MCDATA™

Networking the world's business data™

Fibre Channel Advances Book

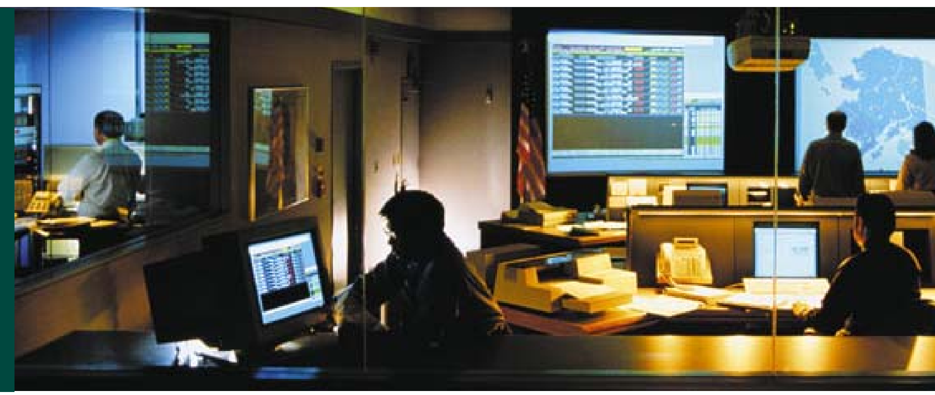

Scott Kipp
12/01/05

The FCIA Trilogy

- **Fibre Channel Storage Area Networks**
  - Focused on the benefits of FC in an overview
  - Showed practical applications of FC
- **Fibre Channel Fundamentals**
  - Focused on the basics of FC
  - Gave details of physical interfaces and Fabric Services
- **Fibre Channel Advances**
  - Will Focus on the advances of FC
  - Will be an overview of new technologies following the similar format of the other books

McDATA

Chapter Outline

Chapters:

- **1: Virtual Overview**
  – Review basic FC
  – Show applications of various new technologies
- **2: Port, Channel and Storage Virtualization**
  – N_Port_ID Virtualization
  – Virtual Channels
  – Fabric Based Virtualization - FAIS
- **3: Virtual Fabrics**
  – All the gory details
- **4: Global Links**
  – FCIP
  – FC-BB – ATM, SONET, GFP

Chapter Outlines

Chapters:

- **5: Inter-Fabric Routing**
  - Now covered in FC-IFR
- **6: Securing SANs**
  - Security update
- **7: Fibre Channel Roadmap**
  - Speeds and Feeds
  - Standards update
- **8: Conclusion**
  - Show large SAN example using techniques
  - Summary

Graphics are completed To here.

I've written to here, So I'm about done.

Schedule

- **Finish first draft and graphics in January**
- **Put draft out for review by February meeting**
  - Horst Truestedt has said he will review it
  - Any others?
- **Finish review by April Meeting**
- **Go to print around June**
- **McDATA has agreed to buy 1,000 books**
- **Following is Chapter 1 without graphics and Chapter 3 with graphics**
  - We’re using the same graphic formats from the first two books

Chapter 1: The Times They Are A-Changin'

The only constant is change. This well-known aphorism isn't completely correct though since the rate of change is not constant. Change is accelerating and people need to adapt to the change to succeed.

The business world is facing dramatic change and is continually evolving to new levels of complexity and scale. As companies tackle global markets with tight budgets and schedules, modern businesses are faced with making more out of less while facing communication challenges as their workers are distributed across multiple sites. Workers need to share unprecedented amounts of data and more sophisticated applications span the global enterprise. Many techniques that used to solve problems are no longer applicable and better ways of planning and acting are required to solve today's problems.

Paralleling the growth of the business, Fibre Channel Storage Area Networks (SAN) have also grown in ways that require new solutions. Old techniques that solved yesterday's problems don't apply to today's problems. Managing tens of Fabrics at various locations makes centralized management virtually impossible. Managing one very large Fabric also became nonsensical and beyond the comprehension level of most administrators. To tackle the complexity problem in large SANs, new concepts and management techniques were needed to keep the SAN manageable and efficient. This book discusses the many advances in Fibre Channel that have been developed to help administrators meet the challenges that they are seeing in today's business world.

Virtualize Everything

While the size of the corporation and their data infrastructure have grown, the size and capabilities of Fibre Channel equipment has grown to keep pace. Instead of a few 32-port Switches, large companies are deploying multiple 256-port Directors in redundant configurations across multiple sites. Storage subsystems have tens of ports and multiple tiers of storage within a cabinet that has petabytes of capacity. Servers host multiple operating systems and even more applications that share common ports across multiple Fabrics. Each of these devices are being virtualized in different ways to increase functionality.

One advance in Fibre Channel, known as N_Port_ID Virtualization, helps multiple applications run on a server over a single physical port. N_Port_ID Virtualization is a technique that enables different applications or users to have their own N_Port_IDs while using the same physical port. Virtual Channels is another technique that can be used by Inter-Switch Links (ISLs) to divide the physical link into multiple logical components to provide Quality of Service (QoS) over the ISL. N_Port_ID Virtualization and Virtual Channels are both explored in Chapter 2.

Another Virtualization technique is known as Fabric-based Virtualization. Also covered in Chapter 2, Fabric Based Virtualization enables multiple advanced replication and caching techniques that increases the performance of applications. From tiered storage implementations to remote backup, Fabric-based Virtualization can save time and money by increasing SAN performance.

Breaking Them Up and Putting Them Back Together

One of the most common problems faced by Information Technology (IT) departments today is the persistent growth of applications. With high growth, employee turnover and multiple acquisitions, administrators have difficulty managing heterogeneous applications, servers and storage. While tens of servers and applications with a few terabytes of storage could be handled by a single administrator, hundreds of servers and applications with hundreds of terabytes of storage takes a team effort for secure, uninterrupted application availability. The simple SAN has often grown into multiple complex SANs spread across multiple time zones.

Instead of building one large Fabric or multiple disparate Fabrics, Fibre Channel has enabled techniques to breakup large Fabrics and selectively put them back together. The large Fabric can be broken up with the concept of Virtual Fabrics that lets multiple Fabrics reside on a physical piece of hardware. Instead of having an 800-port Fabric on a few Directors and Switches as seen in Figure 1-1, Virtual Fabrics lets the administrator carve up the SAN into four 200-port Fabrics on the same infrastructure as shown in Figure 1-2. As with any type of organization, an intelligent division of the resources is required at some point.

Having multiple small Fabrics on the same physical hardware enables the administrator to break the SAN management into comprehensible pieces that can be managed more effectively. For example, instead of having one large Zone Set that many administrators alter, each administrator can be responsible for a limited set of devices and applications in a smaller Zone Set. This divide and conquer mentality of Virtual Fabrics will be explained in Chapter 3.

The story of Virtual Fabrics doesn't end here. After the Fabric is carved up into manageable pieces, selective devices can be configured to communicate between the separate Fabrics via Inter-Fabric Routing. Inter-Fabric Routing brings devices from different Virtual Fabrics back together without merging the Fabrics. As shown in Figure 1-3, an Inter-Fabric Router (IFR) sits above Virtual Fabrics and proxies devices in and out of Fabrics so that the devices can communicate while the Fabrics remain separate. Inter-Fabric Routing is discussed in Chapter 5.

Long Distances and Security

Besides scaling the number of ports in the SAN, the distance between ports has been increasing as companies scale to multiple sites. Partially driven by government mandates that require disaster recovery, corporations are using long distance links to connect remote sites at gigabit speeds. Fibre Channel frames can now be encapsulated and sent over Virtual ISLs on a variety of telecommunication or Internet Protocol networks. The use of non-Fibre Channel networks to transport Fibre Channel will be discussed in Chapter 4.

The extension of Fibre Channel outside of the glass house of the data center was one of the major reasons why the Fibre Channel community began implementing security. Securing Fibre Channel SANs is not limited to long distance links or any other single aspect of the SAN. To have true security, the integrity of the SAN needs to be viewed from a holistic perspective. Security is only feasible when multiple layers of the problem are considered and addressed. Chapter 6 reveals the developments in securing Fibre Channel so that companies can sleep well at night knowing their data is safe.

Putting It All Together

As seen in this quick overview of the book, Fibre Channel is evolving quickly into a highly configurable and secure network. Figure 1-4 shows how a Frame may be converted into various forms while it traverses multiple Fabrics. Multiple layers of virtualization and networks are traversed before the frame from the server reaches the storage device. Before the end of this book, each of these techniques will be considered in greater detail so that users can see the best place to use each capability.

Fibre Channel continues to advance at the physical layer as well. Fibre Channel is continually advancing to new speeds. Unlike Ethernet that jumps by a factor of 10 between generations, Fibre Channel typically advances by a factor of two. Fibre Channel takes evolutionary increases that are easier to deploy and can gradually replace the existing technology if they are backward compatible. The Fibre Channel Roadmap is disclosed in Chapter 7 along with other advances at the physical layer.

This book is intended for advanced users who have a basic understanding of Fibre Channel. The Fibre Channel Industry Association (FCIA) has already released books that cover the basic functionality of Fibre Channel in *Fibre Channel Fundamentals* and *Fibre Channel Storage Area Networks*. These books can be purchased through the FCIA at www.fibrechannel.org.

Chapter 3: Virtual Fabrics

After virtualizing N_Port_IDs, SCSI Ports and channels, it's only logical to virtualize Fabrics (Author's note: This is a hook to the previous chapter). Virtual Fabrics is another concept that divides physical links, devices and switches into multiple logical entities. Virtual Fabrics applies to N_Ports, F_Ports and E_Ports so that multiple logical levels exist within a physical entity. Physical N_Ports may support multiple Logical N_Ports while a physical E_Ports can be virtualized to support multiple Virtual Switches. The logical world allows great flexibility behind the limited physical world.

Virtual Fabrics allows a large Director to be carved into multiple Virtual Switches for isolation and resource sharing. Virtual Switches are isolated from each other but may share physical links to consolidate resources. One chassis may support multiple Virtual Fabrics and may be administered by different people and assigned to different departments and functionality. Virtual Fabrics enables physical resources to be managed at a finer level to consolidate resources and utilize them more effectively.

Virtual Fabric Tagging

Figure 3-1 shows how an Nx_Port can have multiple logical ports and send frames to two N_Ports in different Fabrics over a single link. To do this, the frames from the Logical N_Port must be tagged with the Virtual Fabric Tagging (VFT) Header. The Virtual Fabric Tag format is shown in Figure 3-2. The Virtual Fabric ID (VF_ID) can be seen as a new level of address space that resides above the 3-byte N_Port_ID. The 12-bit VF_ID enables 4096 Virtual Fabrics but most switches will only support tens of Virtual Switches.

The process of tagging frames begins when the ports exchange Virtual Fabric Parameters. N_Ports signify support for Virtual Fabrics by setting a bit in FLOGI exchanges while E_Ports declare support in the Exchange Switch Capabilities SW_ILS. After these signals have been sent, the ports exchange a bitmap that states which VF_IDs are supported. The logical AND of the two bitmaps decides which Virtual Fabrics will be enabled on the link. After the initialization, the Virtual Switches behave like separate physical switches.

Figure 3-1
Virtual Fabric Tagging

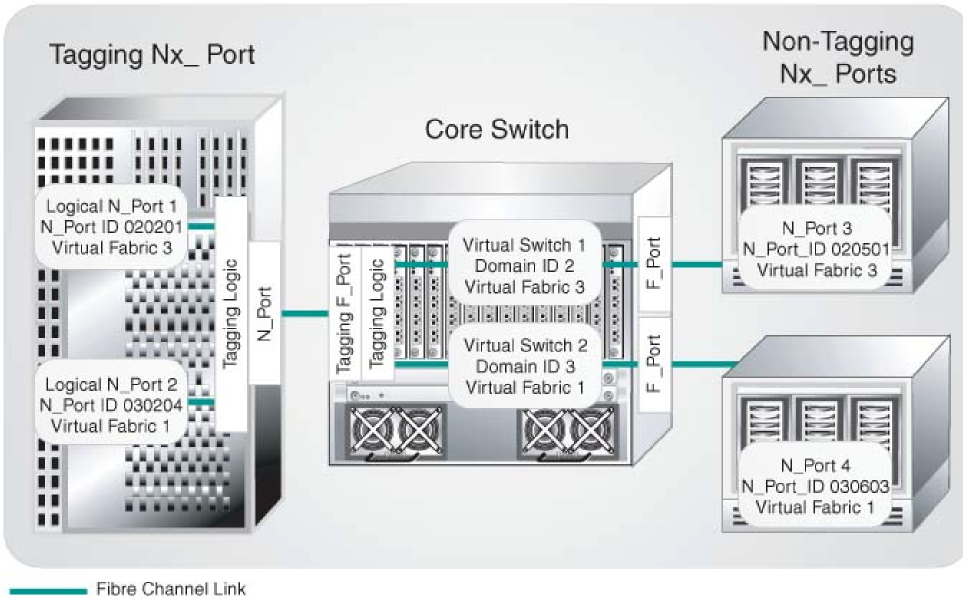

Figure 3-1 gives a good example of how the tagging logic works. When logical N_Port 1 sends a frame to N_Port 3, the Nx_Ports tagging logic add the Virtual Fabric Tag with the VF_ID of 3. The Switch receives the tagged frame and the tagging logic strips off the Virtual Fabric Tag and sends the frame to Virtual Switch 1. Virtual Switch 1 sends the untagged frames out to N_Port 3. The concept of the Virtual Fabric Tagging is transparent from N_Port 3's perspective. The first implementations of VFT were for ISLs that are shared between two Fabrics

Figure 3-2
Virtual Fabric Tagging Header

Start of Frame = 4 Bytes
Virtual Fabric Header = 8 Bytes
Fibre Channel Header = 24 Bytes
Data Field = 0 - 2, 112 Bytes
End of Frame = 4 Bytes

Table 39 - VFT_Header Format

Bits / Word	31..24	23	22	21..18	17	16	15..13	12..01	0
0	R_CTL	Ver		Type	R	R	Priority	VF_ID	R
1	HopCt			Reserved					

R_CTL = Routing Control Field = 50h
Ver = Version = 0
Type = Type of tagged Frame = 0
R = Reserved
VF_ID = Virtual Fabric Identifier
HopCt = Hop Count = 16h by default

Virtual Fabric Isolation

Virtual Fabrics don't need to use VFT Headers to find value. Physical Switches are often divided into multiple Virtual Switches to isolate management and improve scalability. Figure 3-3 shows how a Switch can be divided into multiple Virtual Switches for different departments. This example has the Finance department on a separate server than the Engineering department. Each Switch may have different security policies and different administrators for each Virtual Switch and Fabric.

Figure 3-3
Virtual Fabric Isolation

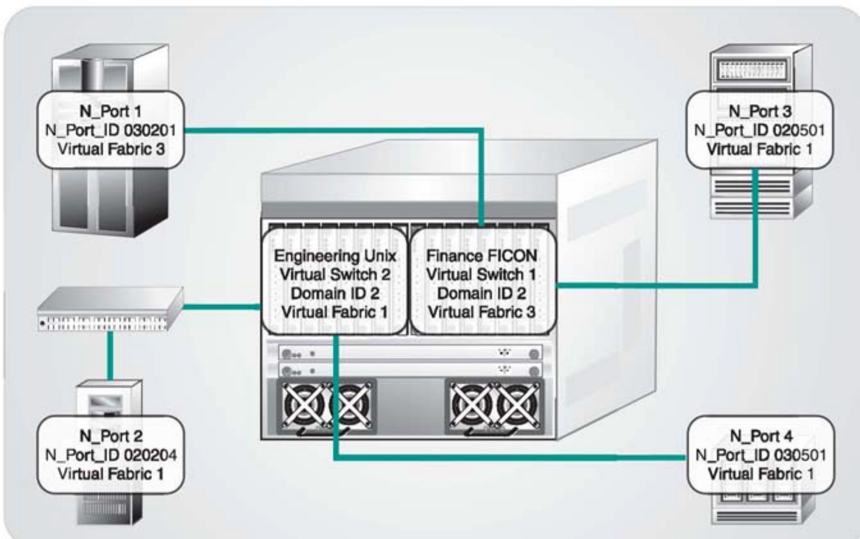

Virtual Fabric Tagging does not occur until ports are assigned to more than one Virtual Switch. Figure 3-4 shows how four ports on each Switch have been assigned to both Virtual Fabrics and the Port VF_ID of 1. The Port VF_ID is the VF_ID of the Fabric that initial frames are assigned to before tagging begins. If the attached port does not support Virtual Fabrics or tagging, then all frames will be associated with the Virtual Fabric of the Port VF_ID. Ports on both sides of the link must support a given VF_ID before traffic can flow between the two Virtual Fabrics.

Figure 3-4
Virtual Fabric Sharing

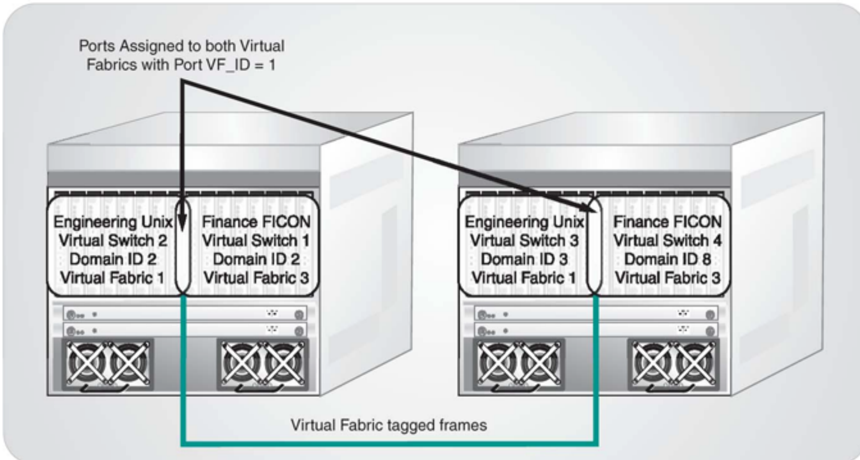

Conclusion

Virtual Fabrics is a handy tool for keeping the size of Fabric Services to a manageable level. Instead of having Zone Sets with thousands of Zones, Virtual Fabrics lets administrators carve physical infrastructure into multiple Fabrics that limits the size of the Zone Set to tens or hundreds of Zones. The scale of other Fabric Services like Name Server and Fabric Configuration Server may also be optimized with Virtual Fabrics.

The idea of carving up a large Fabric into smaller Fabrics and then combining them over tagged links seems counter-intuitive, but it does add benefits. Since Fabrics can be smaller, they are easier to manage but administrators need to manage more of them. Certain valuable resources, like tape libraries and long distance links, may be shared across multiple Virtual Fabrics and achieve higher utilizations.

Virtual Fabrics: The concept of breaking a physical entity into multiple logical entities to improve management and utilization of resources.

Virtual Fabric ID (VF_ID): The 12-bit identifier for a Fabric that is used to route frames to Virtual Switches and Logical N_Ports.

Virtual Fabric Tagging: The use of an additional, pre-pended header that associates the frame to the specified Fabric.

Port VF_ID: The configured parameter that associates untagged frames to the specified Virtual Fabric. The default value of Port VF_ID is 1.