Next Generation Directors, DASD Arrays & Multi-Service, Multi-Protocol Storage Networks

By Steve Guendert

Since early 2003, it's been possible to combine mainframe FICON storage networks and open systems fibre channel Storage Area Networks (SANs) onto a common storage › network. This is known as FICON/FCP Intermix. Although Intermix has been blessed by IBM and FICON director vendors, many end users have been hesitant to try Intermix for a variety of reasons, including security concerns, management concerns, and a lack of understanding over how FICON/FCP Intermix really works. Recent advances in FICON director technology and disk array technology have provided many advanced features and functionality, such as logical domains and storage LPARs, to help alleviate those concerns. This article discusses why FICON/FCP Intermix makes sense, addresses common concerns/objections to implementing it, and details how new technologies, such as director logical domains and storage LPARs, enable the vision of a common mainframe/open systems SAN.

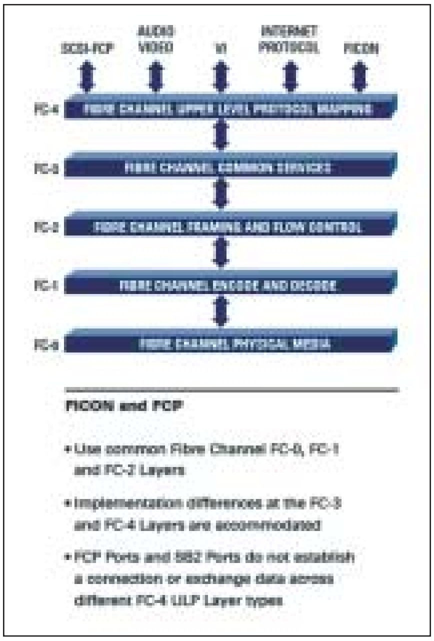

Figure 1: Fibre Channel Standard

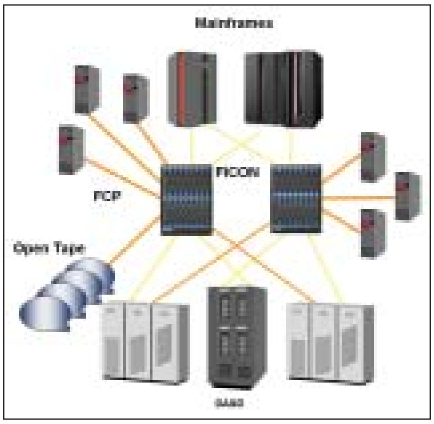

Figure 2: An Intermix Environment

What Is FICON/FCP Intermix?

Let's clarify an important point right off the bat: When the storage networking industry discusses FICON/FCP Intermix, they mean Intermix at the connectivity layer; i.e., on the same directors, switches, and fibre cable infrastructure. Up to this point, FICON/FCP Intermix didn't refer to Intermix of open systems and mainframe disk storage on the same DASD array. The new storage LPAR technologies change that, which we'll discuss later.

First of all, FICON (FC-SB-2), just like the open systems SAN Fibre Channel Protocol (FCP-SCSI-3), is merely a different Upper Layer Protocol (ULP) in the overall Fibre Channel standard (see Figure 1). Essentially, FICON and your open systems fibre channel SAN merely have a different data payload packet at the FC-4 layer. The FC-4 type of Fibre Channel frame is an element of the payload and isn't a part of the routing algorithms. Also, there's really no difference at the hardware level between a director used for FICON and a director used for open systems SAN. If your director vendor were to place two directors in a rack and ask you to determine which director was which, unless you guessed correctly, you wouldn't know which to pick, as the differences are in the software features used on the directors. For example, a director used for FICON (we'll refer to it as a FICON director) would have software features installed for management through the Control Unit Port (CUP) interface to the host mainframe.

Why Would I Want to Intermix FICON and FCP?

Given that open systems and mainframe environments haven't traditionally shared system or staff resources, let alone cultures, you might be tempted to ask why anyone would want to move toward an Intermix environment. When you take into account the added benefits of director logical domains and DASD storage LPARs, FICON/FCP Intermix really warrants serious consideration.

First, consolidating your infrastructure to include fewer fabric devices (directors and switches) to manage plus the added flexibility of being able to add different types of connections can make good ROI and Total Cost of Ownership (TCO) sense. With some exceptions, the new generation of high port count directors is designed with a non-blocking architecture and full throughput to all ports, even when the director port count is fully populated. Many large IT organizations with separate open systems and FICON storage networks are so large they can fully utilize their directors in terms of ports, even when running segregated storage network environments. However, many other organizations that have the two environments segregated aren't as large. These organizations aren't fully utilizing their directors and could realize significant cost savings by consolidating these separate, under-utilized directors into one common storage network. Fewer serial numbers on the floor for a given job is usually a good thing financially. Also, having one common cable plant/fibre infrastructure rather than separate ones for mainframe and open systems may help reduce the cost of your cable infrastructure as well as the costs of managing the cable plant and infrastructure.

Second, but on a similar note, many of the very large organizations referred to previously that are considering migrating from ESCON to FICON may already have a well-established open systems SAN that is delivering reliable, high-performance connectivity. If such an organization had open ports available in their director infrastructure, by looking at FICON/FCP Intermix, they might consider a short-term allocation of some of these unused ports from the existing SAN to test FICON, or to complete an initial deployment once testing is complete. This will help the organization delay the purchase of a separate FICON infrastructure until the connectivity requirements are sufficient to justify purchasing separate FICON directors.

FICON/FCP Intermix also makes good sense for test, development, quality assurance environments, and disaster recovery environments/data centers. These types of environments require both flexibility and resource sharing. Other already established FICON implementations will consider Intermix with the likely evolution of storage subsystem-based DASD mirroring applications from ESCON via channel extension technology to FCP for increased performance between sites. Finally, if you are running, or considering running, Linux on your mainframe, you will essentially be running FICON/FCP Intermix because channels for supporting Linux images will need to be defined as FCP. Even though you'll be using FICON Express channel cards on the mainframe, the channels need to be defined as FCP (see Figure 2).

What Has Hampered the Adoption of FICON/FCP Intermix?

While Intermix has its advantages, being very new, and thus hard to swallow, its adoption rate has been rather slow. Aside from the newness of the concept, there are two primary reasons for this. First, politics and differences in operating systems cultures enter into the picture. Many open systems shops and mainframe shops want nothing to do with each other. They've always been separate and intend to keep it that way.

Second, there are significant management differences between the two protocols. Recall that FICON and FCP are really just different FC4 layer protocols. Since this is true, they don't affect the actual switching of frames. The differences between the two protocols don't manifest themselves until the end user wants to control the connectivity of the switching (such as in zoning). Let's briefly discuss this topic of definition vs. discovery.

Open systems environments have a legacy of performing discovery for storage devices. Open systems communications, therefore, can be thought of as discovery-oriented, fabric-assigned, and name-centric (using the Fibre Channel name server to determine device communication). In most open systems SAN implementations, the devices that are physically connected on a system SCSI or IDE interface are automatically "discovered" at system boot time and are accommodated in the operating system with no pre-definition. These operating systems begin walking through the addresses for all known host bus adapters, looking for devices. The Fibre Channel architecture accounts for this behavior through the use of fabric services with layer FC-3, particularly the name server. Using management techniques such as zoning, FC-3 facilitates the traditional discovery mode while adding some new capability for controlling access to devices (such as different levels of binding).

On the other hand, mainframe host environments have essentially lived forever by the rule that planning is everything. Changes to the system configuration are planned weeks in advance, and every possible preparation is made to ensure a successful and efficient transition to the newly defined environment. This planning focuses heavily on the system IOCP definition file, which details the overall I/O configuration of the system, the resources, the devices, and the connections between them. If all the elements of the link haven't been defined in the IOCP, the connection simply doesn't exist. Therefore, we can view FICON communications as definition-oriented, address-centric, and host-assigned. The Fibre Channel architecture supports this type of device definition through a robust addressing scheme that supports millions of addresses.

Given that open systems and mainframe environments haven't traditionally shared system or staff resources, let alone cultures, you might be tempted to ask why anyone would want to move toward an Intermix environment.

These different approaches to connectivity management "styles" and the limitations of earlier directors are really what limited the adoption of Intermix. Users had to decide if they wanted to manage their directors in a mainframe mode or open systems mode and weigh the advantages and disadvantages of each. FICON devices needed to be zoned separately. When Inter-Switch Links (ISLs) were part of the storage network, how did we determine which ISL was used for FICON and which for open systems? If I needed to perform a firmware update on my directors, and the firmware update was required for open systems environments, but wasn't yet required for the mainframe, what would I do?

With the advent of logical domains/logical partitioning on the new generation of directors and storage LPARs on the new DASD arrays, these hardware management issues have been significantly reduced. In addition, they should help alleviate the political and turf issues discussed earlier.

Logical Domains and Next Generation High Port Count Directors

While high port count directors meet the need for large building blocks as the foundation for a highly scalable Intermix storage network, the industry has been slow to resolve other issues affecting the scalability of storage networks. FICON/FCP Intermix demands scalability in the fabric as well as a robust, well-designed Fibre Channel switching implementation; to date, the lack of stability in large storage networks has been one of the predominant inhibitors to effectively scaling a SAN.

Logical domains in a director allow a single high port count director to be partitioned into multiple, independent logical switches. In effect, logical domains create multiple virtual switches within a single, high port count director chassis. Each of the logical domains has independent and isolated data, control, and management environments. Devices that are attached to director ports in one logical domain aren't allowed to "see" or connect to other ports in other logical domains. This guarantees that frames from one logical domain can't be unintentionally passed to another logical domain. As a result, data traffic isolation between logical domains is guaranteed.

For control isolation, each logical domain typically will have discrete firmware and processor elements. This then allows independent sets of fabric services (e.g., name server, login server, zone server, etc.) to exist within each logical domain. Another name for this is "hard" partitioning. How this is implemented and the terminology currently used vary by switch/director vendor. Standards organizations are starting to work on SAN fabric segmentation and partitioning. Management isolation typically can be provided (specifics depend on the director vendor software) by assigning management privileges in the form of roles to each logical domain. The administrator assigned to a specific role and included in the access list will be entitled to access the logical domain.

Management activities are done on a per logical domain basis via the regulated access list.

Logical domains, therefore, allow you to partition the high port count director into separate logical switches for FICON and multiple separate switches for open systems. Setting up the logical domains by protocol type to provide the necessary isolation helps overcome the management difficulties inherent in Intermix, as well as the political issues of running separate environments in the director infrastructure. Next, we need to look at the implications of storage LPARs on a DASD array.

What Storage LPARS on DASD Arrays Mean to FICON/FCP Intermix

For a variety of extremely sound reasons, the majority of customers—even those running a FICON/FCP Intermix environment—haven't put open systems disk storage and mainframe disk storage on the same physical DASD arrays. Reasons for not doing so have ranged from purely political reasons, such as not wanting to share a frame in an era of large but uncertain growth in disk storage requirements, to purely technical reasons. An example of a common, purely technical reason is a mainframe storage manager's concern over the reliability of open systems LUN masking, and SAN zoning. It is well-known that certain open operating systems' servers attempt to write their signatures to any disk volumes they can "see," effectively destroying mainframe data in the process.

The latest DASD arrays to hit the marketplace now have the capability similar to the new high port count directors to perform logical partitioning. I'll refer to this as a storage LPAR or storage logical partition; I won't go too much into detail here, as this technology is just hitting the market. Storage LPAR technology allows the user to optimize their DASD investment, and helps to segregate workloads and protect them from one another. For example, users can run a test environment in one storage system LPAR, completely segregated from a mission-critical workload running in another storage system LPAR. End users can run dual production workloads, production and development or test partitions, and they can dedicate partition resources to meet new Service Level Agreements or to production or data mining. This allows a variety of capabilities, such as running multiple operating systems' production workloads, testing new operating releases, etc., offering unprecedented flexibility and an opportunity to cut costs. This offers peace of mind that the storage system LPARs are segregated from one another, thereby allowing higher scalability and lower cost per megabyte.

Storage LPARs will make a true FICON/FCP Intermix environment reality—not just Intermix on the connectivity, but Intermix throughout the whole environment, including the physical storage. End users will be able to more fully utilize their DASD arrays, and CIOs can save money in the process by having fewer physical frames on the floor. The user will be able to confidently place mainframe and open systems storage on the same array.

Where Do We Go From Here?

FICON/FCP Intermix is one of the enablers toward developing a Service-Oriented Architecture (SOA) in the data center. This SOA will enable companies to become "on demand" businesses and to grow their infrastructure smartly, rather than in silos dictated by the type of operating system they run.

Logical domains on directors and storage LPARs on DASD arrays solve many of the issues that have hurt the growth in adoption of FICON/FCP Intermix. The political issues that remain to be solved internally in organizations need to be weighed against the cost savings FICON/FCP Intermix will provide and against the capability it will give a business to grow. Once that takes place, FICON/FCP Intermix will truly warrant serious consideration in your data center.

About the Author

Steve Guendert is CNT Technical Solutions Group's Storage principal consultant and is the director of their FICON practice. He is a PMI certified Project Management Professional (PMP), and has M.B.A. and M.S. degrees in Management Information Systems. He is a well-regarded FICON consultant, and is one of the industry's experts on the subject. He lectures and consults regularly on FICON and business case justifications for migrating from ESCON to FICON.

e-Mail: stephen_guendert@cnt.com

Website: www.cnt.com