS-enabled system through the RADIUS server




Figure 2: RADIUS Example

In the example above, the primary database is RADIUS, and the RADIUS servers have been properly configured with user names and passwords. When a management station attempts access, the following scenario occurs:

1. The management server attempts to access the switch (RADIUS client) using a user name/password combination configured on the RADIUS servers
2. The authentication request is sent to the first RADIUS server in the RADIUS configuration (RADIUS Server 1)
3. If the response from RADIUS Server 1 is “accept”, management access is achieved; if the response is “deny”, the management server does not get access.
4. If there is a timeout from RADIUS Server 1, then the authentication request is sent to the second RADIUS server in the RADIUS configuration (RADIUS Server 2)
5. If the response from RADIUS Server 2 is “accept”, management access is achieved; if the response is “deny”, the management server does not get access
6. If there is a timeout from RADIUS Server 2 AND the switch database is configured as a secondary database, then the user name/password is authenticated on the local switch

Error messages related to RADIUS access attempts are displayed at the serial port console. Up to five RADIUS servers can be configured.

Security protocols provide endpoint authentication and communications privacy using cryptography. Typically, you are authenticated to the switch while the switch remains unauthenticated to you. This means that you can be sure with what you are communicating. The next level of security, in which both ends of the conversation are sure with whom they are communicating, is known as two-factor authentication. Two-factor authentication requires public key infrastructure (PKI) deployment to clients. Fabric OS supports:

- HTTPS
 - HTTPS is a Uniform Resource Identifier scheme used to indicate a secure HTTP connection. Web Tools supports the use of hypertext transfer protocol over secure socket layer (HTTPS). HTTPS is used for secure logins.
- LDAPS
 - Lightweight Directory Access Protocol over SSL uses a certificate authority (CA). By default, LDAP traffic is transmitted unsecured. You can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL)/Transport Layer Security (TLS) technology in conjunction with LDAP.
- SCP
 - 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. Configuration upload and download support the use of SCP.
- SNMP
 - Supports SNMPv1 and v3. SNMP is used in network management systems to monitor network-attached devices for conditions that warrant administrative attention.
- SSH
 - Secure Shell (SSH) is a network protocol that allows data to be exchanged over a secure channel between two computers. Encryption provides confidentiality and integrity of data. SSH uses public-key cryptography to authenticate the remote computer and allow the remote computer to authenticate the user, if necessary. Disable telnet before implementing SSH.
- SSL
 - Supports SSLv3, 128-bit encryption by default. Fabric OS uses secure socket layer (SSL) to support HTTPS. A certificate must be generated and installed on each switch to enable SSL.

Policy Database Distribution

Fabric OS lets you manage and enforce the ACL policy database on either a per-switch or fabric-wide basis. The local switch distribution setting and the fabric-wide consistency policy affect the switch ACL policy database and related distribution behavior. The ACL policy database is managed as follows:

Switch database distribution setting - Controls whether or not the switch accepts or rejects databases distributed from other switches in the fabric. The `distribute` command sends the database from one switch to another, overwriting the target switch database with the distributed one. To send or receive a database the setting must be `accept`. Virtual Fabric considerations: FCS, DCC, SCC, and AUTH databases can be distributed using the `distribute` command, but the PWD and IPFILTER databases are blocked from distribution.

Manually distribute an ACL policy database - Run the `distribute` command to push the local database of the specified policy type to target switches.

Fabric-wide consistency policy - Use to ensure that switches in the fabric enforce the same policies. Set a strict or tolerant fabric-wide consistency policy for each ACL policy type to automatically distribute that database when a policy change is activated. If a fabric-wide consistency policy is not set, then the policies are managed on a per switch basis.

Virtual Fabric considerations: Fabric-wide consistency policies are configured on a per logical switch-basis and are applied to the fabrics connected to the logical switches. Automatic policy distribution behavior for DCC, SCC and FCS is the same as that of pre-v6.2.0 releases and are configured on a per logical switch basis.

Managing SAN Data Security

Devices

The Brocade Encryption Switch - 32-port 8 Gbit/sec switch, with 48 to 96 Gbit/sec encryption processing bandwidth.

FS8-18 blade - 18-port 8 Gbit/sec blade with same functionality as the encryption switch. Up to 4 may be installed on a Brocade DCX Backbone or DCX-4S, however they only represent one encryption node.

Usage Limitations

There are usage limitations to be aware of when planning an encryption implementation:

- You cannot host both disk storage and tape storage on the same encryption switch or blade.
- Special redirection zones are created to handle data that is redirected to an encryption switch or blade. Quality of Service (QoS) cannot be applied to a redirection zone. They are part of the defined configuration.
- In order for frame redirection to be applied, regular zones for hosts and targets must be defined in the effective configuration. Hosts and targets must be zoned together by worldwide port name (WWPN) rather than worldwide node name (WWNN) in configurations where frame redirection will be used. If hosts or targets are zoned together using worldwide node name, frame redirection will not occur properly.
- On tapes written in DataFort format, the encryption switch or blade cannot read and decrypt files with a block size of one MB or greater.
- The Top Talker feature is not compatible with redirection zones. The Top Talker feature should not be enabled when an encryption switch or blade is present in the fabric.

Master Key Management

Communications with the RSA Key Management (RKM) and HP Secure Key Manager (SKM) key management systems are encrypted using a master key that is created by the encryption engine on the encryption switch. A master key must be generated by the group leader encryption engine. The master key can be generated once by the group leader, and propagated to the other members of an encryption group. It is essential to back up the master key immediately after it is generated. The master key may be backed up to any of the following:

- To a file as an encrypted key.
- To the key management system as an encrypted key record.
- To a set of recovery smart cards. This option is only available if the switch is managed by the Data Center Fabric Manager (DFCM), and if a card reader is available for attachment to the DCFM workstation.

The use of smart cards provides the highest level of security. When smart cards are used, the key is split and written on up to five cards, and the cards may be kept and stored by up to five individuals, and all are needed to restore the master key.

CryptoTargets

A Crypto LUN is the LUN of a target disk or tape storage device that is enabled for and capable of data-at-rest encryption. Crypto LUN configuration is done on a per-LUN basis. You configure the LUN for encryption by explicitly adding the LUN to the CryptoTarget container and turning on the encryption property and policies on the LUN. Using the CLI is the only way to manage CryptoTargets without disruption to data traffic.

Encryption Group Configuration

An encryption group consists of a set of member nodes that share the same key vault and are managed as a group. At least one node is required to form an encryption group (an encryption group of one would have one member acting as the group leader). An encryption group may include one or more High Availability (HA) clusters and data encryption key (DEK) clusters. An encryption group has the following properties:

- It is identified by a user-defined name
- It is managed from a designated group leader
- All group members share the same key vault, either a NetApp Lifetime Key Management (LKM), RSA Key Manager (RKM) or Hewlett Packard Secure Key Manager (SKM)
- If an opaque key vault is configured (RSA KM or HP SKM) the same master key is used for all encryption operations in the group
- All encryption engines in a chassis are part of the same encryption group
- In Fabric OS v6.2.0, an encryption group may contain up to sixteen encryption engines—up to four nodes with a maximum of four encryption engines per node

The basic encryption group configuration must be completed before you can set up a key vault or configure a storage device.

2 - Management Tools

SAN Management Concepts and Tools

HBA Management

HCM (Host Connectivity Manager) is the Brocade management tool for HBAs

- Has an intuitive and easy to use GUI (Graphical User Interface)
- Provides facilities for installing, configuring, troubleshooting, and monitoring Brocade HBAs
- HCM can be run on the local host or on a remote host for remote management
 - Discovery of remote HBAs requires IP addresses and user credentials
- Automatic LUN discovery displays from the device tree
- Target Rate Limiting requires Fabric OS v6.2

Management tasks can be performed using the BCU (Brocade Command Line Utility)

- BCU can be used to manage HBAs on the local host
- BCU can only be used to manage HBAs installed in the local host
- Can be launched from DCFM and EFCM

Management tasks can also be performed using DCFM Enterprise.

DCFM

Data Center Fabric Manager (DCFM) is the next-generation Brocade fabric management application. It combines the best features from Brocade manageability products, Enterprise Fabric Connectivity Manager (EFCM) and Fabric Manager (FM). DCFM is available in two editions that differ in supported features, hardware platforms, and supported scalability limits

- DCFM Enterprise for enterprise customers
- DCFM Professional for Small and Medium Businesses (SMB)

DCFM centralizes management of Brocade fabrics within and across data centers, including those with different protocols.

- Supports server virtualization
- Reduces expenses and maximizes productivity by automating tasks and providing easy-to-use operations
- Secures fabrics by managing user access controls and ensuring consistent security settings
- Delivers real-time and historical performance monitoring
- Integrates seamlessly with leading third-party automation solutions to provide a holistic approach to data center management

DCFM Professional

Light-weight, fabric management application for SMBs. Available at no charge with standalone switches and embedded products

- A disk is shipped with new switches or can be downloaded

Single fabric management of pure Fabric OS Fabrics

- 10 switches
- 640 switch ports
- 1,000 hosts or storage devices

Local authentication – single client

- Runs on a desktop, laptop, or server
- Easy installation and configuration
- Topology rendering maps showing devices and links
- Last-known-good configurations can be backed up and restored

DCFM Professional is a fabric management application included with Brocade switches that contains a subset of the DCFM Enterprise edition features. DCFM Professional enables centralized management of up to 1000 ports in a single, pure Fabric OS fabric. It performs important functions, including firmware downloads and fabric configuration for groups of switches while accessing device-level information. It is shipped with Brocade 4 Gbit/sec and 8 Gbit/sec hardware products and is also available via web download. DCFM Professional is targeted at the small and medium business segment and does not support any of the enterprise-class products/technologies, such as the DCX Backbone, Fabric-based Encryption products, Fibre Channel Routing, FICON, etc. Customers using these products and technologies should consider DCFM Enterprise. DCFM Professional supports only pure-Fabric OS fabrics and manages only one fabric at a time. DCFM Professional does not support M-EOS platforms. It requires SNMPv3 for FCIP counters.

DCFM Enterprise

Enterprise-class data center fabric management application that provides end-to-end management across multiple fabrics


- Works with Enhanced Group Management (EGM) and Fabric OS
 - EGM is required for firmware downloads
- Designed for Director-class installations, FICON environments, and users leveraging routing and extension
- Encryption configuration of targets for end-to-end security
- Includes support for Brocade HBAs and virtualization
- Support for Virtual Fabrics
 - Displays logical switches denoted as a “V”
 - Requires SNMPv3
- Support for Fibre Channel Routing
- Browser and firewall settings must be correct

Integrates seamlessly with other data center products providing holistic datacenter management, monitoring, and orchestration capabilities:

- Server management products
- Storage management products
- Enterprise monitoring products

Brocade Data Center Fabric Manager (DCFM) Enterprise is a comprehensive network management application that enables end-to-end management of data center fabrics. Brocade DCFM Enterprise manages and secures the flow of data across multiple fabrics. It builds on top of the functionality available with DCFM Professional. DCFM Enterprise provides a comprehensive fabric management framework for the end-to-end management of the data center fabric from the server HBA ports through the fabrics to storage ports, and supports all enterprise-class products and technologies including the DCX Backbone, Fabric-based Encryption products, FICON, Fibre Channel Routing, Fibre Channel over IP (FCIP), etc. DCFM Enterprise integrates seamlessly with other datacenter server management products, storage management products, and enterprise monitoring products to provide holistic datacenter management, monitoring, and orchestration capabilities.

SAN Health



Automates the documentation of a SAN and comes in different forms. SAN Health is a free utility that helps you create:

- Comprehensive documentation
- Historical performance graphs
- Detailed topology diagrams
- Best practice recommendations

SAN Health Professional provides a straightforward, easy-to-navigate user interface for auditing SAN Health data captures it a valuable tool for inventory tracking and change management activities. Organizations can import up to two SAN Health captures to SAN Health Professional for immediate, detailed analysis about any component. In addition to its standard data analysis and search capabilities, the SAN Health Professional framework supports optional add-on modules on a subscription-fee basis. One such module, SAN Health Professional Change Analysis, highlights changes that have occurred to any SAN component between two SAN Health data captures.

SAN Health Professional Change Analysis identifies new components and any changes to the characteristics of components that are common to both data captures. Organizations can view specific types of changes by

using sophisticated search and filter functions. As with the standard SAN Health Professional search function, results can be limited to one category or a combination of categories (devices, ports, switches, directors, fabrics, aliases, zones, configurations) and can be filtered further to remove information on components that do not occur in both data captures—or any other specific values.

SAN Health Expert is a Professional Services engagement subscription service designed for organizations that want additional analysis and advice from a Brocade SAN expert. Live consultations are done quarterly.

Advanced Performance Monitor

The Brocade Advanced Performance Monitoring is a licensed feature. Additional performance monitoring features are provided through Web Tools and DCFM. Based on Brocade Frame Filtering technology and a unique performance counter engine, Advanced Performance Monitoring is a comprehensive tool for monitoring the performance of networked storage resources.

Advanced Performance Monitoring provides the following monitors:

- End-to-End monitors measure the traffic between a host/target pair.
 - Configured using the SID and DID
 - Can be set using either CLI or Web Tools
- Filter-based monitors measure the traffic transmitted through a port with specific values in the first 64 bytes of the frame.
- ISL monitors measure the traffic transmitted through an Inter-Switch Link (ISL) to different destination domains.
- Top Talkers monitors measure the flows that are major consumers of bandwidth on a switch or port.
 - Will conflict with end-to-end monitors
 - Uses CLI, Web Tools or DCFM

Advanced Performance Monitoring is not supported on VE_Ports (virtual FC Ports) and EX_Ports. If you issue commands for any Advanced Performance Monitors on VE_Ports or EX_Ports you will receive error messages.

3 - Troubleshooting

Troubleshooting and Data Gathering

LSANs

An LSAN consists of zones in two or more edge or backbone fabrics that contain the same devices. LSANs essentially provide selective device connectivity between fabrics without forcing you to merge those fabrics. FC routers provide multiple mechanisms to manage interfabric device connectivity through extensions to existing switch management interfaces. You can define and manage LSANs using Advanced Zoning and create LSAN zones using DCFM to push them. FCR harvests the LSAN zones from all administrative domains. If both edge fabrics have the matching LSAN zones and both devices are online, FCR triggers a device import. To support legacy applications, WWNs are reported based on the administrative domain context. As a result, you must not use the network address authority (NAA) field in the WWN to detect an FC Router. LSAN zone enforcement in the local fabric occurs only if the administration domain member list contains both of the devices (local and imported device) specified in the LSAN zone. Use of the PWWN is essential for device sharing.

Using `lsanzoneshow` output to troubleshoot can be helpful. Using the `-s` operand will display state information for the devices.

- `EXIST` - the device exists in this fabric with the zone entry
- `Imported` - the device has been proxy created (imported) into this fabric
- `Configured` - the device is configured to be in the LSAN, but neither exists nor imported
- `Initializing` - the device is in transition and not yet imported into this fabric

Ports, Media and ISLs

The `portperfshow` command can be used to determine port performance and ISL traffic. If you suspect the problem is media-related, other commands to be used would be `portloopbacktest`, `porttest` and `spinfab`. Be aware that some of these commands could cause some disruption while running. They will give an adequate picture of what is happening across ISLs or on individual ports. A SAN Health report could also give a snapshot over time to determine ISL traffic. To get a graphic of the ISL performance, Web Tools and DCFM could be used.

Performance Baselines

Performance baselines can be established with many tools:

- `portperfshow`
- Home-grown applications
- Advanced Performance Monitor
- DCFM
 - DCFM Enterprise can get historical performance
- CLI
- SAN Health

Zoning and Segmentation Issues Across a Fabric

- PID format mismatches
- Zone conflicts due to configuration, content or type mismatches
- Improper usage of the `defzone` command
- Zone database maximum size has been exceeded
- DCFM Enterprise allows you to easily compare two zone databases using a compare option
- Correct any issues by using either CLI, Web Tools or DCFM

Audit Event Logging

- Captures security violations and configuration changes, zoning, switch configuration changes, fabric events, firmware download, and login/logout
- The switch streams AUDIT events to the console and syslog server only
 - AUDIT events can be viewed using `auditdump -s`
 - The log keeps the last 256 messages/events
- For any given event, AUDIT messages capture the following information:
 - User Name - The name of the user who triggered the action
 - User Role - The access level of the user, such as, root or admin
 - Event Name - The name of the event that occurred
 - Status - The status of the event that occurred: success or failure
 - Event Info - Information about the event

Operand	Event Class	Description
1	Zone	You can audit zone event configuration changes, but not the actual values that were changed. For example, you may receive a message that states "Zone configuration has changed," but the message does not display the actual values that were changed.
2	Security	You can audit any user-initiated security event for all management interfaces. For events that have an impact on the entire fabric, an audit is only generated for the switch from which the event was initiated.
3	Configuration	You can audit configuration downloads of existing SNMP configuration parameters. Configuration uploads are not audited.
4	Firmware	You can audit configuration downloads of existing SNMP configuration parameters. Configuration uploads are not audited.
5	Fabric	You can audit Virtual Fabric related changes.

Figure 3: Event Classes

- Event auditing is a configurable feature, filtering is set to off by default
- Enable event auditing by configuring the syslog daemon to send the events to a configured remote host using the `syslogipadd` command
- Set up filters to screen out particular classes of events using the `auditcfg` command

Fabric State

The `supportsave` command is an excellent way to gather the current state of the fabric. It includes trace dump files, RASlog entries, other dumps and core files. You can collect `supportsave` data either with the CLI command or by using DCFM. Sometimes it is useful to have a “before” and “after” `supportsave` so the support provider can make comparisons. The same would be true of SAN Health reports.

Problem Escalation

All too often the resolution of a problem is delayed because proper information is not gathered, and made ready for escalation to a support provider. Sometimes the obvious is overlooked. These are Brocade recommended items:

- 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
- The switch serial number
- The WWN of the switch either by the `wwn` command or from the printed label

4 - SAN Management and Monitoring

Monitoring SAN Performance

Adaptive Networking

The use of Adaptive Networking, which requires a license, can aid in optimizing output across ISLs. Adaptive Networking is a suite of tools and capabilities that enable you to ensure optimized behavior in the SAN. Even under the worst congestion conditions, the Adaptive Networking features can maximize the fabric behavior and provide necessary bandwidth for high-priority, mission-critical applications and connections. Fabric OS v6.0 and higher is required. These are the features in the Adaptive Networking suite:

- Traffic Isolation Zoning
- QoS Ingress Rate Limiting
- QoS SID/DID Traffic Prioritization
- Top Talkers
 - Monitors ingress E_Port traffic

Port Fencing

Port Fencing is not a part of Adaptive Networking. It requires a Fabric Watch license and Fabric OS v6.1 or later. It performs direct discovery using MPI in the M-Series environments running v09.09.09 or later.

SMI-S

The Storage Management Initiative (SMI) was created by the Storage Networking Industry Association (SNIA) to develop a vendor-neutral application programming interface (API) that can be built into networked storage devices and management applications. SMI-S enables a single management application to manage devices from different vendors (arrays, switches, HBAs, tape libraries, hosts). The key SMI-S components are:

- Common Information Model (CIM) - CIM is the data model for WBEM, and provides a common definition of management information for systems, networks, applications and services, and allows for vendor extensions
- Web-based Enterprise Management (WBEM) - WBEM is a set of management and internet standard technologies developed to unify the management of enterprise computing environments
- Service Location Protocol (SLP) - SLP enables computers and other devices to find services in a local area network without prior configuration
- SMI-A Agent for B-Series Switches
 - Provides the topology details of the mixed fabric and will manage all B-Series switches present in mixed fabrics
- Fusion Agent for M-Series Switches
 - Provides the topology details of the mixed fabric and will manage all M-Series switches present in mixed fabrics

The Brocade SMI Agent is a proxy agent that resides on a separate host system. Brocade has released versions for Sun Solaris, Microsoft Windows, and Linux. The Brocade SMI Agent allows applications to

manage Brocade SAN infrastructures from a single access point by communicating with multiple switches and multiple fabrics. The Brocade SMI Agent does not require any modification or upgrade to deployed Brocade fabrics. It supports the widest range of switches in the industry.

Fabric Watch

Optionally licensed per switch, it monitors the performance and status of the switch, including:

- Fabric events – Fabric reconfigs, zone changes, and new logins
- Switch status – Environmental (fans, power supplies, and temperature), SFP (Tx/Rx power, current, & voltage), Security, resource and FRU
- Port status – Monitors E/F/FL_Port signal quality parameters

Fabric Watch maintains a set of counters for each of the monitored conditions

- Tracks the number of occurrences of each condition
- Each counter is compared with an upper boundary and lower boundary
- Can monitor up to 8 logical switches

A default Fabric Watch configuration is available for the purpose of saving setup time. The custom settings available in Fabric Watch provide user flexibility for redefining thresholds and alarm notification methods. Customization is recommended to achieve the following objectives:

- Selecting a message delivery method for critical and non-critical events
- Selecting thresholds and alarm levels relevant to each class element
- Defining the appropriate Time Base event triggering based on the class element traits
- Eliminating message delivery that has little or no practical value to the SAN administrator
- Eliminating multiple messages generated from a single event

Fabric Watch provides the following types of automatic notifications:

- A continuous alarm provides a warning message whenever a threshold is breached; it continues to send alerts until the condition is corrected.
 - For example, if a switch exceeds its temperature threshold, Fabric Watch activates an alarm at every measurement interval until the temperature returns to an acceptable level
- A triggered alarm generates the first warning when a threshold condition is reached and a second alarm when the threshold condition is cleared

Alarms Custom Defined					
	ERROR_LOG	SNMP_TRAP	PORT_LOG_LOCK	RAPI_TRAP	EMAIL_ALERT
Changed	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Below	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Above	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Inbetween	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>



Figure 4: Fabric Watch Alarms

TI Zones

Fabric OS v6.0 introduced the concept of Traffic Isolation zones

- Can create a dedicated route
- Do not modify the routing table
- Are implemented across the entire data path from a single location

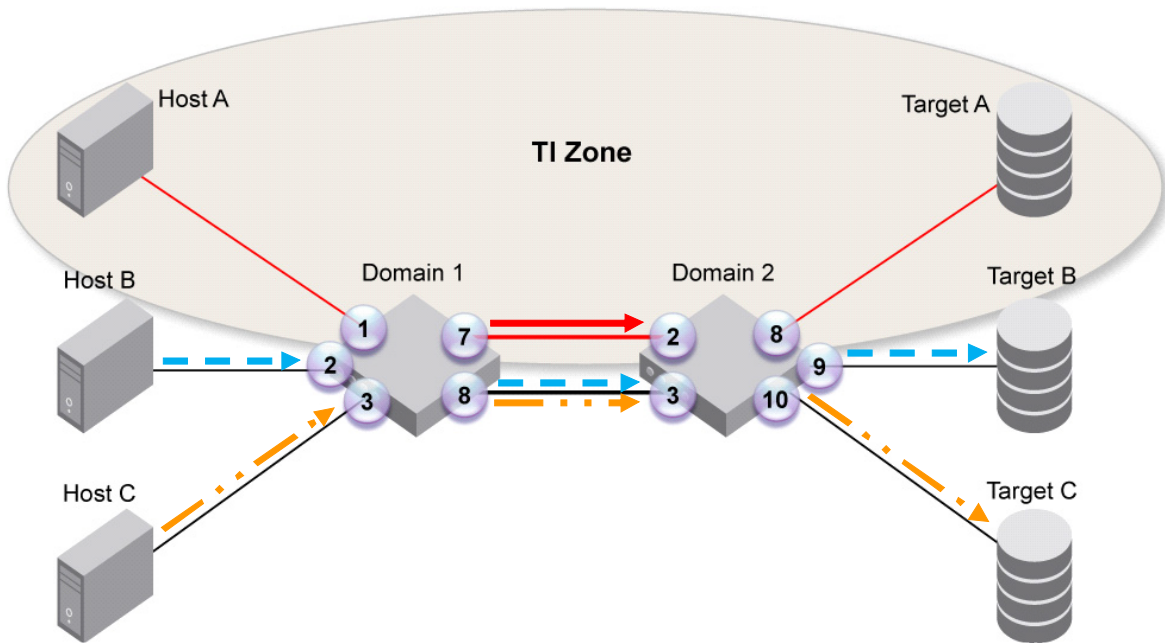


Figure 5: Sample TI Zone

Traffic Isolation zones use a special zoning command, `zone --create`, and are intended to control the routing of frames between zone members, **not** to control access to devices. A normal zone must be in effect granting access between devices before a TI zone will be effective. TI zones will only appear in the *defined* zoning configuration, not in the *effective* zoning configuration. TI zones can **only** be created using D,I (Domain, Index) notation and must include E_Ports and F_Ports in order to create a complete, dedicated, end-to-end route from initiator to target. Ports can only be members of a **single** TI zone.

ISL Trunking

Trunking combines two or more physical ISLs into a single logical link. Trunking goals:

- Reduce individual ISL congestion by providing more bandwidth
- Form a fault-tolerant, high bandwidth logical ISL (called a trunk or trunk group) that withstands the failure of individual ISLs
- Greater resiliency in the fabric

Trunk group characteristics:

- Frames are multiplexed across ISLs in the trunk group
- One port in the trunk group represents the link in the routing data base
- ASICs preserve in-order delivery

Trunking license required for all switches participating in trunking

- Available when the license is installed and ports are re-initialized

Trunking is enabled by default:

- If previously disabled, it must be re-enabled (`portcfgtrunkport`) on the trunk ports
 - Trunk ports must operate at a common port speed
 - Trunk ports must originate and end in a valid port group
 - When trunking criteria is met, the trunk forms automatically




Figure 6: B-Series ISL Trunking

Monitoring and Managing Fabrics

NPIV

- N_Port ID Virtualization (NPIV) enables a single Fibre Channel port to appear as multiple, distinct ports
- Provides separate port identification within the fabric for each operating system image behind the port (as if each operating system image had its own unique physical port)




Figure 7: NPIV Example

- With NPIV support:
 - Each NPIV-enabled port on the switch can support up to 255 unique World Wide Port Names (WWPNs) for use by virtual and physical servers
 - Fabric switches can assign a unique 24-bit address to each virtual server as they login to the fabric
 - Standard fabric zoning and storage LUN masking can be used with virtual machines to isolate storage ports and LUNs to the appropriate virtual server just as they are with physical servers
- Available on all 4 and 8 Gbit/sec switches
- Enabled by default and configurable on a per-port basis

Virtual Fabrics

- Allows a physical switch to be divided up into multiple logical switches
- Logical switches can be interconnected to form logical fabrics
- Each Logical Fibre Channel Switch acts as an independent Fabric Element (Logic Switch) in terms of protocol and management

The benefits include an architecture to virtualize hardware boundaries, allowing for multiple fabrics within the same hardware. This dictates increased flexibility. A switch must be rebooted after Virtual Fabrics have been enabled.

Allows movement of ports from fabric to fabric without re-wire

- Allows easy assignment of unused ports as needed
- Limited support for performance monitoring however
- Having multiple small fabrics on the same physical hardware enables SAN management to be more effective
- Allows bandwidth sharing for better utilization or traffic isolation between physical switches
- Two FICON Virtual Fabrics are supported in a chassis

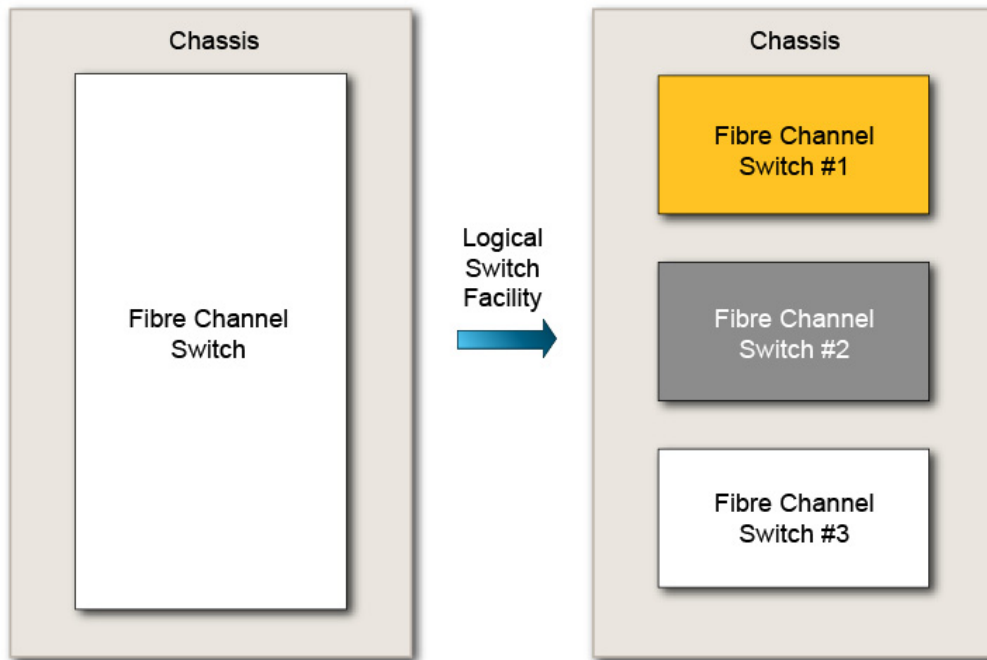


Figure 8: Virtual Fabric Example with Logical Switches

Logical Switches

- A collection of zero or more user ports that act as a single Fibre Channel switch
 - Ports can be F/F/E/VE/EX/VEX
- A Logical Switch is a complete and self-contained FC switch
 - Fabric services (Name Server, Zoning...)
 - Configuration (port, switch, fabric)
 - Fabric characteristics (operating mode, addressing etc)
- Logical switches can be created and deleted without impact to other logical switches
 - If performance monitors are being used, reassigned ports will need new monitors
- Only one version of Fabric OS per chassis (physical switch/Director/Backbone)

Each Logical Switch has its own independent address space (FID). There could be multiple domain IDs in the chassis. The exact same PID can exist in two different Logical switches/fabrics, in the same chassis. All fabric services are architected to manage independent data sets for each logical switch/logical fabric.

- Three types of Brocade Logical Switches:
 - Default Switch (DS)
 - o Initially contains all the ports in the chassis
 - o Default FID is 128
 - o Supports F/FL/E/VE Ports
 - Logical Switch (LS)
 - o Standard Logical Switch
 - o Supports F/FL/E/VE Ports
 - Base Switch (BS)
 - o Used for IFL connections between logical fabrics
 - o Used for ISL connections between logical switches
 - o Supports E/EX/VEX ports **only**

The 256-area addressing mode is available only in a logical switch on the Brocade DCX and DCX-4S platforms. In this mode, only 256 ports are supported and each port receives a unique 8-bit area address. This mode can be used in FICON environments, which have strict requirements for 8-bit area FC addresses.

There are two types of area assignment modes in the 256-area addressing mode: zero-based and port-based. Zero-based mode which assigns areas as ports are added to the partition, beginning at area 0x00. This mode allows FICON customers to make use of the upper ports of a 48-port blade; but this mode may not be compatible with domain,index zoning in InteropMode 2, because M-EOS switches are not capable of handling indexes greater than 255. In the 256-area mode, both zero-based and port-based modes, you can assign from the entire range of 0x0000 - 0xff00 for the PID.

Using DCFM, a port can only be assigned to one logical switch. If you assign a port already assigned to a logical switch, it is removed from the original logical switch and assigned to the new logical switch. To assign ports to a logical switch, perform the following steps.

1. Select **Configure > 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. Select how to display ports on this dialog box from the **Port Display** list

You can display ports by user port number, port address, or slot and port

4. Select the ports you want to include in the logical switch from the **Ports** table

5. Right-click anywhere in the **Existing Logical Switches** table and select **Table > Expand All**

6. Select the logical switch in the **Existing Logical Switches** table

7. Click the right arrow button

The ports display in the selected logical switch node in the **Existing Logical Switches** table.

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

9. Select the **Re-Enable ports after moving them** check box

10. Click **Start** to send these changes to the affected chassis

The status of each change is displayed in the Status column and Status area

11. When the changes are complete, click **Close**

ICLs

- Inter-Chassis Link (ICL) - A licensed feature that provides specialized grouping of ISLs that connects two DCX or two DCX-4S Backbones using special cabling
- When adding an ICL license, all ICL ports must be disabled, and then re-enabled for the license to take effect
- All port parameters associated with ICL ports are static and all portcfg commands are blocked from changing any of the ICL port parameters
- Addition or failure of ICLs is exactly same as ISLs (non-disruptive, no RSCN)
- If using Virtual Fabrics, ICL ports can only be in the base or default switch. If XISL is turned off, you can connect ICLs to other logical switches
- When connecting ICL cables:
 - ICL0 (bottom port) must be connected to ICL1 (top port) on the other DCX
 - ICL1 (top port) must be connected to ICL0 (bottom port) on the other DCX1
 - Cables can be cross connected to the other slot. Example: Cables connected to slot 5 on one DCX can be connected to slot 8 on the other DCX



Figure 9: DCX with ICLs

Access Gateway F_Port Trunking

On edge switches, a masterless trunking feature is called F_Port Trunking because it trunks F_Ports; on the switches running in Access Gateway mode, this feature is called Access Gateway N_Port trunking because it trunks the N_Ports. All switches need Trunking licenses installed, and must be running Fabric OS v6.2 or later.

F_Port Trunking provides a trunk group between a switch (N_Port) in Access Gateway (AG) mode and Condor-based platforms. This feature keeps F_Ports from becoming disabled when they are mapped to an N_Port on a switch in Access Gateway mode. With F_Port trunking, any link within a trunk can go offline or become disabled, but the trunk remains fully functional and there are no reconfiguration requirements. This requires that all ports in the F_Port trunk be in the same trunk port group.

5 - SAN Configuration and Maintenance

Routine Maintenance

QoS Zones

- QoS enables the setting of low and high priorities between specific hosts and targets
- Prioritization is accomplished by the use of QoS zones, which will appear as normal zones
- All normal zoning rules apply
- To distinguish QoS zones from normal zones, special prefixes are used in the zone names:
 - QOSH_ to set high priority
 - QOSL_ to set low priority
- Default setting is medium priority and is used when no QoS zones are specified or when QoS cannot be enforced
- Zones must be created using WWN notation
- Virtual Fabrics can prioritize flows between devices in a logical fabric. The priority is retained for traffic going across ISLs and through the base fabric XISLs.
- The QoS feature only comes into play if there is contention on the link. If there is no congestion on the link QoS will not engage.




Figure 10: QoS Zones

Brocade Branded USB Device

The Brocade 300, 5100, 5300 switches and the Brocade DCX and DCX-4S Backbones support a Brocade branded USB device, used for:

- Uploading a `supportsave` capture
- Uploading or downloading a configuration file
- Downloading firmware

In order to do a firmware download from a Brocade branded USB device, it must be attached to the switch or active CP. Before the USB device can be accessed by the `firmwaredownload` command, it must be enabled and mounted as a file system. The firmware images to be downloaded must be stored under the relative path from `/usb/usbstorage/brocade/firmware` or use the absolute path in the USB file system. Multiple images can be stored under this directory. There is a `firmwarekey` directory where the public key signed firmware is stored. When the `firmwaredownload` command line option, `-U` (upper case), is specified, the `firmwaredownload` command downloads the specified firmware image from the USB device. When specifying a path to a firmware image in the USB device, you can only specify the relative path to `/firmware` or the absolute path.

- `usbstorage -e` - this enables the USB device
- `usbstorage -l` - this lists the contents of the file system




Figure 11: Brocade Branded (supported) USB Device

Buffer-to-Buffer Credits

Buffer-to-Buffer (BB) credit flow control is implemented to limit the amount of data a port may send based on the number and size of the frames sent from that port. Buffer credits represent finite physical port memory. Within a fabric, each port may have a different number of BB credits. Within a connection, each side may have a different number of BB credits. One buffer credit allows a device to send one payload up to 2112 bytes (2148 with headers). Assuming that each payload is 2112, you need one credit per 1 km of link length at 2 Gbit/sec (smaller payloads require additional BB credits to maintain link utilization).

You can allocate buffer credit using the `portcfglongdistance` command, which allows you to allocate sufficient numbers of full-size frame buffers on a particular port or to support a long distance link. Only E_Ports can be configured for extended distances. Changes made by this command are persistent across switch reboots and power cycles.

Level 0 static mode (LO) - LO is the normal (default) mode for a port. Each user port reserves eight buffer credits and competes with other ports for additional buffer credits. No buffer credits are reserved for extended distance ISLs.

Level E static mode (LE) - LE reserves a static number of buffer credits that supports distances up to 10 km. The number reserved depends on the port speed. The baseline for the calculation is one credit per km at 2 Gbit/sec. This yields the following values for 10 km:

- 5 credits per port at 1 Gbit/sec
- 10 credits per port at 2 Gbit/sec
- 20 credits per port at 4 Gbit/sec
- 40 credits per port at 8 Gbit/sec

Dynamic Mode (LD) - LD calculates buffer credits based on the distance measured during port initialization. An upper limit is placed on the calculation by providing a desired distance value. If the measured distance is more than desired distance, the desired distance is used in the calculation; otherwise, the measured distance is used. This is a mechanism for controlling the number of reserved buffer credits ensuring buffer availability for other ports in the same group.

Static long-distance mode (LS) - LS calculates a static number of buffer credits based on a desired distance value.

Web Tools – B-Series Element Manager

- Embedded Java-based GUI that enables administrators to monitor and manage single or small fabrics, manage switches and ports
- May be launched directly from a Web browser. It may also be launched from the DCFM client
- Provides administrative control point for optionally licensed Brocade features, including:
 - ISL Trunking, Advanced Performance Monitoring, Fabric Watch, FCRS, FCIP, and a telnet interface
- SSL can be used to enable secure Web Tools access to the fabric
 - SSL certificates are based on the switch IP address or fully qualified domain name
- The Web Tools access portal into a fabric should be through a switch with the latest Fabric OS version

FC-FC Routing

An EX_Port and VEX_Port function similarly to an E_Port and VE_Port respectively, but terminate at the switch and do not propagate fabric services or routing topology information from one edge fabric to another. IFLs are links between E_Ports and EX_Ports, or VE_Ports and VEX_Ports.

Every EX_Port and VEX_Port uses the fabric ID (FID) to identify the fabric at the opposite end of the interfabric link. The FID for every edge fabric must be unique from the perspective of each backbone fabric.

- If EX_Ports and VEX_Ports are attached to the same edge fabric, they must be configured with the same FID
- If EX_Ports and VEX_Ports are attached to different edge fabrics, they must be configured with a unique FID for each edge fabric

If two different backbone fabrics are connected to the same edge fabric, the backbone fabric IDs must be different, but the edge fabric IDs must be the same. If you configure the same fabric ID for two backbone fabrics that are connected to the same edge fabric, a RASLog message displays a warning about fabric ID overlap.




Figure 12: MetaSAN with Edge-to-Edge, Backbone Fabrics and LSAN Zones

FCIP

- FCIP can be used to create fabric connections across distance:
 - VE_Ports (Virtual E_Ports) are used to create 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 as long as all of the merge criteria are met.
 - VEX_Ports (Virtual EX_Ports) enable FC-FC Routing Service functionality over an FCIP tunnel. VEX_Ports use IFLs. If a VEX_Port is on one end of an FCIP tunnel, the fabrics connected by the tunnel are not merged.
- If using FCIP in your FC-FC Routing configuration, you must:
 - First configure FCIP tunnels
 - Once a tunnel is created, it defaults to a disabled state
 - Configure the VEX_Port
 - After the appropriate ports are configured, enable the tunnel
- After a tunnel is online, using the `portcfg fcipunnel [slot/]ge[port] modify` option to change the FCIP tunnel configuration options and parameters, disrupts traffic on the specified FCIP tunnel.
- Tunnels can carry FCR traffic through VEX_Ports

The WAN tool `ipperf` is an option of the Fabric OS `portcmd` command. This option allows you to specify the slot and port information for displaying performance statistics for a pair of ports. For this basic configuration, you can specify the IP addresses of the endpoints, target bandwidth for the path, and optional parameters such as the length of time to run the test and statistic polling interval. Running the command on an active tunnel will share the available bandwidth.

Only a single `ipperf` session can be active on an FCIP GbE port at any time. Each FCIP port supports a single instance of the WAN tool-embedded client running in only sender or receiver mode. You can, however, use multiple CLI sessions to invoke simultaneous `ipperf` sessions on different FCIP ports.

- The end-to-end IP path performance characteristics that are available using `portcmd --ipperf`
 - Bandwidth
 - Loss
 - Delay
 - Path MTU

The Internet Protocol security (IPSec) uses cryptographic security to ensure private, secure communications over Internet Protocol networks. 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. By default, when creating an FCIP tunnel, IPSec is disabled. IPSec works on FCIP tunnels with or without IP compression (IPComp), FCIP Fastwrite, and Tape Pipelining. IPSec can only be created on tunnels using IPv4 addressing. It requires the High-Performance Extension over FCIP/FC license.

Virtual Ports and FCIP Tunnels

Each extension blade and extension switch platform presents 16 FC ports and 16 virtual ports. Each GbE interface can support up to 8 FCIP Tunnels which are represented as 8 virtual ports on ge0 and 8 virtual ports on ge1.

Gbe Port	Tunnels	Virtual Ports
ge0	0	16
	1	17
	2	18
	3	19
	4	20
	5	21
	6	22
	7	23
ge1	0	24
	1	25
	2	26
	3	27
	4	28
	5	29
	6	30
	7	31

Table 2: Tunnel and Virtual Port Numbering

Taking the Test

After the Introduction Screen, once you click on Next, you will see the non-disclosure agreement:

[000-000]

Section Time Remaining 04:51

IMPORTANT: PLEASE READ THE FOLLOWING BROCADE NON-DISCLOSURE CONFIDENTIALITY AGREEMENT CAREFULLY BEFORE TAKING THIS EXAM.

The following Non-Disclosure Confidentiality Agreement (the "Agreement") sets forth the terms and conditions of your use of the exam materials as defined below.

The Disclosure to you of this Exam and any questions, answers, worksheets, computations, drawings, diagrams, or any communications, including verbal communication by any party, regarding or related to the Exam and such Exam Materials and any derivatives thereof is subject to the Terms and Conditions of this Agreement.

You understand, acknowledge and agree:

- That the questions and answers of the Exam are the exclusive and confidential property of Brocade and are protected by Brocade intellectual property rights;
- That you may not disclose the Exam questions or answers or discuss any of the content of the Exam Materials with any person, without prior approval from Brocade;
- Not to copy or attempt to make copies (written, photocopied, or otherwise) of any Exam Material, including, without limitation, any Exam questions or answers;
- Not to sell, license, distribute, or give away the Exam Materials, questions, or answers.
- You have not purchased, solicited or used unauthorized (non-Brocade sanctioned) Exam Materials, questions, or answers in preparation for this exam.
- That your obligations under this Agreement shall continue in effect after the Exam and, if applicable, after termination of your Certification, regardless of the reason or reasons for terminations, and whether such termination is voluntary or involuntary.

Brocade reserves the right to take all appropriate actions to remedy or prevent disclosure or misuse, including, without limitation, obtaining an immediate injunction. Brocade reserves the right to validate all results and take any appropriate actions as needed. Neither this Agreement nor any right granted hereunder shall be assignable or otherwise transferable by you.

By clicking on the "A" button ("YES, I AGREE"), you are consenting to be bound by the terms and conditions of this agreement and state that you have read this agreement carefully and you understand and accept the obligations which it imposes without reservation. You further state that no promises or representations have been made to induce agreement and that you accept this agreement voluntarily and freely.

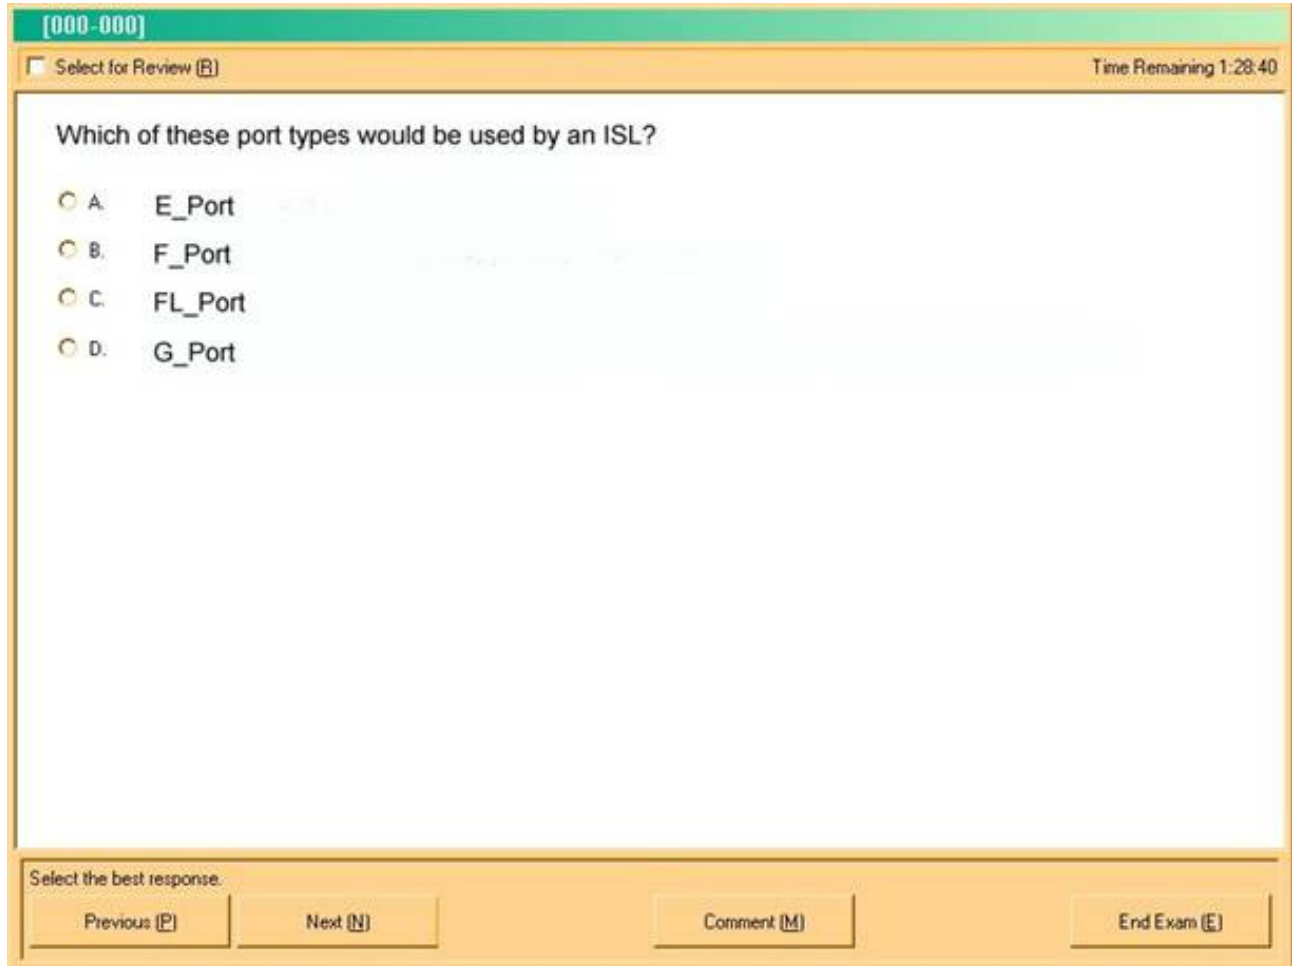
A. YES, I AGREE
 B. NO, I DO NOT AGREE

Please select one of the two options (Yes or No). agree (Scored) Item 1

Previous [P]
Next [N]
Comment [M]
End Exam [E]

Figure 13: Sample NDA

Once you agree to the non-disclosure terms, the timed exam will begin. This is a sample of how the questions will look. In this example, you see a multiple-choice question.




The screenshot shows a software interface for an exam. At the top left, there is a green header with the text "[000-000]". Below this, a light blue bar contains a checkbox labeled "Select for Review (R)" on the left and "Time Remaining 1:28:40" on the right. The main content area is white and contains the question: "Which of these port types would be used by an ISL?". Below the question are four radio button options: "A. E_Port", "B. F_Port", "C. FL_Port", and "D. G_Port". At the bottom of the interface, there is a light blue bar with the text "Select the best response:" on the left. To the right of this text are four buttons: "Previous (P)", "Next (N)", "Comment (M)", and "End Exam (E)".

Figure 14: Sample Question

This is a sample of the score sheet you will see at the end of the exam. You also see the breakdown of how many questions there are in each section of the exam. A hard copy of this will be printed at the testing center. It is **vital** that you obtain and save this hard copy as proof and validation.

[000-000]



Brocade Certified SAN Manager 8 Gbit/sec
Examination Score Report

CANDIDATE:
CANDIDATE ID: 2695281
REGISTRATION NUMBER: 2695281 **DATE:** 14-Aug-2009
VALIDATION NUMBER: 749197037 **SITE:** 1
EXAM: Brocade Certified SAN Manager 8 Gbit/sec
SERIES: 000-000

PASSING SCORE: 60% **YOUR SCORE:** 92% **GRADE:** Pass

SECTION NAME:	SECTION SCORE:
Security	88%
Management Tools	100%
Troubleshooting	80%
SAN Management and Monitoring	94%
SAN Configuration and Maintenance	100%

Brocade thanks you for taking the exam.
 Candidates with passing scores will receive their Certification Packages within 6 weeks of the date of the exam.

Pearson VUE's digital embossing process is available online at: <http://www.pearsonvue.com/authenticate>. Digital embossing will allow you to verify that this score report is authentic. Digital embossing preserves the integrity of this testing program, maintains the value of your certification, and eliminates the possibility of unauthorized embossing of counterfeit score reports.

Previous [P]

Next [N]

Comment [M]

BRSCRPT (Unanswered)
End Exam [E]

Figure 15: Sample Exam Score Report

