WHITE PAPER

Better Performance, Better Insight for Your Mainframe Storage Network with Brocade Gen 6

TABLE OF CONTENTS

Brocade and the IBM z Systems IO product team share a unique history of technical development, which has produced the world's most advanced mainframe computing and storage systems. Brocade's technical heritage can be traced to the late 1980s, with the creation of channel extension technologies to extend data centers beyond the "glass-house." IBM revolutionized the classic "computer room" with the invention of the original ESCON Storage Area Network (SAN) of the 1990s, and, in the 2000s, it facilitated geographically dispersed FICON® storage systems. Today, the most compelling innovations in mainframe storage networking technology are the product of this nearly 30-year partnership between Brocade and IBM z Systems.

As the flash era of the 2010s disrupts the traditional mainframe storage networking mindset, Brocade and IBM have released a series of features that address the demands of the data center. These technologies leverage the fundamental capabilities of Gen 5 and Gen 6 Fibre Channel, and extend them to the applications driving the world's most critical systems. This white paper provides a summary of these unique mainframe features, which reflect Brocade and IBM's dedication to technical excellence.

Overview

Partnership

The Brocade and IBM partnership consists of many teams working together to produce the best-in-class products that mainframe users have come to expect. The close relationship between the business teams allows both organizations to guide the introduction of key technologies to the market place, while the integration between the system test and qualification teams ensures the integrity of those products. Driving these efforts are the deep technical relationships between the Brocade and IBM z Systems IO architecture and development teams, whose close collaboration provides the transformative technology to propel today's ever-changing data center.

Technology Themes

During the past five to seven years, the technology landscape of mainframe storage networking systems has evolved to embrace new levels of scale. Not only are the computational demands increasing, but the physical reach of the systems is also growing. Business continuity and recovery requirements have stretched the limits of yesterday's technology, and the millennial cloud movement is forcing organizations to reexamine transactional processing. These pressures led IBM to produce one of the most innovative and highly functioning processors in history—the z13. At the same time, Brocade provided the networking dexterity necessary to complement the z13's impact.

Feature Summary

The technological leap induced by the z13 is underscored by the inventiveness and integration of the Brocade® Gen 5 and Gen 6 feature set:

- Inter-Chassis Links (ICLs) provide world-class scalability and flexible chassis interconnectivity.
- Multi-hop topologies provide flexibility to geographically dispersed data centers to support unique disaster recovery and business continuity configurations.
- 16 GFC FICON Express16S device connectivity and 32 GFC fabric connectivity provide the fastest data rates in the industry.
- FICON Dynamic Routing utilizes the full capacity of Inter-Switch Link (ISL) and ICL connections in cascaded network topologies, leveraging Brocade Exchange-Based Routing (EBR) technology.
- Fabric I/O Priority capabilities provide a foundation for applications, such as zOS Workload Manager, to prioritize individual IOs for optimal storage network performance.
- Port Decommission is used for easier infrastructure management through automated interaction between the fabric and the mainframe IO subsystem.
- IBM Health Checker for z/OS leverages the internal Control Unit Port (CUP) of the FICON director to evaluate single-point-of-failure conditions as well as path integrity and performance.
- Problem determination and isolation are enhanced through the Read Diagnostic Parameters (RDP) operations.
- Forward Error Correct (FEC) improves connection reliability and performance.

Technology Highlights

Brocade mainframe technology has produced several "first to market" innovations, including Gen 6 32 Gbps Fibre Channel, Gen 5 16 Gbps Fibre Channel with Forward Error Correction, FICON Dynamic Routing, z/OS Health Checker integration, and Read Diagnostic Parameters functionality. Leveraging the IBM z Systems IO team partnership, Brocade has also produced significant mainframe storage networking market exclusives such as ICLs and Port Decommissioning. This shared dedication to the success and growth of the mainframe storage network can also be seen in newer endeavors, including Fabric I/O Priority and FICON multi-hop topologies. These technologies are detailed below to provide insight into Brocade Gen 6/Gen 5 mainframe storage networking functionality.

Availability Note

Included in these descriptions are released, planned, and under-development efforts. Please see your local Brocade representative for details regarding the Brocade Fabric OS® (FOS) releases containing a specific feature or functionality.

Speed, Scale, and Performance

Inter-Chassis Links

Inter-Chassis Links (ICLs) provide short-distance, chassis-to-chassis connectivity for Gen 5 and Gen 6 FICON directors for use with FICON and/or FCP mainframe solutions (see Figure 1). This technology is used to build a powerful storage networking core without sacrificing device ports for Inter-Switch Link (ISL) connectivity.

ICLs minimize the latency between chassis to provide the lowest-latency switching via a backplane relative to using

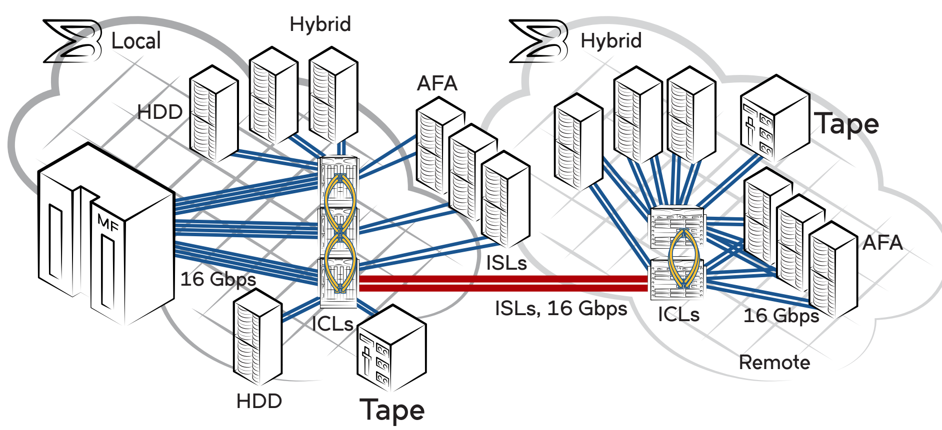

Figure 1: Inter-Chassis Links versus Inter-Switch Links.

comparable ISLs. Frame-based trunking is automatically enabled between four ICLs, which maximizes the load balancing and availability characteristics of the chassis. When used with FICON Dynamic Routing, the IO exchanges are evenly distributed across the ICL infrastructure for optimal performance. Furthermore, since the ICL connection is not considered to be a "hop of concern" in cascaded topologies, it can provide configuration flexibility for large deployments.

Multi-hop Topologies

The movement to disperse mainframe data centers geographically has forced the evolution of classical FICON SAN topologies. Simple cascaded systems have given way to multi-hop variations that exploit Brocade ICL and FCIP channel extension technologies, while maintaining the integrity of the "hop of concern" provision. Advanced diagnostic capabilities, such as the IBM Health Checker for z/OS and the Fibre Channel Read Diagnostic Parameters function, have provided the apparatus necessary to maintain the deterministic reliability required for mainframe storage networks in extended topologies.

Multi-hop topologies include two-, three-, and four-site disaster recovery and business continuity configurations (see Figure 2). Each topology is designed to address specific business requirements as well as regulatory provisions. In addition, the supported topologies are verified to operate with the same level of predictability and deterministic reliability found in today's classical cascaded solutions.

Availability Note

Brocade and IBM are currently validating the most effective topology solutions for mainframe environments to ensure that performance and reliability requirements of z Systems solutions are maintained with each supported configuration. See your Brocade representative for full details.

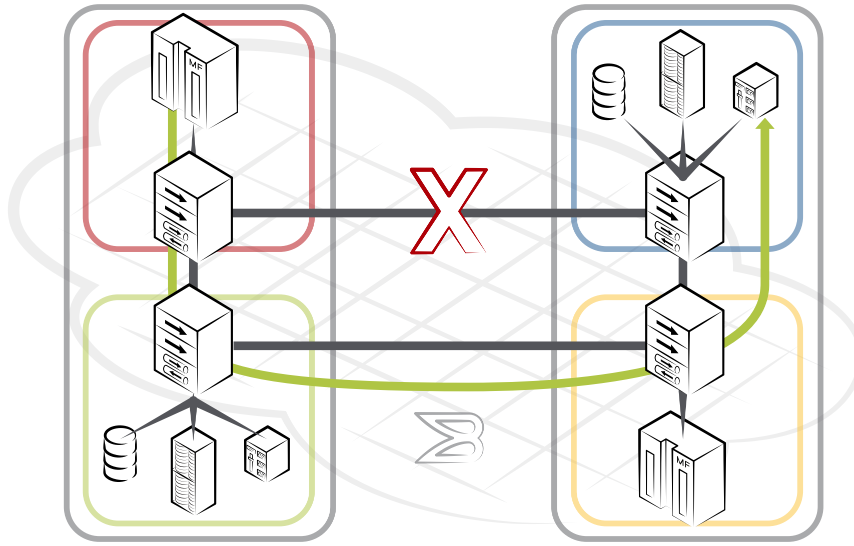

Figure 2: Example of multi-hop topology.

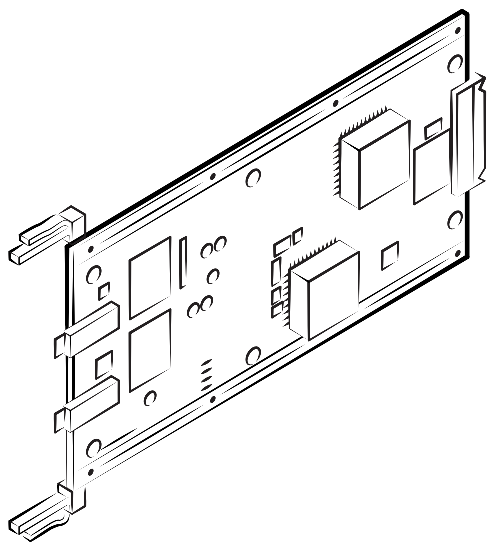

Figure 3: FICON Express16S channel card.

Gen 5/Gen 6 Performance

Brocade Gen 5 and Gen 6 FICON directors, shown in Figure 3, provide optimal interface attachment rates for the FICON Express16S (FX16S) channel as well as storage arrays. The 16 Gbps Fibre Channel device connectivity and 32 Gbps Fibre Channel fabric connectivity provide minimal storage network latency, which significantly enhances data center application operations and performance.

The faster link speeds provide reduced IO latency for critical middleware IO, such as database log writing, which significantly improves DB2 transactional latency. The faster link speed also shrinks batch windows and reduces the elapsed time for IO-bound batch jobs. Since faster link speeds are more sensitive to the quality of the cabling infrastructure, new industry-leading standards are included to provide enhanced error correction on optical connections, ensuring a smooth transition to 16 Gbps and 32 Gbps Fibre Channel technologies.

FICON Dynamic Routing

The introduction of the z13 server included several performance optimizations designed to improve overall throughput in the mainframe storage network. FICON Dynamic Routing allows FICON devices to utilize both static and dynamic SAN routing policies. It is designed to support the dynamic routing policies provided by the FICON director, such as the Brocade Exchange-Based Routing (EBR) protocol.

FICON Dynamic Routing enables organizations to use SAN dynamic routing policies across cascaded FICON directors to simplify configuration and capacity planning, and to provide persistent and repeatable performance and higher resiliency (see Figure 4). In Peer-to-Peer Remote Copy configurations, the sharing of switches is simplified, and hardware costs can be reduced by allowing FICON and FCP to share the same storage network infrastructure.

The benefits of FICON Dynamic Routing are highlighted in the IBM technical reference *FICON Dynamic Routing (FIDR) - Technology and Performance Implications*, co-authored by Brocade (see http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102651).

Fabric I/O Priority

SAN Fabric I/O Priority integration translates application importance into Fabric Priority (Quality of Service), which allows the application to maintain the desired quality of service through the fabric. This allows the priority level to be communicated between the host and the end devices for end-to-end management of application goals (see Figure 5).

Availability Note

Brocade and IBM z Systems are exploring the full integration of Fabric I/O Priority to allow system management functions, such as zOS Workload Manager, to exploit end-to-end, integrated quality of service characteristics. See your Brocade representative for full details.

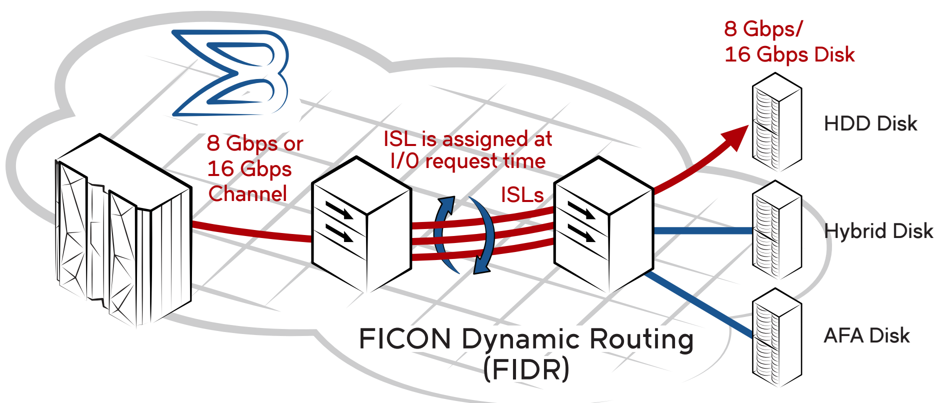

Figure 4: FICON Dynamic Routing.

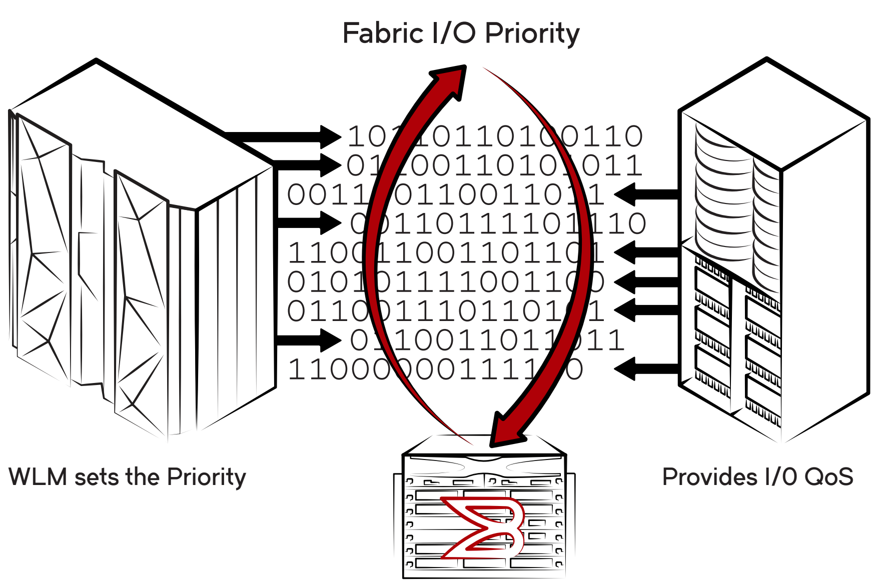

Figure 5: Fabric I/O Priority.

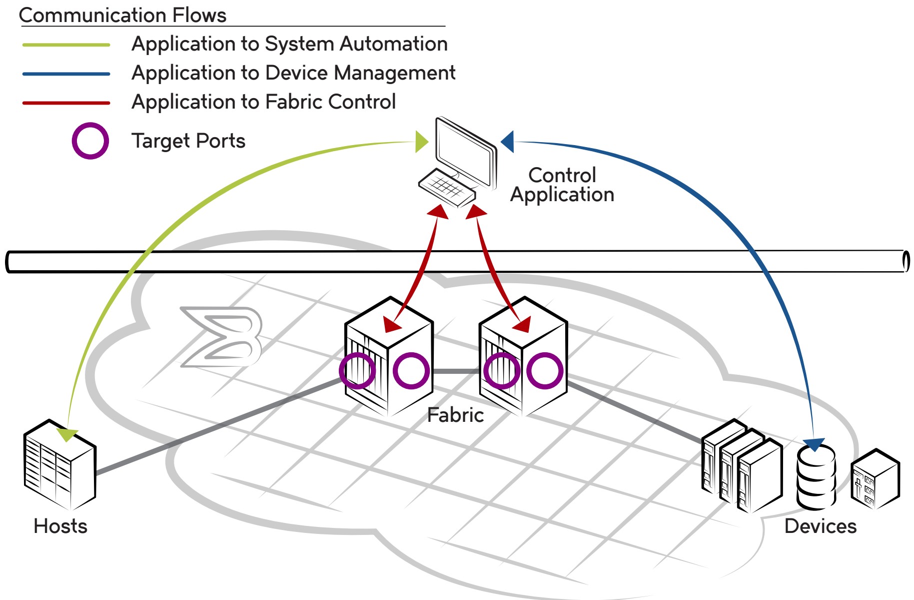

Figure 6: Port Decommission.

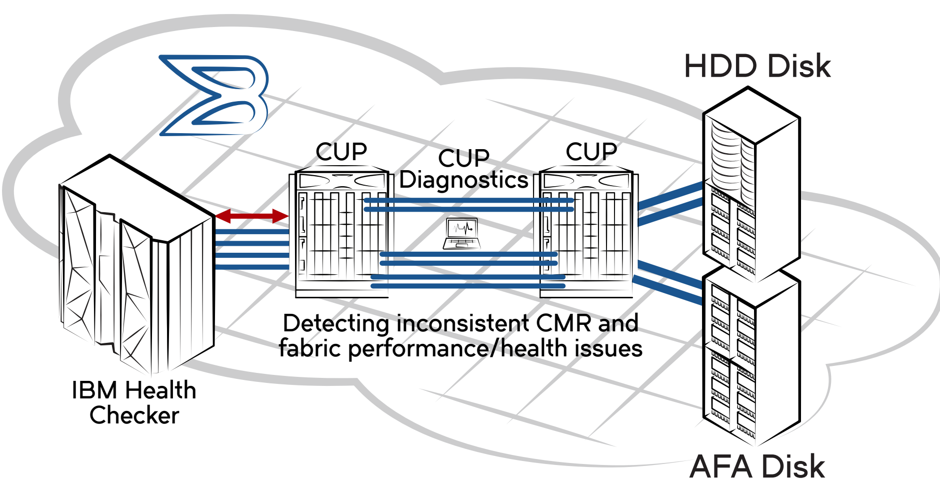

Figure 7: IBM Health Checker for z/OS integration.

Autonomics and Integrity

Port Decommission

Port Decommission is an autonomic, integrated function designed to simplify daily port/path maintenance. Routine port maintenance that once required significant administrative overhead as well as complex operations coordination can now be executed seamlessly without operator intervention. Upon receiving an indication to remove a port or path from active service, Brocade Network Advisor coordinates the activities with both the host IO subsystem and the fabric. After z/OS sends a confirmation that it has migrated the active workload from the designated port/path, Brocade Network Advisor removes the port from operation and indicates it is ready for maintenance (see Figure 6).

Once the maintenance operations have been completed, the port/path is returned to full operational status through the Port Recommission function. As with the decommission operation, Brocade Network Advisor coordinates and communicates the recommission action with z/OS to ensure the operation occurs error-free.

IBM Health Checker for z/OS Integration

The IBM Health Checker for z/OS extracts information from the FICON director through the Brocade CUP interface to the switch (see Figure 7). This information provides the Health Checker with the ability to identify and report the shared components of the mainframe storage network. The resulting path analysis provides system administrators with the insight necessary to ensure members of a path group are not susceptible to single-point-of-failure conditions.

In addition to the single-point-of-failure analysis, IBM Health Checker for z/OS leverages the FICON director CUP interface to assess performance characteristics of the members of a path group. The "flow" information provided includes utilization characteristics, operational conditions, and diagnostic error statistics. This unique solution also incorporates the Brocade Fabric OS® (FOS) bottleneck detection mechanism, which provides alerts to z/OS when fabric anomalies are detected.

Enhanced Diagnostics

Keeping pace with the ever-growing complexity of mainframe storage networks, Brocade and IBM co-developed in-band diagnostic mechanisms that enrich the view of fabric components. The z Systems channel and control unit solutions include functionality that allows the FICON director to read information about the optical transceiver attached to the z Systems device (see Figure 8). This same functionality is utilized by zOS to provide similar information about the transceivers installed in the director.

The new Fibre Channel-standard Read Diagnostic Parameters (RDP) command was created to enhance path evaluation and automatically differentiate between cable hygiene errors and failing components. The RDP function provides the optical characteristics of each end of the link, including optical signal strength and environmental operating metrics, without requiring manual insertion of optical measurement devices. Furthermore, the RDP operation is performed periodically throughout the mainframe storage network to proactively detect marginal conditions that can be addressed before they affect production operations.

Forward Error Correction

Gen 6 Fibre Channel includes an integrated Forward Error Correction (FEC) technology to provide reliable signal integrity at high data rates (that is, 32 GFC/128 GFC). Gen 5 Fibre Channel offers this same functionality as an optional feature to be exploited

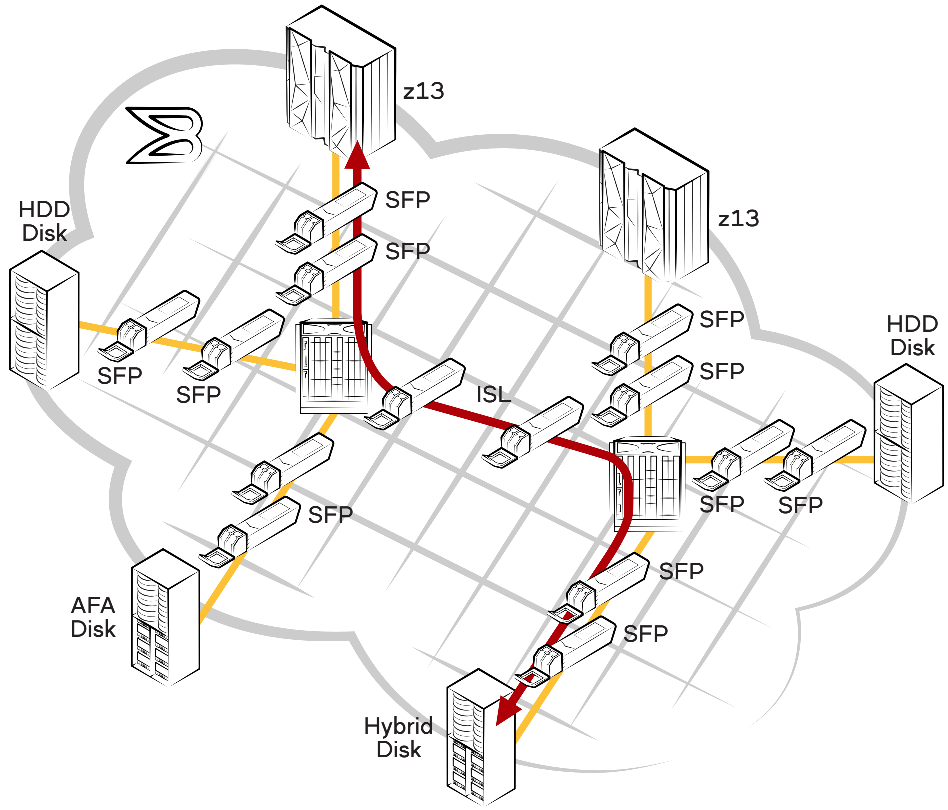

Figure 8: Read Diagnostic Parameters.

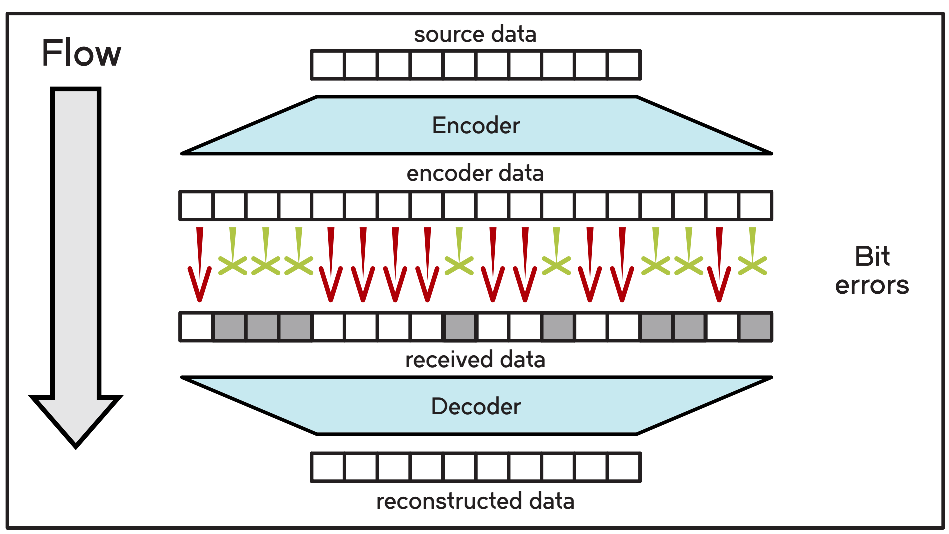

Figure 9: Forward Error Correction.

at the solution provider's discretion. Brocade includes Gen 5 Fibre Channel FEC technology with all Brocade FICON director and switch solutions. The IBM FICON Express16S channel leverages this technology to ensure zOS FICON solutions achieve the highest level of end-to-end reliability and performance in the industry.

Gen 6 Fibre Channel Technology Benefits for z Systems and Flash Storage

While at first glance, given that the FICON Express16S channels were GA in March 2015 on the IBM z13 family of processors, the need for Gen 6 and 32 Gbps-capable FICON/FCP may not be readily obvious. However, considering the reasons Gen 6 Fibre Channel technology was developed in the first place, and the recent z13, Linux ONE, and flash storage array enhancements, it is clear why the Gen 6 technology is a great fit for many organizations today. When Gen 5 Fibre Channel made its debut in 2011, many asked similar questions about the need for it, given that channel and storage connectivity was still all 8 Gbps, and remained so until 2015. Gen 6 Fibre Channel brings three primary, immediate technology benefits to a z13 and flash storage architecture:

- Improved performance
- Improved Virtual Machine (VM) support
- Support for the upcoming NVMe over Fabrics technology

Improved Performance

Enterprise-class storage performance has historically focused on four metrics: latency, response time, throughput (bandwidth), and Input/Output Operations Per Second (IOPS). There are some common misperceptions, including that the terms "response time" and "latency" can be used interchangeably, so it is good to review what each of these metrics actually measures.

Latency

Latency is a measure of delay in a system. In the case of storage, it is the time it takes to respond to an I/O request. In the example of a read from disk, this is the amount of time taken until the device is ready to start reading the block—that is, not including the time taken to complete the read. In the spinning disk (HDD) world, this includes things such as the seek time (moving the actuator arm to the correct track) and the rotational latency (spinning the platter to the correct sector), both of which are mechanical processes, and therefore slow. Latency in a storage system also includes the finite speed for light in the optical fiber, and electronic delay due to capacitive components in buses, Host Bus Adapters (HBAs), switches, and other hardware in the path. The primary advantage flash storage/SSD has over spinning disk technology is the lower latency and the consistency in the latency. SSDs can exhibit latency that is lower by a factor of 100 compared to HDD technology, and this is primarily due to the absence of the mechanical processes described above. Equally important, for random access IO, the SDD with its consistently low latency is able to move data continuously, while the HDD must continuously return for chunks of data, repeating the mechanical process and slowing performance.

Response Time

Latency and response time are often confused with each other and are incorrectly used interchangeably. Response time is the total amount of time it takes to respond to a request for service. It is the sum of the service time and wait time. The service time is the time it takes to do the actual work requested. The wait time is how long the request had to wait in a queue before being serviced, and it can vary greatly depending on the amount of IO. Transmission time gets added to response time when the request and response have to travel over a network, and varies with distance. Mainframe response time metrics for DASD include IOSQ, PEND, CONN, and DISC. Since SSD technology is capable of handling significantly more IOPS than HDD, the time spent waiting in a queue is typically shorter for SSD than HDD, leading to better response times for SSD. When SSD is coupled with Gen 6 Fibre Channel, studies have demonstrated a 71 percent improvement in response time.

Throughput

Throughput (MB/sec) is the data transfer speed. It is the measure of data volume transferred over time—the amount of data that can be pushed or pulled though a system per second. It is simply a product of the number of IOPS and the IO size. As a performance metric, throughput matters most for large block-size data traffic. In a mainframe environment, these would be batch-type processing and data warehousing applications. Batch processing handles requests that are cached (prestored) and then executed all at the same time. An example of batch processing is the reconciliation of a bank's ATM transactions for a given day being compared against all the customers' bank records, and then reporting the ATM transaction records at a time outside normal business hours. Higher throughput capabilities, such as Gen 6 Fibre Channel's 32 Gbps, will lead to shortened batch windows, which can be invaluable to many large z Systems customers (the infamous race to the sunrise). When IBM introduced FICON Express16S channels in 2015, one of the primary cited benefits compared to FICON Express8S channels was an average 32 percent reduction in the elapsed time (wall clock) of the batch window.

IOPS

For a modern mainframe environment, throughput is secondary in importance to a storage device's IOPS capability. Flash storage/SSD technology is capable of significantly more IOPS than spinning disk technology. Mainframes are primarily used for Online Transaction Processing (OLTP) workloads. Transaction processing consists of large numbers of small block-size data IO transactions, often seen in banking, ATMs, order processing, reservation systems, and financial transactions. The most common block-size traffic seen in today's mainframe environment is 4,000 bytes, which coincidentally is used for most storage performance benchmarking tests, as well as z Systems FICON channel performance testing.

The IBM z13 family of mainframes uses the latest industry-standard Peripheral Component Interconnect Express Generation 3 (PCIe Gen3) technology for its IO subsystem (IOS). It also supports 50 percent more IO devices per FICON and FCP channel (32,000 IO devices per channel), with a maximum capacity of 320 FICON channels per z13. A FICON Express16S channel is capable of 93,000 z High-Performance FICON (zHPF) IOPS running 4,000-byte block-size traffic typical of most OLTP workloads. In sum, the z13 processors are IO "monsters." They were designed to handle massive amounts of IO due to the increases in IO demands brought about by what IBM refers to as CAMS (Cloud, Analytics, Mobile, and Social) workloads. For many end users, mobile workloads and their IO requirements alone are driving the need for flash storage attached to z13.

Brocade Gen 6 technology provides a significant leap in IOPS capabilities. The Condor 4 ASIC is capable of handling 33 percent more IOPS than the prior generation Condor 3 ASIC used in Brocade Gen 5 SAN technology. Brocade Gen 5 FICON SAN technology is capable of outstanding performance in a z13 flash storage architecture. Brocade Gen 6 FICON SAN technology takes this performance to a whole new level and will allow flash/SSD to reach its full performance potential.

The enterprise computing space has been at the forefront of the flash storage revolution. Applications that are performance-critical have been/are being moved from the traditional spinning disk DASD arrays to flash storage arrays. This movement started with hybrid arrays, but once mainframe storage array vendors such as EMC, HDS, and IBM brought all flash arrays to the market, end users began migrating to the all-flash array technology. The reduction in latency and improved response times that can be realized by implementing flash storage array technology allow for faster processing of transactions, or more transactions processed for a given time interval. For some industry verticals, this can translate directly into more revenue. A TABB Group study from July 2010 determined that: "A 1-millisecond advantage in a trading application can be worth US$100 million a year to a brokerage firm."

Improved VM Support

The IBM z13 has up to 141 cores using the world's fastest commercial processor, running at 5.0 GHz, to enable high performance coupled with massive scaling. The z13 and the IBM LinuxONE Emperor systems take scalability for VMs to a completely new level and are capable of supporting up to 8,000 virtual servers. When connected to flash storage via FCP channels and Brocade SAN hardware, such as Brocade X6 Directors, the z13 and the LinuxONE Emperor support Node Port ID Virtualization (NPIV). This support of NPIV has increased significantly with z13 and LinuxONE Emperor, doubling to support 64 virtual images per FCP channel from the prior support of 32 virtual images per FCP channel. The IO subsystem of these IBM servers is built on Peripheral Component Interconnect Express (PCIe) Gen 3 technology, which can have up to 40 PCIe Gen 3 fanouts and integrated coupling adapters per system at 16 GBps each (yes, 16 GB, not 16 Gbps). The z13 also supports 10 TB of available RAM (3x its predecessor). This massive amount of memory, as well as server and IO virtualization, allows VMs and applications to run at their full, unconstrained potential. Gen 6 Fibre Channel is ideal for these environments and is capable of supporting the greater virtualization levels at sustained high levels of performance, enabling full utilization of the z13 and LinuxONE Emperor virtualization capabilities.

In the future, zLinux environments connected to storage via Brocade Fibre Channel technology will be able to gain insight into application performance using Brocade Fabric Vision® VM Insight capabilities.

Support for Non-Volatile Memory Express over Fabrics using Fibre Channel

Non-Volatile Memory Express (NVMe) is a new and innovative method of accessing storage media, and has emerged as the new storage connectivity platform that will drive massive performance gains. It is ideal for flash/SSD. Applications will see better random and sequential performance by reducing latency and enabling much more parallelism through an optimized PCIe interface purpose-built for solid-state storage. The momentum behind NVMe has been increasing since it was introduced in 2011. In fact, NVMe technology is expected to improve along two dimensions over the next couple of years: improvements in latency and the scaling up of the number of NVMe devices in large solutions.

In 2014, the Fibre Channel Industry Association (FCIA) announced a new working group within the INCITS T11 committee (responsible for Fibre Channel standards) to align NVMe with Fibre Channel as part of the NVM Express over Fabrics initiative. This was an important evolution because it kept Fibre Channel at the forefront of storage innovation. FC-NVMe defines a common architecture that supports a range of storage networking fabrics for NVMe block storage protocol over a storage networking fabric. This includes enabling a front-side interface into storage systems, scaling out to large numbers of NVMe devices, and extending the distance within a data center over which NVMe devices and NVMe subsystems can be accessed.

FC-NVMe offers compelling performance that is synergistic with the SSD and mainframe technologies that demand it. Gen 6 Fibre Channel is NVMe-ready and offers the ideal connectivity for high-performance mainframe and flash storage.

Summary

The enterprise computing community represents the unique sector of the storage networking market that propels the most critical systems in the world. These systems must operate cohesively and under extremely stringent conditions in order to provide the everyday services and operations enjoyed by billions.

Brocade and the IBM z Systems IO team share an extraordinary relationship that has yielded countless enhancements to the world of enterprise computing. This nearly 30-year bond has stood the test of time, as well as the turmoil of the high-tech marketplace, to produce reliable solutions for these critical systems. At a time when most companies are boasting of alliances around the technology "flavor of the day," Brocade and IBM z Systems have delivered the technology for the ages.

Brocade builds on that 30 years of mainframe leadership to deliver the industry's highest performance and most reliable and scalable FICON infrastructure. With seamless FICON connectivity and support for innovative features that only Brocade can offer—including ClearLink® D_Port, CUP Diagnostics, Port Decommissioning, and ISL encryption—organizations can achieve the full potential from new flash storage and IBM z13 and future mainframe technology.

Learn more at www.brocade.com.

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