BCFP in a Nutshell Study Guide for Exam 143-080

Brocade University

Revision 1111

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

© 2011 Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, the Brocade B-weave logo, Fabric OS, File Lifecycle Manager, MyView, Secure Fabric OS, SilkWorm, and StorageX are registered trademarks and the Brocade B-wing symbol and Tapestry are trademarks of Brocade Communications Systems, Inc., in the United States and/ or in other countries. FICON is a registered trademark of IBM Corporation in the U.S. and 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 1111

BCFP in a Nutshell 16 Gbps Edition

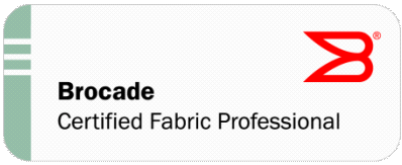

Objective: The BCFP Nutshell guide is designed to help you prepare for the BCFP Certification, exam number 143-080.

Audience: The BCFP Nutshell self-study guide is intended for those who have successfully completed the CFP 300 Brocade Certified Fabric Professional 16 Gbps Training course, and who wish to undertake self-study or review activities before taking the actual BCFP exam. The BCFP guide is not intended as a substitute for classroom training or hands-on time with Brocade products.

How to make the most of the BCFP **guide:** The BCFP guide summarizes the key topics on the BCFP 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. To benefit from the BCFP guide, we strongly recommend you have successfully completed the CFP 300 Brocade Certified Fabric Professional 16 Gbps Training course.

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

Joe Cannata
Certification Manager

©2011 Brocade Communications

Table of Contents

Chapter 1 — Product Features 1
Access Gateway 1
Adaptive Networking 2
Host Bus Adapters 2
Brocade 1860 Fabric Adapters 2
Brocade 415, 425, 815, and 825 Fibre Channel HBAs 2
Brocade 1010 and 1020 CNAs 2
HBA Target Rate Limiting 3
F_Port Trunking 3
FC-FC Routing 4
EX_Port Trunking 4
Brocade Backbone Switches 5
E_Port Trunking 5
Licensed Features 6
Time-Based Licensing 7
Slot-Based Licensing 7
ICL Licensing 7
Chapter 2 — Implementation and Configuration 8
Fibre Channel over IP 8
FC-FC Routing 9
Fabric OS Security 10
Access Gateway 11
Access Gateway Policies 12
Access Gateway Failover 12
Trunking in Access Gateway Mode 12
Chapter 3 — SAN Management 13
Brocade Network Advisor 13
SMI Agent Configuration 14
Client and Server Support Save 14
Performance Data 15
Simple Network Management Protocol (SNMP) 15
Bottleneck Detection 16
Types of Bottlenecks 16
Supported Configurations 17
Advanced Performance Monitoring (APM) 18
Types of Monitors 18
End-to-End Performance Monitors 19
Chapter 4 — Adaptive Networking 20
Adaptive Networking Overview 20
QoS Zones 20

Ingress Rate Limiting 21
Traffic Isolation Zones 21
Enhanced TI Zones *23*
Chapter 5 — FC-FC Routing *24*
FC-FC Routing Terminology 25
LSAN Tagging *26*
Device Sharing Considerations *26*
FC-FC Routing With FCIP *27*
Chapter 6 — Extension Features and Implementation *28*
FCIP 28
FCIP Hardware Offerings *29*
Tunnels *30*
Circuits *30*
FCIP Adaptive Rate Limiting *30*
Long Distance Fabrics 31
Long Distance Modes *31*
Configuring Fill Words *32*
Considerations for Using IR_RDY on TDM Devices *32*
Chapter 7— Advanced Troubleshooting *33*
Diagnostic Port 33
Diagnostic Port Uses *33*
Top Talkers 33
Top Talker Modes *34*
Top Talker Usage *34*
Top Talker Limitations *34*
Miscellaneous Troubleshooting Commands 35
porterrshow *35*
framelog *36*
fabriclog *36*
errshow *36*
portcmd *37*
zone --show tierrors *37*
trunkshow *37*
lsanzoneshow *37*
fcrphydevshow *38*
Taking the Test *39*

©2011 Brocade Communications

List of Figures

Access Gateway Diagram 1
FC-FC Routing Diagram 4
Trunking Port Groups 6
FC-FC Routing Port Types 9
Brocade 300 Access Gateway Default Port Map 11
Traffic Flow Without TI Zones 22
Traffic Flow With TI Zones 22
Enhanced TI Zones Example 23
FC-FC Routing with LSAN Zones 24
Fabric IDs 25
FC-FC Routing Showing Front and Translate Domains 26
Device Sharing Considerations 27
FCIP Example 28
FCIP Tunnel Example 30
FCIP Circuit Example 30

©2011 Brocade Communications

Chapter 1 – Product Features

After reviewing this section be sure you can perform the following:

- Identify advanced product features
- Demonstrate a fundamental knowledge of DCB/FCoE
- Given a scenario, demonstrate knowledge of interopmode
- Given a scenario, demonstrate advanced knowledge of Brocade Trunking
- Describe licensing options and features

Access Gateway

Brocade Access Gateway (AG) is a Fabric OS feature that lets you configure your enterprise fabric to handle additional N_Ports instead of domains.

Switches in AG mode are logically transparent to the host and the fabric. It increases the number of hosts that have access to the fabric without increasing the number of switches in the fabric. This greatly increases the scalability of a fabric and simplifies configuration and management in a large fabric by reducing the number of domain IDs and ports.

- Access Gateway (AG) is a Fabric OS feature that enables:
  - Seamless connectivity to any fabric that supports NPIV
  - Enhanced scalability
  - Simplified manageability
- Designed to connect numerous servers with minimal impact to any existing fabric
  - Allows a switch in AG mode to be added to a fabric without consuming a domain ID
- Focus is connectivity, bandwidth is shared
- Included in the base Fabric OS
  - Needs all available Ports-On-Demand licenses installed to enable AG
- Attached F_Port devices must be Fibre Channel Protocol (FCP) initiators or targets
  - Not supported: loop devices, FICON, virtual iSCSI initiators

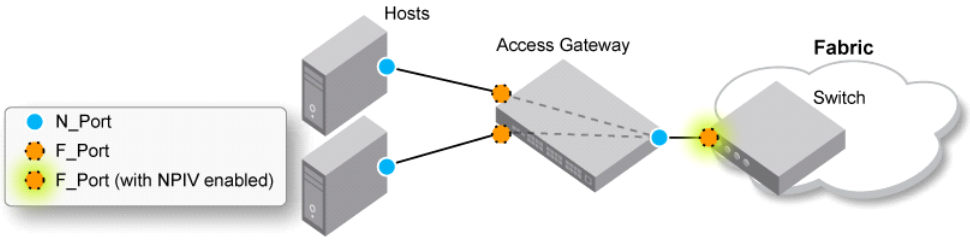

Figure 1: Access Gateway Diagram

Adaptive Networking

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. The Adaptive Networking framework includes the following features:

- Top Talkers
- Traffic Isolation Zoning
- QoS Ingress Rate Limiting
- HBA QoS Target Rate Limiting
- QoS SID/DID Traffic Prioritization
- QoS HBA Traffic Prioritization
- Bottleneck Detection

The Adaptive Networking license only activates these features:

- QoS Ingress Rate Limiting
- QoS SID/DID Traffic Prioritization

Host Bus Adapters

Brocade 1860 Fabric Adapters

Available in single- and dual-port models, the Brocade 1860 features Brocade AnyIO technology, making it the industry's first multiprotocol server connectivity product designed for highly virtualized, cloud-enabled data centers. The Brocade 1860 is a new class of adapter that combines a 16 Gbps Fibre Channel Host Bus Adapter (HBA) and a 10 Gigabit Ethernet (GbE) Converged Network Adapter (CNA) and Network Interface Card (NIC), supporting native Fibre Channel, Data Center Bridging (DCB), Fibre Channel over Ethernet (FCoE), and Internet Small Computer System Interface (iSCSI) running simultaneously in a single product.

Brocade 415, 425, 815, and 825 Fibre Channel HBAs

Available in 4 Gbps and 8 Gbps, single- and dual-port models, the Brocade line of HBAs provide unmatched hardware capabilities and unique software features. This new class of HBAs is designed to help IT organizations deploy and manage end-to-end Storage Area Network (SAN) services across next-generation data centers. Brocade HBAs support NPIV, QoS at the application level, and boot from SAN capabilities.

Brocade 1010 and 1020 CNAs

The Brocade 1010 (single port) and Brocade 1020 (dual port) Converged Network Adapters (CNAs) integrate 10 Gbps Ethernet Network Interface Card (NIC) functionality with Fibre Channel technology—enabling transport over a 10 Gigabit Ethernet (GbE) connection through the new Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) protocols, providing best-in-class LAN connectivity and I/O consolidation to help reduce cost and complexity in next-generation data center environments.

©2011 Brocade Communications

HBA Target Rate Limiting

- Can be enabled on Brocade 16 and 8 Gbps capable HBAs
  - Not supported on CNAs
  - Licensed feature
- Managed only through the Brocade Configuration Utility (BCU)
  - Host-based CLI
  - Not configured through the switch
- Used to minimize congestion at the HBA port due to a slow drain device operating in the fabric at a slower speed
- Traffic destined to the remote port is limited to its current operating speed
- The default rate limit is 1 Gbps

F_Port Trunking

The FC trunking feature works in conjunction with the trunking feature on Brocade switches, whereby the Fabric Operating System provides a mechanism to trunk different switch ports of the same port group into one routing address to aggregate available bandwidth. When FC trunking is enabled, two physical ports belonging to the same Brocade dual-port HBA are trunked together to form a logical Fibre Channel port. Both HBA ports must be operating at the same speed while in trunk mode.

The following licenses must be installed on the switch connected to the HBA port:

- Server Application Optimization (SAO) license
- Trunking license

You can configure F_Port trunking in the following scenarios:

- Between F_Ports on a Fabric OS switch and N_Ports on an Access Gateway module.
- Between F_Ports on a Fabric OS switch and N_Ports on a Brocade adapter.

Before enabling trunking, consider the following requirements:

- When trunking is enabled, a trunked logical port (Port 0) is created and reported per adapter. Most BCU commands are applicable in this logical port's context only.
- Before enabling trunking on the adapter, you must first enable trunking on the connected switch and assign the trunk area. Follow the steps in “Enabling FC trunking on Brocade switches and adapters” in the Fabric OS Administrators Guide for configuring trunking on the switch and adapter.
- Both adapter ports should be connected to the same port group on the switch.
- Only two ports on the same HBA can participate in trunking and both ports should be operating at the same speed.
- FC Trunking is supported on the dual-port cards only.
- FC Trunking is supported on 8 Gbps ports installed on switch models using Fabric OS v6.4.1 or later.

FC-FC Routing

Fabric OS provides Layer 3 Fibre Channel-to-Fibre Channel routing (FC-FC Routing) between fabrics and allows device access between two or more fabrics without merging the fabrics.

FC-FC Routing is supported between the following fabric types:

- Fabric OS-to-Fabric OS
- Fabric OS-to-M-EOS
- M-EOS-to-M-EOS

Physical connectivity is accomplished through the use of a Fibre Channel router (FC router). Logical connectivity is accomplished through the use of Logical Storage Area Networks (LSANs), by creating uniquely named zones called “LSAN zones”

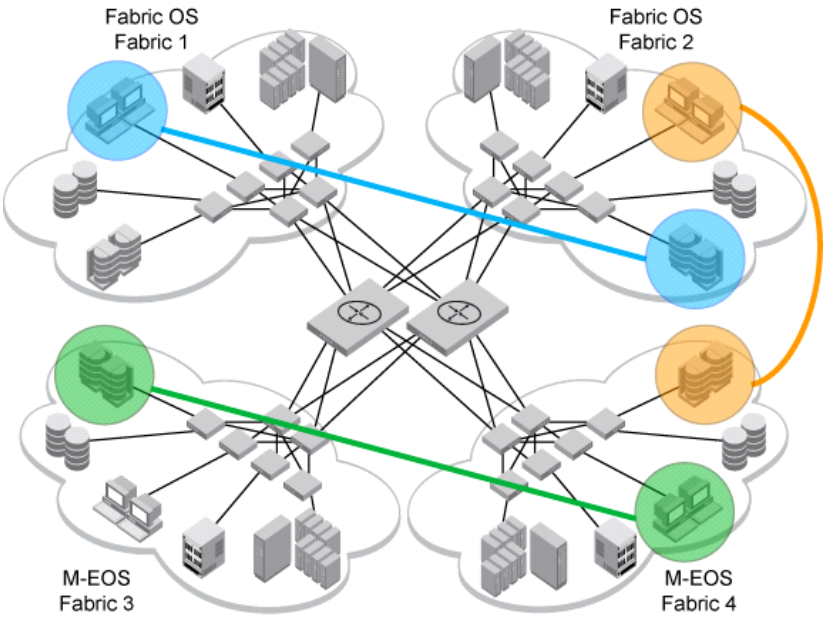

Figure 2: FC-FC Routing Diagram

EX_Port Trunking

You can configure EX_Ports to use trunking just as you do regular E_Ports. EX_Port trunking support is designed to provide the best utilization and balance of frames transmitted on each link between the FC router and the edge fabric. You should trunk all ports connected to the same edge fabrics.

EX_Port trunking is masterless except for EX_Ports on enterprise-class platforms. For the enterprise-class platforms, Virtual Fabrics must be enabled for masterless EX_Port trunking to take effect. For the fixed-port switches, Virtual Fabrics can be enabled or disabled. If masterless EX_Port trunking is not in effect and the master port goes offline, the entire EX_Port-based trunk re-forms and is taken offline for a short period of time. If there are no other links to the edge fabric from the backbone, the master port going offline may cause a traffic disruption in the backbone.

©2011 Brocade Communications

Brocade Backbone Switches

ICL trunking is configured on an inter-chassis link (ICL) between two enterprise-class platforms and is applicable only to ports on the core blades. ICL trunks automatically form on the ICLs when you install the Trunking license on each platform.

The Brocade DCX 8510-8 has 4 port groups on the CR16-8 core blade. The Brocade DCX 8510-4 has two port groups on the CR16-4 core blade. Each port group has four QSFP connectors, and each QSFP connector maps to four user ports. Each of the four user ports in a QSFP terminates on a different ASIC, so a trunk cannot be formed among these ports. To establish ICL trunking between platforms in the Brocade DCX 8510 family, follow these configuration rules:

- You need at least two ICLs between the platforms. A single ICL does not enable trunking.
- Each QSFP has four ports. However, these ports cannot form a trunk with each other, but can form trunks only with corresponding ports on another QSFP.
- You can have a maximum of four ports in an ICL trunk.
- You can have a maximum of eight 4-port trunks to a neighboring domain. Each core blade can have a maximum of four ICLs to a neighboring domain.
- The QSFP cables must be connected to the same trunk group, or cage, on each individual platform.

E_Port Trunking

The trunking feature optimizes the use of bandwidth by allowing a group of links to merge into a single logical link, called a trunk group. Traffic is distributed dynamically and in order over this trunk group, achieving greater performance with fewer links. Within the trunk group, multiple physical ports appear as a single port, thus simplifying management. Trunking also improves system reliability by maintaining in-order delivery of data and avoiding I/O retries if one link within the trunk group fails.

Trunking is frame-based instead of exchange-based. Since a frame is much smaller than an exchange, this means that frame-based trunks are more granular and better balanced than exchange-based trunks and provide maximum utilization of links.

Trunking is supported in the following configurations:

- Trunk links can be 2 Gbps, 4 Gbps, 8 Gbps, 10 Gbps, or 16 Gbps depending on the Brocade platform.
- The maximum number of ports per trunk and trunks per switch depends on the Brocade platform.
- You can have up to eight ports in one trunk group to create high performance ISL trunks between switches with up to 128 Gbps (based on 16 Gbps port speed).
- If in-flight encryption/compression is enabled, you can have a maximum of only two ports per trunk.
- An E_Port or EX_Port trunk can include up to eight ports from an octet. All the ports must be adjacent to each other using the clearly marked groups on the front of the product.

The following requirements apply to all types of trunking:

- The Trunking license must be installed on every switch that participates in trunking.
- All of the ports in a trunk group must belong to the same port group.
- All of the ports in a trunk group must be running at the same speed.
- All of the ports in a trunk group must be configured for the same distance.
- All of the ports in a trunk group must have the same encryption, compression, QoS, and Forward Error Correction (FEC) settings.

- Trunk groups must be between Brocade switches (or Brocade adapters, in the case of F_Port trunking). Brocade trunking is proprietary and not supported on M-EOS or third-party switches.
- There must be a direct connection between participating switches.
- Trunking cannot be done if ports are in ISL R_RDY mode. (You can disable this mode using the `portcfgislmode` command.)
- Trunking is supported only on FC ports. Virtual FC ports (VE_ or VEX_Ports) do not support trunking.

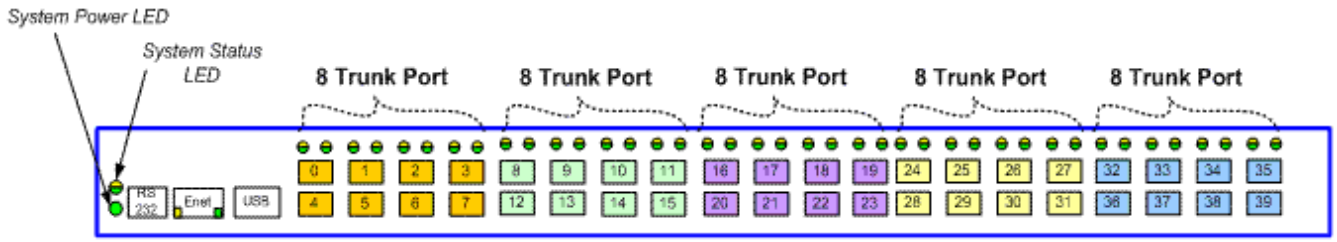

Figure 3: Trunking Port Groups

Licensed Features

Feature licenses are often part of the licensed paper pack supplied with your switch software; if not, they can be purchased separately from your switch vendor, who provides the transaction keys to activate the associated feature or features. Each product, each feature, and each individual switch within a fabric requires its own license key.

Licences might be associated with a feature version. If a feature has a version-based license, that license is valid only for a particular version of the feature. If you want a newer version of the feature, you must purchase a new license. If a license is not version-based, then it is valid for all versions of the feature. Likewise, if you downgrade Fabric OS to an earlier version, some licenses associated with specific features of the version you are downgrading might not work.

As of Fabric OS v7.0.0 the following optional licenses are available. For a complete description of each license see the published release notes for the version of Fabric OS you are running.

- Brocade Ports on DemandBrocade Extended Fabrics
- Brocade ISL Trunking
- Brocade Advanced Performance Monitoring
- Brocade Fabric Watch
- High Performance Extension over FCIP/FC
- Brocade Accelerator for FICON
- FICON Management Server
- Adaptive Networking with QoS
- Server Application Optimization
- Integrated Routing
- Encryption Performance Upgrade
- DataFort Compatibility
- Brocade 8000 FC Ports on Demand
- Advanced Extension

©2011 Brocade Communications

- 10GbE FCIP/10G Fibre Channel
- Advanced FICON Acceleration
- 7800 Port Upgrade
- ICL 16-link, or Inter Chassis Links
- ICL 8-Link
- ICL POD License

Time-Based Licensing

A time-based license applies a try-before-you-buy approach to certain features so that you can experience the feature and its capabilities prior to buying the license. Once you have installed the license, you are given a time limit to use the feature. These are the licenses that are time-based:

- 10 Gigabit FCIP/Fibre Channel license
- Advanced Extension
- Advanced FICON Acceleration license
- Adaptive Networking
- Advanced Performance Monitoring
- Fabric
- Fabric Watch
- Extended Fabric
- High Performance Extension over FCIP/FC
- Integrated Routing
- Trunking

Slot-Based Licensing

Slot-based licensing is used on the Brocade DCX 8510 family, DCX, and DCX-4S platforms to support the FX8-24 blade and on the Brocade DCX 8510 family to support the 16 Gbps FC port blades (FC16-24 and FC16-48). License capacity is equal to the number of slots. These licenses allow you to select the slots that the license will enable up to the capacity purchased and to increase the capacity without disrupting slots that already have licensed features running. Each slot-based license key is for a single feature.

ICL Licensing

Brocade ICL links operate between the core blades of the DCX 8510 family of enterprise-class platforms, or between the core blades of the DCX and DCX-4S platforms. Typically, if both core blades are installed then they are active on the DCX and DCX-4S (or DCX 8510 family) enterprise-class platforms.

- ICL ports on core blades of a DCX 8510-8 can be used only with an ICL (1st or 2nd) POD license.
- ICL ports on core blades of a DCX 8510-4 can be used only with an ICL 1st POD license.
- ICL ports on core blades of a DCX can be used only with an ICL 16-link or 8-link license. ICL ports on core blades of a DCX-4S can be used only with an ICL 8-link license.

ICLs can only be connected between like ASICs. DCX to DCX, or DCX-4S or DCX 8510-8 to DCX 8510-8 or 8510-4.

Chapter 2 – Implementation and Configuration

After reviewing this section be sure you can perform the following:

- Demonstrate knowledge of the different types of Brocade connectivity
- Demonstrate knowledge of implementing fabric security
- Describe implementation of Access Gateway features

Fibre Channel over IP

The Fibre Channel-over-IP (FCIP) protocol connects Fibre Channel switches over an IP network. IP packets generated by an FCIP-compliant port navigate the IP network to reach the destination end point. Implementation uses standards-based TCP and it interoperates with regular network equipment.

- An FCIP tunnel is represented in a Brocade fabric as a Virtual E_Port (VE_Port)
  - Just like an E_Port, except underlying transport is IP not FC
- The VE_Port emulates an E_Port on either end of the FCIP tunnel:
  - The FCIP platforms at both ends of the links merge to form a single fabric
  - VE_Ports do not use FC flow control mechanisms (BB credits); they utilize TCP flow control mechanisms
  - VE_Ports do not support FC ISL trunking, but they do support exchange-based routing (Dynamic Path Selection) and DLS (Dynamic Load Sharing)
  - Note that with FCIP Trunking, it is recommended to implement a multiple circuit trunk instead of having multiple VE_Ports to the same fabric
  - VE_Ports can be used as XISLs
- Fabric OS supports FC-FC routing over an FCIP tunnel, creating a Virtual EX_Port (VEX_Port)
  - Allows long-distance FCIP connections with fabric-to-fabric isolation
  - VEX_Ports are no different from EX_Ports, except the underlying transport is IP rather than FC
  - Both EX_Ports and VEX_Ports are supported in the same backbone fabric
- Connectivity rules with VEX_Ports:
  - A VEX_Port connects only to a VE_Port – it may not connect to another VEX_Port
  - There can be multiple VEX-to-VE port connections between a Backbone fabric and an Edge fabric
  - VEX_Ports are not supported with FICON
- FastWrite and OSTP (Open Systems Tape Pipelining) provide a means to mitigate latency thereby increasing performance of certain storage applications such as:
  - Remote Data Recovery (RDR)
  - Backup/Recovery/Archive (BURA)
- FastWrite accelerates SCSI writes over FCIP
- OSTP accelerates SCSI tape reads and writes over FCIP

©2011 Brocade Communications

- FW and OSTP reduce the effects of latency by
  - Initiating data transfer sooner
  - Reducing the number of round trips across the WAN
- FastWrite and OSTP are supported over VE and VEX_Ports

FC-FC Routing

In FC-FC routing EX_Ports in the backbone fabric are connected to E_Ports in the edge fabrics as shown in the figure below.

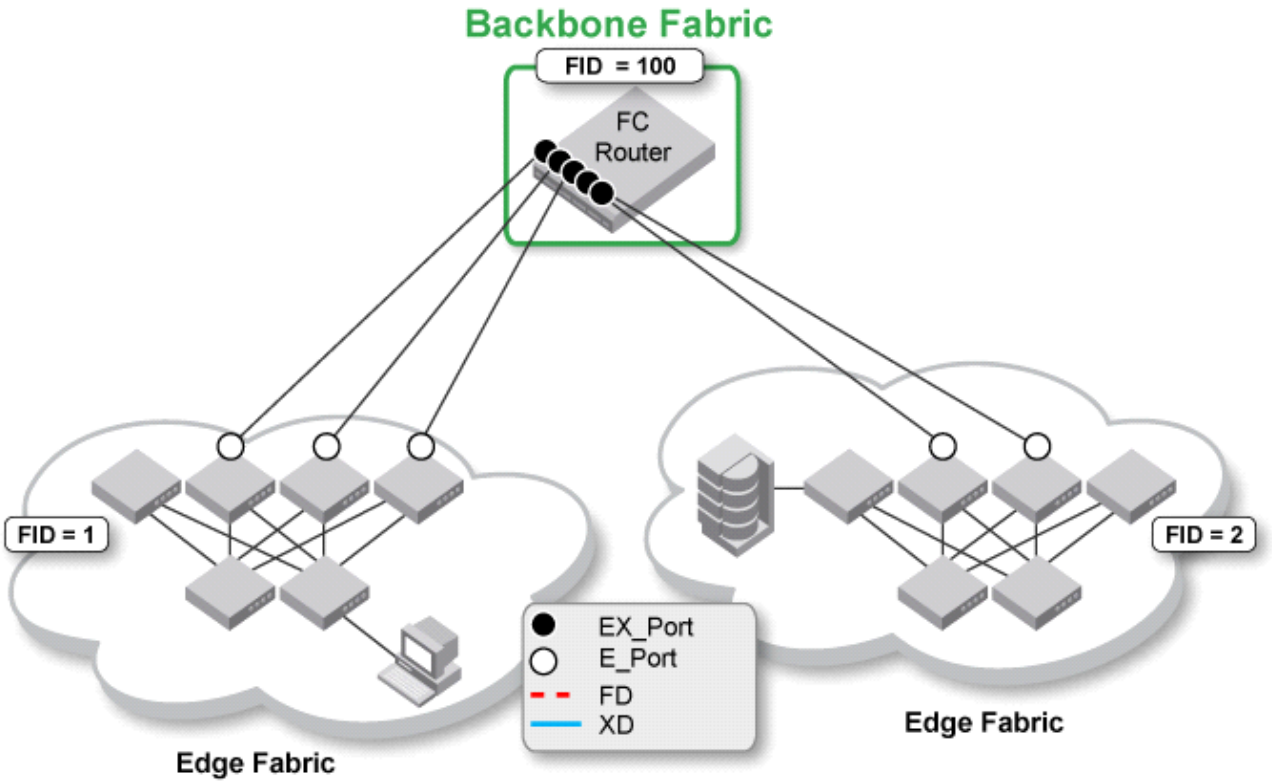

Figure 4: FC-FC Routing Port Types

Fabric OS Security

Fabric OS v7.0.0 supports the following Access Control List (ACL) security policies: (no license required)

- Fabric Configuration Server (FCS): Restricts which switches can be used to change the configuration of the fabric
- Switch Connection Control (SCC): Restricts which switches can join a fabric
- Device Connection Control (DCC): Restricts which Fibre Channel devices can connect to which Fibre Channel switch ports
- Advanced Device Security (ADS): Policy-based Device Connection Control (DCC) used to restrict device access on the Access Gateway
- IP Filter Policy (IPFILTER): Filters IP management interface traffic
- Authentication policy for fabric elements (AUTH): Public Key Infrastructure (PKI)
- Password database and user policy (PWD): Enables users and passwords configured on one switch to be distributed to other switches

©2011 Brocade Communications

Access Gateway

When you first enable a switch for AG mode, the F_Ports are mapped to a set of predefined N_Ports by default. The following figure shows an example port mapping for the Brocade 300. By default, Failover and Failback policies are enabled on all N_Ports.

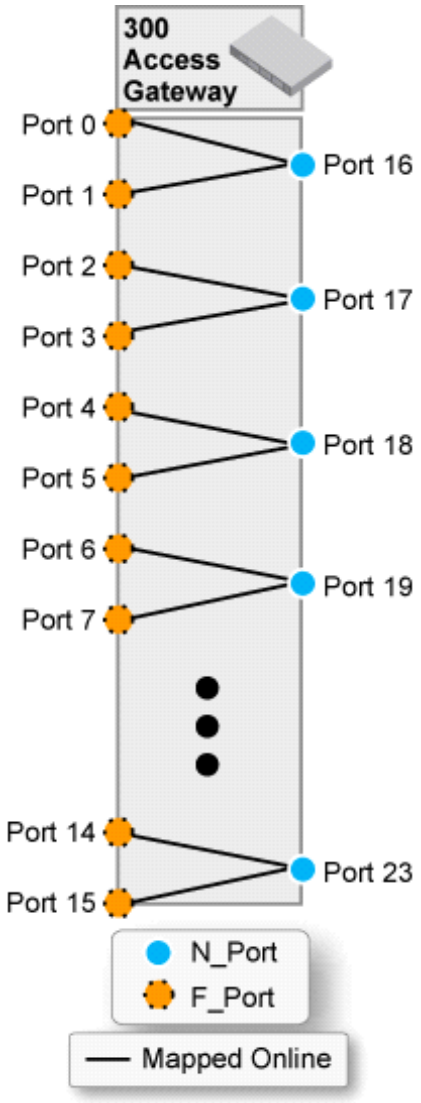

Figure 5: Brocade 300 Access Gateway Default Port Map

Access Gateway Policies

- Port Grouping (PG)
  - Limits N_Port failover to occur only within a user-defined group of ports
  - Enabled by default
- Auto Port Configuration (APC)
  - Automatically detects F_Ports and N_Ports
  - Automatically maps the ports
  - Disabled by default
- Advanced Device Security (ADS)
  - Discussed in the Security eLearning module
  - WWN Load Balance
  - Device logins are sent to the least loaded port in the port group to which they are mapped
- Each policy must be in an enabled state before it can be used

Access Gateway Failover

If an N_Port goes offline, the Access Gateway N_Port failover policy allows hosts to be automatically remapped to another online N_Port

- F_Ports connected to failed N_Port are evenly distributed across N_Ports connected to the same fabric
- F_Ports receive a new FC address based on the new N_Port
- Enforced at N_Port initialization as well (cold failover)
- The default configuration requires all N_Ports to be connected to the same fabric
- Enabled by default; managed on a per-N_Port basis

Trunking in Access Gateway Mode

Trunking in Access Gateway mode provides a trunk group between N_Ports on the AG module and F_Ports on the Edge switch module. With trunking, any link within a trunk group can go offline or become disabled, but the trunk remains fully functional and no reconfiguration is required. Trunking prevents reassignments of the port ID when N_Ports go offline.

Because AG Trunking configuration is mostly on the edge switch, information in this section is applicable to the edge switch module and not the AG module. On the AG module, you only need to ensure that the trunking license is applied and enabled. On the edge switch, you must first configure an F_Port trunk group and statically assign an Area_ID to the trunk group. Assigning a Trunk Area (TA) to a port or trunk group enables F_Port masterless trunking on that port or trunk group. On switches running in Access Gateway mode, the masterless trunking feature trunks N_Ports because these are the only ports that connect to the enterprise fabric. When a TA is assigned to a port or trunk group, the ports will immediately acquire the TA as the area of its port IDs (PIDs). When a TA is removed from a port or trunk group, the port reverts to the default area as its PID.

©2011 Brocade Communications

Chapter 3 — SAN Management

After reviewing this section be sure you can perform the following:

- Demonstrate knowledge of using advanced management interfaces
- Demonstrate knowledge of how to gather performance data
- Describe implementation of Access Gateway features

Brocade Network Advisor

Brocade Network Advisor is the industry's first unified network management solution for data, storage, application delivery, wireless, and converged networks. It supports Fibre Channel SANs, FCoE, IP switching and routing (including Ethernet fabrics), and MPLS networks, providing end–to–end visibility across different network types through a seamless and unified user experience.

TABLE 1 Network Advisor Packages

Package	Versions
SAN with SMI Agent +IP	Licensed Version — Enterprise • SAN — Support for 24 fabrics, 10,000 devices, and 9,000 switch ports • IP — Support for 5,050 devices, 252,500 ports, and 100 MPLS devices
	Professional • SAN — Support for 1 fabric, 1,000 devices, and 1,000 switch ports • IP — Support for 20 devices and 1,000 ports **Note:** SMI Agent is not supported on Professional.
SAN with SMI Agent	Licensed Version — Enterprise • SAN — Support for 24 fabrics, 10,000 devices, and 9,000 switch ports
	Licensed Version — Professional Plus • SAN — Support for 4 fabrics, 4,096 devices, and,2560 switch ports
	Professional • SAN — Support for 1 fabric, 1,000 devices, and 1,000 ports **Note:** SMI Agent is not supported on Professional.
IP	Licensed Version — Enterprise • IP — Support for 5,050 devices, 252,500 ports, and 100 MPLS devices
	Professional • IP — Support for 20 devices and 1,000 ports
SMI Agent only	N/A

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.

Professional edition can discover, but not manage M-EOS devices or 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.

SMI Agent Configuration

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. The SMIA Configuration Tool consists of the following tabs:

- **Home**—enables you to launch the following Management application dialog boxes: Fabric Discovery, Host Discovery, Users, Options, Server, and About.
- **Authentication**—enables you to configure mutual authentication for client, CIMMOM server, and Indication using a secure protocol.
- **CIMOM**—enables you to configure the CIMOM server port, the CIMOM Bind Network Address, and the CIMOM log.
- **Certificate Management**—enables you to import client and Indication certificates, export server certificates, as well as view and delete current certificates.
- **Summary**—enables you to view the CIMOM server configuration and current configuration.

Client and Server Support Save

You can use Technical Support to collect Support Save 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)

Client Support save data includes:

- Client log files
- Client data model log

©2011 Brocade Communications

Performance Data

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 monitoring to indicate the devices that create the most traffic and to identify the ports that are most congested. Performance monitoring allows you to monitor your SAN using the following methods (requires a licensed version of Brocade Network Advisor):

- 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 (FC ports, ISL ports, Device ports, GbE ports, FCIP tunnels, managed HBA ports, managed CNA ports, E_Port trunks, and 10 GbE ports). The Professional version only allows you to monitor your SAN by gathering and displaying real-time performance data (FC ports, ISL ports, device ports, GbE ports, FCIP tunnels, managed HBA ports, managed CNA ports, E_Port trunks, and 10 GbE ports).
- Persist and display historical performance data (FC ports, ISL ports, device ports, FCIP tunnels, and 10 GbE 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 percentage utilization for FC and FCIP links.

Simple Network Management Protocol (SNMP)

The Simple Network Management Protocol (SNMP) is a standard method for monitoring and managing network devices. Using SNMP components, you can program tools to view, browse, and manipulate Brocade switch variables and set up enterprise-level management processes.

You can configure SNMPv3 and SNMPv1 for the automatic transmission of SNMP information to management stations. The configuration process involves configuring the SNMP agent and configuring SNMP traps. Use the `snmpconfig` command to configure the SNMP agent and traps for SNMPv3 or SNMPv1 configurations, and the security level. You can specify no security, authentication only, or authentication and privacy.

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 enables you to do the following:

- Prevent degradation of throughput in the fabric.
  The Bottleneck Detection feature alerts you to the existence and locations of devices that are causing latency. If you receive alerts for one or more F_Ports, use the CLI to check whether these F_Ports have a history of bottlenecks.
- Reduce the time it takes to troubleshoot network problems.
  If you notice one or more applications slowing down, you can determine whether any latency devices are attached to the fabric and where. You can use the CLI to display a history of bottleneck conditions on a port. If the CLI shows above-threshold bottleneck severity, you can narrow the problem down to device latency rather than problems in the fabric.

You can use the Bottleneck Detection feature with other Adaptive Networking features to optimize the performance of your fabric. For example, you can do the following:

- If the Bottleneck Detection feature detects a latency bottleneck, you can use TI zones or QoS SID/DID traffic prioritization to isolate latency device traffic from high priority application traffic.
- If the Bottleneck Detection feature detects ISL congestion, you can use ingress rate limiting to slow down low priority application traffic, if it is contributing to the congestion.

You configure Bottleneck Detection on a per-switch basis, with optional per-port exclusions. Bottleneck Detection does not require a license.

Types of Bottlenecks

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.

By default, Bottleneck Detection detects latency bottlenecks that are severe enough that they cause 98% loss of throughput. This default value can be modified to a different percentage.

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.

©2011 Brocade Communications

Supported Configurations

- Bottleneck Detection is supported only on Fibre Channel ports and FCoE F_Ports
- Bottleneck Detection is supported only 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, 8, and 16 Gbps platforms, including 10 Gbps speeds
- 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

Advanced Performance Monitoring (APM)

Advanced Performance Monitoring is a licensed feature that provides a comprehensive tool for monitoring the performance of networked storage resources. Additional performance monitoring features, such as CRC error reports, are provided through Web Tools and Brocade Network Advisor.

Advanced Performance Monitor commands are available only to users with admin permissions. Use the `perfhelp` command to display a list of commands associated with Advanced Performance Monitoring.

Types of Monitors

Advanced Performance Monitoring provides the following monitors:

- End-to-End monitors (EE monitors) measure the traffic between a host/target pair.
- Frame monitors measure the traffic transmitted through a port with specific values in the first 64 bytes of the frame.
- Top Talker monitors measure the flows that are major consumers of bandwidth on a switch or port.
  - Port mode Top Talker
    - A Top Talker monitor can be installed on a port to measure the traffic originating from the port and flowing to different destinations.
    - You can configure Top Talker monitors on F_Ports and, depending on the switch model, on E_Ports. The following platforms support Top Talker monitors on E_Ports:
    - Brocade 6510
    - Brocade DCX 8510 family
  - Fabric mode Top Talker
    - In fabric mode, Top Talker monitors are installed on all E_Ports in the fabric and measure the data rate of all the possible flows in the fabric (ingress E_Port traffic only). In fabric mode, Top Talker monitors can determine the top *n* bandwidth users on a given switch.

Note

When Top Talkers is installed in fabric mode all end-to-end monitors must be removed from all switches.

©2011 Brocade Communications

End-to-End Performance Monitors

Use end-to-end monitoring when you want to monitor throughput between a pair of devices. End-to-end performance monitoring counts the number of words in Fibre Channel frames for a specified Source ID (SID) and Destination ID (DID) pair.

To enable end-to-end performance monitoring, you must configure an EE monitor on a port, specifying the SID-DID pair (in hexadecimal). The monitor counts only those frames with matching SID and DID.

When end-to-end monitoring is configured on a trunk it is associated with the trunk master port. If the trunk master should go offline the end-to-end monitor will automatically move to the newly designated master.

Note

Monitor data storage is limited to prevent exceeding the flash memory on the switch.

Chapter 4 — Adaptive Networking

After reviewing this section be sure you can perform the following:

- Given a scenario, describe how to implement Adaptive Networking in a fabric

Adaptive Networking Overview

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.

The Adaptive Networking suite includes the following features:

- Bottleneck Detection: The Bottleneck Detection feature identifies devices attached to the fabric that are slowing down traffic. Bottleneck Detection does not require a license.
- Top Talkers: The Top Talkers feature provides real-time information about the top "n" bandwidth-consuming flows passing through a specific port in the network. Top Talkers requires an Advanced Performance Monitoring license.
- Traffic Isolation Zoning: Traffic Isolation Zoning (TI zoning) allows you to control the flow of inter-switch traffic by creating a dedicated path for traffic flowing from a specific set of source ports (F_Ports). Traffic Isolation Zoning does not require a license.
- Ingress Rate Limiting: Ingress rate limiting restricts the speed of traffic from a particular device to the switch port. Ingress rate limiting requires an Adaptive Networking license.
- Quality of Service (QoS) SID/DID Traffic Prioritization: SID/DID traffic prioritization allows you to categorize the traffic flow between a host and target has having a high or low priority. QoS SID/DID traffic prioritization requires an Adaptive Networking license for 8 Gbps platforms, but does not require a license for 4 Gbps platforms.

QoS Zones

You assign high or low priority (QoS level) using a QoS zone. A QoS zone is a special zone that indicates the priority of the traffic flow between a given host/target pair. The members of a QoS zone are the host/target pairs. QoS zones can contain WWN members (WWNN or WWPN) or Domain, Index (D,I) members.

A QoS zone has a special name to differentiate it from a regular zone. The format of the QoS zone name is as follows:

- For high priority: QOSH*id_xxxxx*
- For low priority: QOSL*id_xxxxx*
- The default priority for all traffic is medium, this does not have to be specified

Where id is a flow identifier that designates a specific virtual channel for the traffic flow and *xxxxx* is the user-defined portion of the name. For example, the following are valid QoS zone names:

- QOSH3_HighPriorityTraffic
- QOSL1_LowPriorityZone

It is important to note that traffic prioritization is only enforced on the egress port and not on the ingress port. QoS only takes affect where there is contention on the link and takes priority over the Ingress Rate Limiting (IRL) feature.

 ©2011 Brocade Communications

Ingress Rate Limiting

IRL is a licensed feature that requires the Adaptive Networking license. IRL restricts the speed of traffic from a particular device to the switch port. Use ingress rate limiting for the following situations:

- To reduce existing congestion in the network or proactively avoid congestion.
- To enable you to offer flexible bandwidth limit services based on requirements.
- To enable more important devices to use the network bandwidth during specific services, such as network backup.

IRL enforcement is needed only if the port can run at a speed higher than the rate limit. For example, if the rate limit is 4 Gbps and the port is only a 2 Gbps port, then IRL is not enforced.

When enabling IRL you must specify a rate limit for the port. Additionally enabling IRL on a port is a non-disruptive operation, there is no need to disable the port.

Traffic Isolation Zones

The Traffic Isolation Zoning feature allows you to control the flow of inter-switch traffic by creating a dedicated path for traffic flowing from a specific set of source ports (N_Ports). For example, you might use Traffic Isolation Zoning for the following scenarios:

- To dedicate an ISL to high priority, host-to-target traffic.
- To force high volume, low priority traffic onto a given ISL to limit the effect on the fabric of this high traffic pattern.
- To ensure that requests and responses of FCIP-based applications such as tape pipelining use the same VE_Port tunnel across a metaSAN.

Traffic isolation is implemented using a special zone, called a Traffic Isolation zone (TI zone). A TI zone indicates the set of N_Ports and E_Ports to be used for a specific traffic flow. When a TI zone is activated, the fabric attempts to isolate all inter-switch traffic entering from a member of the zone 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 does not require a license.

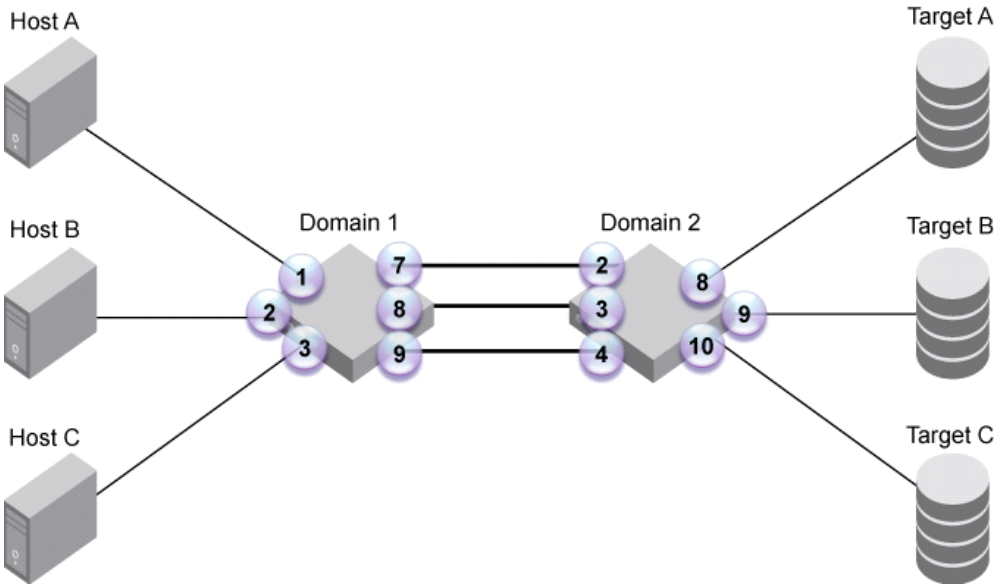

Figure 6: Traffic Flow Without TI Zones

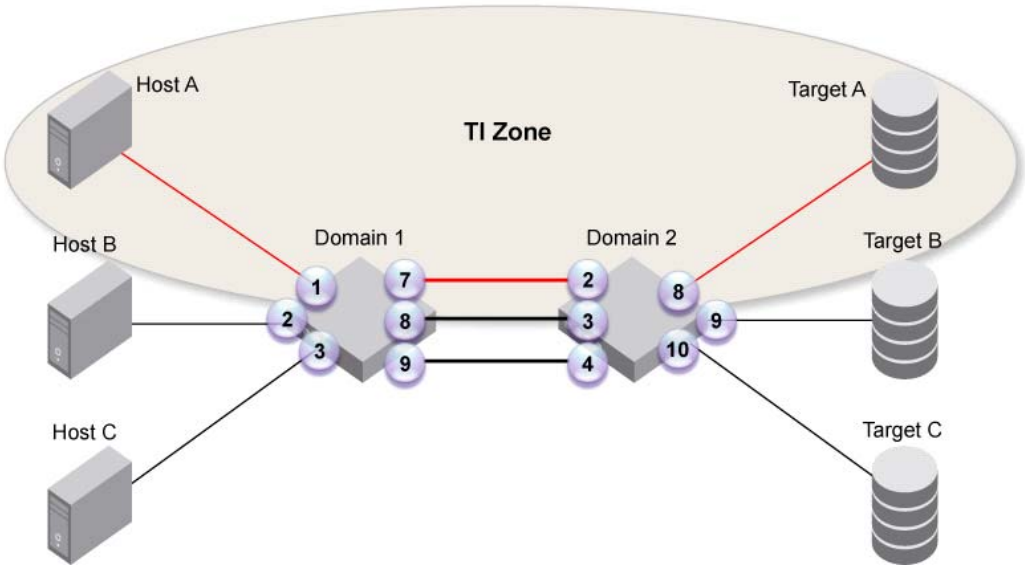

Figure 7: Traffic Flow With TI Zones

As shown in Figure 7, TI zones must include E_Ports and F_ and FL_Ports in order to create a complete, dedicated, end-to-end route from initiator to target. Both VE_Ports and VEX_Ports are supported in TI zones.

©2011 Brocade Communications

Enhanced TI Zones

Prior to Fabric OS v6.4.0, a port could be in only one TI zone at a time. Starting in Fabric OS v6.4.0, ports can be in multiple TI zones at the same time. Zones with overlapping port members are called enhanced TI zones.

Figure 8 shows an example of two TI zones. Because these TI zones have an overlapping port (3,8), they are enhanced TI zones.

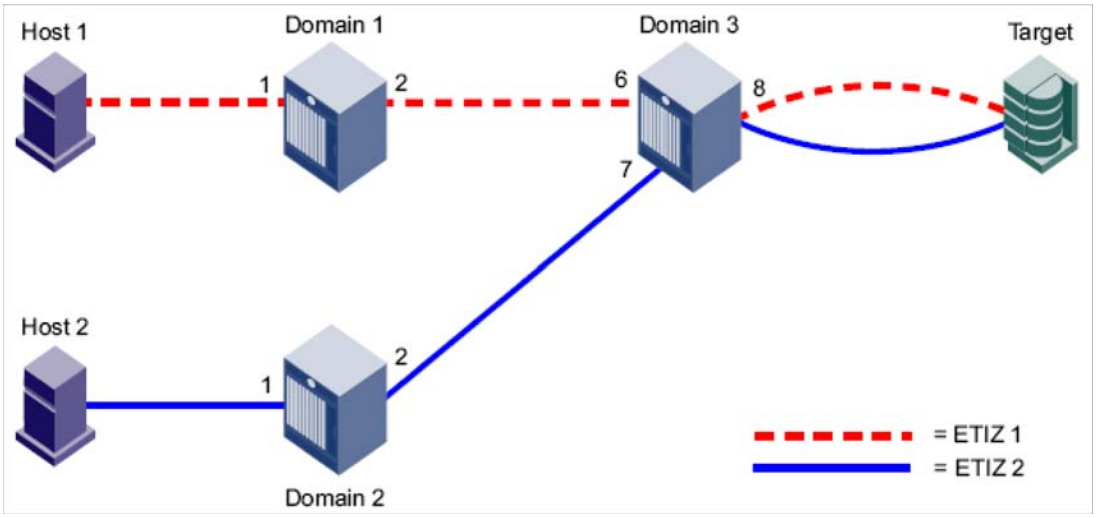

Figure 8: Enhanced TI Zones Example

Chapter 5 – FC-FC Routing

After reviewing this section be sure you can perform the following:

- Describe Fibre Channel routed fabrics
- Demonstrate knowledge of how to implement FC-FC routing

The FC-FC routing service provides Fibre Channel routing (FCR) between two or more fabrics without merging those fabrics. For example, using FCR you can share tape drives across multiple fabrics without the administrative problems, such as change management, network management, scalability, reliability, availability, and serviceability, that might result from merging the fabrics.

A Fibre Channel router (FC router) is a switch running the FC-FC routing service. The FC-FC routing service can be simultaneously used as an FC router and as a SAN extension over wide area networks (WANs) using FCIP.

Configuration of FC-FC Routing requires the installation of the Integrated Routing license.

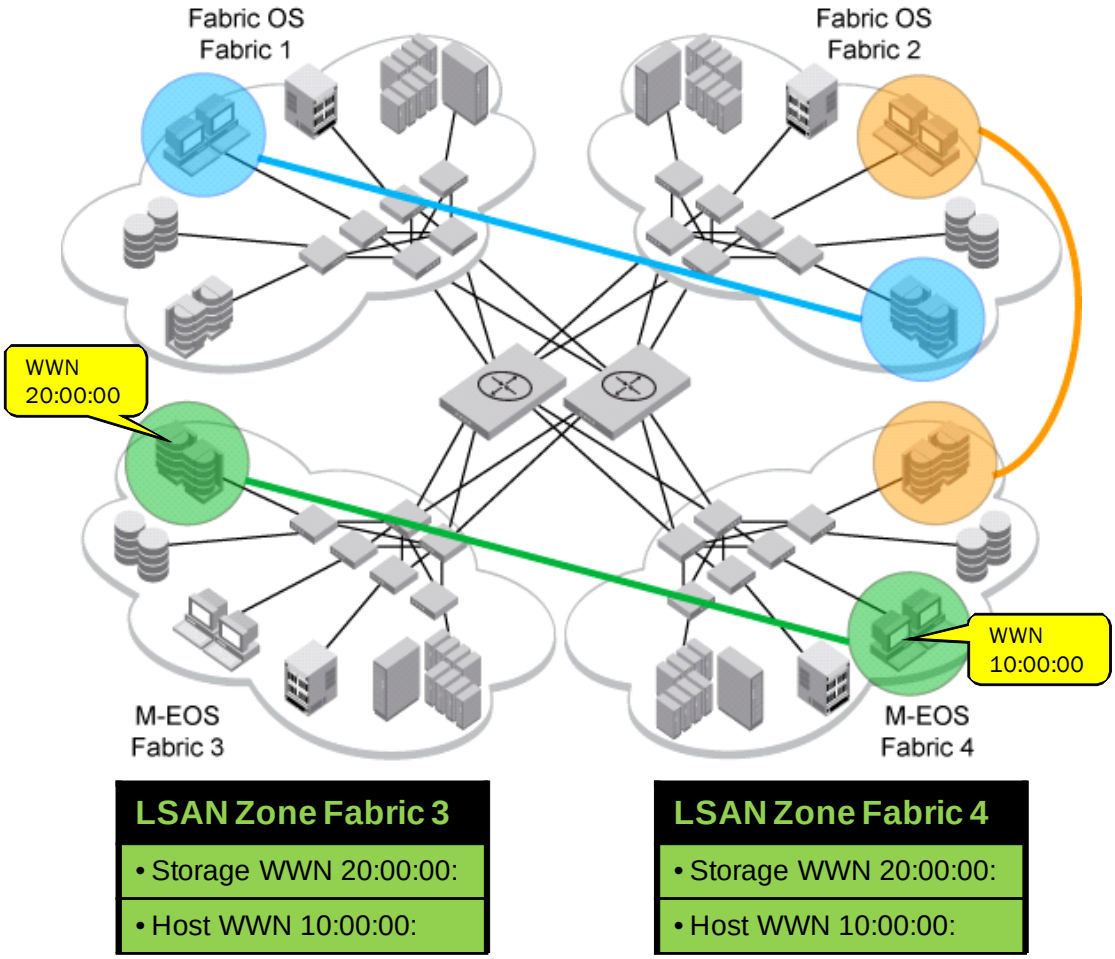

Figure 9: FC-FC Routing with LSAN Zones

©2011 Brocade Communications

FC-FC Routing Terminology

Edge Fabric: A fabric that is attached to one or more FC router EX_Ports.

Backbone Fabric (BB): The interconnection point for edge fabrics, containing at least one FC router.

Fabric ID (FID): Uniquely identifies each fabric participating in routed fabrics. Configuring the FID on a switch requires disabling the switch.

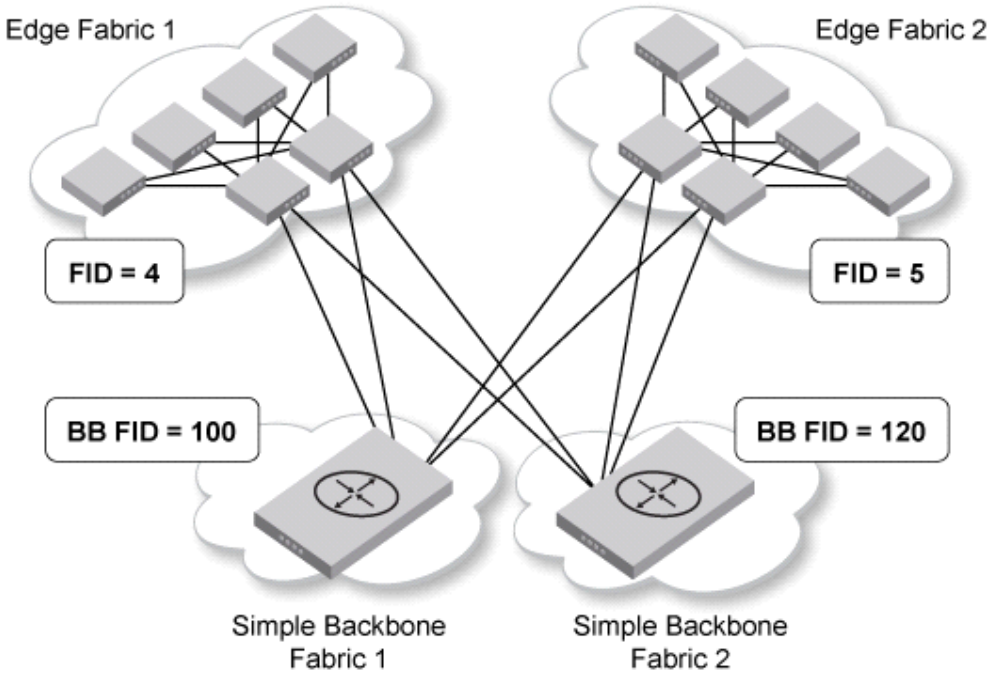

Figure 10: Fabric IDs

EX_Port: A type of E_Port used to connect an FC router port to an edge fabric without merging the two.

Inter-Fabric Link (IFL): The connection between a backbone fabric EX_Port and an edge fabric E_Port.

LSAN: A logical storage area network that spans multiple physical fabrics.

LSAN Zone: Zones that define which devices are to be shared between fabrics.

Front Domain (fd): A logical domain created in the edge fabric when edge fabrics are connected to backbone fabrics. A front domain represents the router in an edge fabric. Front domains are not created in a backbone fabric. Instead, they are a tier domain between the translate domains (xd) and the edge fabric. Imported devices are NOT attached to front domains, they are attached to translate domains.

Translate Domain (xd): A logical domain created when routed fabrics share devices. They are created in edge or backbone fabrics, but only created when physical devices in both fabrics requiring an xd are online and are part of an LSAN zone in two or more fabrics. Only one xd exists for each remote routed fabric.

Exported Device: A physical device defined in an LSAN zone that the router shares out of a fabric (edge or backbone).

Imported Device: A logical device defined in an LSAN zone that acts as a proxy device for the physical device in a different routed fabric (edge or backbone). Imported, or proxy, devices are logically connected to translate domains.

Fibre Channel Router Protocol (FCRP): A Brocade-authored Layer 3 routing protocol with two distinct components.

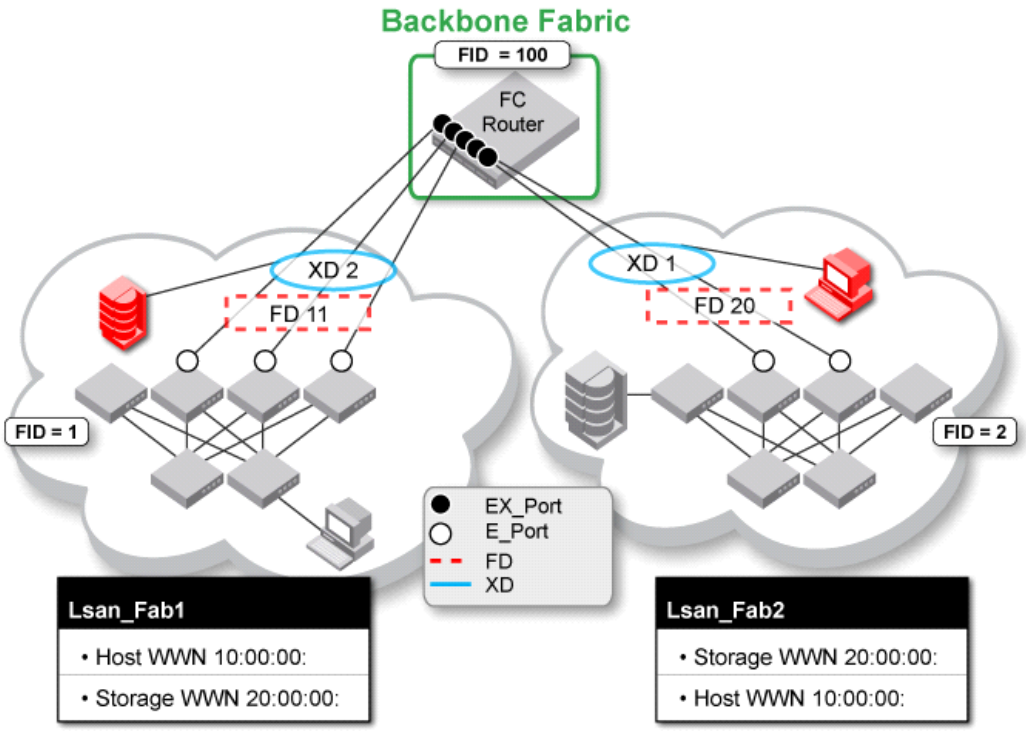

Figure 11: FC-FC Routing Showing Front and Translate Domains

LSAN Tagging

FC routers with Fabric OS v6.2 and greater support two types of optional LSAN tags used for special processing:

- Enforce tags: Only LSAN zones with names containing a configured enforce tag name are imported by the FC router
- Speed tag: Used when FC router needs to import targets to a routed fabric before the host is online (e.g. boot over SAN)

Device Sharing Considerations

- In configurations with two or more backbone fabrics connected to the same edge fabric, routing is not supported between edge fabrics that are not directly attached to the same backbone fabric.
- Routing over multiple backbone fabrics is a multi-hop topology and is not allowed.
- Sharing devices between two or more backbone fabrics is also not supported.

©2011 Brocade Communications

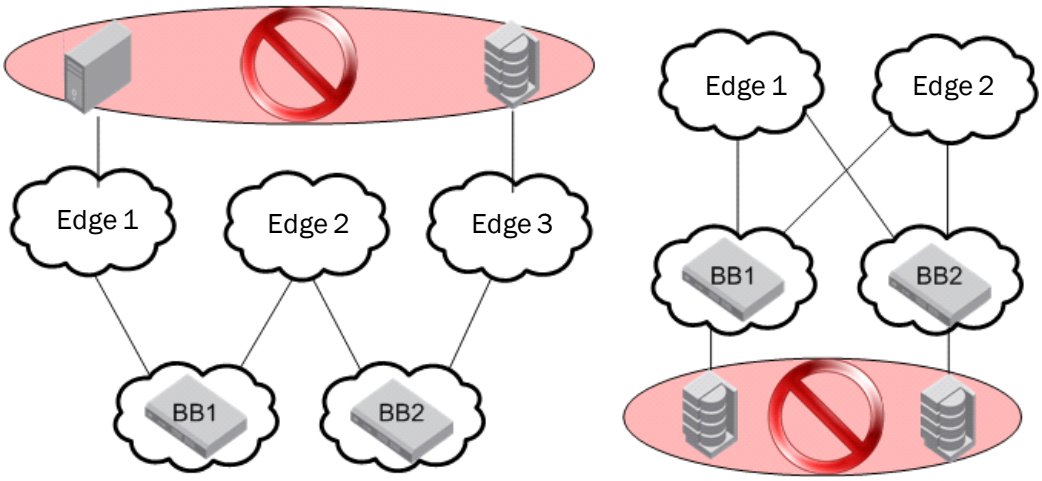

Figure 12: Device Sharing Considerations

FC-FC Routing With FCIP

Fabric OS supports FC-FC routing over an FCIP tunnel, creating a Virtual EX_Port (VEX_Port). This allows for long-distance FCIP connections with fabric-to-fabric isolation.

VEX_Ports are no different from EX_Ports, except the underlying transport is IP rather than Fibre Channel. Both EX_Ports and VEX_Ports are supported in the same backbone fabric.

VEX_Port connectivity rules:

- A VEX_Port connects only to a VE_Port – it may not connect to another VEX_Port
- There can be multiple VEX-to-VE port connections between a backbone fabric and an edge fabric
- VEX_Ports are not supported with FICON

Chapter 6 – Extension Features and Implementation

After reviewing this section be sure you can perform the following:

- Demonstrate knowledge of how to manage FCIP/FC distributed external solutions
- Demonstrate knowledge of advanced long distance Fibre Channel solutions

Brocade has a number of extension and distance solutions available. Brocade extension solutions can be used to extend a fabric over dark fiber using Time Division Multiplexing (TDM) or Dense Wave Division Multiplexing (DWDM). Long Wave Length SFPs can be used to extent native Fibre Channel across large campuses and between buildings, and FCIP can be used to extend a fabric through an IP network to a remove location.

FCIP

The Fibre Channel-over-IP (FCIP) protocol connects Fibre Channel switches over an IP network. IP packets generated by an FCIP-compliant port navigate the IP network to reach the destination end point. Implementation uses standards-based TCP, it interoperates with regular network equipment. FCIP is a tunneling protocol that allows transparent interconnection of remote locations through an IP-based network. From the fabric view, an FCIP link is an ISL, transporting all needed FC control and data frames between switches. The IP network is invisible to the fabric. The FC fabric and protocols are invisible to the TCP/IP network.

FCIP allows customers to extend their data centers over large distances by leveraging the existing IP infrastructure. This leads to less fiber cable being used in extension projects.

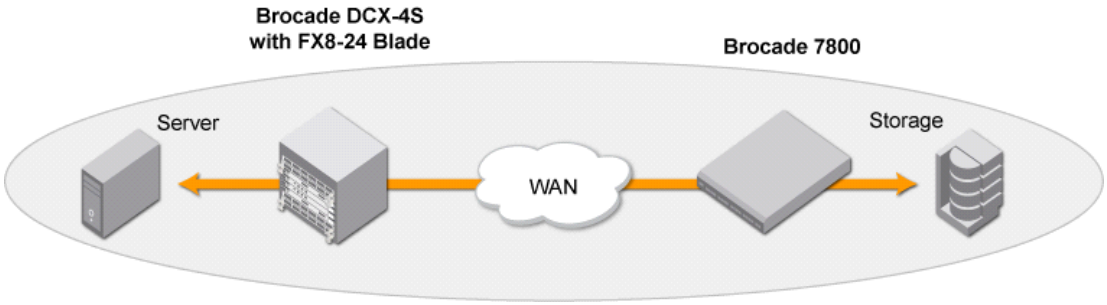

Figure 13: FCIP Example

©2011 Brocade Communications

FCIP Hardware Offerings

Brocade 7800 Extension Switch

- The Brocade 7800 Extension Switch has two configurations
  - Brocade 7800 4/2 – 4 x 8 Gbps FC Ports and 2 x 1 GbE Ports
  - Brocade 7800 16/6 – 16 x 8 Gbps FC Ports and 6 x 1 GbE Ports

Note

An Upgrade License is required to enable all 16 FC ports and 6 GbE ports.

- Auto-sensing FC ports at 8, 4, 2 or 1 Gbps
- The Inband Management feature allows a customer to define one or more IP addresses on the nonmanagement GbE ports and will provide a management path from the WAN network to the switch CP. Customers can then use the IP addresses to access SNMP events, telnet, etc.
- Support for FICON
- One GoldenEye2 ASIC for FC ports
- One FCIP subsystem

FX8-24 Extension Blade

- 2 x 10 GbE ports (license required)
  - Enabling either of the two 10 GbE ports requires a 10GbE FCIP license, which is a slot-based license. The supported operational modes are:
    - 10 x 1 GbE port
    - 10 x 1 GbE ports and 1 x 10 GbE port
    - 2 x 10 GbE ports
- 10 x 1 GbE ports
- 12 x 8 Gbps FC ports
- Max of 4 blades supported in the DCX products as of Fabric OS v7.0

Note

Check release notes for FOS version support. Only two blades are supported in a chassis with FOS v6.3.x.

- Auto-senses FC link speed at 8, 4, 2 or 1 Gbps
- One Condor2 ASIC for FC ports
- Two FCIP complexes

Tunnels

FCIP tunnels are a collection of one or more FCIP circuits that create one logical connection between two FCIP devices. Each FCIP tunnel presents a VE_Port to the FC fabric. Tunnels can span multiple physical ports. Multiple FCIP circuits from different physical port interfaces added to a VE/FCIP tunnel increases the bandwidth available to an FCIP tunnel. Configuring a tunnel with more than one circuit requires an Advanced Extension license. Without a license present, a second circuit will not be allowed to be configured. The administrator will receive a message stating as such.

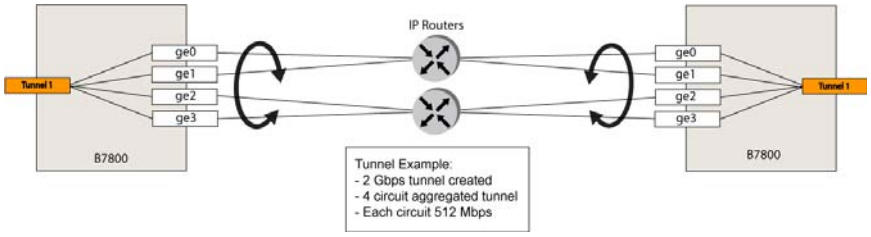

Figure 14: FCIP Tunnel Example

Circuits

An FCIP circuit is a logical connection created between two IP address end points. When created, a committed rate can be configured. Each 1 GigE circuit supports a rate of 10 Mbps to 1 Gbps. Each 10 GigE circuit supports a rate of 10 Mbps to 10 Gbps.

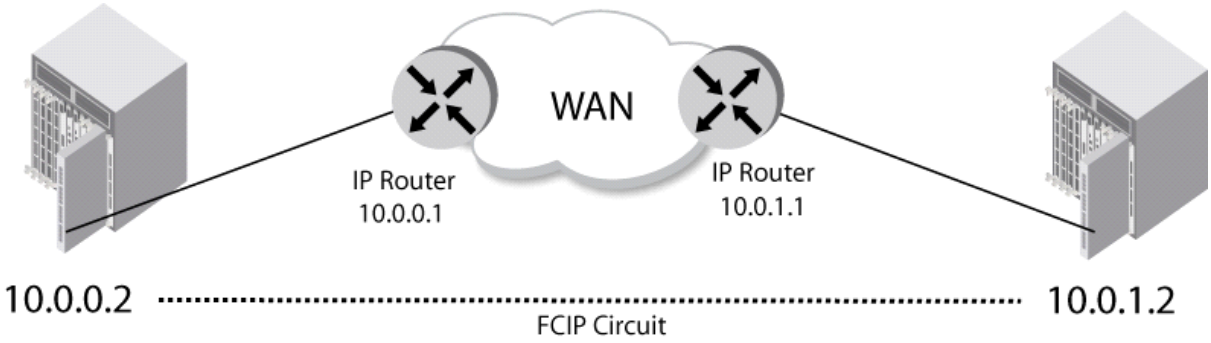

Figure 15: FCIP Circuit Example

FCIP Adaptive Rate Limiting

Adaptive Rate Limiting (ARL) provides for an adaptive committed rate configuration on an FCIP circuit. This is implemented by configuring a minimum and a maximum committed rate. ARL allows the traffic rate on a circuit to float between the minimum and maximum.

If there is traffic demand from FCIP and the network connection is clean (no retransmits) then the rate will increase up to the maximum. If TCP retransmits are detected, the rate will immediately retreat back to the minimum. Growth is accomplished by testing the ceiling and stepping the available bandwidth up gradually. This provides more efficient bandwidth sharing between applications using the same network infrastructure.

Use of the FCIP ARL feature requires the Advanced Extension license.

©2011 Brocade Communications

Long Distance Fabrics

The most effective configuration for implementing long-distance SAN fabrics is to deploy Fibre Channel switches at each location in the SAN. Each switch handles local interconnectivity and multiplexes traffic across long-distance dark fiber or wave division multiplexing (WDM) links while the Brocade Extended Fabrics software enables SAN management over long distances. Brocade Extended Fabrics is an optional licensed feature for Brocade SAN deployment over distance beyond 10 km. A Brocade Extended Fabrics license is required before you can implement long distance dynamic (LD) and long distance static (LS) distance levels. The LD and LS settings are necessary to achieve maximum performance results over Inter-Switch Links (ISLs) that are greater than 10 km.

Long Distance Modes

Use the `portcfglongdistance` command to configure long distance links and to allocate sufficient numbers of full size frame buffers on a particular port. Changes made by this command are persistent across switch reboots and power cycles. This command supports the following long-distance link modes:

- Static Mode (LO) - L0 is the normal (default) mode for a port. It configures the port as a regular port. A total of 20 full-size frame buffers are reserved for data traffic, regardless of the port operating speed; therefore, the maximum supported link distance is up to 5 km at 2 Gbps, up to 2 km at 4 Gbps, and up to 1 km at 8, 10, and 16 Gbps.
- Static Mode (LE) - LE configures an E_Ports distance greater than 5 km and up to 10 km. LE does not require an Extended Fabrics license. The baseline for the buffer credit calculation is one buffer credit per km at 2 Gbps. This yields the following values for 10 km:
  - 10 buffer credits per port at 2 Gbps.
  - 20 buffer credits per port at 4 Gbps.
  - 40 buffer credits per port at 8 Gbps.
  - 50 buffer credits per port at 10 Gbps
  - 80 buffer credits per port at 16 Gbps
- Dynamic Mode (LD) - LD calculates BB credits based on the distance measured during port initialization. Brocade switches use a proprietary algorithm to estimate distance across an ISL. The estimated distance is used to determine the BB credits required in LD (Dynamic) extended link mode based on a maximum Fibre Channel payload size of 2,112. You can place an upper limit on the calculation by providing a desired_distance value. Fabric OS confines user entries to no larger than what it has estimated the distance to be. When the measured distance is more than desired_distance, the desired_distance (the smaller value) is used in the calculation.
- Static Long-Distance Mode (LS) - LS calculates a static number of BB credits based only on a user-defined desired_distance value. LS mode also assumes that all FC payloads are 2112 bytes. Specify LS mode to configure a static long distance link with a fixed buffer allocation greater than 10 km. Up to a total of 1452 full-size frame buffers are reserved for data traffic, depending on the specified desired_distance value.

Configuring Fill Words

Brocade Fibre Channel switches support the use of two different fill words on long distance links; ARB (default) and IDLE. On Condor2/GoldenEye2-based platforms the fill word can be modified using the `portcfgfillword` command. Condor3 based platforms are able to auto detect the required fill word and the use of the `portcfgfillword` command has been deprecated.

Below is a list of Condor2/GoldenEye2-based switches:

- Brocade 300
- Brocade 5100
- Brocade 5300
- Brocade 7800
- All 8 Gbps DCX and DCX-4S blades
  - FC8-16
  - FC8-32
  - FC8-48
  - FC8-64

Considerations for Using IR_RDY on TDM Devices

Some TDM devices require the use of IDLE fill words instead of the default ARB. When connecting to these devices the following considerations must be made:

- Disable QoS on the port using the `portcfgqos` command
- Disable credit recovery on the port using the `portcfgcreditrecovery` command

©2011 Brocade Communications

Chapter 7— Advanced Troubleshooting

After reviewing this section be sure you can perform the following:

- Describe which tools can be used to troubleshoot a complex fabric problem
- Demonstrate knowledge of how to troubleshoot extension, FC routing, and FC network scenarios

Diagnostic Port

Fabric OS v7.0.0 allows you to configure a Fibre Channel port, including ISLs and loop back ports, into a Diagnostic Port (D_Port). This port lets you isolate the inter-switch link (ISL) to diagnose link level faults. The D_Port does not carry any fabric traffic, and is designated to run only specific diagnostics tests on it. The creation of a D_Port is subject to Virtual Fabric restrictions that may be in place. The ports must be 10G or 16G Brocade-branded SFPs on a Brocade DCX 8510, and running Fabric OS v7.0.0 or later.

You must configure both ends of the link between a given pair of switches, and you must disable the port before you can configure a D_Port. Re-enabling the D_Ports automatically starts the diagnostics when the ports come online, and includes the following tests:

- Electrical loop back (16G SFPs only)
- Optical loop back (16G SFPs only)
- Link traffic (16G SFPs and 10G SFPs)
- Link latency and distance measurement (16G SFPs and 10G SFPs)

Diagnostic Port Uses

You can use the D_Port to test the following:

- Testing a new ISL before adding it to the fabric
- Testing a trunk member before joining it with the trunk
- Testing long distance cables and SFPs
  - Determining the distance of a FC cable
- Testing loopback ports

Top Talkers

The Top Talker feature is based on the Advanced Performance Monitor (APM) feature. The Top Talker feature determines the largest users of F_Port bandwidth by monitoring all flows (SID-DID pairs) through one or more switch F_Ports on any switch in the fabric.

Top Talkers is an enhancement to Advanced Performance Monitor (APM) end-to-end monitors. When enabled, these monitors determine which SID-DID pairs are the major users of switch F_Port bandwidth Can be enabled on specific switch E_Ports or F_Ports in the fabric. Top Talkers can determine the flows (SID-DID pairs) that are the major users of bandwidth, measure bandwidth usage data in real-time and relative to the port on which the monitor is installed. Top Talkers also requires APM license.

Top Talker Modes

Top Talker can be configured in two modes:

- **Port Mode:** Enabled on an F_Port to measure the traffic between the F_Port and all other devices with which it can communicate. You can configure this on E_Ports on the on the following 16 Gbps platforms: Brocade 6510 and Brocade DCX 8510 family.
- **Fabric Mode:** Enabled on all E_Ports in the fabric to measure the data rate of all the possible flows in the fabric (ingress E_Port traffic only). In Fabric Mode, Top Talker monitors can determine the top *n* bandwidth users on a given switch.

Note

Can be configured in Port Mode or Fabric Mode only, not both simultaneously

Top Talker Usage

- **Traffic monitoring:** Determine the largest flows so traffic can be rerouted through a switch or fabric to balance bandwidth utilization.
- **Virtual servers:** Determine the largest flow virtual servers through a given HBA by taking NPIV devices into consideration when calculating the top talking flows.
- **Service tracking:** Identify when the largest flows exceed a maximum standard set by a service agreement.
- **Adaptive Networking:** Identify flows that would gain the greatest benefits from Adaptive Networking features used to optimize fabric behaviors.

Top Talker Limitations

- Top Talker monitors cannot detect transient surges in traffic through a given flow
- Top Talker monitors cannot be installed on a mirrored port
- Can monitor up to 10,000 flows
- Cannot detect spikes in traffic on a given flow

©2011 Brocade Communications

Miscellaneous Troubleshooting Commands

porterrshow

- The `frames tx` and `rx` are the number of frames being transmitted and received.
- The `crc_err` counter are frames with CRC errors. If this counter goes up, then the physical path should be inspected. Check the cables to and from the switch, patch panel, and other devices. Check the SFP by swapping it with a known good working SFP. If you see this issue on an 8 Gbps blade, use the `portcfgfillword` command to reduce EMI.
- The `crc_g_eof` counter are frames with CRC errors and a good EOF. The first port detecting a CRC error marks the frame with a bad EOF and passes the frame on to its destination. Subsequent ports in the path also detect the CRC error and the `crc_err` counter increments on these ports. However, since the first port marked the frame with a bad EOF, the good EOF counter on the subsequent ports does not increment. The marginal link associated with the port with an increasing good EOF counter is the marginal link and the source of the errors.
- The `enc_out` counter are errors that occur outside the frame and usually indicating a bad primitive. To determine if you are having a cable problem, take snapshots of the port errors by using the `porterrshow` command in increments of 5 to 10 minutes. If you notice the `crc_err` counter go up, you have a bad or damaged cable, or a bad or damaged device in the path. ICLs see `enc_out` errors when ports on one side of the link are disabled.
- The `disc_c3` errors are discarded class 3 errors, which means that the switch is holding onto the frame longer than the hold time allows. One problem this could be related to is ISL oversubscription.

```
switch:admin> porterrshow
   frames    enc crc   crc     too    too    bad   enc disc link loss loss frjt fbsy
   tx    rx  in  err   g_eof   shrt   long   eof   out  c3  fail sync sig
==============================================================================
0: 665k 7.0k 0  0     0       0      0      0     6   0    0    1    2    0    0
1: 0    0    0  0     0       0      0      0     0   0    0    0    2    0    0
2: 0    0    0  0     0       0      0      0     0   0    0    0    1    0    0
3: 0    0    0  0     0       0      0      0     0   0    0    0    1    0    0
4: 0    0    0  0     0       0      0      0     0   0    0    0    1    0    0
5: 0    0    0  0     0       0      0      0     0   0    0    0    1    0    0
6: 0    0    0  0     0       0      0      0     0   0    0    0    1    0    0
7: 0    0    0  0     0       0      0      0     0   0    0    0    1    0    0
8: 78   60   0  0     0       0      0      0     7   0    0    3    6    0    0
9: 12   4    0  0     0       0      0      0     3   0    0    1    2    0    0
10: 0   0    0  0     0       0      0      0     0   0    0    0    1    0    0
11: 0   0    0  0     0       0      0      0     0   0    0    0    1    0    0
12: 0   0    0  0     0       0      0      0     0   0    0    0    1    0    0
13: 0   0    0  0     0       0      0      0     0   0    0    0    1    0    0
14: 0   0    0  0     0       0      0      0     0   0    0    0    1    0    0
15: 0   0    0  0     0       0      0      0     0   0    0    0    1    0    0
16: 665k 7.4k0  0     0       0      0      0     6   0    0    1    2    0    0
<output truncated>
```

framelog

When a frame is unable to reach its destination due to timeout, it is discarded. You can use Frame Viewer to find out which flows contained the dropped frames, which in turn can help you determine which applications might be impacted. Using Frame Viewer, you can see exactly what time the frames were dropped, timestamps are accurate to within one second. This also assists in the debug process.

You can view and filter up to 20 discarded frames per chip per second for 1,200 seconds using a number of fields with the `framelog` command.

```
switch:user> framelog --show -mode summary -n 8
=====================================================================
              Mon Jan 31 23:54:59 UTC 2011
=====================================================================
Log              TX   RX
timestamp        port port SID      DID      SFID DFID Type    Count
=====================================================================
Jan 31 23:49:37 2    2    0xfffffd 0xfffffd 1    1    timeout 1
Jan 31 23:49:37 2    1    0x051500 0x060100 1    1    timeout 4
Jan 31 23:49:37 2    0    0x051700 0x060000 1    1    timeout 4
Jan 31 23:49:36 2    1    0x051500 0x060100 1    1    timeout 3
Jan 31 23:49:36 2    0    0x051700 0x060000 1    1    timeout 3
Jan 31 23:49:35 2    1    0x051500 0x060100 1    1    timeout 2
```

fabriclog

Use this command to display, clear, disable, enable, or resize the fabric log. When used with the `--show` option, this command displays additional information. Among the information provided is link resets, build fabrics, principal switch information, and trunk port information.

errshow

Use this command to display external error log messages one at a time. When executed without operands, this command prints the error messages for the logical switch context in which the command is executed. When used with the `-a` option, the command prints the error messages for the entire chassis. The messages are displayed with page breaks. The `-r` operand displays the messages in reversed order.

The output of this command is unique for each CP. On dual CPs this command must be executed on each CP to obtain a complete record.

```
switch:admin> errshow -r
Fabric OS: v7.0.0a

2011/11/09-18:49:58, [SEC-1203], 103, FID 128, INFO, switch, Login
information: Login successful via TELNET/SSH/RSH. IP Addr: 10.255.255.25
```


©2011 Brocade Communications

portcmd

Use this command to invoke the end-to-end IP path performance (IPperf) characterization feature, to ping or trace a route to a destination IP host from an intelligent GbE port, or to determine the path characteristics between a local data source and a remote data sink.

The following portCmd features are platform-specific:

- IPperf - Supported only on the Brocade FR4-18i.
- Tperf - Supported only on the Brocade 7800/FX8-24.
- VLAN tagging - Supported on the Brocade 7800/FX8-24 and FR4-18i.
- IPv6 - Supported on the Brocade FX8-24/7800 and FR4-18i.

When issued with the `--tperf` option, this command determines the path characteristics to a remote host or tunnel destination. Tperf generates statistics every 30 seconds by default unless you specify a different value for `-interval`.

zone --show tierrors

Analyzes real and potential routing problems with the activated TI zoning set and prints a report. This command must be executed in the local domain and analyzes only that domain.

trunkshow

Use this command to display trunking information of both E_Ports and EX_Ports. You can use the `-perf` option to display the total bandwidth, throughput, and percentage of link utilization information for the trunk group (Rx, Tx, and combined total for Tx+Rx).

```
switch:admin> trunkshow
  1:  8->  8 10:00:00:05:33:7f:05:b1   3 deskew 15 MASTER
      9->  9 10:00:00:05:33:7f:05:b1   3 deskew 15

  2: 15-> 15 10:00:00:05:1e:0b:55:3f   1 deskew 15 MASTER
     14-> 14 10:00:00:05:1e:0b:55:3f   1 deskew 15
```

lsanzoneshow

Use this command to display the inter-fabric zones or LSAN zones. These zones are normal WWN zones created in FC Router EX_Port-connected fabrics and backbone fabrics. The LSAN zones are identified by the text string "lsan_" in the zone name. Note that the string is case insensitive so "LSAN_" also is valid. The FC Router uses these zones to establish the inter-fabric device import and export policy. The LSAN zones are established by zoning administration in each EX_Port-attached fabric and backbone fabric. Inter-fabric device sharing is allowed between two devices if the LSAN zones defined in their respective fabrics both allow the two devices to communicate; for example, the intersection of LSAN zones in two fabrics define the device sharing policy.

The LSAN zones are listed by fabric. Zone membership information (information about the devices in the zone) is provided for each LSAN zone. The default output displays only WWNs of the zone members. The `-s` option can be used to show device states.

fcrphydevshow

Use this command to display the physical (real) devices that are configured to be exported to other fabrics. A device is considered to be configured to be exported to another fabric if it is a member of an LSAN zone. The device is displayed only if it is discovered in the EX_Port-attached fabric and backbone fabric's name server (for instance, the device is online).

Physical device information is available only for physical devices that exist in fabrics attached to EX_Ports of FC routers on the same backbone fabric as the current FC router.

The default output displays only physical device information relevant to the current FC router. Relevant physical devices include physical devices that are configured to be exported from fabrics attached to the current FC router's EX_Ports.

The physical devices are listed by fabric.

©2011 Brocade Communications

Taking the Test

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

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 credential, 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. Brocade also reserves the right to use any technologies and methods for verifying the identity of candidates. Such technology may include, without limitation, personally identifiable information, challenge questions, identification numbers, photographic information, and other measures to protect against fraud and abuse.

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

©2011 Brocade Communications