

Metro Cloud Connectivity: Integrated Metro SAN Connectivity in Gen 5 Fibre Channel Switches

Table of Contents

Introduction	3
Overview	4
Brocade Gen 5 Fibre Channel—Metro Connectivity Features	4
Native Gen 5 Fibre Channel Long-Distance Support	5
Fibre Channel Buffer-to-Buffer Flow Control.....	5
Native Gen 5 Fibre Channel Long-Distance Support	8
Integrated 10 Gbps Fibre Channel Speed Support	9
The Evolution of Optical Network Standards	9
10 Gbps Fibre Channel Speed Support.....	10
Integrated Inter-Switch Link (ISL) Compression	12
Integrated Inter-Switch Link (ISL) Encryption	13
Enhanced Diagnostic and Error Recovery Technology	14
Forward Error Correction, (FEC)	15
Brocade ClearLink Diagnostic Port, (D_Port)	15
Buffer Credit Loss Detection and Recovery	15
Monitoring and Alerting Policy Suite (MAPS).....	16
Flow Vision.....	16
Conclusion	17

Introduction

As customers look to Fibre Channel Storage Area Networks (FC SANs) for building private storage cloud services for their enterprise data centers, some of the key requirements are:

- Consolidated and highly virtualized pools of compute, storage, and network resources
- Secure and efficient use of inter-fabric connectivity
- Lower capital and operational costs, higher asset utilization
- On-demand and efficient provisioning of application resources through automated management

Brocade has developed solutions to address these key requirements, leveraging the Brocade® Application-Specific Integrated Circuit (ASIC) and Brocade Fabric OS® (FOS) version 7.x in the Gen 5 Fibre Channel platforms, Brocade Network Advisor, and Brocade Host Bus Adapter (HBA) and Converged Network Adapter (CNA) technology (as well as working with transceiver vendors). To enable customers to achieve their goals, Brocade is delivering key technologies (see Figure 1) that allow customers to:

- Scale up/out based on business growth
- Secure data and optimize inter-data center bandwidth utilization
- Reduce CapEx and OpEx cost with built-in diagnostics and SAN management tools
- Optimize both bandwidth and Input/Output Operations Per Second (IOPS), while being energy efficient




Figure 1: Enabling private storage clouds.

Overview

The Fibre Channel network architecture enabled the deployment of the SAN within IT infrastructures beginning in the mid-1990s. It also became the catalyzing point in the evolution of the dynamic data center. As the connection between server and storage changed from a direct connection to a network connection, IT architects began building what are understood today as cloud data center infrastructures. Specifically, with the introduction of FC and SANs, IT architects began building what we know today as the storage cloud.

Almost from the initial deployments of the SAN in the data center, there has been a clear understanding of the need to protect this critical infrastructure. Following the basic high-availability precept of designing a system with no single point of failure, IT architects realized that regardless of how well-designed, the local SAN infrastructure would eventually require protection from a total site failure. This realization led to the development and implementation of site-to-site SAN infrastructures, alternately referred to as SAN extension, wide-area SAN, or metro-wide-area SAN networks.

Today there are two general methods for extending the reach of a SAN from one physical site to another. The first method is by transporting the native FC protocol, without any kind of protocol conversion, over a physical link between sites, through some type of fiber-optic connectivity. The second method transports the FC protocol by encapsulating it within a transport protocol, which is then sent over the physical link between sites. The most common implementation of this method utilizes the Fibre Channel over IP (FCIP) protocol to transport FC frames by encapsulating them within an IP frame, which is then sent over a standard IP network that links both sites. The main difference between these two methods is that the native FC method offers better performance but at shorter distances, while the encapsulation method offers longer distances but at lower performance rates. Both methods have their place in today's wide-area storage cloud infrastructures.

Brocade GEN 5 Fibre Channel–Metro Connectivity Features

Since delivering the first FC SAN switch in the industry in March 1997, Brocade has been a leader in developing SAN extension solutions to meet the distance connectivity needs of IT SAN architects, utilizing native as well as encapsulated extension solutions. However, with the introduction of the Brocade Gen 5 Fibre Channel SAN switching solutions, Brocade has again raised the bar for SAN distance extension solutions.

Powered by the Brocade engineered Application-Specific Integrated Circuit (ASIC), Brocade Gen 5 Fibre Channel SAN switches support a number of new features that allow IT architects to build larger, more reliable, wide-area “stretched” SAN storage clouds.

With the new class of Gen 5 Fibre Channel switches, Brocade has introduced a number of ASIC integrated extension capabilities that allow IT architects to design and implement a variety of metro-area SAN extension solutions without additional hardware. The metro area is defined as the distance from site to site of no more than 100 km (62 miles), and with a Round-Trip Time (RTT) latency of no more than 5 milliseconds. In the event that SAN extension solutions are required over greater distances, with longer RTT times, Brocade provides dedicated SAN extension solutions that satisfy the most rigorous of these long-distance requirements.

The Brocade Gen 5 Fibre Channel metro connectivity features include:

- **Native Gen 5 Fibre Channel Long-Distance Support:** In addition to doubling the overall throughput, Gen 5 Fibre Channel switches are able to utilize a buffer credit pool

of 8,192 buffers, which quadruples the Brocade 8 Gbps buffer credit pool of 2,048 buffers. A variety of distance extension architectures are also supported, utilizing native Fibre Channel connectivity.

- **Integrated 10 Gbps Fibre Channel Speed Support:** Brocade Gen 5 Fibre Channel switches support port operating speeds of not only 16, 8, 4, and 2 Gbps, but they also support a port operating speed of 10 Gbps. In addition to native fiber, operating the port at 10 Gbps is also supported over Wave Division Multiplexing (WDM) solutions, such as Dense Wave Division Multiplexing (DWDM) and Coarse Wave Division Multiplexing (CWDM).
- **Integrated Inter-Switch Link (ISL) Compression:** Brocade Gen 5 Fibre Channel switches provide the capability to compress all data in flight, over an ISL. This requires a Brocade Gen 5 Fibre Channel switch on both sides of the ISL, and a maximum of 4 ports per Brocade DCX® 8510 Backbone blade, 8 ports per Brocade 6520 Switch, or 2 ports per Brocade 6510 Switch can be utilized for this data compression. With Brocade FOS v7.1 or later, additional ports can be utilized for data compression if running at lower than 16 Gbps speeds. The compression rate is typically 2:1.
- **Integrated Inter-Switch Link (ISL) Encryption:** Brocade Gen 5 Fibre Channel switches provide the capability to encrypt all data in flight, over an ISL. This requires a Brocade Gen 5 Fibre Channel switch on both sides of the ISL, and a maximum of 4 ports per Brocade DCX 8510 blade, 8 ports per Brocade 6520, or 2 ports per Brocade 6510 can be utilized for this data encryption. With Brocade FOS v7.1 or later, additional ports can be utilized for data encryption if running at lower than 16 Gbps speeds. Both encryption and compression can be enabled on the ISL link simultaneously.
- **Enhanced Diagnostic and Error Recovery Technology:** Brocade Fabric Vision technology, an extension of Gen 5 Fibre Channel, is a breakthrough hardware and software technology that offers innovative diagnostic, monitoring, and management capabilities to help administrators avoid problems, maximize application performance, and reduce operational costs. Fabric Vision technology includes diagnostic and recovery features that enable smooth operation of metro-connected SAN clouds. These features also include the ability to measure, verify, and saturate the ISL links between switch Expansion Ports (E_Ports), as well as between a Brocade Gen 5 Fibre Channel adapter and Gen 5 Fibre Channel switch, utilizing a Brocade-unique ClearLink diagnostic port (D_Port) feature. Additionally, Gen 5 Fibre Channel switches are able to detect and recover from buffer credit loss situations, in some cases without traffic disruption. Finally, SAN architects have the ability to enable Forward Error Correction (FEC) on ISL E_Ports and EX_Ports to improve error detection and correction.

Native GEN 5 Fibre Channel Long-Distance Support

Fibre Channel Buffer-to-Buffer Flow Control

The Fibre Channel network flow control mechanism is described as a buffer-to-buffer credit system. To transmit a frame in this system, the sending port must know beforehand that a buffer is available at the receiving port. This forward flow control model is one of the key technologies that ensure that no frames are ever dropped or lost in a normally operating Fibre Channel network.

The mechanism that allows an FC port to know, ahead of time, how many frames they can send is described as a buffer credit pool. This information is determined upon port initialization, when port pairs exchange information about how many buffer credits each port has available and how many buffer credits each port requires.

When the port knows this information, it can begin sending frames. The key transaction is that when a port sends a frame, it decrements its buffer credit pool by one. When the receiving port receives the frame, it sends an acknowledgement primitive, called Receiver Ready or R_RDY back to the sending port, which then allows it to increment its buffer credit pool by one.

This way, frames can be sent over very complex FC network architectures with guaranteed forward flow control and ensure lossless delivery. Furthermore, in the Brocade FC architecture, if congestion does develop in the network because of a multiplexed ISL traffic architecture, all sending traffic slows down gradually in response to the congestion. (Brocade refers to a multiplexed ISL traffic architecture as Virtual Channels [VCs].) Without an architecture that provides fair access to the wire for all traffic streams, congestion could impose a condition where one traffic stream blocks other streams, generally referred to as Head-of-Line (HoL) blocking.

Buffer-to-buffer credits are generally allocated and managed transparently within the FC SAN fabric. In most cases, architects do not need to overtly manage them. However, when architecting a SAN design that incorporates a long-distance link (>500 meters), it is important to understand the buffer credit requirements and allocations that are available across these long-distance links.

The challenge with long-distance links is that even travelling at the speed of light, frames take time to get from one end of an optical cable to the other. If the cable is long enough and the link speed is fast enough, the result is a situation where multiple frames are in transit on the wire at any point in time. With each frame that is sent, the sending port is decreasing its buffer credit pool by one and, if this goes on long enough without receiving the R_RDY primitives returning across the long link, its buffer credits will reach zero and it is forced to wait for its buffer credits to be replenished before it can send frames again.

In the case of long-distance FC links, the sending port needs to have enough buffer credits available to fill the link with as many frames as is required so that this condition does not occur. The SAN architect must ensure that there are enough credits to keep the ISLs "full" at all times.

As a result, SAN architects understand that the relationship between the length of the link, the speed of the link, and the frame size that is being transmitted across the link determines the correct number of buffer credits that are required for the link. A general formula for determining the correct number of buffer credits that are required for a given port-to-port link in an FC network is as follows:

$$\text{(link speed in Gbps * distance in kilometers) / frame size in kilobytes} = \text{the number of buffer credits that are required}$$

This means that for a 16 Gbps link of 500 km moving frames of 2 KB in size would require:

$$(16 * 500) / 2 = 4,000$$

or approximately 4,000 buffer credits to be available to the ports on both sides of the extended SAN link in order to operate at full speed.

The problematic parameter in this equation is, of course, the frame size. Fibre Channel defines a variable length frame consisting of 36 bytes of overhead and up to 2,112 bytes of payload for a total maximum size of 2,148 bytes. Devices such as HBAs and storage arrays negotiate to 2 KB frame sizes for payload. Though this means that a majority

of frames are full size (2 KB), lower frame sizes are used for various Small Computer Systems Interface (SCSI) commands and FC-class F traffic (such as zoning updates, Registered State Change Notifications [RSCNs], and name server information). In many instances, SAN architects typically assume the maximum 2 KB Fibre Channel frame size for these buffer credit calculations. However, if the actual frame size in flight is only, for example, 1 KB, the amount of buffer credits that are required is doubled. A more accurate parameter for these formulas would be the average frame size.

It is important to understand that traffic can still flow in an FC network when insufficient buffers are available—they will just operate at slower line rates. This condition is known as buffer credit starvation and traffic continues to flow as long as buffer credits are not lost. If, in fact, a buffer credit pool reaches zero, the port can no longer send frames until the credits are replenished or until the link is reset. So, using our earlier example, if only 2,000 buffer credits were available, our 16 Gbps 500 km link with an average frame size of 2 KB would be capped at an 8 Gbps line rate.

It is also important to understand that, assuming an average 2 KB frame size, if our example ISL of 16 Gbps at 500 km is assigned 4,500 buffer credits (more than the required 4,000), there is no performance improvement. In other words, assigning more credits does not yield better line-rate performance.

In the Brocade FC SAN fabric, to enable advanced buffer credit configurations, the ports must be configured in one of several long-distance modes, two of which are enabled via the Extended Fabrics license. This table lists the Brocade Fabric long-distance modes.

Table 1: Brocade Fabric Long-Distance Modes

Distance Mode	Distance	License Required
LO	Local Data Center	No
LE	10 km	No
LD	Auto-Discovery	Yes
LS	Static Assignment	Yes

LO designates local or “normal” buffer credit allocations. This is the default port configuration and, as

previously mentioned, it provides sufficient buffer credits within the normal data center link distances (less than 500 meters).

LE is used to support distance up to 10 km and does not require a license. The 10 km limit is not dependent on speed because if the ports negotiate to higher speeds, more credits are automatically assigned to support the higher line speed.

LD (dynamic distance discovery mode) is the most user-friendly mode. It automatically probes the link and, uses a sophisticated algorithm to calculate the amount of credits that are required, based on the distance and speed set for the link.

LS is a statically configured mode that was added to Brocade FOS v5.1. This is the most flexible mode for the advanced user. It allows complete control of buffer credit assignments, for long-distance requirements.

Brocade FOS v7.1 provides users additional control when configuring a port of an LD or LS link. Users can specify the buffers required or the average frame size for a long-distance port. Using the frame size option, the number of buffer credits required for a port is automatically calculated. These options give users additional flexibility to optimize performance on long-distance links.

In addition, Brocade FOS v7.1 provides users better insight into long-distance link traffic patterns by displaying the average buffer usage and average frame size via Command Line Interface (CLI). Brocade FOS v7.1 also provides a new CLI that automatically calculates the number of buffers required per port given the distance, speed, and frame size.

Native Gen 5 Fibre Channel Long-Distance Support

With the introduction of Gen 5 Fibre Channel, Brocade has significantly expanded the buffer architecture that is available for SAN storage cloud architects, providing a massive 8,192 buffers. This is a four-fold increase over the 2,048 buffers that are available with the previous generation platforms.

When you calculate the maximum amount of buffer credits that are available, it is important to note that the Brocade FC ASIC architecture reserves some buffers. The buffers in reserve are not available for long-distance configuration. For example, the Gen 5 Fibre Channel ASIC reserves eight buffers for every front-end Fibre Channel port, 48 buffers for communicating with the switch control processor, and 48 buffers for additional internal functions. Additional buffers are also reserved internally in multi-ASIC switches, such as the Brocade DCX, to support the internal communications between port blade and core blade. This table lists the total number of reserved and available buffers, per ASIC, within the specific Brocade DCX port blade or standalone switch.

Table 2: Reserved and Available Buffers for Gen 5 Fibre Channel Platforms

Gen 5 Fibre Channel Switch Type	Total Reserved Buffers	Total Available Buffers
Brocade 6520 Switch	3456	4736
Brocade 6510 Switch	480	7,712
Brocade 6505 Switch	288	7904
Brocade DCX FC16-32 Port Blade	2,784	5,408
Brocade DCX FC16-48 Port Blade	3,232	4,960

Using the buffer credit formula that was provided earlier, and assuming an average frame size of 2 KB, the specific long-distance link distance support of Gen 5 Fibre Channel switches are listed in this table.

Table 3: Distance Support at Various Link Speeds for Gen 5 Fibre Channel Platforms

Link Speed	Brocade 6520	Brocade 6510	Brocade 6505	Brocade DCX	Brocade DCX
				FC16-32 Port Blade	FC16-48 Port Blade
				Single ASIC	Single ASIC
2 Gbps	4,736 km	7,712 km	7,904 km	5,408 km	4,960 km
4 Gbps	2,368 km	3,856 km	3,952 km	2,704 km	2,480 km
8 Gbps	1,184 km	1,928 km	1,976 km	1,352 km	1,240 km
16 Gbps	592 km	964 km	988 km	676 km	620 km

The Gen 5 Fibre Channel long-distance buffer architecture supports distances that are far in excess of what the current Small Form-Factor Pluggable (SFP) optical technology can support, (that is, light), today. These calculations offer dramatic examples of the expanded capability of Gen 5 Fibre Channel extended distance support. However, the practical benefit of this feature could mean that instead of an increase in distance between any two SANs, IT architects can increase the number of SANs that are connected within a metro wide area. SAN architects who utilize Gen 5 Fibre Channel switches can design larger, classic "hub and spoke" metro-wide-area SAN storage clouds.

Lastly, because optical network distances must be designed with the appropriate port optics technology (as pointed out earlier), Brocade provides Short Wavelength Laser (SWL), Long Wavelength Laser (LWL) and Extended Long Wavelength Laser (ELWL), 16 Gbps Small Form-Factor Pluggable Plus (SFP+) optics for the Gen 5 Fibre Channel switches. Brocade LWL SFP+ optics support distances of up to 10 km and ELWL SFP+ optics support distances of up to 25 km.

Integrated 10 Gbps Fibre Channel Speed Support

The Evolution of Optical Network Standards

One of the remarkable achievements of the Fibre Channel network standards was that it introduced a network architecture that could evolve through successive generations without imposing significant infrastructure upgrade costs. For this reason, after the optical network industry switched from the larger 1 Gbps Gigabit Interface Converter (GBIC) pluggable optics to the smaller SFP optics technology, the FC network upgrade from 2 Gbps to 4 Gbps, and finally to 8 Gbps, was achieved fairly easily and cost-efficiently.

Each successive advance in speed meant that as long as the distance limits were sufficient (2 Gbps at 500 meters, 4 Gbps at 380 meters, and 8 Gbps at 150 meters), an IT architect could--for the first time--double the speed of their network without changing the wiring plant. Additionally, because the same conversion layer encoding is used (the 8b/10b scheme), the SFP optics technology required only small changes to support the increasing line rates.

Minimizing change was a goal of the ANSI T11.3 subcommittee, which is the standards group responsible for the main body of Fibre Channel network standards. However, changes were introduced into the network conversion layer when the network speed advanced to 10 Gbps. The increased speed also required changes to the optics technology that was deployed.

With the advent of high-speed optical networking, the relevant network standards organizations made a concerted effort to harmonize the physical and conversion layer standards for different network architectures, where possible. By standardizing, IT architects could benefit from a physical layer that could support multiple network architectures. The result was that the ANSI T11.3 subcommittee and the IEEE 802.3 working group (the standards group that is responsible for the main body of Ethernet standards) settled on the same optical physical and conversion layer standards developed in the T11.3 subcommittee for 1 Gbps FC, specifically standard FC-1. This standardization was mirrored in the IEEE 802.3 Clause 38 Physical Coding Sublayer (PCS) standard. Additionally, because Ethernet did not use the interim 2, 4, and 8 Gbps speeds, the T11.3 subcommittee further developed the physical layer standards for 2, 4, and 8 Gbps FC as extensions to the original FC-1 standard. This compatibility enabled the “easy to upgrade” FC network architecture that exists today.

However, when network speeds advanced to 10 Gbps, the 802.3 group was first to define new conversion layer standards that were different than previous standards. These new standards resulted in the IEEE 802.3 Clause 49 PCS standard, which changed the conversion layer encoding scheme from the 8b/10b scheme to the 64b/66b scheme. This change required new optical technology and standards that, initially resulted in a new, larger, 10 Gbps SFP optical module, referred to as an XFP. The ANSI T11.3 subcommittee also defined the 10 Gbps FC conversion layer standards as utilizing the new 64b/66b encoding conversion scheme. The T11.3 subcommittee further developed new optics standards that provided better backwards compatibility with existing SFP technology, which resulted in the development of the SFP+ standards. The resulting SFP+ optics utilize the same form factor as earlier SFP optics, and draws less power than XFP optical modules. Today, SFP+ is the most popular optical socket technology deployed for 10 Gbps Ethernet as well as for Fibre Channel.

Because of the changes in conversion layer, a noticeable shift occurred in Fibre Channel network architectures with the segregation of 10 Gbps ports and their related technology from the legacy 1, 2, 4, and 8 Gbps FC ports. Some switch vendors began building dedicated 10 Gbps FC ports, which could typically only be used as ISL ports. However, if the port was not needed, customers still had to pay for the port capacity, which forced a “stranded capacity” situation.

Brocade took a different approach, developing a special 10 Gbps blade, called the FC10-6, for the Brocade 48000 4 Gbps Director, which provided 6-10 Gbps FC ports, for use as ISL ports between similar FC10-6-equipped Brocade 48000 Directors. The FC10-6 blade design incorporated two sets of FC ASICs—the 4 Gbps ASIC and the 10 Gbps ASIC.

10 Gbps Fibre Channel Speed Support

The ANSI T11.3 subcommittee also defined the conversion layer encoding scheme for 16 Gbps and 32 Gbps FC speeds to utilize the 64b/66b encoding. Additionally, SFP+ optical technology can be utilized for not only 10 Gbps line rates, but also for higher speed line rates, such as Brocade Gen 5 Fibre Channel switches.

However, Brocade developed a further integrated enhancement, which provides the SAN architect with new capabilities in designing a metro-wide-area SAN. The Gen 5 Fibre Channel ASIC, in addition to supporting FC line rates of 2, 4, 8, and 16 Gbps, also supports FC line rates of 10 Gbps. More specifically, it can do this without specialized hardware and without forcing “stranded capacity.”

In comparison to long-distance fiber-optic links between Brocade Gen 5 Fibre Channel switches, which can run natively at 16 Gbps, the ability to run ports at 10 Gbps might not seem like a benefit. However, if the physical link between SANs is provided through alternate service providers, this capability gives SAN architects the required flexibility in designing a metro-area SAN architecture by providing compatibility with other wide-area network technology.

Today, IT architects can link SANs in a metro area for native FC protocol transmission. They can link the SANs either by directly utilizing a fiber-optic cable between sites or by creating multiple channels on an optical cable between sites, utilizing WDM technology. WDM is a technique for providing multiple channels across a single strand of fiber-optic cable. The optical signals are sent at different wavelengths of light, also called lambda circuits. The two most common WDM technologies are DWDM and CWDM. The main benefit of deploying WDM technology is that you can increase the amount of traffic and the types of traffic over a single optical cable.

Additionally, both types of metro-area SAN connectivity links, either direct cable, or WDM, can be deployed directly, or can be purchased as a service, from a service provider. Even in the case where an IT architect either outright owns or leases the entire fiber optic cable that links two data center sites together, competing interests within the organization might require dividing the cables into multiple channels with WDM technology.

The value driving port speeds on the Brocade Gen 5 Fibre Channel switches at the 10 Gbps rate is because most WDM technology does not currently support 16 Gbps rates. Rather than having to throttle down to either 8 Gbps or 4 Gbps line rates, and waste additional lambda circuits to support required bandwidth, Brocade Gen 5 Fibre Channel switches can drive a given lambda circuit at a 10 Gbps line rate, optimizing the link. Brocade has successfully tested this configuration with CWDM/DWDM solutions from Adva, in the form of the Adva Optical FSP 3000, and Ciena, in the form of the Ciena ActivSpan 4200. Brocade will continue to test additional CWDM/DWDM solutions in the future to ensure compatibility with a wide variety of technology providers. Refer to the Brocade Compatibility Matrix for the full list of compatible third-party CWDM/DWDM solutions.

The actual configuration of the 10 Gbps FC line rate on Gen 5 Fibre Channel switches is done by configuring the speed for an 8-port group, called an octet. This table lists octet speed combination options available on the Brocade Gen 5 Fibre Channel switches.

Table 4: Port Speeds Supported by Speed Mode

Speed Mode	Port Speed Supported
1	16 Gbps, 8 Gbps, 4 Gbps, 2 Gbps
2	10 Gbps, 8 Gbps, 4 Gbps, 2 Gbps
3	16 Gbps, 10 Gbps

The default speed mode is 1, which means any port in the eight-port group octet can operate at 16, 8, or 4 Gbps, utilizing 16 Gbps SFP+ optics, or at 8, 4, or 2 Gbps, utilizing 8 Gbps SFP+ optics. Speed combination modes 2 and 3 enable any port in the octet to operate at a 10 Gbps line rate, but also specifically require 10 Gbps SFP+ optics. The 10 Gbps SFP+ optics are available in SWL and LWL models.

Note that the changing of the octet speed mode is a disruptive event and, prior to Brocade FOS v7.2, is supported only on the first eight ports on any blade in the Brocade DCX 8510 (4-slot and 8-slot) and the first eight ports on the Brocade 6520 or 6510. The maximum configuration supported provides 64 ports of 10 Gbps across all eight port blades on the Brocade DCX 8510-8, 32 ports of 10 Gbps across all four port blades on the Brocade DCX 8510-4, or eight ports of 10 Gbps on the Brocade 6520 or 6510. With Brocade FOS v7.2 and later, this restriction is lifted, allowing any port on a 16 Gbps blade/switch that also supports 10 Gbps Fibre Channel functionality to be configured for 10 Gbps speed.

Implementing the 10 Gbps FC line-rate feature does not require any additional hardware, but does require that a Brocade Extended Fabrics license be enabled on both switches. Additionally, the 10 Gbps FC line-rate capability of the Brocade Gen 5 Fibre Channel switches is not compatible with the prior Brocade FC10-6 10 Gbps FC port blade. So, to establish a 10 Gbps ISL between sites, both sites must be connected utilizing Brocade Gen 5 Fibre Channel switches.

Integrated ISL Compression

Brocade Gen 5 Fibre Channel switches introduce a new capability for metro-area SAN architects: ASIC-integrated, ISL, in-flight, data compression. Each ASIC can provide up to 32 Gbps of compression, via a maximum of 2 - 16 Gbps FC ports, which can be combined and load-balanced, utilizing Brocade ISL Trunking.

Because the Brocade DCX 32-port and 48-port 16 Gbps port blades are equipped with two ASICs, a single port blade in the Brocade DCX 8510 can provide up to 64 Gbps of ISL data compression, utilizing four ports at 16 Gbps speed, six ports at 10 Gbps speed, or eight ports at 8 Gbps speed. The maximum Brocade DCX configuration supported provides 512 Gbps of compression across all eight port blades in the Brocade DCX 8510-8, or 256 Gbps of compression across all four port blades in the Brocade DCX 8510-4. The Brocade 6520 provides up to 128 Gbps of compression, utilizing up to eight FC ports at 16 Gbps speed, 12 ports at 10 Gbps speed, or 16 ports at 8 Gbps speed. The Brocade 6510 is limited to providing up to 32 Gbps of compression, utilizing up to two ports at 16 Gbps speed, three ports at 10 Gbps speed, or four ports at 8 Gbps speed. Integrated data compression is not available on the Brocade 6505.

This compression technology is described as "in-flight" because this ASIC feature is enabled only between E_Ports and EX_Ports, which allows ISL links to have the data compressed as it is sent from the Gen 5 Fibre Channel switch on one side of an ISL and then decompressed as it is received by the Gen 5 Fibre Channel switch that is connected to the other side of the ISL. As mentioned earlier, in-flight ISL data compression is supported across trunked ISLs, as well as multiple ISLs and long-distance ISLs. Brocade Fabric Quality of Service (QoS) parameters are also honored across these ISL configurations.

The compression technology utilized is a Brocade-developed implementation that utilizes a Lempel-Ziv-Oberhumer (LZO) lossless data compression algorithm. The compression algorithm provides an average compression ratio of 2:1, and all Fibre Channel Protocol (FCP) and non-FCP frames that transit the ISL are compressed. The frames that are not compressed are Basic Link Services (BLS) as defined in the ANSI T11.3 FC-FS standard, and Extended Link Services (ELS) as defined in the ANSI T11.3 FC-LS standard.

When enabling the in-flight ISL data compression capability, the ISL port must be configured with additional buffers, requiring the switch port to be configured in an LD mode. Enabling in-flight ISL data compression also increases the time it takes for the ASIC to move the frame. This time is described as latency and should be understood by SAN architects. Normally the transit time for a 2 KB frame to move from one port to another port on a single ASIC is approximately 700 nanoseconds, a nanosecond representing one-billionth (10⁻⁹) of a second. Adding in-flight data compression increases the overall latency by approximately 5.5 microseconds, a microsecond representing one-millionth (10⁻⁶) of a second.

The approximate latency time is 6.2 microseconds for a 2 KB frame to move from a source port, be compressed, and then move to the destination port on a single ASIC. Of course, calculating the total latency across an ISL link means including the latency calculations for both ends. For example, compressing a 2 KB frame and sending it from one Gen 5 Fibre Channel switch to another would result in a total latency of 12.4 microseconds, (6.2 * 2), not counting the link transit time.

One of the use cases for utilizing ASIC-integrated ISL data compression is when a metro-area SAN infrastructure includes an ISL for which there are either bandwidth caps or bandwidth usage charges.

Finally, implementing the ISL compression capability requires no additional hardware and no additional licensing.

Integrated ISL Encryption

Brocade Gen 5 Fibre Channel switches, in addition to ISL data compression, also introduce the new capability of implementing ASIC-integrated ISL, in-flight, data encryption. Each ASIC can provide up to 32 Gbps of encryption, via a maximum of 2 - 16 Gbps FC ports, which can, again, be combined and load-balanced, utilizing Brocade ISL Trunking.

The two ASICs on the Brocade DCX 32-port and 48-port 16 Gbps port blades enable a single port blade in the Brocade DCX 8510 to provide up to 64 Gbps of ISL data encryption, utilizing four ports at 16 Gbps speed, six ports at 10 Gbps speed, or eight ports at 8 Gbps speed. The maximum Brocade DCX configuration supported provides 512 Gbps of encryption across all eight port blades in the Brocade DCX 8510-8, or 256 Gbps of encryption across all four port blades in the Brocade DCX 8510-4. The Brocade 6520 provides up to 128 Gbps of encryption, utilizing up to eight FC ports at 16 Gbps speed, 12 ports at 10 Gbps speed, or 16 ports at 8 Gbps speed. The Brocade 6510 is limited to providing up to 32 Gbps of encryption, utilizing up to two FC ports at 16 Gbps speed, three ports at 10 Gbps speed, or four ports at 8 Gbps speed. Integrated data encryption is not available on the Brocade 6505.

As with integrated compression, the integrated encryption is supported in-flight, exclusively for ISL E_Ports and EX_Ports, linking Gen 5 Fibre Channel switches. Enabling ISL encryption results in all data being encrypted as it is sent from the Gen 5 Fibre Channel switch on one side of an ISL and then decrypted as it is received by the Gen 5 Fibre Channel switch connected to the other side of the ISL. As with integrated ISL compression, this integrated ISL encryption capability is supported across trunked ISLs, as well as multiple ISLs and long-distance ISLs. Brocade Fabric QoS parameters are also honored across these ISL configurations. It is important to note that, when implementing ISL encryption, using multiple ISLs between the same switch pair requires that all ISLs be configured for encryption or none at all.

Compression and encryption can be enabled, utilizing the integrated features of the Brocade Gen 5 Fibre Channel switches. As is the case with integrated data compression, enabling integrated data encryption adds approximately 5.5 microseconds to the overall latency. An approximate latency time of 6.2 microseconds is added for a 2 KB frame to move from a source port, be encrypted, and then move to the destination port on a single ASIC. Also, when you calculate the total latency across an ISL link, include the ASIC latency calculations for both ends. Encrypting a 2 KB frame and sending it from one Gen 5 Fibre Channel switch to another results in a total latency of 12.4 microseconds ($6.2 * 2$), not counting the link transit time. If both encryption and compression are enabled, those latency times are not cumulative. For example, compressing and then encrypting a 2 KB frame incurs approximately 6.2 microseconds of latency on the sending switch and incurs approximately 6.2 microseconds of latency at the receiving switch to decrypt and uncompress the frame. The total latency time is 12.4 microseconds, again, not counting the link transit time.

The encryption method utilized for the integrated ISL encryption is the Advanced Encryption Standard (AES) AES-256 algorithm using 256 bit keys, and uses the Galois Counter Mode (GCM) of operation. AES-GCM was developed to support high-throughput Message Authentication Codes (MACs) for high data rate applications such as high-speed networking. In AES-GCM, the MACs are produced using special structures called Galois field multipliers, which are multipliers that use Galois field operations to produce their results. The key is that they are scalable and can be selected to match the throughput requirement of the data.

As with integrated ISL data compression, when enabling integrated ISL encryption, all FCP and non-FCP frames that transit the ISL are encrypted, with the exception of BLS and ELS frames. To enable integrated ISL encryption, port-level authentication is required and Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) must be enabled. The Internet Key Exchange (IKE) protocol is used for key generation and exchange. The key size is 256 bits, the Initialization Vector (IV) size is 64 bits, and the Salt size is 32 bits. Traditional encryption systems require a key management system for creating and managing the encryption keys. The integrated ISL encryption capability is implemented with a simpler design, utilizing a non-expiring set of keys that are reused. While non-expiring and reused keys can represent a security concern, these keys also allow this integrated ISL encryption to be implemented with very little management impact.

One use case for utilizing integrated ISL encryption is to enable a further layer of security for a metro-area SAN infrastructure. Implementing the ISL encryption capability also requires no additional hardware and no additional licensing.

Advanced Diagnostics and Analytics

Brocade Gen 5 Fibre Channel switches include Fabric Vision technology, offering several innovative diagnostic, monitoring, and management technologies. These technologies are designed to enable the IT SAN architect to monitor, manage, diagnose, and troubleshoot larger and more complex SAN storage clouds. Some of these technologies have increased relevance when planning a metro- or wide-area SAN network infrastructure.

The following features are integrated into the Gen 5 Fibre Channel ASIC, and require no additional hardware or licenses.

Forward Error Correction

The Brocade Gen 5 Fibre Channel ASIC includes integrated Forward Error Correction (FEC) technology, which can be enabled on E_Ports connecting ISLs between switches. FEC is a system of error control for data transmissions, whereby the sender adds systematically generated Error-Correcting Code (ECC) to its transmission. This system allows the receiver to detect and correct errors without the need to ask the sender for additional data.

The Brocade implementation of FEC enables the ASIC to recover bit errors in both 16 Gbps and 10 Gbps data streams. The FEC implementation can enable corrections of up to 11 error bits in every 2,112-bit transmission. This effectively enhances the reliability of data transmissions and is enabled by default on Gen 5 Fibre Channel switch E_Ports.

Enabling FEC does increase the latency of FC frame transmission by approximately 400 nanoseconds, which means that the time it takes for a frame to move from a source port to a destination port on a single ASIC with FEC enabled is approximately .7 to 1.2 microseconds. SAN administrators also have the option of disabling FEC on E-Ports.

Brocade ClearLink Diagnostic Port

Brocade Gen 5 Fibre Channel switches deliver enhanced diagnostic capabilities in the form of a unique port type, called the ClearLink diagnostic port (D_Port). The ClearLink D_Port is designed to diagnose optics and cables before they are put into production. ClearLink performs electrical and optical loopback tests, as well as link-distance measurement and link saturation testing.

The ClearLink diagnostic capability provides an opportunity to measure and thoroughly test links before they are put into production. ClearLink can also be used to test active links. However, the link must first be taken offline to enable the ClearLink D_Port configuration and tests.

Initially supported only on Gen 5 Fibre Channel E_Ports, Brocade FOS v7.1 extends the ClearLink diagnostic functionality to Brocade fabric adapters running at 16 Gbps speeds and Gen 5 Fibre Channel switches running in Access Gateway mode.

Brocade Gen 5 Fibre Channel SFP+ optics support all ClearLink D_Port tests, including loopback and link tests. The accuracy of the 16 Gbps SFP+ link measurement is within 5 meters. 10 Gbps SFP+ optics do not currently support the loopback tests, but they do support the link measurement as well as link saturation tests, and provide link measurement accuracy to within 50 meters.

Buffer Credit Loss Detection and Recovery

Another category of Brocade Gen 5 Fibre Channel diagnostic and error recovery technologies is in the area of buffer credit loss detection and recovery. It should be evident by now that the management of buffer credits in wide-area SAN architectures is critically important. Furthermore, many issues can arise in the SAN network whenever buffer credit starvation or buffer credit loss occurs.

Conditions where a particular link may be starved of buffer credits could include either incorrect long-distance buffer credit allocations or links where buffer credits are being lost. Lost buffer credits can be attributed to error conditions such as a faulty physical layer component or misbehaving end node devices. If this condition persists untreated, it can result in a "stuck" link condition whereby the link is left without buffer credits for an extended time period, (for example, 600 milliseconds), stopping all communications across the link.

These problem conditions are only exacerbated when they exist in wide-area SAN architectures. The Brocade Gen 5 Fibre Channel ASIC includes a number of new features that are designed to diagnose, troubleshoot, and recover from these types of conditions.

As previously mentioned, the Brocade Fibre Channel network implements a multiplexed ISL architecture called Virtual Channels (VCs), which enables efficient utilization of E_Port to E_Port ISL links. However, being able to diagnose and troubleshoot buffer credit issues at the VC level is very important.

While the previous generation Brocade ASIC and Brocade FOS provide the ability to detect buffer credit loss and recover buffer credits at the port level, the Brocade Gen 5 Fibre Channel ASIC diagnostic and error recovery feature set includes these features:

- The ability to detect and recover from buffer credit loss at the VC level
- The ability to detect and recover “stuck” links at the VC level

Brocade Gen 5 Fibre Channel switches can actually detect buffer credit loss at the VC level and, if the ASICs detect only a single buffer credit lost, can restore the buffer credit without interrupting the ISL data flow. If the ASICs detect more than one buffer credit lost or if they detect a “stuck” VC, they can recover from the condition by resetting the link, which would require retransmission of frames that were in transit across the link at the time of the link reset.

Brocade Fabric Vision technology also offers several breakthrough capabilities that dramatically simplify day-to-day SAN administration and provide unprecedented visibility across the extended storage network. The following advanced technologies and capabilities are optionally available.

Monitoring and Alerting Policy Suite (MAPS)

MAPS is a new, easy-to-use, policy-based threshold monitoring and alerting suite that proactively monitors the health and performance of the SAN infrastructure to ensure application uptime and availability. By leveraging pre-built rule/policy-based templates, MAPS simplifies threshold configuration, monitoring, and alerting. Using Brocade Network Advisor, organizations can configure the entire fabric (or multiple fabrics) at one time using common rules and policies. They also can customize policies for specific ports or switch elements, all through a single dialog. The integrated dashboard displays an overall switch health report, along with details on out-of-policy conditions, to help administrators quickly pinpoint potential issues and easily identify trends and other behaviors that occur on a switch or fabric.

Flow Vision

Brocade Flow Vision enables administrators to identify, monitor, and analyze specific application and data flows in order to maximize performance, avoid congestion, and optimize resources. Flow Vision includes the following applications:

- **Flow Monitor:** Provides comprehensive Logical Unit Number (LUN)-level visibility into flows in the fabric, including the ability to automatically learn (discover) flows and non-disruptively monitor specific flows or frame types to identify resource contention or congestion that impacts application performance.
- **Flow Generator:** Provides a built-in test traffic generator for pretesting and validating the SAN infrastructure, including the internal connections within a switch, for robustness before deploying applications.
- **Flow Mirror:** Provides the ability to non-disruptively create copies of specific flows or frame types that can be captured for in-depth analysis.

Conclusion

Developed in partnership with our customers, Brocade Gen 5 Fibre Channel SAN platforms with Fabric Vision technology unleash the full potential of cloud architectures and Flash storage. The powerful Brocade Gen 5 Fibre Channel ASIC, Brocade FOS, and Brocade Network Advisor provide integrated technologies that enable building and managing metro-area SAN storage clouds more efficiently and cost-effectively than ever before.

Being able to optimize performance and protect SAN storage clouds from single points of failure, including site failures, is a requirement that continues to evolve. As Brocade works to integrate more metro- and wide-area SAN technologies into the Brocade Gen 5 Fibre Channel platforms, the expectation is that Brocade customers will continue to guide Brocade forward with their requirements as their SAN storage clouds evolve.

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



© 2015 Brocade Communications Systems, Inc. All Rights Reserved. 05/15 GA-WP-374-04

ADX, Brocade, Brocade Assurance, the B-wing symbol, DCX, Fabric OS, HyperEdge, ICX, MLX, MyBrocade, OpenScript, The Effortless Network, VCS, VDX, Vplane, and Vyatta are registered trademarks, and Fabric Vision and vADX are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of others.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment features, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This information document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

BROCADE 