WHITE PAPER

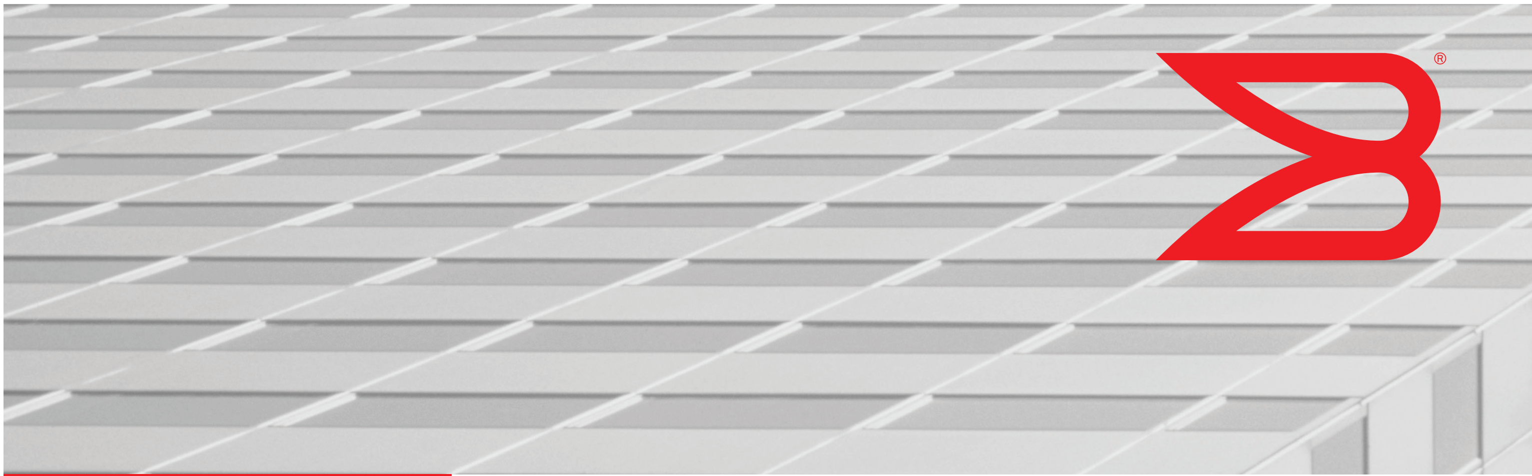

MAINFRAME

Why Switched FICON? (Switched FICON vs. Direct-Attached FICON)

Organizations of all sizes are looking to simplify their IT infrastructures to reduce costs. Some might consider implementing a direct-attached FICON architecture, but what are they giving up? Will it actually result in efficiencies and cost savings?

BROCADE

With the many enhancements and improvements in mainframe I/O technology, the old question "Do I need FICON® switching technology, or should I go direct-attached?" is frequently asked. With up to 288 FICON Express8 channels supported on an IBM System z196, why not just direct-attach the control units? The short answer is that with all of the I/O improvements, switching technology is needed—now more than ever. In fact, there are more reasons to switch FICON than there were to switch ESCON. Some of these reasons are purely technical; others are more business-related. This paper will explore both the technical and business reasons for implementing a switched FICON architecture instead of a direct-attached storage FICON architecture.

TECHNOLOGY BACKGROUND

The raw bandwidth of 8 Gbps FICON Express8 is 40 times greater than ESCON's capabilities—and FICON Express8S on zEnterprise (System z196 and z114) effectively boosts channel performance another 50 percent. The raw I/Os Per Second (IOPS) capacity of FICON Express8 and FICON Express8S channels are even more impressive, particularly when a channel program utilizes the z High Performance FICON (zHPF) protocol. To utilize these tremendous improvements, the FICON protocol is packet-switched and—unlike ESCON—is capable of having multiple I/Os occupy the same channel simultaneously.

FICON/FICON Express8 and FICON Express8S channels on zEnterprise processors can have up to 64 concurrent I/Os (open exchanges) to different devices. FICON Express8 and Express8S channels running zHPF can have up to 750 concurrent I/Os on the System zEnterprise processor family. Only when a director or switch is used between the host and storage device can the true performance potential inherent in these channel bandwidth and I/O processing gains be fully exploited.

TECHNICAL REASONS FOR A SWITCHED FICON ARCHITECTURE

Why is it a best practice to implement switched FICON rather than use point-to point (direct-attaching) FICON for connecting storage control units? Following are five key technical reasons:

- To overcome buffer credit limitations on FICON Express8 channels
- To build fan-in, fan-out architecture designs for maximizing resource utilization
- To localize failures for improved availability
- To increase scalability and enable flexible connectivity for continued growth
- To leverage new FICON technologies

FICON Channel Buffer Credits

When IBM introduced the availability of FICON Express8 channels, one very important change was the number of buffer credits available on each port per 4-port FICON Express8 channel card. While FICON Express4 channels had 200 buffer credits per port on a 4-port FICON Express4 channel card, this changed to 40 buffer credits per port on a FICON Express8 channel card. Organizations familiar with buffer credits will recall that the number of buffer credits required for a given distance varies directly in a linear relationship with link speed. In other words, doubling the link speed would double the number of buffer credits required to achieve the same performance at the same distance.

Also, organizations might recall the IBM System z10 Statement of Direction concerning buffer credits:

> "The FICON Express4 features are intended to be the last features to support extended distance without performance degradation. IBM intends to not offer FICON features with buffer credits for performance at extended distances. Future FICON features are intended to support up to 10 km without performance degradation. Extended distance solutions may include FICON directors or switches (for buffer credit provision) or Dense Wave Division Multiplexers (for buffer credit simulation)."

IBM held true to its statement, and the 40 buffer credits per port on a 4-port FICON Express8 channel card can support up to 10 km of distance for full-frame size I/Os (2 KB frames). What happens if organizations have I/Os with smaller than full-size frames? The distance supported by the 40 buffer credits would increase. It is also likely that at faster future link speeds, the distance supported will decrease to 5 km or less.

A switched architecture allows organizations to overcome the buffer credit limitations on the FICON Express8/FICON Express8S channel card. Depending upon the specific model, FICON directors and switches can have more than 1300 buffer credits available per port for long-distance connectivity.

Fan-In, Fan-Out Architecture Designs

In the late 1990s, the open systems world started to implement Fibre Channel Storage Area Networks (SANs) to overcome the low utilization of resources inherent in a direct-attached architecture. SANs addressed this issue through the use of fan-in and fan-out storage network designs. These same principles apply to a FICON storage network.

As a general rule, FICON Express8 and FICON Express8S channels offer better performance, in terms of IOPS or bandwidth, than the storage host adapter ports to which they are connected. Therefore, a direct-attached FICON storage architecture will typically see very low channel utilization rates. To overcome this issue, fan-in and fan-out storage network designs are used.

A switched FICON architecture allows a single channel to fan-out to multiple storage devices via switching, improving overall resource utilization. This can be especially valuable if an organization's environment has newer FICON channels, such as FICON Express8/8S, but older tape drive technology. Figure 1 illustrates how a single FICON channel can concurrently keep several tape drives running at full-rated speeds. The actual fan-out numbers based on tape drives will, of course, depend on the specific tape drive and control unit; however, it is not unusual to see a FICON Express8 channel fan-out from a switch to five to six tape drives (a 1:5 or 1:6 fan-out ratio). The same principles apply for fan-out to DASD arrays. The exact fan-out ratio is dependent on the DASD array model and host adapter capabilities for IOPS and/or bandwidth.

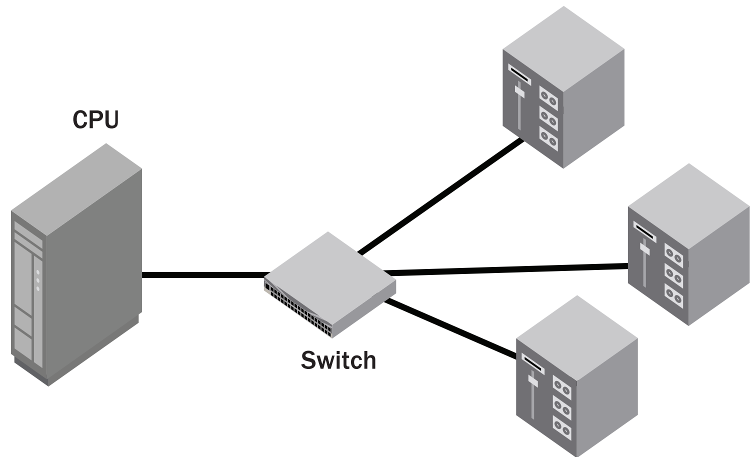

Figure 1.
Switched FICON allows one channel to keep multiple tape drives fully utilized.

Keeping Failures Localized

In a direct-attached architecture, a failure anywhere in the path renders both the channel interface and the control unit port inoperable. This could be the failure of an entire FICON channel card, a port on the channel card, a failure of the cable, a failure of the entire storage host adapter card, or a failure of an individual port on the storage host adapter card. In other words, a failure on any of these components will affect both the mainframe connection and the storage connection. The worst possible reliability, availability, and serviceability for FICON-attached storage are provided by a direct-attached architecture.

With a switched architecture, failures are localized to only the affected FICON channel interface or control unit interface—not both. The non-failing side remains available, and if the storage side has not failed, other FICON channels can still access that host adapter port via the switch or director (Figure 2). This failure isolation, combined with fan-in and fan-out architectures, allows the most robust storage architectures, minimizing downtime and maximizing availability.

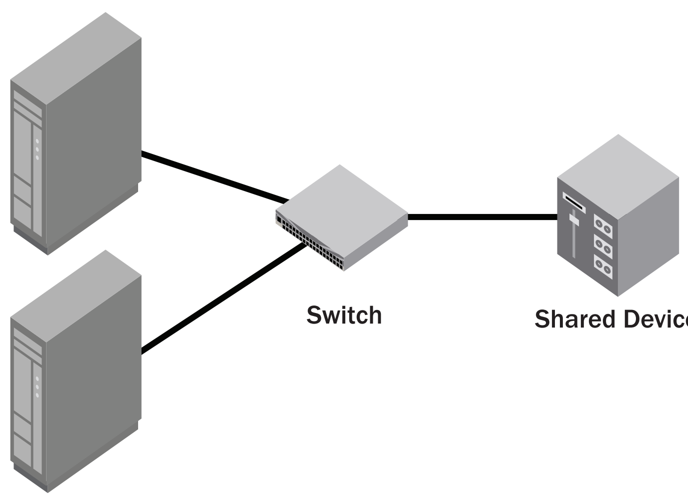

Figure 2.
A FICON director isolates faults, and improves availability.

Scale and Flexible Connectivity

Direct-attached FICON does not easily allow for dynamic growth and scalability, since one FICON channel card port is tied to one storage host adapter port. In such an architecture, there is a 1:1 relationship (no fan-in or fan-out), and since there is a finite number of FICON channels available (dependent on the mainframe model/machine type), growth in a mainframe storage environment can pose a problem. What happens if an organization needs more FICON connectivity, but has run out of FICON channels? Use of switching and proper usage of fan-in and fan-out in the storage architecture design will go a long way toward improving scalability.

In addition, best-practice storage architecture designs include room for growth. With a switched FICON architecture, adding devices such as tape is much easier—simply connect the new control unit to the switch. This eliminates the need to open the channel cage in the mainframe to add new channel interfaces, reducing both capital and operational costs. This also gives managers more flexible planning options when upgrades are necessary, since the urgency of upgrades is lessened.

And what about the next generation of channels? The bandwidth capabilities of channels are growing at a much faster rate than those of storage devices. As channel speeds increase, switches will allow data center managers to take advantage of new technology as it becomes available, while protecting investments and minimizing costs.

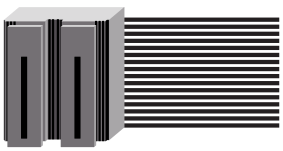

Mainframes have a limited number of FICON CHPIDs.

z800	32 FICON Express
z900	96 FICON Express
z890	40 FICON Express
z890	80 FICON Express2
z990	120 FICON Express2
z990	240 FICON Express2
z9BC	112 FICON Express4
z9EC	336 FICON Express4
z10BC	112 FICON Express8
z10EC	336 FICON Express8
z114	128 FICON Express8
z196	288 FICON Express8

Figure 3.
Direct-attached architectures are limited by the number of channels available.

Also, it is an IBM best-practice recommendation to use single-mode long-wave connections for FICON channels. Storage vendors, however, often recommend using multi-mode short-wave connections on their host adapter ports—usually due to cost. Some organizations' existing storage devices have a mix of single-mode and multi-mode connections. Since they cannot directly connect a single-mode FICON channel to a multi-mode storage host adapter, this could pose a problem. With a FICON director or switch in the middle, however, organizations do not need to change the storage host adapter ports to comply with the single-mode best-practice recommendation for the FICON channels. The FICON switching device can have both types of connectivity—single-mode long-wave ports for attaching the FICON channels, and multi-mode short-wave ports for attaching the storage.

Furthermore, FICON switching elements at two different locations can be interconnected by fiber at distances up to 100 km or more—typically used in disaster recovery and business continuance architectures. As previously discussed, FICON switching allows resources to be shared. With cascaded FICON switching, those resources can be shared between geographically separated locations, allowing data to be replicated or tape backups to be done at the alternate site from the primary site—with no performance loss. Often, workloads will be distributed such that both the local and remote sites are primary production sites, and each site uses the other as its backup.

While the fiber itself is relatively inexpensive, laying new fiber may require a construction project. Dense Wave Division Multiplexing (DWDM) helps get more out of the fiber, but with FICON switching vendors now offering Inter-Switch Links (ISLs) with up to 16 Gbps of bandwidth, the expense of DWDM can be reduced, or even eliminated. FICON switches maximize utilization of this valuable inter-site fiber by allowing multiple environments to share the same fiber. In addition, FICON switching devices offer unique storage network management features—such as ISL trunking and preferred pathing—that are not available with DWDM equipment.

FICON switches allow data center managers to further exploit inter-site fiber sharing by enabling them to intermix FICON and native Fibre Channel Protocol (FCP) traffic—known as Protocol Intermix Mode, or PIM. Even in data centers where there is enough fiber to separate FICON and open systems traffic, preferred pathing features on a FICON switch can be a great cost saver. With preferred paths established, certain cross-site fiber can be allocated for the mainframe environment, while other fiber can be allocated for open systems. The ISLs can be configured such that in the event of a failure—and only in the event of an ISL failure—the links would be shared by both open systems and mainframe traffic.

Leveraging New Technologies

Over the past 24 to 36 months, IBM has announced a series of technology enhancements that require the use of switched FICON. These include:

- NPIV support for z Linux
- Dynamic Channel Path Management (DCM)
- z/OS FICON Discovery and Auto-Configuration (zDAC)

IBM announced support for NPIV on z Linux in 2005. Today, NPIV is supported on the System z9, z10, z196, and z114. Until NPIV was supported on System z, adoption of Linux on System z had been relatively slow. NPIV allows for full support of LUN masking and zoning by virtualizing the Fibre Channel identifiers. This in turn allows each z Linux operating system image to appear as if it has its own individual Host Bus Adapter (HBA)—when those images are in fact sharing FCP channels. Since IBM began supporting NPIV on System z, adoption of z Linux has grown significantly—to the point where IBM believes approximately 19 percent of MIPS shipping on new z196s are for z Linux implementations. Implementation of NPIV on System z requires a switched architecture.

Dynamic Channel Path Management (DCM) is another feature that requires a switched FICON architecture. DCM provides the ability to have System z automatically manage FICON I/O paths connected to DASD subsystems in response to changing workload demands. Use of DCM helps simplify I/O configuration planning and definition, reduces the complexity of managing I/O, dynamically balances I/O channel resources, and enhances availability. DCM can best be summarized as a feature that allows for more flexible channel configurations—by designating channels as "managed"—and proactive performance management. DCM requires a switched FICON architecture because topology information is communicated via the switch or director. The FICON switch must have a Control Unit Port (CUP) license, and be configured/defined as a control unit in the Hardware Configuration Definition (HCD) Sysgen.

z/OS FICON Discovery and Auto-Configuration (zDAC) is the latest technology enhancement for FICON. IBM introduced zDAC as a follow-on to an earlier enhancement in which the FICON channels log into the Fibre Channel name server on a FICON director. zDAC enables the automatic discovery and configuration of FICON-attached DASD and tape devices. Essentially, zDAC automates a portion of the HCD Sysgen process. zDAC uses intelligent analysis to help validate the System z and storage definitions' compatibility, and uses built-in best practices to help configure for high availability and avoid single points of failure. zDAC is transparent to existing configurations and settings. It is invoked and integrated with the z/OS HCD and z/OS Hardware Configuration Manager (HCM). zDAC also requires a switched FICON architecture.

IBM also introduced support for transport-mode FICON (known as z High Performance FICON, or zHPF) in October 2008, with recent enhancements announced in July 2011. While not required for zHPF, a switched architecture is recommended.

BUSINESS REASONS FOR A SWITCHED FICON ARCHITECTURE

In addition to the technical reasons described earlier, the following business reasons support implementing a switched FICON architecture:

- To enable massive consolidation in order to reduce capital and operating expenses
- To improve application performance at long distances
- To support growth and enable effective resource sharing

Massive Consolidation

With NPIV support on System z, server and I/O consolidation is very compelling. IBM undertook a well-publicized project at its internal data centers (Project Big Green) and consolidated 3900 open systems servers onto 30 System z mainframes running Linux. IBM's Total Cost of Ownership (TCO) savings—taking into account footprint reductions, power and cooling, and management simplification costs—was nearly 80 percent for a five-year period. These types of TCO savings are why 19 percent of new IBM mainframe processor shipments are now being used for Linux.

Implementation of NPIV requires connectivity from the FICON (FCP) channel to a switching device (director or smaller port-count switch) that supports NPIV. A special microcode load is installed on the FICON channel to enable it to function as an FCP channel. NPIV allows the consolidation of up to 255 z Linux images ("servers") behind each FCP channel, using one port on a channel card and one port on the attached switching device for connecting these virtual servers. This enables massive consolidation of many HBAs, each attached to its own switch port in the SAN.

IBM currently recommends as a best practice to configure no more than 32 Linux images per FCP channel. Although this level of I/O consolidation was possible prior to NPIV support on System z, implementing LUN masking and zoning in the same manner as with open systems servers/SAN/storage was not possible prior to the support for NPIV with Linux System z.

NPIV implementation on System z has also been driving consolidation and adoption of a common SAN for distributed/open systems and mainframe (FICON), commonly known as Protocol Intermix Mode (PIM). While IBM has supported PIM in System z environments since 2003, adoption rates were low until NPIV implementations for z Linux picked up with the introduction of System z10 in 2008, and enhanced segregation/security beyond simple zoning was possible through switch partitioning or virtual fabrics/logical switches. With 19 percent of new mainframes being shipped for use with Linux on System z, it is safe to say that at least 19 percent of mainframe environments are now running a shared PIM environment.

Leveraging enhancements in switching technology, performance, and management, PIM users can now fully populate the latest high-density directors with minimal or no over-subscription, and use management capabilities such as virtual fabrics/logical switches to fully isolate open systems ports and FICON ports in the same physical director chassis. Rather than having more, partially populated switching platforms that are dedicated to either open systems or mainframe/FICON, PIM allows for consolidation onto fewer switching devices, reducing management complexity and improving resource utilization. This in turn leads to lower operating costs, and a lower TCO for the storage network. It also allows for a consolidated, simplified cabling infrastructure.

> Organizations that implement NPIV with a switched FICON architecture can realize massive consolidation benefits in their z Linux implementation. They can realize even greater cost savings by implementing a PIM SAN.

Figure 4.
Organizations implement NPIV to consolidate I/O in z Linux environments.

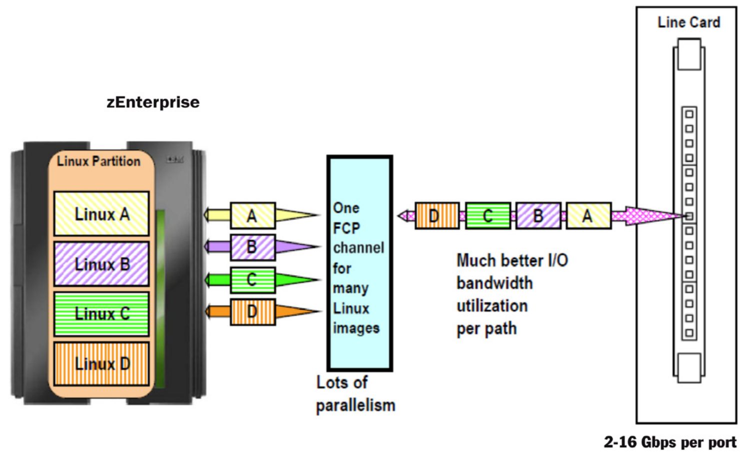

Application Performance over Distance

As previously discussed, the number of buffer credits per port on a 4-port FICON Express8 channel has been reduced to 40, supporting up to 10 km without performance degradation. What happens if an organization needs to go beyond 10 km for a direct-attached storage configuration? They will likely see performance degradation due to insufficient buffer credits. Insufficient quantities of buffer credits do not keep the “pipe” full with streaming frames of data.

Switched FICON avoids this problem. FICON directors and switches have sufficient buffer credits available on ports to allow them to stream frames at full-line performance rates with no bandwidth degradation. IT organizations that implement a cascaded FICON configuration between sites can, with the latest FICON director platforms, stream frames at 16 Gbps rates with no performance degradation for sites that are 100 km apart. This data traffic can also be compressed—and even encrypted—while traversing the network between sites, allowing IT to securely move more data, faster.

Switched FICON technology also allows organizations to take advantage of hardware-based FICON protocol acceleration/emulation techniques for tape (reads and writes), as well as with zGM (z/OS Global Mirror, formerly known as XRC, or Extended Remote Copy). This emulation technology—available on standalone extension switches or as a blade in FICON directors—allows the channel programs to be acknowledged locally at each site and avoids the back-and-forth protocol handshakes that normally travel between remote sites. It also reduces the impact of latency on application performance and delivers local-like performance over unlimited distances. In addition, this acceleration/emulation technology optimizes bandwidth utilization.

Why is bandwidth efficiency so important? Because it is typically the most expensive budget component in an organization’s multisite disaster recovery/business continuity architecture. Anything that can be done to improve the utilization and/or reduce the bandwidth requirements between sites would likely lead to significant TCO savings.

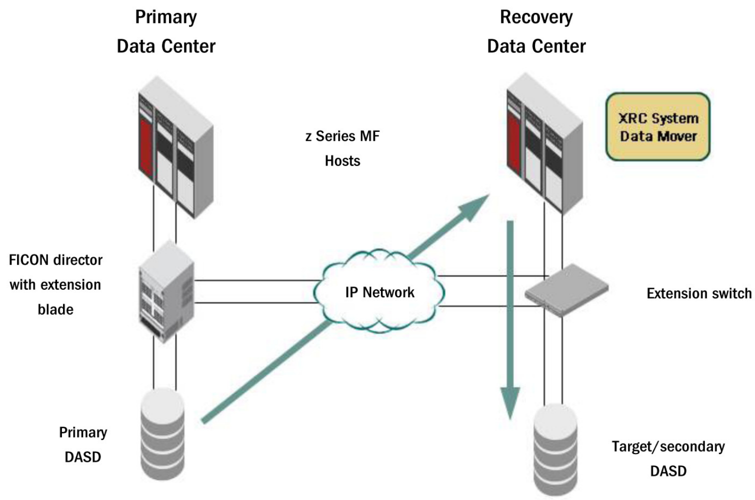

Figure 5.
Switched FICON with emulation allows optimized performance and bandwidth utilization over extended distance.

Enabling Growth and Resource Sharing

Direct-attached storage forces a 1:1 relationship between host connectivity and storage connectivity. In other words, each storage port on a DASD array host adapter requires its own physical port connection on a FICON Express8 channel card. These channel cards are typically very expensive on a per-port basis—typically four to six times the cost of a FICON director port. Also, there is a finite number of FICON Express8 channels available on a System z196 (a maximum of 288), as well as a finite number of host adapter ports on a DASD array. If an organization has a large configuration and a direct-attached FICON storage architecture, how does it plan to scale its environment? What happens if an organization acquires a company and needs additional channel ports? A switched FICON infrastructure allows cost-effective, seamless expansion to meet growth requirements.

Direct-attached FICON storage also typically results in underutilized channel card ports and host adapter ports on DASD arrays. FICON Express8 and FICON Express8S channels can comfortably perform at high-channel utilization rates, and direct-attached storage architecture typically sees channel utilization rates of 10 percent or less. As illustrated in Figure 6, leveraging FICON directors or switches allows organizations to maximize channel utilization.

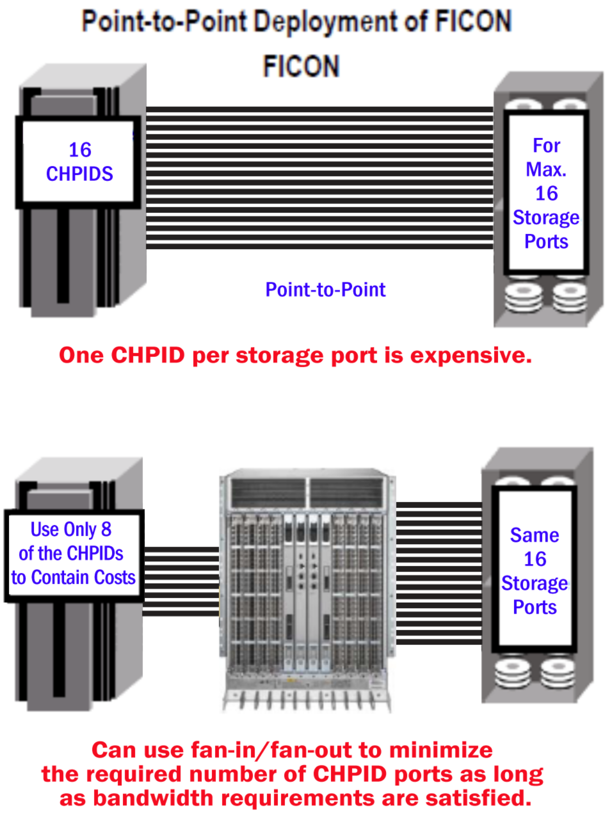

Figure 6.
Switched FICON drives improved channel utilization, while preserving Channel Path Identifiers (CHPIDs) for growth.

It also is very important to keep traffic for tape drives streaming, and avoid stopping and starting the tape drives, as this leads to unwanted wear and tear of tape heads, cartridges, and the tape media itself. This is accomplished by using FICON acceleration/emulation techniques as described earlier. A configuration similar to the one shown in Figure 7 can also be implemented. Such a configuration requires solid analysis and planning, but it will pay dividends for an organization's FICON tape environment.

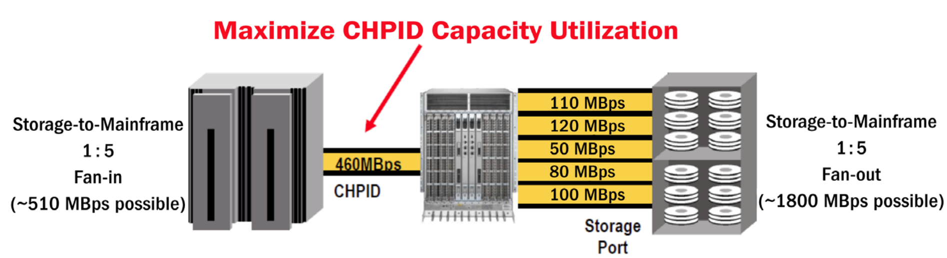

Figure 7.
A well-planned configuration can maximize CHPID capacity utilization for FICON tape efficiency.

Finally, switches facilitate fan-in—allowing different hosts and LPARs whose I/O subsystems are not shared to share the same assets. While some benefit may be realized immediately, the potential for value in future equipment planning can be even greater. With the ability to share assets, equipment that would be too expensive for a single environment can be deployed in a cost-saving manner.

The most common example is to replace tape farms with virtual tape systems. By reducing the number of individual tape drives, maintenance (service contracts), floor space, power, tape handling, and cooling costs are reduced. Virtual tape also improves reliable data recovery, allows for significantly shorter Recovery Time Objectives (RTO) and nearer Recovery Point Objectives (RPO), and offers features such as peer-to-peer copies. However, without the ability to share these systems, it may be difficult to amass sufficient cost savings to justify the initial cost of virtual tape. And the only practical way to share these standalone tape systems or tape libraries is through a switch.

With disk subsystems, in addition to sharing the asset, it is sometimes desirable to share the data across multiple systems. The port limitations on DASD may prohibit or limit this capability using direct-attached (point-to-point) FICON channels. Again, the switch can provide a solution to this issue.

Even when there is no need to share devices during normal production, this capability can be very valuable in the event of a failure. Data sets stored on tape can quickly be read by CPUs picking up workload that is already attached to the same switch as the tape drives. Similarly, data stored on DASD can be available as soon as a fault is determined.

Switch features such as preconfigured port prohibit/allow matrix tables can ensure that access intended only for a disaster scenario is prohibited during normal production.

SUMMARY

Direct-attached FICON might appear to be a great way to take advantage of FICON technology's advances over ESCON. However, a closer examination shows that switched FICON, similar to switched ESCON, is a better, more robust architecture for enterprise data centers. Switched FICON offers:

- Better utilization of host channels and their performance capabilities
- Scalability to meet growth requirements
- Improved reliability, problem isolation, and availability
- Flexible connectivity to support evolving infrastructures
- More robust business continuity implementations via cascaded FICON
- Improved distance connectivity, with improved performance over extended distances
- New mainframe I/O technology enhancements such as NPIV, FICON DCM, zDAC, and zHPF

Switched FICON also provides many business advantages and potential cost savings, including:

- The ability to perform massive server, I/O, and SAN consolidation, dramatically reducing capital and operating expenses
- Local-like application performance over any distance, allowing host and storage resources to reside wherever business dictates
- More effective resource sharing, improving utilization and reducing costs

With the growing trend toward increased usage of Linux on System z, and the cost advantages of NPIV implementations and PIM SAN architectures, direct-attached storage in a mainframe environment is becoming a thing of the past. The advantages of a switched FICON infrastructure are simply too great to ignore.

For more information about Brocade® solutions for FICON, visit www.brocade.com.

> Investments made in switches for disaster recovery and business continuance are likely to pay the largest dividends. Having access to alternative resources and multiple paths to those resources can result in significant savings in the event of a failure.

Corporate Headquarters
San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters
Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters
Singapore
T: +65-6538-4700
apac-info@brocade.com

BROCADE