

8GFC Deployment

Fibre channel's next generation has arrived and now supports 8-gigabit per second (Gb/s) transfer rates over one fibre channel (8GFC) link, effectively doubling available bandwidth to 840 megabytes per second (MB/s). However, 8GFC technology is priced at a significant premium over 4GFC and depending on the configuration, deployment of 8GFC may be up to 2X the cost of using 4GFC. In contrast, one appealing factor of 8GFC links is its downward compatibility with 4GFC and 2GFC links thus allowing a gradual move to this new technology.

Increased bandwidth capability, the large price differential, and the downward compatibility should be the critical factors in determining an imminent proposed move to 8GFC. Aside from these, other less objective factors must also influence the ultimate 4GFC versus 8GFC decision. Indeed, since 4GFC technology was introduced in 2004, major enhancements in storage area networking, server, application, and storage subsystem performance have been made, making the switch to 8GFC even more complicated.

8GFC and storage area network switching

Storage area networking, particularly switch technology, has seen numerous changes in recent years. Specifically, switches have necessarily become more complex to match increasing fabric functionality requirements. In addition, layers of switches and the number of ports have been proliferating to support more complex fabrics. Today, core and core to edge switching uses 4GFC, 10GFC or inter-chassis copper channels for inter-switch links (ISLs). One area particularly suited to 8GFC technology are the ISLs such as between the core networking and/or edge storage switch components thereby reducing port counts and cabling by at least a factor of two.

Additionally, 10GFC ISLs could conceivably be economically replaced with a more general purpose, 8GFC solution. SAN extensions over metro area networks use 10GFC but contrary to the ISL use above, 8GFC metro area networks with FC clocking cost significantly more than 10GFC with Ethernet clocking and thus continues to justify the status quo.

8GFC and servers

Server configurations have experienced several performance-enhancing changes in recent years. One significant change has been the emergence of blade servers to manage multiple commodity servers running similar, high-intensity workloads. With blade servers, separate infrastructure hardware is now shared across a number of server blades within a single cabinet, thereby significantly reducing necessary equipment and corresponding power requirements. Even so, the addition of 8GFC to these blade servers may be of questionable benefit depending on bandwidth requirements to meet the overall I/O demand.

Another major change over recent years in server technology has been the increased adoption of server virtualization. With the emergence of VMware ESX, Microsoft Hyper-V, XenSource, Virtual Iron and others it's now easier to use one high-powered physical server to run multiple virtual machines (VMs), each executing their own application. However, the applications typically virtualized do not require high I/O bandwidth and therefore, virtualized servers are not prime candidates for 8GFC.

One place where 8GFCs might be justified is with high-end proprietary Unix servers from HP, IBM, Sun and others. These high-end Unix servers can support I/O throughput in the range of many hundreds of MB/s requiring the superior bandwidth capabilities provided by 8GFC. In addition, IBM's recently introduced Z10 mainframes also can generate high I/O activity. In this case 8Gbit FICON, based on 8GFC technology, may provide a boost in performance.

8GFC and applications

As hardware has advanced, applications have also advanced. Application IO demand can run the gamut of bandwidth needs from a few MB/s to many hundreds of MB/s depending on the type of application. Regardless of the actual server hardware, only those applications requiring high bandwidth would benefit from 8GFC performance.

Application I/O requirements are often stated using two measures, I/O operations per second (IOPs) and I/O bandwidth also called SAN throughput. The more usual metric is IOPs rate, defined as how many read, and/or write operations per second are required by an application. On the other hand, I/O bandwidth or storage throughput is defined as the number of MB/s or gigabytes per second (GB/s) that an application needs to process. Commonly, an application may need high IOPs but low I/O bandwidth or vice versa.

In a recent Sun Microsystems white paper, application categories were identified and I/O requirements were specified as follows:

- **Transaction processing (OLTP) and Ecommerce applications.** These applications can require the highest IOPs rate with minimal latency but generally do not need high storage throughput. For OLTP and Ecommerce applications storage throughput requirements are usually less than 25 MB/s.
- **Decision support, data warehousing, and visual database applications.** These applications can require high IOPs and need approximately 75 MB/s of storage throughput to operate.
- **Scientific computing and imaging applications.** These applications require medium IOPs and need approximately 125 MB/sec of storage throughput.
- **Audio and video media streaming applications.** These applications require low IOPs but depending on compression and the number of

8GFC Deployment

concurrent streams may need as high as 600 MB/s of I/O bandwidth to sustain adequate streaming performance.

In the above examples, only the video streaming application is a good candidate for 8GFC host connectivity. The other applications have low bandwidth requirements and are unlikely to require 8GFC.

These published application requirements are further substantiated at least as to decision support by Transaction Processing Performance Council (TPC) benchmark results analysis. This analysis indicates that TPC-H (TPC's decision support benchmark) can attain storage bandwidth requirements approaching 350 MB/s under maximum server transaction processing. Under more normal operation conditions much lower bandwidth requirements would be expected.

On the other hand, application workloads aren't the only thing driving data center bandwidth requirements. Storage backups are often some of the worst I/O bandwidth users today. Presently with snapshot technology, which allows for near-instantaneous replication of data on the storage subsystem, the backup window is no longer such a critical data center metric but storage capacity and usage continues to increase. Consequently, more and more storage needs to be backed up for each cycle, consuming significant bandwidth to complete in a timely fashion. For instance, a recently published Microsoft Exchange Solution Review Program (ESRP) result for the EMC Symmetrix DMX-4 4500, a high-end storage subsystem supporting 60,000 mailboxes, reported an aggregate database backup throughput of ~3.9 GB/s using 14 Exchange servers each of which averaged approximately 280 MB/s of storage throughput. As such, this high-end storage subsystem could easily use five or more 8GFC links while each Exchange server could still use one 4GFC for performance or two 4GFC links for availability.

8GFC and storage subsystems

Demand for increased performance of the entire IT department's infrastructure has been felt even at the storage subsystem level. With the move to 8GFC by server and storage area networking links, storage connection link migration to 8GFC is expected. However, the move is still justifiable only by a need for additional bandwidth above 4GFC.

A relatively new benchmark from the Storage Performance Consortium (SPC) reveals that approximately one-third of storage subsystems evaluated can exceed storage throughput of 800 MB/s. For example, one benchmark result for the IBM DS8300 indicates about a 3.2 GB/s storage throughput for the subsystem. This bandwidth usage could justify the deployment of four or more 8GFC links.

Even though a perusal of benchmark results would suggest that few subsystems could effectively use 8GFC links, an individual data center may benefit from additional bandwidth provided by 8GFC. For example, many switch vendors suggest a storage link fan-in of 5 or more to 1, i.e., 5 or more host fibre channel links per 1 storage subsystem link. However, fan-in ratios can be much higher than 5:1 depending on the applications

being served. Anytime aggregate link bandwidth needs are in excess of 800 MB/s then 8GFC should be considered.

Tools for determining any 8GFC advantage

Determining throughput or bandwidth requirements for a data center can be a particularly daunting task. However, vendors are providing tools to help monitor a data center's storage fabric. For example, Cisco Fabric Manager Performance Monitoring can provide detailed bandwidth usage information by drilling down to an individual port, link or ISL level. Such detailed information would provide definitive guidance in determining where additional bandwidth might enhance application throughput or overall data center performance.

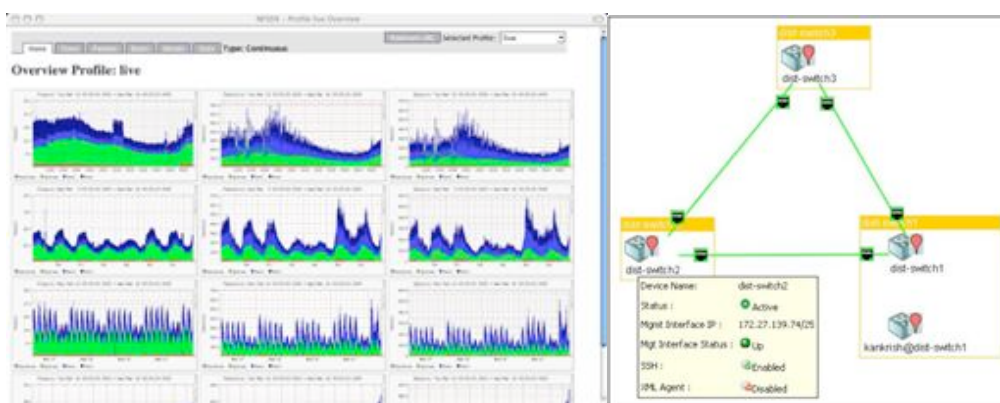


Figure 1 Screenshots from Cisco Fabric Manager

Conclusion

In the relentless quest for superior performance, 8GFC technology has emerged. However, the ultimate decision to adopt this technology should be based on technological and economic reality. That is, does the data center, regardless of its specific configuration need the expanded throughput capabilities of 8GFC links? Is the significant price differential between 4GFC and 8GFC justifiable? Can the new technology be gradually introduced?

Given the current state of SANs, servers, applications, and storage subsystems only modest true need exists for 8GFC today and is perhaps reflected in the current support of this technology. However, future development of these products and new applications will rapidly drive the mainstream adoption of this technology.

Silverton Consulting, Inc. is a Strategy, Storage & Systems consulting services company, based in the USA offering products and services to the data storage community.

<http://www.silvertonconsulting.com>