

**DATA CENTER**

## **Building Private Cloud Storage Infrastructure With The Brocade DCX 8510 Backbone**

A best-in-class hardware architecture enables highly virtualized and private cloud storage data center networks.

**BROCADE**

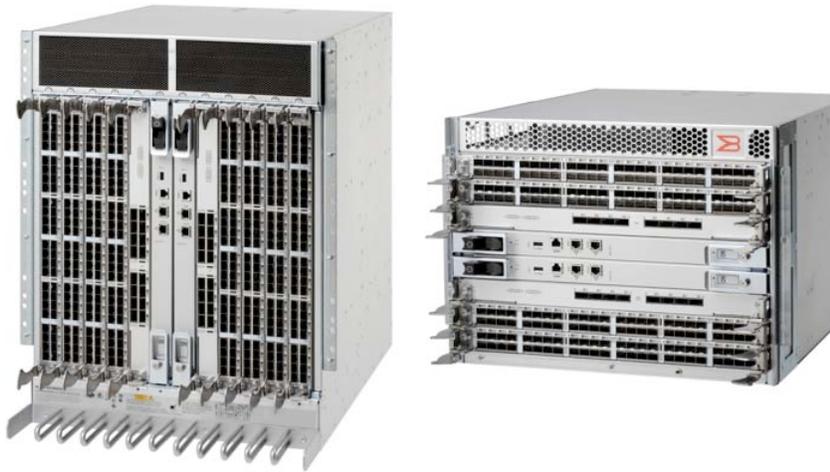
**The Brocade® DCX® 8510 Backbone is the most powerful Fibre Channel switching infrastructure in the industry. It provides the most reliable, scalable, high-performance foundation for private cloud storage and highly virtualized environments. With its intelligent 7th-generation ASICs and new hardware and software capabilities, the Brocade DCX 8510 is designed to increase business agility while providing nonstop access to information and reduction of infrastructure and administrative costs.**

**This paper explains the architectural advantages of the Brocade DCX 8510. It describes how IT organizations can leverage the performance capabilities, modular flexibility, and “five-nines” (99.999 percent) reliability of this Storage Area Network (SAN) platform to meet its most pressing business requirements with a future-proof solution that is built on proven data center technology.**

#### **OVERVIEW**

In May 2011, Brocade introduced the Brocade DCX 8510 Backbone Family (see Figure 1) and Brocade Fabric OS® (Brocade FOS) 7.0, which is the first platform in the industry to provide 16 gigabit-per-second (Gbps) Fibre Channel (FC) capabilities. This new product family builds upon the success and adoption of the 8 Gbps Brocade DCX Backbone Family by data centers across the globe.

Although this paper focuses on the Brocade DCX 8510-8 (8-slot chassis), some information is provided about the Brocade DCX 8510-4 (4-slot chassis). For more details on these two backbone platforms, see the Brocade DCX 8510 Backbone Family Data Sheet on [www.brocade.com](http://www.brocade.com).



**Figure 1.**  
Brocade DCX 8510-8 (left)  
and Brocade DCX 8510-4 (right)

Compared to competitive offerings, the Brocade DCX 8510 Backbone is the fastest and most advanced SAN backbone in the industry. It provides the following advantages:

- It scales non-disruptively to as many as 384 concurrently active 16 Gbps full-duplex ports in a single domain (IBM “FICON” and open systems).
- It scales non-disruptively to as many as 512 concurrently active 8 Gbps full-duplex ports in a single domain (open systems).
- It enables simultaneous uncongested switching on all ports as long as simple best practices are followed.
- It can provide 8.2 terabits per second (Tbps) (Brocade DCX 8510-8) or 4.1 Tbps (Brocade DCX 8510-4) utilizing 16 Gbps blades, ICLs, and local switching.

In addition to providing the highest levels of performance, the Brocade DCX 8510 Backbone features a modular, high-availability architecture that supports mission-critical environments. Moreover, the industry-leading power and cooling efficiency of the platform helps reduce ownership costs while maximizing rack density.

The Brocade DCX 8510 Backbone uses only 5.5 watts AC per port and 0.2 watts per Gbps at its maximum 16 Gbps 384-port configuration. The Brocade DCX 8510-4 uses only 6.0 watts AC per port and 0.3 watts per Gbps at its maximum 16 Gbps 192-port configuration. Both are more efficient than their predecessors and are up to 15 to 16 times more efficient than competitive products. This efficiency not only reduces data center power bills; it also reduces cooling requirements and minimizes or eliminates the need for data center infrastructure upgrades, such as new power distribution units (PDUs), power circuits, and larger heating, ventilation, and air conditioning (HVAC) units. In addition, the highly integrated architecture uses fewer active electric components boarding the chassis, which improves key reliability metrics such as mean time between failures (MTBF).

The Brocade DCX 8510 Backbone leverages a highly flexible multiprotocol architecture, supporting Fibre Channel, Fibre Connectivity (IBM “FICON”), Fibre Channel over IP (FCIP), and IP over Fibre Channel (IPFC). IT organizations can also easily mix FC port blades with advanced functionality blades for SAN encryption and SAN extension to build an infrastructure that optimizes functionality, price, and performance. Ease of setup enables data center administrators to quickly maximize its performance and availability.

This paper describes the internal architecture of the Brocade DCX 8510 Backbone and how best to leverage its industry-leading performance and blade flexibility to meet business requirements.

### **How Is Fibre Channel Bandwidth Measured?**

Fibre Channel is a lossless, low-latency, full-duplex network protocol, meaning that data can be transmitted and received simultaneously. The name of a specific Fibre Channel standard, for example “16 Gbps FC,” refers to how fast an application payload can move in one direction. This is called “data rate.” Vendors sometimes state data rates followed by the words “full duplex”—for example, 16 Gbps full duplex—although that is not necessary when referring to Fibre Channel speeds. The term “aggregate data rate” is the sum of the application payloads moving in each direction (full duplex) and is equal to twice the data rate.

## BROCADE DCX ASIC FEATURES

The Brocade DCX Backbone core routing blades (CR16-8 and CR16-4) and Fibre Channel port blades feature Brocade “Condor3” ASICs, each capable of 768 Gbps of bandwidth. Each Condor3 ASIC has 48 x 16 Gbps ports, which can be combined into trunk groups of multiple sizes. The Brocade DCX 8510 architecture leverages the same Fibre Channel protocols as the front-end ports, enabling back-end ports to avoid latency due to protocol conversion overhead.

When a frame enters the ASIC, the destination address is read from the header, which enables routing decisions to be made before the entire frame has been received. This is known within the industry as “cut-through routing” and means that a frame can begin transmission out of the correct destination port on the ASIC before the frame has finished entering the ingress port. Local latency on the same ASIC is 700 nanoseconds (ns) and blade-to-blade latency is 2.1 microsecond (µsec). Forward error correction (FEC) adds 400 nanoseconds (ns) between 16 or 10 Gbps E\_Ports, and is enabled by default. Integrated encryption or compression each adds less than 6 µsec per node. The Brocade DCX 8510 has the lowest switching latency and highest throughput combination of any Fibre Channel backbone platform in the industry.

Each Condor3 ASIC on a port blade can act as an independent switching engine to provide local switching between port groups in addition to switching across the backplane. Local

Unlike competitive offerings, frames that are switched within port groups are always capable of full port speed.

switching traffic does not cross the backplane, nor does it consume any slot bandwidth. This enables every port on high-density blades to communicate at full 16 Gbps. That is just 700 ns—at least 20 times faster than the next-fastest SAN enterprise platform on the market. On the 32-port blades, local switching is performed within 16-port groups. On 48-port blades, local switching is performed within 24-port groups. Only Brocade offers an enterprise architecture that can make these types of switching decisions at

the port level, enabling local switching and the ability to deliver up to 8.2 Tbps (Brocade DCX 8510-8) and 4.1 Tbps (Brocade DCX 8510-4) of aggregate bandwidth per backbone.

To support long-distance configurations, Condor3 ASICs include a robust set of features that include the following:

- 8,192 buffer-to-buffer credits per 16-port group on 32-port blades, and per 24-port group on 48-port blades
- integrated 10 Gbps Fibre Channel speed for seamless integration with dense wavelength-division multiplexing (DWDM) metro networking
- integrated in-flight compression for link optimization
- integrated in-flight encryption for link security

Condor3 ASICs also support Brocade frame-based Inter-Switch Link (ISL) Trunking across up to eight 16 Gbps ISLs. Brocade has significantly improved frame-level trunking through a “masterless link” in a trunk group. If an ISL trunk link ever fails, the ISL trunk seamlessly reforms with the remaining links, enabling higher overall data availability.

Dynamic Path Selection (DPS) provides exchange-level routing between individual ISLs, ICLs, ICL Trunking groups, and ISL Trunking groups. Exchange-based DPS automatically optimizes fabric-wide performance by automatically routing data to the most efficient available path in the fabric. DPS augments ICL and ISL Trunking to provide more effective load balancing in certain configurations, such as routing data between multiple trunk groups.

Preventing frame loss during an event such as the addition or removal of an ISL while the fabric is active is a critical customer requirement. Lossless Dynamic Load Sharing (DLS) and DPS enable optimal utilization of ISLs by performing traffic rebalancing operations during fabric events such as E\_Port up/down, F\_Port down, and so on. Typically, when a port goes down or comes back up, frames can be dropped, frames can arrive out of order, or a traffic imbalance can occur. Lossless DLS and DPS architecture available from Brocade rebalances traffic at the frame and exchange level, delivering in-order traffic without dropping frames, thus preventing application timeouts or Small Computer Systems Interface (SCSI) retries.

## **BROCADE DCX 8510 PLATFORM ARCHITECTURE**

In the Brocade DCX 8510, each port blade has Condor3 ASICs that expose some ports for front-end server or storage connectivity, as well as some back-end ports for core routing connectivity with ASICs via the backplane. The backbone uses a multistage ASIC layout analogous to a “fat-tree” core-edge topology. The fat-tree layout is symmetrical; that is, all ports have equal access to all other ports.

Brocade DCX 8510 platforms can switch frames locally if the destination port is on the same ASIC as the source. This is an important feature for high-density environments because it allows blades that are oversubscribed to achieve full, uncongested performance when switching on the same ASIC. No other backbone offers local switching. With competitive offerings, traffic must traverse the crossbar ASIC and backplane even when traveling to a neighboring port, which significantly degrades performance.

The flexible Brocade DCX 8510 architecture uses a wide variety of blades for increasing port density, multiprotocol capabilities, and fabric-based applications. Data center administrators can easily mix the blades in the Brocade DCX 8510 to address specific business requirements and optimize cost/performance ratios.

The following blades are currently available (as of May 2011):

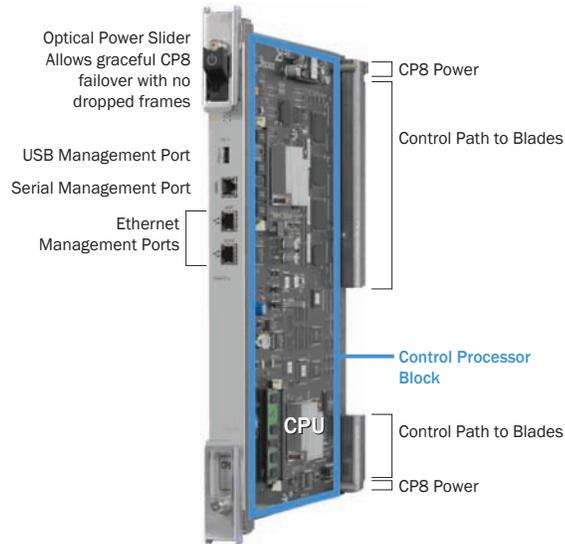
- CP8—Control Processor—manages all the port blades and the two core blades, and monitors power supplies, fans, and other components
- CR16-8—Core Routing Blade—3,074 Tbps per blade (introduced with FOS 7.0)
  - ICL ports on every CR16-8 blade turned on via an optional license
- CR16-4—Core Routing Blade—1,536 Tbps per blade (introduced with FOS 7.0)
  - ICL ports on every CR16-4 blade turned on via an optional license
- FC16-32—32 ports, 16 Gbps FC blade (introduced with FOS 7.0)
- FC16-48—48 ports, 16 Gbps FC blade (introduced with FOS 7.0)
- FC8-64—64 ports, 8 Gbps FC blade (introduced with FOS 6.4)
- FS8-18—Encryption Blade—16 ports, 8 Gbps line-speed encryption of data at rest, (introduced with FOS 6.1.1 enc.)
- FX8-24—Extension Blade—24 ports, 12 x 8 Gbps FC ports, 10 x 1 Gigabit Ethernet (GbE) ports, and two optional 10 GbE ports for long-distance extension of FCIP blade (introduced with FOS 6.3)

## **CORE BLADES AND INTER-CHASSIS LINK (ICL) PORT**

### **Control Processor Blades**

The Brocade DCX 8510 has two Control Processor (CP8) blades (see Figure 2.) that manage the overall functionality of the chassis. The two redundant control processors are highly available and run the Brocade FOS. The control processor functions are redundant active-passive (hot-standby). The blade with the active control processor is known as the “active

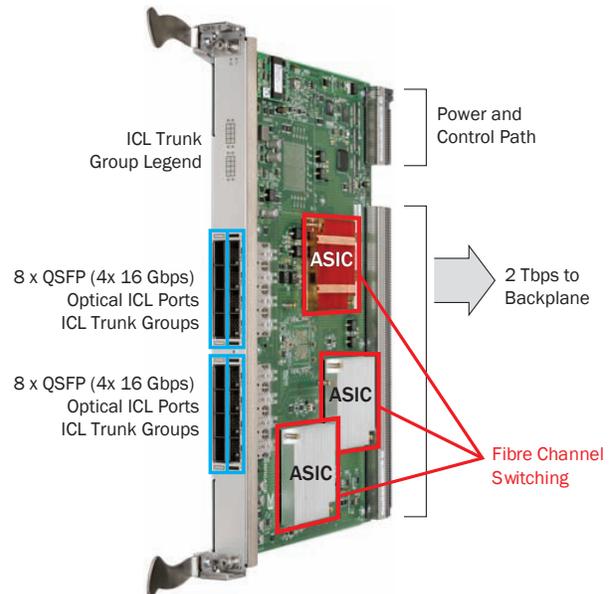
control processor blade,” but either could be active or standby. Additionally, on each processor there is a USB port and two network ports. The USB port is only for use with a USB storage device that is branded by Brocade. The dual IP ports allow a customer to potentially fail over internally on the same control processor without the loss of an IP connection, rather than fail over to the standby control processor blade.



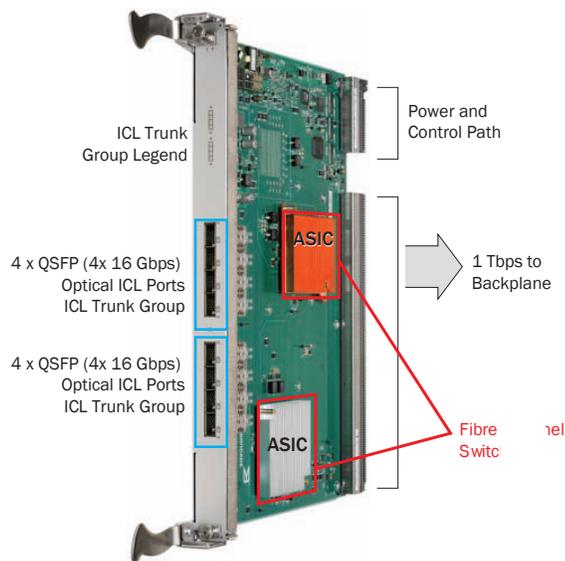
**Figure 2.**  
CP8 Core Processor Blade

### Core Routing Blades

The Brocade DCX 8510 includes two core routing blades for Brocade DCX 8510-8 (CR16-8) or Brocade 8510-4 (CR16-4) (see figures 3 and 4). These blades provide core switching and routing of the frames either from blade to blade or from the Brocade DCX 8510 chassis through the ICL ports. The CR16-8 and CR16-4 blades are active-active in each Brocade DCX 8510 chassis. The CR16-8 has four Condor3 ASICs. The CR16-4 has two Condor3 ASICs. Each ASIC has dual connection to each ASIC group on each line card.

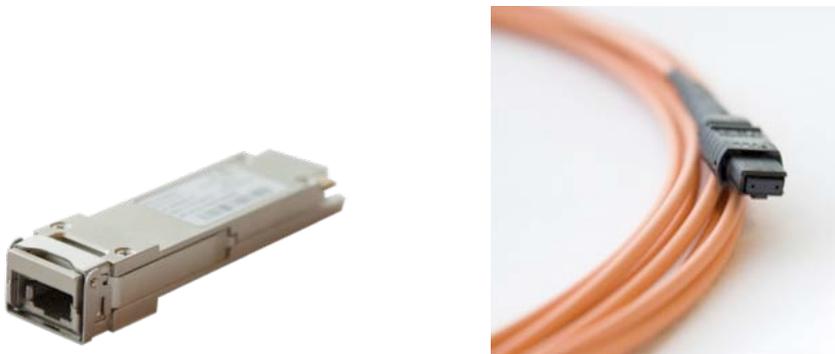


**Figure 3.**  
CR16-8 Core Routing Blade



**Figure 4.**  
CR16-4 Core Routing Blade

There are 16 x Quad Small Form-factor Pluggable (QSFP) ICL connections on each CR16-8 blade (32 QSFP ICL connections per Brocade DCX 8510-8). There are 8 x QSFP ICL connections on each CR16-4 blade (16 QSFP ICL connections per Brocade DCX 8510-4). Each ICL port provides 64 Gbps of bandwidth over a QSFP (4 x 16 Gbps) link (see figure 5). Frame-based trunk groups are automatically formed from up to four ICLs.



**Figure 5.**  
QSFP ICL Optic and Cable

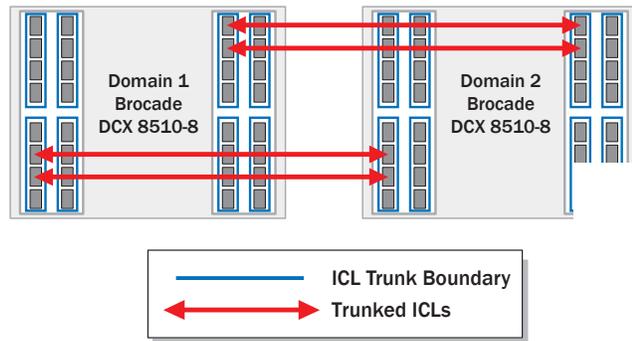
### Multi-Chassis Configuration

Networks are evolving in order to adapt to rapid growth and change in the server and storage infrastructure. New optical ICLs can connect up to six Brocade DCX 8510 Backbones (up to 2,304 x 16 Gbps ports), enabling flatter, faster, and simpler fabrics that increase consolidation while reducing network complexity and costs.

These high-density chassis topologies reduce inter-switch cabling by 75 percent and free up to 33 percent of ports for server and storage. This maximizes overall port density in the lowest amount of rack space.

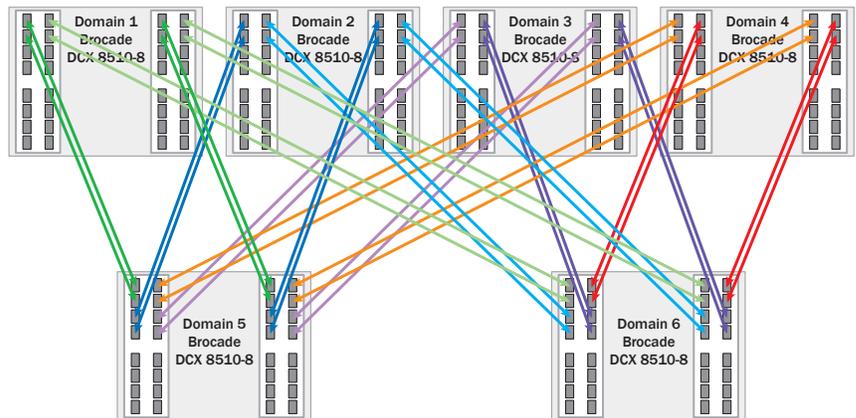
Each QSFP ICL port consists of four separate 16 Gbps links, with each terminating on different ASICs within each core routing blade. Standard QSFP optical cables are supported up to 164 feet (50 meters). Brocade brand QSFPs (4 x 16 Gbps) are required. A minimum of four ICL ports must be connected between chassis for performance and redundancy. A minimum of two ICLs ports from same trunk boundary in one core blade must be connected two ICL ports from the same trunk boundary on the core blade in another chassis (see figure 6). This will provide 256 Gbps (the equivalent to 16 x 16 Gbps ISLs).

**Figure 6.**  
ICL Trunking between two Brocade DCX 8510 Chassis



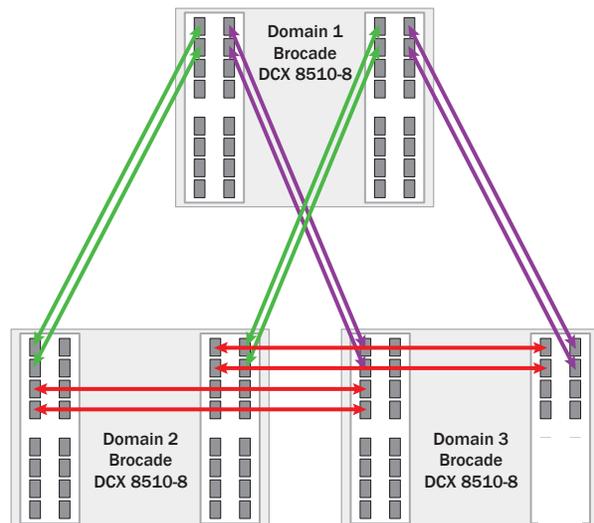
ICLs enable scalable core-edge and active-active mesh chassis topologies. Up to six chassis are supported in a core-edge configuration (see figure 7). This topology will support up to 2304 16 Gbps ports and provide up to 50 Tbps total bandwidth.

**Figure 7.**  
Six chassis core-edge topology



Up to three chassis are supported in a full-mesh configuration (see figure 8). This topology will support up to 1152 16 Gbps ports and provide up to 50 Tbps total bandwidth.

**Figure 8.**  
Three chassis full-mesh topology



For more details on ICLs, refer to Building Hyper-Scale Fabrics with Brocade DCX 8510 on [www.brocade.com](http://www.brocade.com).

## 16 AND 8 GBPS PORT BLADES

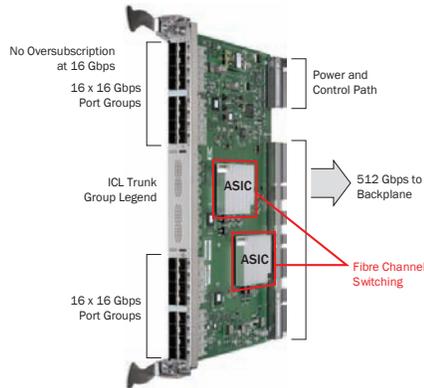
Brocade offers 32- and 48-port 16 Gbps blades as well as the 64-port 8 Gbps blade to connect to servers, storage, or switches. All of the port blades can leverage local switching to ensure full 16 or 8 Gbps performance on all ports. Each CR16-8 blade contains four ASICs that switch data over the backplane between port blade ASICs. A total of 512 Gbps of aggregate bandwidth per blade is available for switching through the backplane. Mixing switching over the backplane with local switching delivers performance of up to 768 Gbps per blade using 48-port 16 Gbps blades.

For distance of more than 10 Gbps DWDM or dark fiber using small form-factor pluggables (SFPs) that are branded by Brocade, the Condor3 ASIC has approximately four times the buffer credits as the Condor2 ASIC (8192 vs. 2048). This enables 2, 4, 8, 10, or 16 Gbps ISLs and more long-wave connections over greater distances.

When connecting a large number of devices that need sustained 16 Gbps transmission line rates, IT organizations can leverage local switching to avoid congestion. Local switching on Fibre Channel port blades reduces port-to-port latency. Frames cross the backplane in 2.1  $\mu$ sec, and locally switched frames cross the blade in only 700 ns. The latency from crossing the backplane is still more than 40 times faster than disk access times and is much faster than any competing product.

All 16 Gbps ports on the FC16-32 blade operate at full line rate through the backplane or with local switching.

Figure 9 shows a photograph and functional diagram of the FC16-32 blade.



### Switching Speed Defined

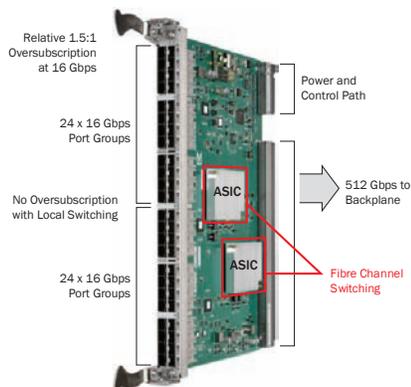
When describing SAN switching speed, vendors typically use the following measurements:

- **Milliseconds (ms):**  
One thousandth of a second
- **Microseconds ( $\mu$ sec):**  
One millionth of a second
- **Nanoseconds (ns):**  
One billionth of a second

**Figure 9.**  
FC16-32 Port Blade

The FC16-48 blade has a higher backplane oversubscription ratio at 16 Gbps but it has larger port groups to take advantage of local switching. Although the backplane connectivity of this blade is identical to the FC16-32 blade, the FC16-48 blade exposes 24 user-facing ports per ASIC rather than 16. Oversubscription occurs only when the first 32 ports are fully utilized.

Figure 10 shows a photograph and functional diagram of the FC16-48 blade.

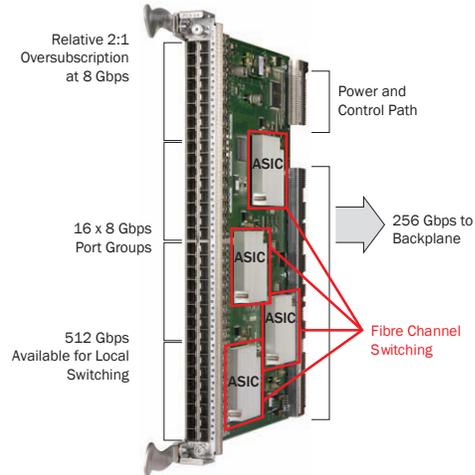


**Figure 10.**  
FC16-48 Port Blade

The FC8-64 blade has a 2:1 oversubscription ratio at 8 Gbps switching through the backplane and no oversubscription with local switching. The FC8-64 blade exposes 16 user-facing ports per ASIC, and up to eight 8-port trunk groups can be created with the 64-port blade.

Figure 11 shows a photograph and functional diagram of the FC8-64 blade.

**Figure 11.**  
FC8-64 Port Blade



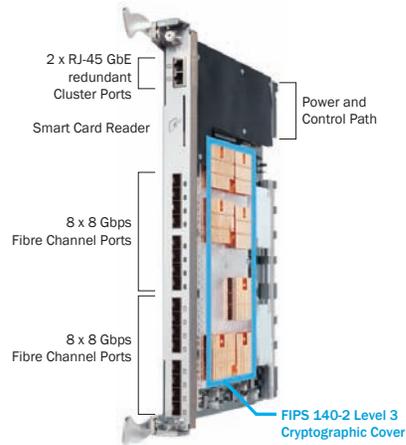
## SPECIALTY BLADES

### SAN Encryption Blade

The Brocade FS8-18 Encryption Blade provides 16 x 8 Gbps Fibre Channel ports, 2 x RJ-45 GbE ports, and a smart card reader. The Brocade FS8-18 is a high-speed, highly reliable Federal Information Processing Standard (FIPS) 140-2 Level 3 validated blade, which provides fabric-based encryption and compression services to secure data assets either selectively or on a comprehensive basis. The blade scales non-disruptively from 48 up to 96 Gbps of disk encryption processing power. It also provides encryption and compression services at speeds up to 48 Gbps for data on tape storage media. Moreover, the Brocade FS8-18 is tightly integrated with four industry-leading, enterprise-class key management systems that can scale to support key lifecycle services across distributed environments.

Figure 12 shows a photograph and functional diagram of the FS8-18.

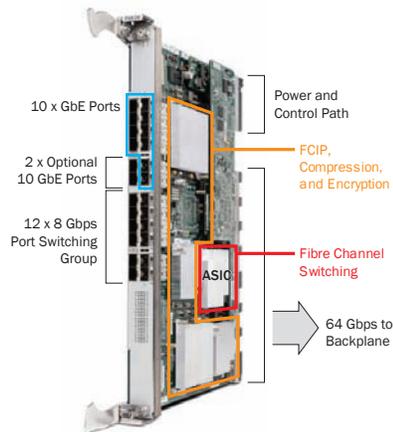
**Figure 12.**  
Brocade FS8-18 Encryption Blade



## SAN Extension Blade

Brocade FX8-24 Extension Blades accelerate and optimize replication, backup, and migration over any distance with the fastest, most reliable, and most cost-effective network infrastructure. Twelve 8 Gbps Fibre Channel ports, ten 1 GbE ports, and up to two optional 10 GbE ports provide unmatched Fibre Channel and FCIP bandwidth, port density, and throughput for maximum application performance over IP WAN links. Whether it is deployed in large data centers or multisite environments, the Brocade FX8-24 enables replication and backup applications to move more data faster and further than ever before to address the most demanding disaster recovery, compliance, and data mobility requirements.

Figure 13 shows a photograph and functional diagram of the FX8-24.



**Figure 13.**  
Brocade FX8-24 Extension Blade

## FLEXIBLE CORE-EDGE AND MESH CHASSIS TOPOLOGIES

Optical ICLs enable large-scale, highly available, high-performance core-edge chassis SANs. Ports and bandwidth can be added to the core or edge to optimize overall fabric performance. ICLs enable practical active-active mesh topologies for simple any-to-any fabric connectivity. This innovative chassis connectivity technology eliminates the cabling complexity and port consumption cost of connecting multiple chassis together through traditional ISLs. Brocade DCX 8510 Backbones that are connected by ICLs are treated as a single hop in a Fibre Channel and IBM “FICON” network.

Although IT organizations could build a network of 48-port switches with similar performance characteristics to the Brocade DCX 8510 Backbone, it would require more than a dozen 48-port switches that are connected in a “fat-tree” fashion. This network would require complex cabling, management of more than 12 discrete switching elements, support for higher power and cooling, and more SFPs to support ISLs. In contrast, the Brocade DCX 8510 delivers the same high level of performance without the associated disadvantages of a large multiswitch network, and brings high-density chassis performance to IT organizations that could previously not justify the investment or overhead costs.

The Brocade DCX 8510 Backbone architecture enables the entire backbone to be a single domain and a single hop in a Fibre Channel network.

In comparison to a multiswitch, fat-tree network, the Brocade DCX 8510 Backbone provides the following advantages:

- It is easier to deploy and manage.
- It simplifies the cable plant by eliminating ISLs and additional SFP media.
- It is far more scalable than a large network of independent domains.
- It is lower in both initial capital expenditures (CapEx) and ongoing operational expenditures (OpEx).
- It has fewer active components and more component redundancy for higher reliability.
- It provides multiprotocol support and routing within a single chassis.

### **PERFORMANCE IMPACT OF CONTROL PROCESSOR OR CORE ROUTING FAILURE**

Any type of failure on the Brocade DCX 8510—whether a control processor or a core ASIC—is extremely rare. However, in the unusual event of a failure, the Brocade DCX 8510 is designed for fast and easy control processor replacement. This section describes potential (albeit unlikely) failure scenarios and how the Brocade DCX 8510 is designed to minimize the impact on performance and provide the highest level of system availability.

#### **Control Processor Failure in a CP8 Blade**

If the CPU of the active control processor blade fails, it affects only the management plane. Data traffic between end devices continues to flow uninterrupted. With two IP management ports, the risk of having to fail over to a standby control processor blade is potentially minimized.

A control processor failure has no effect on the data plane. The standby control processor automatically takes over and the backbone continues to operate without dropping any data frames.

#### **Core Routing Failure in a CR16-8 or CR16-4 Blade**

The potential impact of a core routing blade failure to overall system performance is straightforward. If half of the core elements were to go offline due to a hardware failure, half of the aggregate switching capacity over the backplane would be offline until the condition was corrected. A Brocade DCX 8510 Backbone with just one core routing blade can still provide 2,048 Gbps aggregate slot-to-slot bandwidth, or 256 Gbps to every backbone slot.

## **SUMMARY**

With aggregate chassis bandwidth far greater than competitive offerings, Brocade DCX 8510 Backbones are architected to deliver highly virtualized and private cloud storage environments. As demonstrated by Brocade testing, the Brocade DCX 8510 offers the following features:

- It delivers 16 or 8 Gbps Fibre Channel and IBM “FICON” line-rate connectivity on all ports simultaneously.
- It provides local switching to maximize bandwidth for high-demand applications.
- It enables high-performance chassis scalability through ICLs.
- It offers port blade flexibility to meet specific connectivity, performance, and budget needs.
- It provides investment protection by supporting data security, inter-fabric routing, and SAN extension.
- It delivers “five-nines” availability.

For further details on the capabilities of the Brocade DCX 8510 Backbone Family, refer to <http://www.brocade.com/products-solutions/products/dcx-backbone/index.page>. There you will find the Brocade DCX 8510 Backbone Family Data Sheet and relevant Technical Briefs and White Papers.

**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

© 2011 Brocade Communications Systems, Inc. All Rights Reserved. 07/11 GA-WP-1583-01

Brocade, the B-wing symbol, DCX, Fabric OS, and SAN Health are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, CloudPlex, MLX, VCS, VDX, and When the Mission Is Critical, the Network Is Brocade are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, 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 informational 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**