Fibre Channel over Internet Protocol (FCIP)

EMC Proven™ Professional Knowledge Sharing 2008

Joseph Holbrook
Brocade
jholbroo@brocade.com

Table of Contents

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect EMC Corporation's views, processes or methodologies.

Introduction

Maintaining application availability is a top priority for most IT organizations. Datacenters now run 24x7, and global operations may come to a grinding halt if even one critical application goes down. As a result, IT departments must ensure that operations can continue even in the face of disasters that could disable a site or an entire region with a single blow. Business Continuity, Disaster Tolerance, Disaster Recovery, and similar solutions are being deployed more frequently. Indeed, regulatory requirements are mandating this category of solution for many industries. All solutions in this category must be able to move large amounts of data in a reliable and repeatable manner. As a result, they are generally built on top of a Fibre Channel (FC) Storage Area Network (SAN) infrastructure.

An effective solution requires locations separated by great distances; there is no point in copying data between sites if they are close enough that they may be "hit" by the same disaster. You may need to extend a SAN over hundreds of kilometers or across a continent. You must sometimes use an IP network to transport the FC SAN data to be effective over such distances. Fibre Channel over IP (FCIP) is the standard. It allows transparent tunneling of FC switch-to-switch links across IP networks. In addition to offering a full range of Fibre Channel switching and routing products, Brocade now offers several FCIP gateway solutions that fully complement EMC solutions.

In this article, we will discuss the following topical overview to help both non-technical and technical professionals understand the benefits of using Fibre Channel over IP (FCIP) technology.

TOPIC OVERVIEW

- Meta-SAN Overview
- Fibre Channel over IP (FCIP) Overview
- EMC FCIP products overview
- EMC/Brocade FCIP Design
- FCIP Implementation
- FCIP Performance Considerations
- Brief enterprise level FCIP customer implementation (case study)
- Article Closing
- Glossary

Meta–SAN Overview

When single SAN switches are connected to each other, they form a SAN "fabric". When they are connected through an FCR approach they form a different SAN. This results in a "Routed" SAN; the result is called a "Meta SAN".

The "Meta SAN" can be greatly expanded when SANs are connected over extended long distances using FCIP.

A **Meta-SAN** is a collection of SAN devices, switches, edge fabrics, Logical Storage Area Networks (LSANs), and Routers that comprise a physically connected but logically partitioned storage network. A simple Meta-SAN can be constructed using an EMC Connectrix® MP-7500B, Connectrix ED-48000B with a Connectrix PB-48K-18I blade to connect two or more separate fabrics. Additional EMC Connectrix MP-7500Bs, Connectrix ED-48000B with Connectrix PB-48K-18I blades, can be used to increase the available bandwidth between fabrics, and for redundancy as well as to create a distance extension solution using FCIP with the GE ports on extension products.

A simple FCR/FCIP configuration is illustrated below.

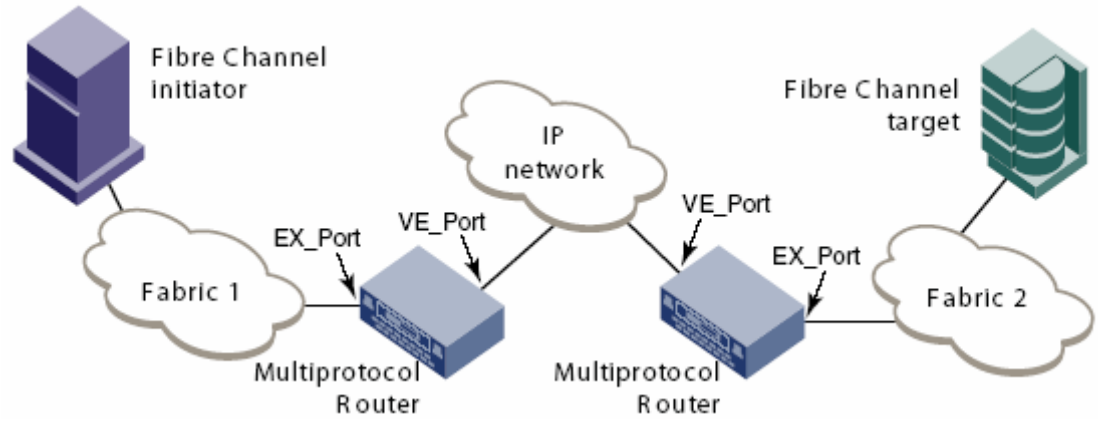

Each edge fabric in the Meta-SAN will remain a SAN island and will be referred to an edge fabric in this case. Edge fabrics will maintain separate fabric services such as name services, zoning databases, routing tables, domain ID spaces, etc.

This ability to maintain separate services greatly reduces management problems like domain and zoning conflicts that would otherwise be a concern when merging a fabric.

The need to increase port counts and share resources across geographical and functional boundaries will not disappear anytime soon. Moreover, the need to ensure secure connectivity for selected resources—especially in heterogeneous environments—further compounds troubleshooting, fault isolation, and management challenges posed by large SANs.

FC routing addresses these issues by enabling organizations to connect devices in different SANs without merging the fabrics. Using this capability, organizations can share resources across multiple SANs and scale beyond current SAN port count support constraints. In addition, they can more easily support multiple firmware revisions and connect devices between SANs purchased from and supported by different OEMs, or from different SAN platform vendors, as business needs dictate.

This level of SAN connectivity gives organizations a powerful tool for reducing or even eliminating disruptions associated with many common operational events, such as:

- Migration to new SAN infrastructure, where FC routing can be used to migrate SANs and devices from 1 or 2 Gbit/sec to 4 Gbit/sec systems.
- Data center consolidation, where the Connectrix PB-48K-18I blade and EMC Connectrix MP-7500B Router can enable data movement between locations, adding value when FC routing is combined with FCIP.

- Mutual data center backup or mirroring, where the Connectrix PB-48K-18I blade and EMC Connectrix MP-7500B Router enable two major data centers to act as backup facilities for each other. This application often involves FCIP, and the tight integration of FCIP and FC routing is usually a value-add for the customer.
- Storage and application rebalancing between fabrics, where the Connectrix PB-48K-18I blade and EMC Connectrix MP-7500B Router can connect devices in previously isolated SANs
- Data migration between test/development SANs and production SANs, enabling data movement between physically or logically separated environments.
- Data migration or sharing between SANs containing different firmware releases.
- Data migration between Brocade and McDATA legacy fabrics.

Benefits of META-SAN Connectivity

- Scalability - Growing SAN environments much more efficiently and maximize the value of high-end resources.
- Security - Connecting devices in different SANs without merging the SANs, organizations can support secure, selective resource sharing through LSANs.
- Centralization- Centralization makes it easier to manage equipment from multiple vendors increasing operational flexibility and enabling the deployment of a best-in-class environment.

EXAMPLE OF EMC META-SAN DISTANCE CONNECTIVTY

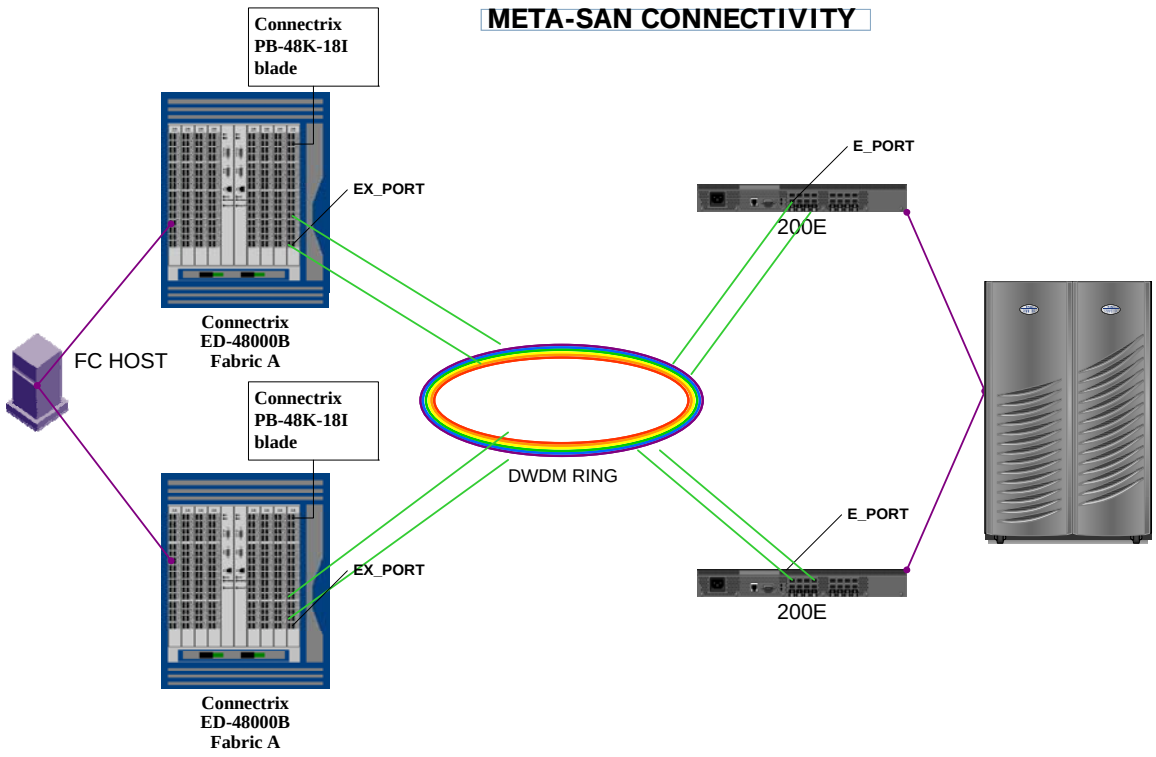

Notice in the EMC Meta-SAN diagram above:

1. Fabric is HA (Redundant Fabric with Dual IFLs)

2. Symmetrically organized by port number and slot number

3. Host and Storage are in different locations

The need to create Fibre Channel SANs that can grow in a scalable, cost-effective manner is one of organizations' basic requirements. For example, organizations with multiple SAN islands would like to connect them into a more unified, centrally managed

SAN environment. Unfortunately, many organizations have avoided merging their SAN islands for fear that the administrative workload, risk, and expense would not justify the benefit of enhanced connectivity.

However, the ability to seamlessly integrate Brocade FC routing into existing SAN fabrics, and realize the benefits of centralized management while maintaining fabric separation, is changing that perception. Running on the Connectrix PB-48K-18I blade and EMC Connectrix MP-7500B switch, FC routing enables devices located on separate SAN fabrics to communicate without the need to merge the fabrics into a large SAN environment. By using this service, organizations can interconnect devices without having to redesign and reconfigure their entire environment, thereby eliminating the potential risks and costs of downtime.

The resulting routed network would consist of multiple individual SAN fabrics that form one storage network connectivity model, known as a "Meta SAN." In this way, FC routing offers key strategic advantages such as:

- Simplifying SAN design, implementation, and management through centralization
- Providing a seamless and secure way to share resources across multiple SANs without the complexity of physically merging those SANs
- Creating a more unified SAN environment with easier interconnection and support for SANs and SAN resources purchased from different storage vendors
- Reducing disruptions created by events such as data migration, storage or server consolidation, migration to production environments, and application rebalancing between fabrics

When devices on different fabrics are allowed to communicate through FC routing, the resulting connectivity group is known as a Logical SAN (LSAN). LSANs enable selective and secure resource sharing across multiple SANs by leveraging existing zoning tools and methodologies.

In addition to optimizing resource utilization, this approach improves scalability by:

- Minimizing the risk and complexity of large fabrics
- Right-sizing SANs based on application and business requirements
- Simplifying management and fault isolation
- Protecting and extending current technology investments since LSANs require no changes to existing SAN switches or attached edge devices but do leverage existing zoning tools

The Brocade FC-FC routing service (FCRS) provides connectivity between two or more fabrics without the need to merge the fabrics. This service allows the creation of Logical Storage Area Networks (LSAN) that can provide connectivity that can span fabrics. LSAN is a zone that spans fabrics and allows connectivity without actually merging the fabrics. This service is implemented on the router as an "EX" port. The fact that FCR can connect fabrics without merging the fabrics has advantages in terms of scalability, network management, change management, availability and serviceability.

It is import to note that the act of projecting a node into another fabric is called *exporting*. When a host is exported from Fabric 1 into Fabric 2, it also must be imported into Fabric 2 from Fabric 1.

To create an LSAN, both exporting and importing must occur, so these statements are functionally equivalent in normal cases.

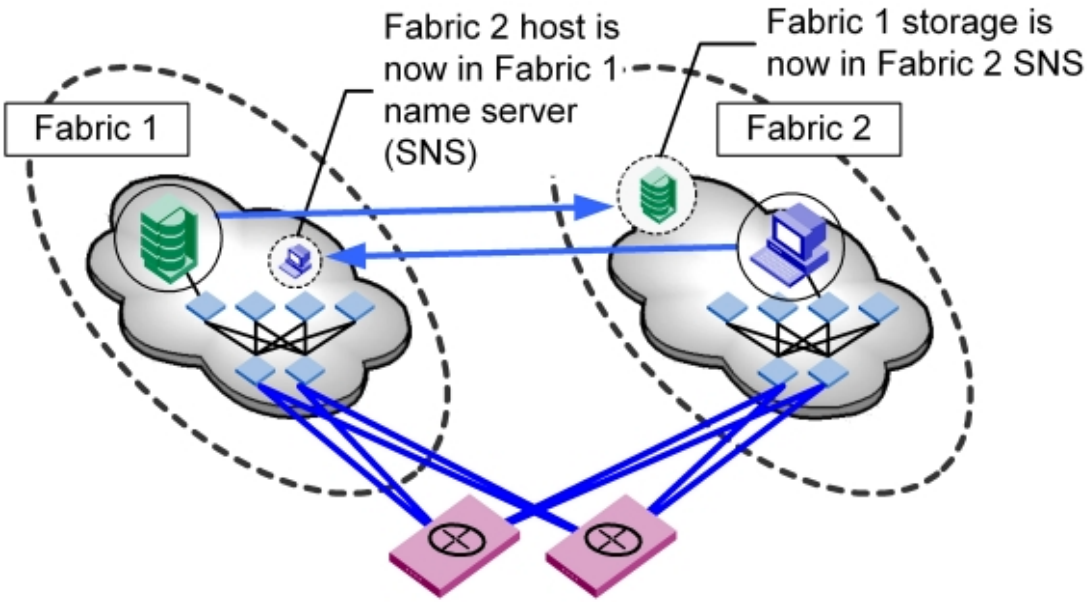

You must create an LSAN on both sides of the router. LSAN zones are indistinguishable to an edge fabric from any other kind of zone, which is why they are compatible with previous Brocade Fabric Operating System (Fabric OS) versions.

There are just two distinguishing features of an LSAN zone:

1. First, they must begin with the prefix "LSAN_" so that routers will recognize them.
2. Second, they must contain only port WWNS or aliases of devices intended for inter-fabric sharing.

This is because Fibre Channel Port IDs (PIDs) are not unique identifiers in a Meta SAN.

The same PID can exist in multiple edge fabrics, so a router would not know what the administrator wanted to do if PIDs were used to create LSAN zones. This constraint,

however, does not in any way preclude the use of PID based zoning in edge fabrics for other zones. Those zones continue to work as usual.

When a set of devices on different edge fabrics are allowed to communicate through an FC router in this way, the resulting connectivity group is an LSAN.

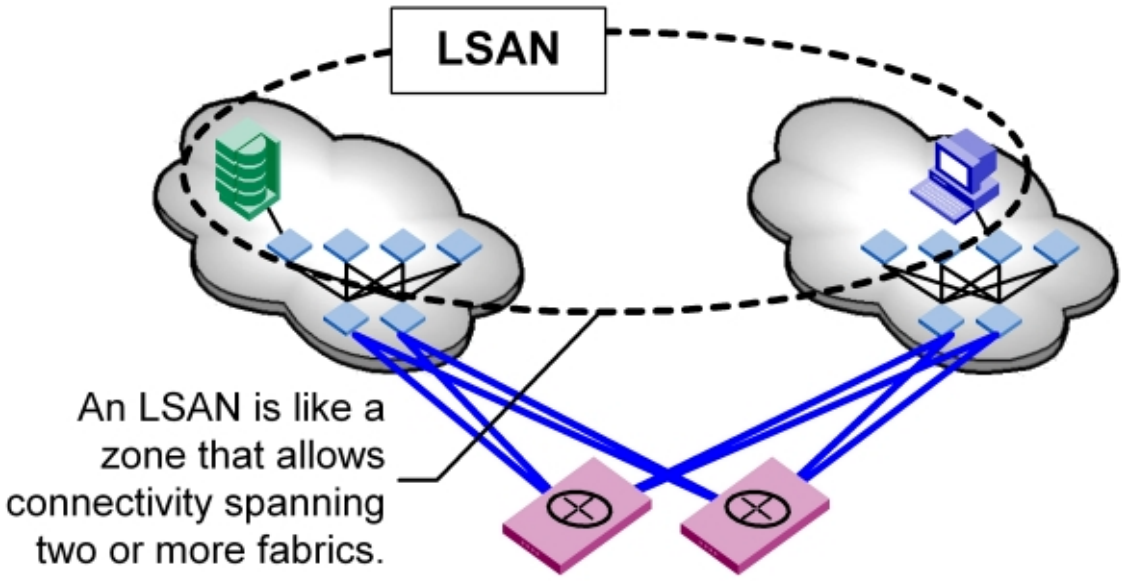

LSAN EXAMPLE

Many different LSANs can exist in a Meta SAN in the same way that many zones can exist within a single fabric. Indeed, many different LSANs can exist between any given set of fabrics. Devices can be members of multiple LSANs, and LSANs can overlap with traditional zoning mechanisms on local fabrics as well.

Fiber Channel Routing is somewhat equivalent to an IP router like a firewall since it uses a DENY all approach when the FCR services are on it. The SAN administrator must use an ACL approach by zoning to allow traffic to flow between fabrics.

It is important that we clarify some terminology since FCR has some significant new terminology that SAN Administrators and SAN Engineers may not be familiar with. *Please see the FCR Terminology Glossary in the back of this article.*

PHANTOM DOMAINS AND DEVICES

The router introduces a new model where a phantom topology is presented. It consists of phantom domains and phantom devices and does not correlate directly to physical entities. The phantom topology is presented via protocols such as Fabric Shortest Path First, Name Server, Management Server, etc. It is important to understand that this is automatically done by the routers FC-NAT protocols.

All devices that are translated between fabrics are "hung off of" xlate domains. To maximize FC NAT address space, translated devices are given FC PIDs using both the area and port bytes, so to the human eye xlate addresses may look like NL_Port devices in destination fabrics even if they are really N_Ports in their source fabrics.

PIDs are "made up" by the router, and whatever PID is used for any given device exported into any given edge fabric is persistent. Even simultaneously rebooting every host, storage, and network device in the Meta SAN - including all routers - will not cause any xlate PIDs to change.

NOTE: Translation table can be saved and loaded using the configUpload / configDownload commands

META-SAN REVIEW

The Brocade FC-FC routing service (FCRS) provides connectivity between two or more fabrics without merging them. This service allows you to create Logical Storage Area Networks (LSAN) that can provide connectivity to span fabrics. LSAN is a zone that spans fabrics and allows connectivity without merging the fabrics. This service is implemented on the router as an "EX" port.

The fact that FCR can connect fabrics without merging the fabrics has advantages like:

- scalability
- network management
- change management
- availability
- serviceability
- economical

FCR routing is:

- Logically connected to SAN islands and shares resources across multiple fabrics
- Maintained through administration and fault isolation of separately managed fabrics
- Supported by any tool that supports zoning

FCR Solves problems like:

- Scaling SANs for number of ports
- Scaling SAN's over distance
- Allowing the combining of SANs that do not have direct E_port interoperability

FCIP with FCR provides a significant Meta-SAN benefit in scalability and distance capabilities.

***FCIP is the technology that can greatly enable Meta-SAN extension and scalability when used with* FC-FC routing service (FCR).**

FCIP Overview

Fibre Channel over IP is a complex combination of transport technologies that address the dual requirements of storage networking and networking over distance. A mature technology optimized for storage data movement within the campus and data center, Fibre Channel represents a major investment in software compatibility, interoperability, and proven applications for campus-based storage networking. Likewise, IP is a mature technology optimized for data movement across WAN distances. It represents a major investment in software compatibility, equipment interoperability, and proven applications for WAN-based data networking.

Today's Fibre Channel-over-IP solutions encapsulate Fibre Channel and transport it over a TCP socket. Performance can vary based on the types of switches and routers, the number of hops the packets must traverse, and the level of congestion in the network. Today, storage transport performance over IP networks, especially over public networks, is limited due to the variable latency of service provider networks.

As IP and Ethernet equipment continues to evolve, higher levels of Quality of Service (QoS), Cost of Service (CoS), provisioning, and circuit emulation should provide the latency guarantees required by synchronous storage applications. In controlled environments, these technologies might even improve the performance of IP networks. Regardless, Fibre Channel over IP is currently a cost-effective technology for asynchronous applications such as remote data backup.

The True Benefits of Fibre Channel Over IP (FCIP)

While multiple technologies are capable of interconnecting SANs, very few can be widely deployed cost-effectively. Because most organizations already have IP connections and significant experience with Ethernet and IP networks, they can usually

leverage this equipment and expertise to manage data in conjunction with Fibre Channel SANs. For example, IP connectivity provides the greatest flexibility at the lowest cost for latency-tolerant applications. As a result, it can be used to back up data across a campus network, Metropolitan Area Network (MAN), or WAN. Moreover, this flexible technology can be deployed within a single enterprise or in an SSP multi-tenant environment.

Today, many types of organizations are beginning to transport SAN storage over IP, especially for non–real-time data transfer. By combining the better of two mature technologies, FCIP solutions provide a more standardized and lower-cost way to increase SAN interconnectivity for a variety of applications. In fact, FCIP includes full support for the Fibre Channel set of equipment and software. Organizations can seamlessly extend existing and planned Fibre Channel SANs over long distances through IP networks, thereby protecting significant investments in both technologies. In addition, FCIP provides a cost-effective way to achieve business protection (enabling solutions such as remote tape archiving).

FCIP can transport existing Fibre Channel services across the IP network such that two or more interconnected SANs can appear as a single large SAN and can be managed by traditional SAN management applications. In addition, FCIP enables SAN applications to support additional protocols without modification. These applications might include disk mirroring between buildings in a campus network or remote replication over the WAN. The type of applications utilized is based on the distance the data must travel, the network bandwidth, and the QoS requirements and/or abilities of the network connection.

While some FCIP implementations are point-to-point "tunnels," the protocol does not require that the "gateways" support only point-to-point tunneling. The FCIP standard supports all Fibre Channel services, including FSPF routing algorithms, so that multiple logical links created from a single gateway can route Fibre Channel packets over the IP

infrastructure. Not only is FCIP routable, but IP networks do not need to know anything about the packets being routed.

The Fibre Channel services handle all routing between logical links, while the TCP protocol handles the delivery of packets to the specific gateway device.

Defining FCIP

FCIP is defined as a tunneling protocol for connecting geographically distributed Fibre Channel SANs transparently over IP networks. It uses TCP/IP as the transport while keeping Fibre Channel fabric services intact. The FCIP standard is a cost-effective technology that uses widely deployed IP standards while enabling management of remote SANs through existing SAN management tools. This extremely powerful standard is designed to leverage the rapid advances in Gigabit Ethernet and emerging 10 Gigabit Ethernet technologies as well as tested local-loop and long-haul technologies. As bandwidth increases in local networks and prices continue to fall, FCIP becomes even more cost-effective.

A simple tunneling Configuration is illustrated below.

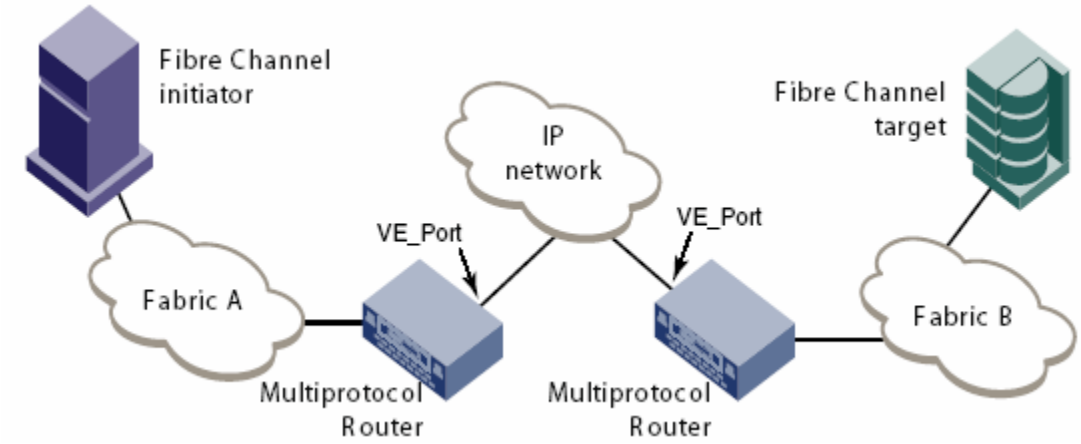

FCIP Tunneling introduces the following concepts:

Tunnel

An FCIP tunnel carries Fibre Channel traffic (frames) over IP networks such that the Fibre Channel fabric and all Fibre Channel devices in the fabric are unaware of the IP network's presence. Fibre Channel frames "tunnel" through IP networks by dividing frames, encapsulating the result in IP packets upon entering the tunnel, and then reconstructing them as they leave the tunnel.

VE_Port

Special types of ports, called VE_Ports (virtual E_Port), function somewhat like an E_Port. The link between a VE_Port and a VE_Port is called an interswitch link (ISL). You can configure multiple ISLs from a Connextrix MP-7500B or Connextrix ED-48000B with a PB-48K-18i blade. After you configure the VE_Ports on either two Connextrix MP-7500Bs or Connextrix ED-48000B s with the PB-48K-18i blade, an FCIP connection is established between them. VE_Ports do not prevent fabric merging. Using an EX_Port is one way to prevent fabrics from merging.

VEX_Port

A VEX_Port enables routing functionality through an FCIP tunnel. VEX_Ports are virtual FC_Ports that are exposed by FCIP tunnels connecting to either the Connextrix MP-7500B or with Connextrix ED-48000B and a PB-48K-18i blade; they run interfabric links (IFLs) as EX_Ports to enable Fibre Channel router capability. You can have up to 8 VEX_Ports per GbE on the Connextrix ED-48000B with a PB-48K-18i blade.

GbE

Gigabit Ethernet ports are available on the Connextrix MP-7500B and Connextrix ED-48000B with a PB-48K-18i blade. These ports support FCIP with link speeds up to 1-Gbit/sec. Each GbE port (ge0, ge1) supports up to 8 FCIP tunnels.

Combining FCR/FCIP

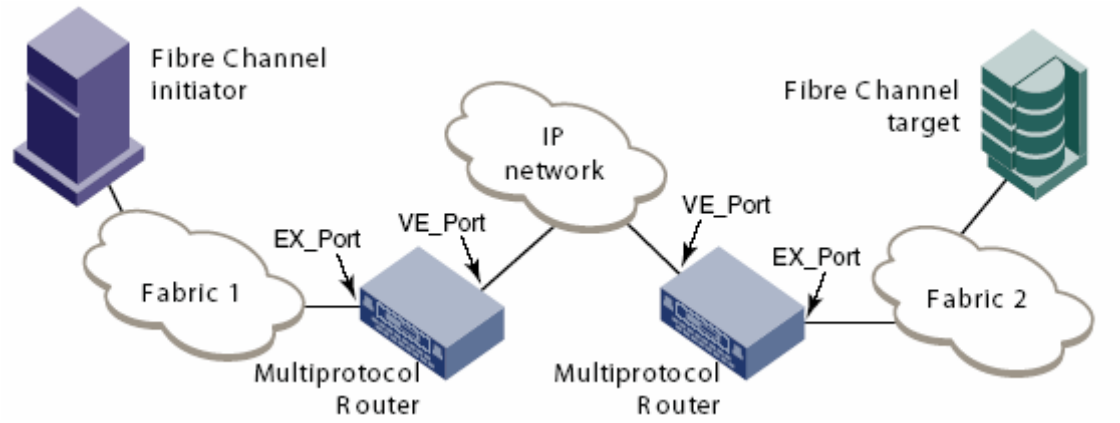

FCIP for Distance Connectivity

Solution

- Ability to share resources and move data between geographies
- Leverage IP MAN/WAN in conjunction with SAN resources
- Integrate FCIP tunneling within Brocade framework
- Minimize protocol conversion events
- Standards-based solution

Benefits

- Simplified management, integrated within Brocade framework
- Reduced management, training and troubleshooting costs when compared to multi-vendor solutions
- Extend remote replication and backup functions to long distances
- Reduce cost by leveraging IP networks
- Flexible, cost effective foundation for utility computing architecture

FCIP Overview Summary

With FCIP, FC frames are encapsulated in TCP/IP packets. It is a point-to-point tunneling protocol which transparently interconnects two or more fabrics across intermediate IP networks.

Brocade's integrated FCIP products have valuable features, including but not limited to:

- Full integration with Brocade switch and network management software
- Full integration with the Brocade Fibre Channel routing feature (LSANs)
- Traffic Shaping to support efficient performance in WAN environments
- Jumbo Ethernet frames to support more efficient throughput
- Load balancing for FCIP links to provide greater aggregate bandwidth

Brocade does not recommend FCIP for use in every possible distance extension scenario as no technical solution can be "all things to all people". FCIP has inherent performance, reliability, data integrity, and manageability limitations when compared to native FC solutions, but native FC extension may not always be available. FCIP *does* have its place, and it *can* support very long distances and high performance, as long as the carrier network is extremely high performance and reliable.

EMC FCIP Hardware Product Overview

The industry-leading Brocade family of fabric switches connects servers and storage devices through Fibre Channel SAN fabrics. These high-speed, robust storage networks enable organizations to access and share data in a high-performance, highly available, manageable, and scalable manner.

To help protect existing investments, the EMC Connectrix® family of products is fully forward and backward compatible. This capability enables organizations to migrate from 1 and 2 Gbit/sec to 4 Gbit/sec SAN environments and deploy a highly scalable core-to-edge storage networking infrastructure.

The **Connectrix PB-48K-18I** FC Routing and FCIP blade coupled with the **EMC Connectrix MP-7500B FC Router** and FCIP switch augments this family by providing a solid range of storage networking capabilities uniquely constructed to maximize the value of SANs. They provide new connectivity options by supporting FC routing and FCIP services to increase SAN functionality and versatility within the data center and across geographies. The primary advantage is the ability to connect devices between two or more fabrics without merging them, whether that is across native FC or FCIP—thereby providing a more flexible storage networking foundation for implementing value-added services across the data center infrastructure.

Product Review

EMC Connectrix MP-7500B Switch, Connectrix PB-48K-18I Blade (For Connectrix ED-48000B Director)

Highlights

- Unified SAN architecture
- Dual GigE FCIP ports for connecting to IP networks

- 1, 2, or 4 Gbit/sec Fibre Channel routing ports to connect Fibre Channel SAN fabrics or storage devices

- **Superior performance:**
  - Industry-first 4 Gbit/sec Fibre Channel routing
  - Optimizations for high-latency, low-bandwidth WANs
  - Line-rate performance for high-speed WANs
  - Hardware-based compression
- **Efficiency:**
  - Efficient encapsulation of Fibre Channel into IP
- **High scalability:**
  - 16 4 Gbit/sec Fibre Channel routing ports, 2 GigE FCIP ports
  - 8 FCIP tunnels per port

Connectrix PB-48K-18I
Director Multi Protocol Blade

EMC Connextrix MP 7500B
SAN Router

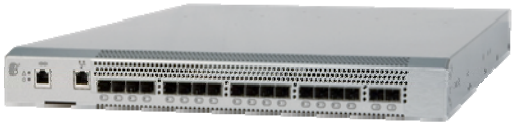

PHOTOS ABOVE

Connectrix PB-48K-18I Blade

(For Connectrix ED-48000B Director)

EMC Connectrix MP-7500B Switch

HARDWARE OVERVIEW

Connectrix PB-48K-18I Blade (Connectrix ED-48000B Director)

- Enterprise solution for SAN director:
  - Sixteen 1, 2, and 4 Gbit/sec Fibre Channel routing ports
  - Two Gigabit Ethernet FCIP ports
  - Onboard blade processor
  - High availability
  - Hot-swappable
  - Hot code activation for Fibre Channel routing
    (FCIP may have disruption)

- FCIP distance extension

- Hardware compatibility:
  - Supported with the following:
    - Connectrix ED-48000B Director and control processors
    - Connectrix ED-48000B blades:
      - FC4-16, 16-port 1, 2, and 4 Gbit/sec blade
      - FC4-32, 32-port 4 Gbit/sec blade

- Two blades per chassis:
  - 32 ports of Fibre Channel routing
  - 4 ports of FCIP (32 virtual FCIP tunnels)

- High availability:

– A standby control processor takes over for a failed active control processor and becomes the active one (traffic forwarding continues during failover)

EMC Connectrix MP-7500B Switch

- Standalone solution for enterprise and midrange SAN environments:
  - Sixteen 1, 2, and 4 Gbit/sec Fibre Channel routing ports
  - Two Gigabit Ethernet FCIP ports
  - Fixed configuration
  - Onboard processor
  - Redundant fans
  - Redundant power supply
  - Hot code activation for Fibre Channel routing (FCIP may have disruption)

- FCIP distance extension (License required)

EMC CONNECTRIX 7500 HARDWARE OVERVIEW (FRONT)

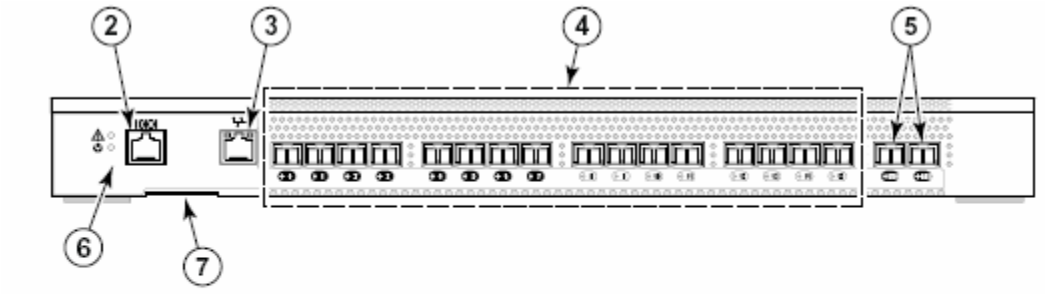

1	SilkWorm 7500	5	GbE ports (2)
2	Console Managment Port	6	System Status LED (top)
3	Ethernet Management Port		System Power LED (bottom)
4	Fibre Channel Ports (16)	7	IP Address pull out tab

EMC CONNECTRIX PORT OVERVIEW (FRONT)

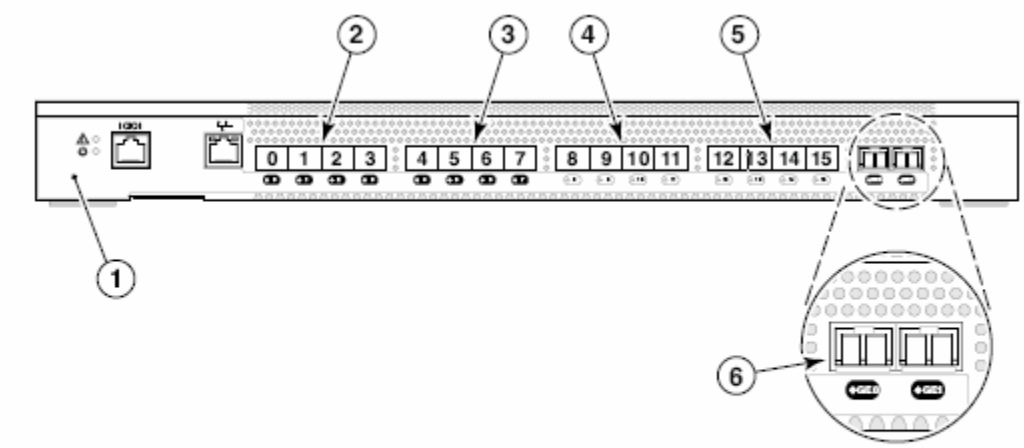

1	SilkWorm 7500	4	Fibre Channel Ports 8 through 11
2	Fibre Channel Ports 0 through 3	5	Fibre Channel Ports 12 through 15
3	Fibre Channel Ports 4 through 7	6	GbE ports (2)

EMC CONNECTRIX HARDWARE OVERVIEW (BACK)

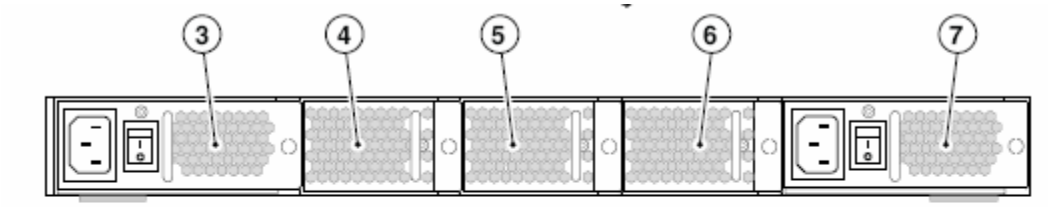

1	SilkWorm 7500	5	Fan Assembly #2
2	Nonport Side of Switch	6	Fan Assembly #1
3	Power Supply #2	7	Power Supply #1
4	Fan Assembly #3		

EMC CONNECTRIX PB-48K-18i

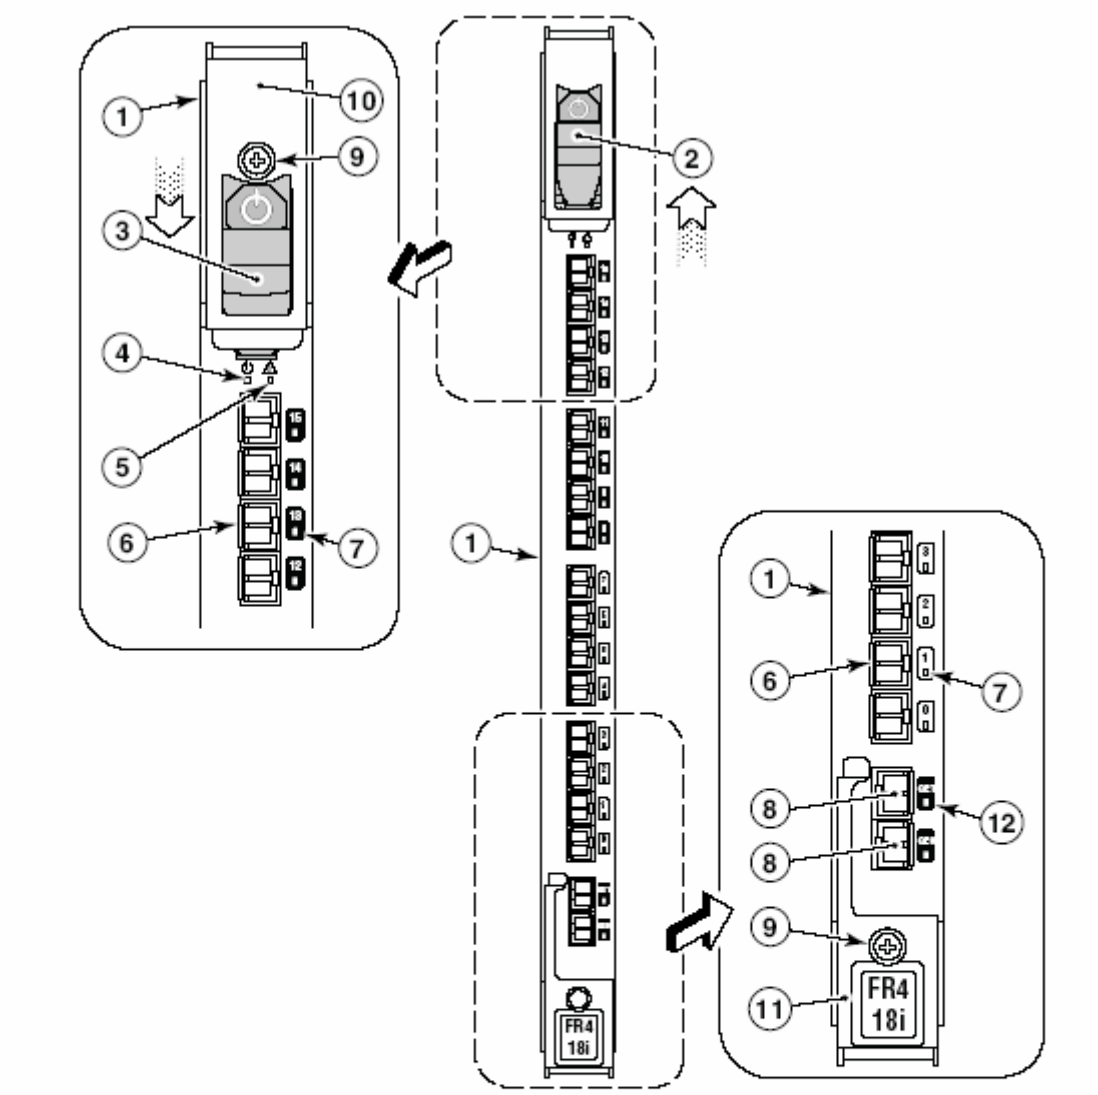

1	FR4-18i Blade	7	Fibre Channel Port Status LED
2	ON/OFF Slider Switch (in the ON position)	8	GbE Ports (x2)
3	ON/OFF Slider Switch (in the OFF position)	9	Thumb Screw
4	Blade Power LED	10	Upper Ejector
5	Blade Status LED	11	Lower Ejector
6	Fibre Channel Port	12	GbE Port Status LED

The Fibre Channel ports are numbered from bottom to top, in eight-port groups, and are numbered on the faceplate.

It is a best practice to review the current EMC Topology Guide and any affiliated software and hardware release notes such as Connectrix manuals and Fabric OS Guides when designing an EMC FCIP configuration. Also, it is critical to validate the current EMC tested Fabric OS before upgrading the operating system to ensure supportability and compatibility.

Below are some important notes from the EMC Topology Guide.

Note: E-Lab™ Navigator describes the latest supported configurations and minimum code requirements.

Symmetrix® setup: Symmetrix SRDF® ports should be configured as standard Fibre Channel SRDF ports. In a Fibre Channel environment, the MP-7500B provides all the services of a Fibre Channel switch, similar to those provided by any other Fibre Channel switch.

CLARiiON® setup: CLARiiON MirrorView™ ports should be configured as standard Fibre Channel MirrorView ports.

Note these configuration rules:

- The MP-7500B can be used as part of a DR (disaster recovery) and/or data migration SAN.
- FCIP can be used only as a backbone connection between the routers.
- SRDF, MirrorView, and SAN Copy™ are supported.
- Host I/O across the FCIP link can be supported if the application can tolerate the latency incurred due to the FCIP link.

- N_Port devices can be connected only to the Fibre Channel switches in the fabric connected to the router.

FCIP Implementation

FCIP Setup – FOS 5.1 or greater

Ver 1.0

Please read these notes before you start.

IMPORTANT NOTES:

FCIP TCP PORTS IMPORTANT NOTES:

TCP Ports 3225 & 3226 need to open for FCIP Connectivity. If a listener endpoint is configured, port 3227 is used.

Before creating a TCP connection to a peer FCIP Entity the FCIP_LEP needs a static IP address, a TCP port (TCP port 3225 is used for FCIP COS F traffic and TCP port 3226 is used for COS 2,3 traffic), the expected WWN of the other end of the link, and TCP parameter and Quality of Service (QoS) information.

FCIP / FCR IMPORTANT NOTES:

NOTE: Fabrics connected through FCIP merge if the ports are configured as VE_Ports, and do not merge if they are configured as VEX_Ports. If VE_Ports are used in a Fibre Channel Routing Services backbone fabric configuration, then the backbone fabric merges but the EX_Port attached to edge fabrics do not merge.

The port types for FCIP tunneling are either VE_Port or VEX_Port.

- An FCIP tunnel using VE_Ports will merge the two fabrics.
- An FCIP tunnel using a VEX_Port will not merge the fabrics. A VEX_Port can connect only to a VE_Port.

The IP encapsulation of the Fibre Channel frame on one port and the reconstruction of Fibre Channel frames on the other port are transparent to the initiator and target.

NOTE: If using FCIP in your FC-FC Routing configuration, you must first configure FCIP tunnels.

Valid IFL Connections (when using FC-FC Routing Services)

FC Routing Port Type	Port Type at Destination IFL
VEX_Port	VE_Port
EX_Port	E_Port

Once a tunnel is created, it defaults to a disabled state. Configure the VE_Port or VEX_Port. After the appropriate ports are configured, enable the tunnel.

Secure Fabric OS, Management Server Platform services, and interopmode are *not* supported in the backbone fabric. FOS 5.2 and above support FC-FC connections to McDATA Fabrics. Fabric OS v5.2.0 furnishes the FC router with the ability to connect to McDATA fabrics in both McDATA Open or McDATA Fabric mode.

Configure the FCIP ports on the FR4-18i blades as VEX_Ports, and the FCIP ports on the SilkWorm 7500 routers as VE_Ports. This allows the tunneled FC links to form as IFLs instead of ISLs: a routed FC SAN, or “Meta SAN” topology.

Once the switch (SW7500 & FR4-18i) is configured, the **switchShow** command displays 32 Fibre Channel ports (port numbers 0 through 31) and 2 GbE ports. The first 16 Fibre Channel ports are physical ports on the SilkWorm 7500, Ports 16-23 are virtual ports associated with the GE0 physical GbE link and ports 24-31 are virtual ports associated with GE1 physical GbE link. The GbE ports are displayed as **ge0** and **ge1** and are not assigned port numbers or area numbers.

To Obtain SW7500 SN & Uptime information enter the chassisshow command.

Overview for FCIP Setup

Review VE, VEX & EX port concepts before designing FCIP connectivity

- Simplicity vs. Isolation
- Simplicity vs. Scale
- Verify Backbone-Edge Device Requirements/Constraints

Review Application requirements for traffic across FCIP connections

- Is the Storage Application qualified with the Brocade FCIP?
- Is FastWrite desired (and will the application be doing writes?)
- Is Tape Pipelining desired? (multiple Pipeline support is limited!)

Confirm IP Network capabilities regarding:

- Firewalls that could block TCP ports 3225/3226/3227 (and possibly 4112/4113)
- Support for Jumbo Frames (max MTU 2284)
- Connectivity: will optical or Copper SFPs be used?
- What kind of bandwidth is available?
- What kind of latency is on the line?
- What kind of packet loss is on the line?
- Is the line predictable or variable in performance?

Review IP performance

- Use the PortCmd ipperf tool (FOS 5.2 and greater) to provide WAN performance analysis once GE ports are configured (see below)

Setup LSAN zoning if using EX or VEX ports.

- Use LSAN zones to share devices across routed fabrics

Configuring FCIP Tunnels

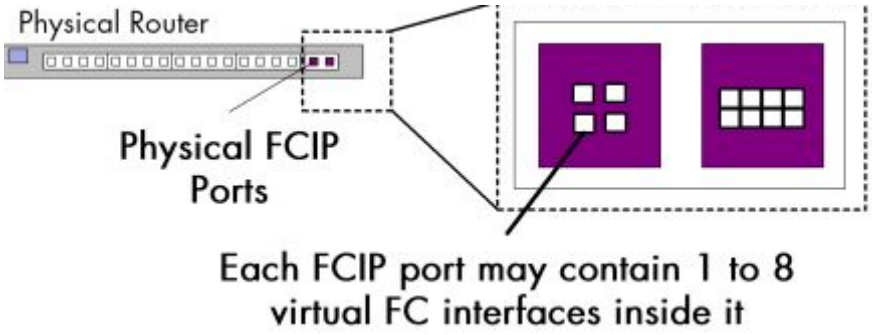

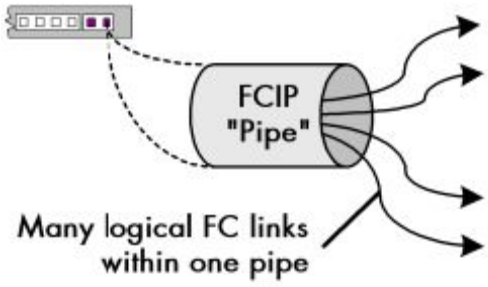

FCIP Up to 8 Logical instances per GE port example

Following are the steps for configuring an FCIP tunnel:

1. Define the IP Interface of the GbE Port
2. Add IP Routes on a GbE Port (Optional)
3. Verify IP Connectivity
4. Configure FCIP Tunnels
5. Verify the FCIP Tunnel Configuration
6. View and Enable Persistently Disabled Ports
7. Verify E_Port is Online

1. Define the IP Interface of the GbE Port

switch:admin> **portcfg ipif 8/ge0 create xxx.xxx.xxx.40 255.255.255.0 1500**
switch:admin> **portcfg ipif 8/ge0 create xxx.xxx.xxx.41 255.255.255.0 1500**
switch:admin> **portshow ipif 8/ge0**

portcfg *action* [*slot/*][**ge**]*port args*
Action: ipif [*slot/*][**ge**]*port args*
args for **ipif** include:
create *ipaddr netmask mtu_size*
Creates IP interfaces.
delete *ipaddr*
Deletes IP interfaces.
Action: **fciptunnel** [*slot/*][**ge**]*port args* [*optional_args*]
args for **fciptunnel** include:
create *tunnel_id remote_ipaddr local_ipaddr comm_rate*
Creates FCIP tunnels. *optional_args* for **create** include:

-c Enables compression on the tunnel specified.

-k *timeout*

Specifies the keep alive timeout, in seconds. *timeout* values are 8 to 7,200; default is 10.

-m *time* Specifies the minimum retransmit time, in milliseconds. *time* values are 20 to 5,000; default is 100.

-n *remote_wwn*

Specifies the remote-side FC entity WWN.

-r *retransmissions*

Specifies the maximum retransmissions. *retransmissions* values are 1 to 8; default is 8.

-s Disables selective acknowledgement code (SACK) on the tunnel specified.

-w Enables WAN TOV on the tunnel specified.

delete *tunnel_id*

Deletes FCIP tunnels.

2. Add IP Routes on a GbE Port (Optional)

switch:admin> **portcfg iproute 8/ge0 create 192.168.101.0 255.255.255.0 xxx.xxx.xxx.1 1**

switch:admin> **portcfg iproute 8/ge0 create 192.168.102.0 255.255.255.0 xxx.xxx.xxx.1 1**

portshow iproute 8/ge0

portcfg *action* [*slot/*][**ge**]*port args*

Action: iproute [*slot/*][**ge**]*port args*

args for **iproute** include:

create *ipaddr netmask gateway_router metric*

Creates IP routes.

delete *ipaddr netmask*

Deletes IP routes.

3. Verify IP Connectivity

portcmd ping 8/ge0 -s xxx.xxx.xxx.40 -d xxx.xxx.xxx.50

portcmd *action* [*slot*/]**ge***port args*

Action: ping [*slot*/]**ge***port* **-s** *source_ip* **-d** *destination_ip* [**-z** *size* **-n** *num_requests*]

Pings a destination IP address from one of the source IP interfaces on the GbE

port. Valid arguments include:

-s *source_ip*

Specifies the source IP interface that originates the ping request.

-d *destination_ip*

Specifies the destination IP address to which to target the ping request.

-z *size* Overrides the default packet size to some fixed size in bytes. The size of

the ping request must be less than the configured MTU size on the IP interface (see **portcfg** for details on setting the MTU size).

-n *num_requests*

Generates specified number of ping requests.

slot For bladed systems only, specifies the slot number.

ge*port* Specifies the port number of the GbE port on the blade.

4. Configure FCIP Tunnels

switch:admin> **portcfg fciptunnel 8/ge0 create 2 xxx.xxx.xxx.50 xxx.xxx.xxx.40 100000**

switch:admin> **portcfg fciptunnel 8/ge0 create 3 xxx.xxx.xxx.51 xxx.xxx.xxx.41 100000**

portcfg *action* [*slot/*][**ge**]*port args*

Action: **fciptunnel** [*slot/*][**ge**]*port args* [*optional_args*]

args for **fciptunnel** include:

create *tunnel_id remote_ipaddr local_ipaddr comm_rate*

Creates FCIP tunnels. *optional_args* for **create** include:

-c Enables compression on the tunnel specified.

-k *timeout*

Specifies the keep alive timeout, in seconds. *timeout* values are 8 to 7,200; default is 10.

-m *time* Specifies the minimum retransmit time, in milliseconds. *time* values are 20 to 5,000; default is 100.

-n *remote_wwn*

Specifies the remote-side FC entity WWN.

-r *retransmissions*

Specifies the maximum retransmissions. *retransmissions* values are 1 to 8; default is 8.

-s Disables selective acknowledgement code (SACK) on the tunnel specified.

-w Enables WAN TOV on the tunnel specified.

delete *tunnel_id*

Deletes FCIP tunnels.

5. Verify the FCIP Tunnel Configuration

portshow fciptunnel 8/ge0 all

portshow ipif 8/ge0

6. View and Enable Persistently Disabled Ports

portenable 8/ge0

switch:admin> **portcfgshow 8/ge0**

Area Number: 8
Speed Level: AUTO
Trunk Port ON
Long Distance OFF
VC Link Init OFF
Locked L_Port OFF
Locked G_Port OFF
Disabled E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable ON
NPIV capability ON
EX Port ON

switch:admin> **portcfgpersistentenable 8/ge0**

switch:admin> **portcfgshow 8/ge0**

Area Number: 8
Speed Level: AUTO
Trunk Port ON
Long Distance OFF
VC Link Init OFF
Locked L_Port OFF
Locked G_Port OFF
Disabled E_Port OFF
ISL R_RDY Mode OFF
RSCN Suppressed OFF
Persistent Disable OFF

NPIV capability ON
EX Port ON
switch:admin>

7. Verify E_Port is Online

switch:admin> **switchshow**

WAN Analysis of the IP Network (end-to-end)

1) Configure GigE ports with IP interfaces
2) Run the 'portCmd' tools with the --ipperf option (FOS 5.2 and above)

portcmd ipperf 4/ge0 -s xxx.xxx.xxx.50 -d xxx.xxx.xxx.40 -S
ipperf to xxx.xxx.xxx.40 from IP interface xxx.xxx.xxx.50 on 4/0:3227
30s: BW:115.67MBps WBW(30s): 56.70MBps Loss(%):0.0 Delay(ms):2
PMTU:1500

portcmd --ipperf [*slot*/]**ge***port* **-s** *src_ip* **-d** *dst_ip* **-S** | **-R** [**-i** *interval*] [**-p** *port*] [**-r** *committed_rate*] [**-**
t *running_time*] [**-z** *size*]

--ipperf output displays:
Sampling frequency(s)
This is the interval specified with the **--ipperf** command with **-i** option or the default (30s).
BW
This is the bandwidth measured in the last interval. Bandwidth is defined

as the total packets and bytes sent. Note: BW represents what the FCIP tunnel / FC application sees for throughput rather than the Ethernet on-the-wire bytes.

WBW

This is the weighted bandwidth currently with a gain of 50%.

Loss (%)

This is the number of TCP retransmits. This number is an average rate over the last display interval.

Delay (ms)

This is the TCP smoothed RTT and variance estimate in milliseconds.

PMTU

Path MTU; This is the largest IP-layer datagram that can be transmitted over the end-to-end path without fragmentation. This value is measured in bytes and includes the IP header and payload. Note: There is limited support for black hole PMTU detection. If the Jumbo PMTU (anything over 1500) does not work, **--ipperf** will try 1500 bytes (the minimum PMTU supported for FCIP tunnels). If 1500 PMTU fails, **--ipperf** will give up. There is no support for aging. During black hole PMTU detection the BW, WBW, Loss and PMTU values printed may not be accurate.

FCR

Backbone Configuration

1) Disable all active EX and VEX ports with “portdisable [<slot>/]<port number>”
2) Disable the FCR service on the switch with “fcrdisable”
3) Configure the backbone Fabric ID with the command “fcrconfigure”
4) Enable the FCR service on the switch with “fcrenable”
5) Enable the ports that were disabled in step 1 with “portenable [<slot>/]<port number>”

EX<->E

1) Disable all ports that will become an EX_port with “portdisable [<slot>/]<port number>”

2) Configure the EX_Port “portcfgexport <port> -a <1|2> -f <Fabric ID> -d <domain id> -p <pidformat>”

a. (1-enable, 2-disable)

Note: The following optional parameters may be used: [-r r_a_tov] [-e e_d_tov] [-d domain] [-p 0-native 1-core 2-extended edge] [-t 1-Enable 2-Disable]. The –t will negotiate fabric parameters, making other parameters listed here unnecessary.

3) Enable the port for the EX_Port “portenable [<slot>/]<port number>”

4) Connect EX_Port to edge fabric if not already connected

5) Check fabric with fcrfabricshow on backbone fabric and fabricShow on edge fabric to confirm connections.

6) Repeat 1 – 5 as necessary for each edge fabric.

VEX<->VE (*This will not work with the SilkWorm 7420***)**

1) Configure ports to be VE ports first as mentioned in the FCIP section

2) Disable all ports that will become a VEX_port with “portdisable [<slot>/]<port number>”

3) Configure the VEX_Port “portcfgvexport <port> -a <1|2> -f <Fabric ID> -d <domain id> -p <pidformat>”

a)(1-enable, 2-disable)

Note: The following optional parameters may be used: [-r r_a_tov] [-e e_d_tov] [-d domain] [-p 0-native 1-core 2-extended edge] [-t 1-Enable 2-Disable]. The –t will negotiate fabric parameters, making other parameters listed here unnecessary.

4) Enable the port for the VEX_Port "portenable [<slot>/]<port number>"

5) Connect VEX_Port to edge fabric if not already connected

6) Check fabric with fcrfabricshow on backbone fabric and fabricShow on edge fabric to confirm connections.

7) Repeat 1 – 5 as necessary for each VEX <-> VE edge fabric connection.

Configure LSAN Zones

1) Identify devices to share between fabrics

2) Create aliases using **Port WWN**

3) Create zones on edge fabrics that begin with “lsan_” or “LSAN_” and add to configuration

Note: All zone members of an LSAN zone must be Port WWN.

4) Run the following commands on the router

b. fcrfabricshow

c. fcrphydevshow

d. fcrproxydevshow

e. fcrrouteshow

f. fcrxlateconfig [-r remove] <edgeFabricId> <remoteFabricId> <preferred-DomainId>

i. The preferred Domain ID will represent the remote fabric in the edge fabric. This will be needed if specific domain ids need to be set for xlate domain.

g. lsanzoneshow

FCIP PERFORMANCE TIPS

FC ports run at 4Gbits and FCIP ports run at 1Gbit. Theoretically, FC is 4x faster. However, line rate is not the only variable in determining application performance, and it is important for SAN designers to keep this in mind when designing FCIP solutions.

First, consider the line rates of all devices between the FCIP endpoints. Picture a SAN in which four 1Gbit FCIP links enter a Gigabit Ethernet switch, then traverse a T1 link to another Gigabit Ethernet switch, where they connect to four more FCIP links. Even though the solution provides 4Gbits of FCIP "line rate" performance, the end-to-end throughput will be 1.5Mbps or less: the speed of the T1 link.

The "or less" part of that statement comes from several areas. If other devices (e.g. IP applications like ftp or http) are running over the same T1, then they will be taking up some of the bandwidth. If the distance between the sites it great enough, latency might create application-level effects. And even if nothing else is using the link at all, and the distance is short, FCIP requires using up a portion of the bandwidth for network headers. While FCIP is much more efficient than iSCSI, it does still require adding quite a bit of protocol overhead when compared to native FC.

This is not intended to be an indictment of FCIP. It simply means that FCIP solutions require more performance analysis and tuning than their native FC counterparts. SAN designers should look beyond mere interface bandwidth metrics before deciding on FCIP as a transport, and certainly should do so before provisioning specific numbers of FCIP ports and allocating WAN resources.

Specific parameters to consider before an implementation follow.

Fast Write

In an effort to work around the latency of long-distance links, Fast Write (also known as Write Acceleration) removes one of the two round-trips on a write-request by pre-acknowledging the SCSI write commit. Upon the first write request, the data is cached and the host is sent an acknowledgement. However, it does not actually transfer the data until the second part of the write request (scsi xfer_rdy) is received. While Fast Write has benefits in dealing with high-latency distance interconnects, it belongs in the higher end of the network stack, not on the switch. Consider that if the network acknowledges a write, but the write does not actually complete, the host is not aware of the situation and may not be able to recover. Also, many data applications have Fast Write built in, making it an unneeded and expensive feature to implement in the switch.

Supported Fastwrite & Tape Pipelining Configurations

Consider the configurations shown in the two figures below to understand the supported configurations. In both cases, there are no multiple equal-cost paths.

In the first figure on this page, there is a single tunnel with fastwrite and tape pipelining enabled. In the second figure (following page), there are multiple tunnels, but none of them create a multiple equal-cost path.

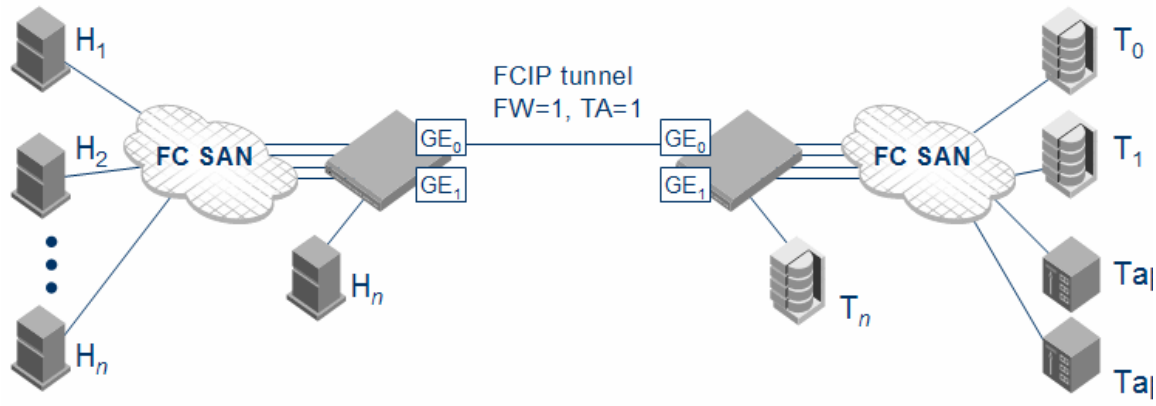

Single Tunnel, Fastwrite & Tape Pipelining Enabled

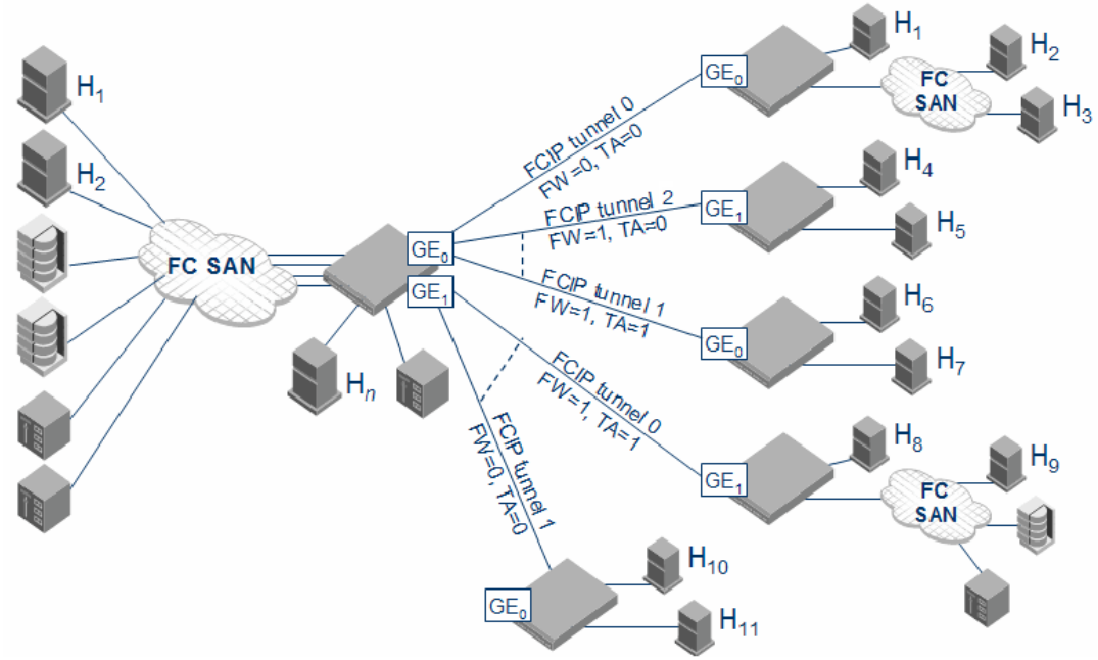

******IMPORTANT NOTE FOR SRDF USERS:**

SRDF already has a Fast Write feature built into it. Fast Write on the switch can interfere with Fast Write on the SRDF application, slowing throughput considerably. The use of fast write will not benefit you if you are using any of the extension products with SRDF.

Tape Pipelining

Tape Acceleration (also known as tape pipelining) refers to not only pre-acknowledging the Write Commit, but the entire Write Operation as well. The local switch will cache the data and immediately return 'Status=good'. It will then negotiate with its remote partner, buffer the data and stream it to the remote switch, which in turn will stream it to its local tape device.

Tape appears to be local regardless of distance. (possible because tape uses synchronization points making it easier to "spoof" or pre-acknowledge I/O without risking data corruption, as could happen if a Fast Write were inappropriately acknowledged).

Since I/O is cached on both ends of a long-distance link, a steady flow of data to and from the remote tape device is ensured by the maximum utilization of the FCIP link. The tape is kept busy by this steady flow of data over the network, avoiding the “shoe-shine” or unnecessary physical wear of the tape head.

Fastwrite	Tape Pipelining
Each GbE port supports up to 2048 simultaneous accelerated exchanges, which means *a total of 2048 simultaneous exchanges combined* for fastwrite and tape pipelining.	Each GbE port supports up to 2048 simultaneous accelerated exchanges, which means *a total of 2048 simultaneous exchanges combined* for fastwrite and tape pipelining.
Does not affect FICON traffic	Does not affect FICON traffic

Does not support multiple equal-cost path configurations (see "Supported Configurations").	Does not support multiple equal-cost path configurations or multiple non-equal-cost path configurations (see "Supported Configurations").
Class 3 traffic is accelerated with fastwrite.	Class 3 traffic is accelerated between host and sequential device.
	With sequential devices (tape drives), there are 1024 initiator-tape (IT) pairs per GbE port, but 2048 initiator-tape-LUN (ITL) pairs per GbE port. The ITL pairs are shared among the IT pairs. For example: • 2 ITL pairs for each IT pair as long as the target has two LUNs. • If a target has 32 LUNs, 32 ITL pairs for IT pairs. In this case, only 64 IT pairs are associated with ITL pairs. The rest of the IT pairs are not associated to any ITL pairs, so no tape pipelining is performed for those pairs. By default, only fastwrite-based acceleration is performed on the unassociated pairs.
	Does not support multiple non-equal-cost path between host and sequential device

Traffic Shaping

Users can optimize their existing network connections, avoiding the well-known TCP slow-start issue where packets are dropped due to threshold congestion. Users can set the maximum transmission rate, and traffic will never exceed the available bandwidth. This also provides Quality of Service (QoS) by ensuring storage traffic does not overrun other IP traffic on the WAN. Flow control/traffic shaping is key in avoiding FC vs. IP line speed mismatches, or network congestion in general.

Jumbo Packets

A basic Ethernet MTU is 1518 bytes. This forces an FC frame to be broken into two Ethernet packets when travelling via FCIP. A 2250 byte jumbo packet can accept an entire standard FC frame (2148 bytes) leading to improved FCIP performance and more

efficient network utilization. Many IP service providers can support jumbo packets throughout their network.

Exchange Based Trunking

Up to four links can be used to create an FCIP trunk to provide load balancing and failover for multiple exchanges. An FCIP trunk looks and feels like a single ISL from a Fibre Channel perspective.

(continue to next page)

FCIP Enterprise Customer Implementation

Case Study Overview:

Customer : Telcom

SAN Usage : DR/Replication

Ports : 8,000 +

Directors: 10 +

Switches : 100 +

Routers : 4

Devices : 1000+

Users : 600+

Problems to be addressed:

1. No DR & Meta-SAN Connectivity
2. Replication capabilities were behind the curve and exposed the customer to a high risk of downtime.
3. Technical expertise varied from department to department
4. Customers' growth became unmanageable
5. Performance issues were routine
6. No unified standards with SAN management concerns like zoning, switch naming, port assignments, etc
7. No significant documentation of the Meta-SAN

Solutions to the Problems:

1. Implemented Meta-SAN connectivity
   (FCIP through ***Connectrix PB-48K-18I and EMC Connectrix MP-7500B***)
2. Replication capabilities were tested and increased.
3. Departments to "hand over" SAN management to "centralized" SAN group.
4. Growth became manageable through sound SAN management
5. Ensure SAN interoptability.
6. Connectrix ED-48000B directors are the most efficient SAN directors in industry
7. Provided documentation and training on FCIP

Customer replication and DR design

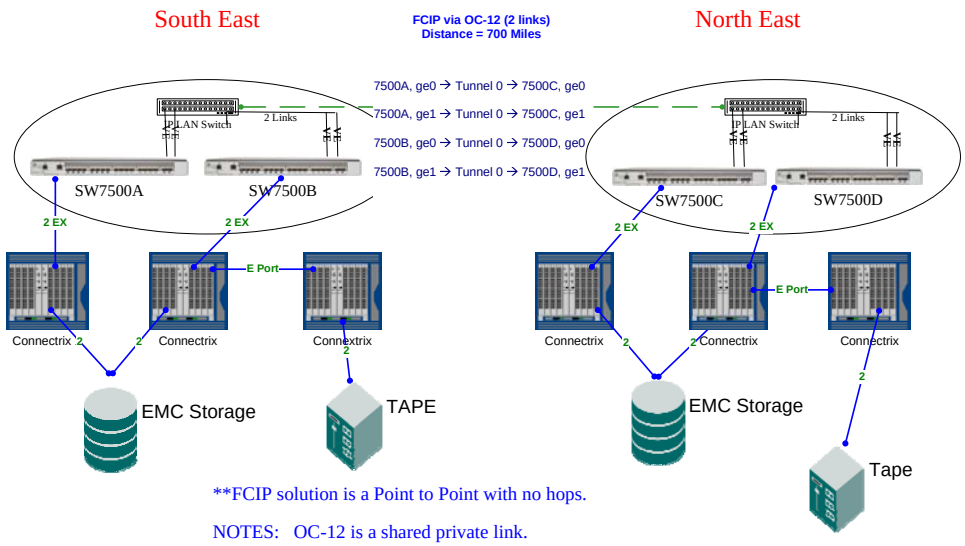

The Customer now has a well designed, scalable Meta-SAN that will allow them flexibility and room for growth in the foreseeable future.

Article Closing

As one of the first applications for storage networking based on TCP/IP, extending connectivity for Fibre Channel SANs over long distances is an essential component of disaster recovery, remote backup and other critical business solutions. Since IP networks have no inherent distance or speed limitations, moving storage data over TCP/IP enables a more flexible and robust alternative to direct Fibre Channel-to-Fibre Channel linkage.

Until more recently, business-continuance and disaster-recovery plans that transfer critical data to diverse locations have been feasible only for the largest enterprises. The costly WAN connections and equipment-transferring storage area network (SAN) traffic across significant distances have traditionally been accomplished with technology that puts the SAN traffic directly on a SONET system, directly onto dense wavelength-division multiplexing (DWDM) systems, or even onto dark fibers.

FCIP technology avoids all these expensive alternatives and makes use of the low-cost and ubiquitous IP network to transfer SAN data. This one technology brings true business continuance and disaster recovery within reach of the small-business and midsized-business (SMB) customer. It also provides a much lower-cost alternative for the large enterprise.

FCIP is a tunneling protocol that transports all FC ISL traffic. Similarly, FCIP uses TCP/IP as the transport protocol and IPSec for security.

A FCIP link tunnels all ISL traffic between a pair of FC switches, and may have one or more TCP connections between a pair of IP nodes for the tunnel end points. From the FC fabric view, an FCIP link is an ISL transporting all FC control and data frames between switches, with the IP network and protocols invisible. One can configure one

or more ISLs (using FCIP links) between FC switches using FCIP links.

Brocade solutions for FCIP are an industry first in the 4 Gig/sec market. Brocade solutions are fully complementary with EMC SAN solutions including the industry leading EMC Clariion and Symmetrix platforms.

FCIP Implementation is relatively non-complex if implemented correctly and requires only basic SAN planning. The main concern is validation of the distance extension and the correct analysis of the IP network configuration.

The major benefits of **Fibre Channel over IP (FCIP) are:**

- Full integration with Brocade switch and network management software
- Full integration with the Brocade Fibre Channel routing feature (LSANs)
- Traffic Shaping to support efficient performance in WAN environments
- Jumbo Ethernet frames to support more efficient throughput
- Load balancing for FCIP links to provide greater aggregate bandwidth
- Increased Meta-SAN distance and capacity

Backbone Fabric: A capability that enables scalable Meta SANs by allowing the networking of multiple routers that connect to the backbone fabric via E_Port interfaces. Devices attached to routers via F_Port or FL_Port, or imported via the iSCSI Gateway Service, are also considered part of the backbone. A backbone fabric is an intermediate network that connects two or more edge fabrics. It consists of at least one EMC Connectrix MP-7500B, Connectrix ED-48000B with a Connectrix PB-48K-18I blade and possibly a number of Fabric OS-based Fibre Channel switches. It also enables hosts and targets in one edge fabric to communicate with devices in the other edge fabrics. A backbone fabric enables hosts and targets in one edge fabric to communicate with devices in other edge or backbone fabrics.

Backbone-to-Edge Routing - Fibre Channel routers can connect to a common fabric–known as a *backbone fabric*–via E_Ports. A backbone fabric can be used as a transport fabric that interconnects edge fabrics. Fibre Channel routers also enable hosts and targets in edge fabrics to communicate with devices in the backbone fabric–this is known as *backbone-to-edge routing*. From the edge fabric's perspective, the backbone fabric is just like any other edge fabric. For the edge fabric and backbone fabric devices to communicate, the shared devices need to be presented to each other's native fabric. To do so, at least one translate phantom domain (switch) is projected into the backbone fabric. This translate phantom switch represents the entire edge fabric. The shared physical device in the edge has a corresponding proxy device on the translate phantom domain switch. Each edge fabric has one and only one xlate switch to the backbone fabric. The backbone fabric device communicates with the proxy devices whenever it needs to contact the shared physical device in the edge. The FC-FC routing service receives the frames from the backbone switches destined to the proxy device,
and redirects the frame to the actual physical device.

E_Port: A standard Fibre Channel mechanism that enables switches to network with each other.

Edge Fabric: A Fibre Channel fabric connected to a router via one or more EX_Ports. This is where hosts and storage are typically attached in a Meta-SAN.

Edge-to-Edge Routing - Occurs when devices in one edge fabric communicate with devices in another edge fabric through one or more Fiber Channel routers.

EX_Port: The type of E_Port used to connect a router to an edge fabric. An EX_Port follows standard E_Port protocols and supports FC-NAT but does not allow fabric merging across EX_Ports.

Exported Device: A device that has been mapped between fabrics. A host or storage port in one edge fabric can be exported to any other fabric through LSAN zoning.

Fabric: A collection of Fibre Channel switches and devices, such as hosts and storage.

Fabric ID (FID): Unique identifier of a fabric in a Meta-SAN. Every EX_Port and VEX_Port uses the FID property to identify the fabric at the opposite end of the IFL. You should configure all of the EX_Ports and VEX_Ports attached to the same edge fabric with the same FID. The FID for every edge fabric must be unique from each backbone fabric's perspective.

FCIP Tunneling Service: A service that enables SANs to span longer distances than could be supported with native Fibre Channel links. FCIP is a TCP/IP-based tunneling

protocol that allows the transparent interconnection of geographically distributed SAN islands through an IP-based network.

Fibre Channel: The primary protocol for building SANs. Unlike IP and Ethernet, Fibre Channel is designed to support the needs of storage devices of all types.

Fibre Channel Network Address Translation (FC-NAT): A capability that allows devices in different fabrics to communicate when those fabrics have addressing conflicts. This is similar to the "hide-behind" NAT used in firewalls.

Fibre Channel Router Protocol (**FCRP):** A Brocade-authored standards-track protocol that enables LSAN switches to perform routing between different edge fabrics, optionally across a backbone fabric.

FC-FC Routing Service: A service that extends hierarchical networking capabilities to Fibre Channel fabrics. It enables devices located on separate fabrics to communicate without merging the fabrics. It also enables the creation of LSANs.

Inter-Fabric Link (IFL): A connection between a router and an edge fabric. Architecturally, these can be of type EX_Port-to-E_Port or EX_Port-to-EX_Port. The former method is supported in the first release.

Logical Storage Area Network (LSAN): A logical network that spans multiple fabrics. The path between devices in an LSAN can be local to an edge fabric or cross one or more Routers and up to one intermediate backbone fabric. LSANs are administered through LSAN zones in each edge fabric.

LSAN Zone: The mechanism by which LSANs are administered. A Router attached to two fabrics "listens" for the creation of matching LSAN zones on both fabrics. If this occurs, it creates phantom domains and FC-NAT entries as appropriate, and inserts

entries for them into the nameservers on the fabrics. LSAN zones are compatible with standard zoning mechanisms.

Meta-SAN: The collection of all devices, switches, edge and backbone fabrics, LSANs, and Routers that make up a physically connected but logically partitioned storage network. This would simply be called "the network." However, an additional term is required to specify the difference between a single-fabric network ("SAN"), a multifabric network *without* cross-fabric connectivity (for example, a "dual-redundant fabric SAN"), and a multifabric network *with* connectivity ("Meta SAN").

Phantom Domains - The Fibre Channel Router emulates two levels of phantom domains. The first set of are front phantom domains. There is one front phantom domain from FCR to an edge Fabric.

The second level is a "translate phantom domain." The EX_Ports also present translate phantom domains in edge fabrics as being topologically behind the front domains; if the translate phantom domain is in a backbone fabric, then it is topologically present behind the Fibre Channel router because there is no front domain in a backbone fabric. The translate phantom domain is a router virtual domain that represents an entire fabric. You can achieve device connectivity from one fabric to another over the backbone or edge fabric through this virtual domain without merging the two fabrics.

Translate phantom domains are sometimes called "translate domains," or "xlate domains." A Connectrix PB-48K-18I blade is attached to an edge fabric using an EX_Port, will create translate phantom domains in the fabric corresponding to the imported edge fabrics with active LSANs defined. If you import devices into the backbone fabric, a translate phantom domain is created in the backbone device (in addition to the one in the edge fabric).

Proxy Devices:

An EMC Connectrix MP-7500B or Connectrix ED-48000B with a Connectrix PB-48K-18I blade achieves interfabric device connectivity by creating proxy devices (hosts and targets) in attached fabrics that represent real devices in other fabrics.

For example, a host in Fabric 1 can communicate with a target in Fabric 2 as follows:

- A proxy target in Fabric 1 represents the real target in Fabric 2.
- Likewise, a proxy host in Fabric 2 represents the real host in Fabric 1.

The host discovers and sends Fibre Channel frames to the proxy target. The EMC Connectrix MP-7500B or Connectrix ED-48000B with a Connectrix PB-48K-18I blade receives these frames, translates them, and delivers them to the destination fabric for delivery to the target.

The target responds by sending frames to the proxy host. Hosts and targets are exported from the edge SAN to which they are attached and, correspondingly, imported into the edge SAN reached through Fibre Channel routing.

Proxy ID - The port ID of the proxy device. A proxy device is a virtual device presented into a fabric by a Fibre Channel router. It represents a real device on another fabric. When a proxy device is created in a fabric, the real Fibre Channel device is considered to be imported into this fabric. The presence of a proxy device is required for inter-fabric device communication. The proxy device appears to the fabric as a real Fibre Channel device, has a name server entry, and is assigned a valid port ID. The port ID is only relevant on the fabric in which the proxy device has been created.

Router: A device that enables Brocade routing services.

Multiprotocol routing services: Available on the Connectrix PB-48K-18I blade and the EMC Connectrix MP-7500B Router, that includes the FC-FC Routing Service, and the FCIP Tunneling Service.

NR_Port: A port used as a source and destination address for frames traversing a backbone fabric. A normal E_Port (not an EX_Port) is used to connect a Router to a

backbone. An NR_Port appears to the rest of the backbone as a standard N_Port connected to the Router domain.

VE_Port: Virtual E_Port; an FCIP tunnel without routing is a VE_Port.

VEx_Port: The type of VE_Port used to connect a router to an edge fabric. A VEx_Port follows standard E_Port protocols and supports FC-NAT but does not allow fabric merging across VEX_Ports.

Author Biography

Joe Holbrook has been in the computer field since 1993 when he was exposed to several UNIX systems on board a US Navy ship (USS JFK, CV-67) while on active duty.

He has updated his career from the UNIX world to specialize in EMC SANS and the Disaster Recovery specializations. Joe has worked as a consultant/employee for numerous companies like Hewlett Packard, EMC, Northrup Grumman, CSC, Ibasis.net, Chematch.com, SAIC and Siemens Nixdorf.

He currently works for Brocade Communications as a Brocade Solutions Consultant where he is specialist in Fiber Channel SANS implementation. He also recently worked on a long term engagement for Brocade Communications as an onsite engineer at a large US Government site that had over 17,000 SAN ports with a large FC routing architecture.

He attended Central Texas University while in the Navy and received and AA in Electronics Technology. He received a Certificate in Total Quality Management from the United States International University (USIU) in San Diego. He received several Certificates in Information Systems, Project Management and a BSIS from the University of Massachusetts Lowell.