



**STORAGE AREA  
NETWORK**

## **Achieving Enterprise SAN Performance with the Brocade 48000 Director**

A best-in-class architecture enables optimum performance, flexibility, and reliability for enterprise data center networks.

**BROCADE**

**The Brocade® 48000 Director is the industry’s highest-performing director platform for supporting enterprise-class Storage Area Network (SAN) operations. With its intelligent sixth-generation ASICs and new hardware and software capabilities, the Brocade 48000 provides a reliable foundation for fully connected multiprotocol SAN fabrics, FICON® solutions, and Meta SANs capable of supporting thousands of servers and storage devices.**

**The Brocade 48000 also provides industry-leading power and cooling efficiency, helping to reduce the Total Cost of Ownership (TCO).**

**This paper outlines the architectural advantages of the Brocade 48000 and describes how IT organizations can leverage the performance capabilities, modular flexibility, and “five-nines” (99.999 percent) reliability of this SAN director to achieve specific business requirements.**

## OVERVIEW

In May 2005, Brocade introduced the Brocade 48000 Director (see Figure 1), a third-generation SAN director and the first in the industry to provide 4 Gbit/sec (Gb) Fibre Channel (FC) capabilities. Since that time, the Brocade 48000 has become a key component in thousands of data centers around the world.

With the release of Fabric OS® (FOS) 6.0 in January 2008, the Brocade 48000 adds 8 Gbit/sec Fibre Channel and FICON performance for data-intensive storage applications.



**Figure 1.**

The Brocade 48000 Director in a 384-port configuration.

Compared to competitive offerings, the Brocade 48000 is the industry's fastest and most advanced SAN director, providing numerous advantages:

- The platform scales non-disruptively from 16 to as many as 384 concurrently active 4Gb or 8Gb full-duplex ports in a single domain.
- The product design enables simultaneous uncongested switching on all ports as long as simple best practices are followed.
- The platform can provide 1.536 Tbit/sec aggregate switching bandwidth utilizing 4Gb blades and Local Switching between two thirds or more of all ports, and 3.072 Tbit/sec utilizing 8Gb blades and Local Switching between approximately five sixths or more of all ports.

In addition to providing the highest levels of performance, the Brocade 48000 features a modular, high-availability architecture that supports mission-critical environments. Moreover, the platform's industry-leading power and cooling efficiency help reduce ownership costs while maximizing rack density.

The Brocade 48000 uses just 3.26 watts AC per port and 0.41 watts per gigabit at its maximum 8Gb 384-port configuration. This is twice as efficient as its predecessor and up to ten times more efficient than competitive products. This efficiency not only reduces data center power bills—it 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 Failure (MTBF).

## How Is Fibre Channel Bandwidth Measured?

Fibre Channel is a full-duplex network protocol, meaning that data can be transmitted and received simultaneously. The name of a specific Fibre Channel standard, for example "4 Gbit/sec 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, "4 Gbit/sec full duplex," although it is not necessary to do so 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..

The Brocade 48000 is also highly flexible, supporting Fibre Channel, Fibre Connectivity (FICON), FICON Cascading, FICON Control Unit Port (CUP), Brocade Accelerator for FICON, FCIP with IP Security (IPSec), and iSCSI. IT organizations can easily mix Fibre Channel blade options to build an architecture that has the optimal price/performance ratio to meet the requirements of specific SAN environments. And its easy setup characteristics enable data center administrators to maximize its performance and availability using a few simple guidelines.

This paper describes the internal architecture of the Brocade 48000 Director and how best to leverage the director's industry-leading performance and blade flexibility to achieve business requirements.

## **BROCADE 48000 PLATFORM ASIC FEATURES**

The Brocade 48000 Control Processors (CP4s) feature Brocade "Condor" ASICs each capable of switching at 128 Gbit/sec. Each Brocade Condor ASIC has thirty-two 4Gb ports, which can be combined into trunk groups of multiple sizes. The Brocade 48000 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 whole frame has been received. This allows the ASICs to perform cut-through routing, which means that a frame can begin transmission out of the correct destination port on the ASIC even before the frame has finished entering the ingress port. Local latency on the same ASIC is 0.8  $\mu$ s and blade-to-blade latency is 2.4  $\mu$ s. As a result, the Brocade 48000 has the lowest switching latency and highest throughput of any Fibre Channel director in the industry.

Because the FC8 port blade Condor 2 (8Gb) and the FC4 port blade Condor (4Gb) ASICs can act as independent switching engines, the Brocade 48000 can leverage localized switching within a port group in addition to switching over the backplane. On 16- and 32-port blades, Local Switching is performed within 16-port groups. On 48-port blades, Local Switching is performed within 24-port groups. Unlike competitive offerings, frames being switched within port groups do not need to traverse the backplane. This enables every port on high-density blades to communicate at full 8 Gbit/sec or 4 Gbit/sec full-duplex speed with port-to-port latency of just 800 ns—25 times faster than the next-fastest SAN director on the market. Only Brocade offers a director architecture that can make these types of switching decisions at the port level, thereby enabling Local Switching and the ability to deliver up to 3.072 Tbit/sec of aggregate bandwidth per Brocade 48000 system.

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

To support long-distance configurations, 8Gb blades have Condor 2 ASICs, which provide 2,048 buffer-to-buffer credits per 16-port group on 16- and 32-port blades, and per 24-port group on 48-port blades; 4Gb blades with Condor ASICs have 1,024 buffer-to-buffer credits per port group.

The Condor 2 and Condor ASICs also enable Brocade Inter-Switch Link (ISL) Trunking with up to 64 Gbit/sec full-duplex, frame-level trunks (up to eight 8Gb links in a trunk) and Dynamic Path

Selection (DPS) for exchange-level routing between individual ISLs or ISL Trunking groups. Up to eight trunks can be balanced to achieve a total throughput of 512 Gbit/sec. Furthermore, 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 will seamlessly reform with the remaining links, enabling higher overall data availability.

### **Switching Speed Defined**

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

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

## BROCADE 48000 PLATFORM ARCHITECTURE

In the Brocade 48000, each port blade has Condor 2 or Condor ASICs that expose some ports for user connectivity and some ports to the control processors core switching ASICs via the backplane. The director uses a multi-stage 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. The director 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 when switching between blade ASICs to achieve full uncongested performance when switching on the same ASIC. No other director offers Local Switching: with competing offerings, traffic must traverse the crossbar ASIC and backplane even when traveling to a neighboring port—a trait that significantly degrades performance.

The flexible Brocade 48000 architecture utilizes 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 48000 to address specific business requirements and optimize cost/performance ratios. The following blades are currently available (as of mid-2008).

Blade Name	Description	Introduced with
FC8-16	16 ports, 8Gb FC blade	FOS 6.0
FC8-32	32 ports, 8Gb FC blade	FOS 6.1
FC8-48	48 ports, 8Gb FC blade	FOS 6.1
FC4-16	16 ports, 4Gb FC blade	FOS 5.1
FC4-32	32 ports, 4Gb FC blade	FOS 5.1
FC4-48	48 ports, 4Gb FC blade	FOS 5.2
FR4-18i Extension Blade	FC Routing and FCIP blade with FICON support	FOS 5.2
FC4-16IP iSCSI Blade	iSCSI-to-FC gateway blade	FOS 5.2
FC10-6	6 ports, 10Gb FC blade	FOS 5.3
FA4-18 Fabric Application Blade	18 ports, 4Gb FC application blade	FOS 5.3
CP4	Control Processor with core switching: at 256 Gbit/sec per CP4 blade	FOS 5.1

### 8Gb Fibre Channel Blades

Brocade 16-, 32-, and 48-port 8Gb blades are the right choice for 8Gb ISLs to a Brocade DCX Backbone or an 8Gb switch, including the Brocade 300, 5100, and 5300 Switches. Compared with 4Gb port blades, 8Gb blades require half the number of ISL connections. Connecting storage and hosts to the same blade leverages Local Switching to ensure full 8 Gbit/sec performance. Mixing switching over the backplane with Local Switching delivers performance of between 64 Gbit/sec and 384 Gbit/sec per blade.

For distance over dark fiber using Brocade Small Form Factor Pluggables (SFPs), the Condor 2 ASIC has approximately twice the buffer credits as the Condor ASIC—enabling 1Gb, 2Gb, 4Gb, or 8Gb ISLs and more long-wave connections over greater distances.

Figure 2 shows a photograph and functional diagram of the 8Gb 16-port blade.

**Figure 2.**  
FC8-16 blade design.




Figure 3 shows how the blade positions in the Brocade 48000 are connected to each other using FC8-16 blades in a 128-port configuration. Eight FC8-16 port blades support up to 8 x 8 Gbit/sec full-duplex flows per blade over the backplane, utilizing a total of 64 ports. The remaining 64 user-facing ports on the eight FC8-16 blades can switch locally at 8 Gbit/sec full duplex.

While Local Switching on the FC8-16 blade reduces port-to-port latency (frames cross the backplane in 2.2  $\mu$ s, whereas locally switched frames cross the blade in only 700 ns), the latency from crossing the backplane is still more than 50 times faster than disk access times and is much faster than any competing product. Local latency on the same ASIC is 0.7 us (8Gb blades) and 0.8 us (4Gb blades), and blade-to-blade latency is between 2.2 and 2.4  $\mu$ s.

**Figure 3.**  
Overview of a Brocade 48000 128-port configuration using FC8-16 blades. Numbers are all data rate.


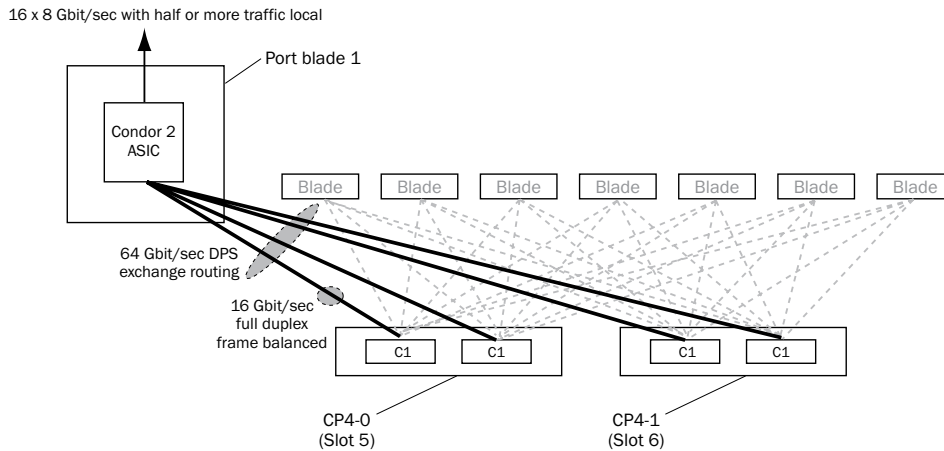


Figure 4 illustrates the internal connectivity between FC8-16 ports blades and the Control Processor blades (CP4). Each CP4 blade contains two ASICs that switch over the backplane between the ASICs. The thick line represents 16 Gbit/sec of internal links (consisting of four individual 4 Gbit/sec links) between the port blade ASIC and each ASIC on the CP4 blades. As each port blade is connected to both control processors, a total of 64 Gbit/sec of aggregate bandwidth per blade is available for internal switching.




**Figure 4.**  
FC8-16 blade internal connectivity.

### 32-port 8Gb Fibre Channel Blade

The FC8-32 blade operates at full 8 Gbit/sec speed per port for Local Switching and up to 4:1 oversubscribed for non-local switching.

Figure 5 shows a photograph and functional diagram of the FC8-32 blade.




**Figure 5.**  
FC8-32 blade design.

### 48-port 8Gb Fibre Channel Blade

The FC8-48 blade has a higher backplane oversubscription ratio but larger port groups to take advantage of Local Switching. While the backplane connectivity of this blade is identical to the FC8-32 blade, the FC8-48 blade exposes 24 user-facing ports per ASIC rather than 16.

Figure 6 shows a photograph and functional diagram of the FC8-48 blade.


**Figure 6.**  
FC8-48 blade design.



## SAN Extension Blade

The Brocade FR4-18i Extension Blade consists of sixteen 4Gb FC ports with Fibre Channel routing capability and two Gigabit Ethernet (GbE) ports for FCIP. Each FC port can provide Fibre Channel routing or conventional Fibre Channel node and ISL connectivity. Each GbE port supports up to eight FCIP tunnels. Up to two FR4-18i blades and 32 FCIP tunnels are supported in a Brocade 48000. Additionally, the Brocade FR4-18i supports full 1 Gbit/sec performance per GbE port, FastWrite, compression, IPSec encryption, tape pipelining, and Brocade Accelerator for FICON. The Local Switching groups on the Brocade FR4-18i are FC ports 0 to 7 and ports 8 to 15.

Figure 7 shows a photograph and functional diagram of this blade.




**Figure 7.**  
FR4-18i FC Routing and  
Extension blade design.

### iSCSI Blade

The Brocade FC4-16IP iSCSI blade consists of eight 4Gb Fibre Channel ports and eight iSCSI-over-Gigabit Ethernet ports. All ports switch locally within the 8-port group. The iSCSI ports act as a gateway with any other Fibre Channel ports in a Brocade 48000 chassis, enabling iSCSI hosts to access Fibre Channel storage. Because each port supports up to 64 iSCSI initiators, one blade can support up to 512 servers. Populated with four blades, a single Brocade 48000 can fan in 2048 servers. The iSCSI hosts can be mapped to any storage target in the Brocade 48000 or the fabric to which it is connected. The eight FC ports on the FC4-16IP blade can be used for regular FC connectivity.

Figure 8 shows a photograph and functional diagram of this blade.

**Figure 8.**  
FC4-16IP iSCSI  
blade design.



### 6-port 10 Gbit/sec Fibre Channel Blade

The Brocade FC10-6 blade consists of six 10Gb Fibre Channel ports that use 10 Gigabit Small Form Factor Pluggable (XFP) optical transceivers. The primary use for the FC10-6 blade is for long-distance extension over dark fiber. The ports on the FC10-6 blade operate only in E\_Port mode to create ISLs. The FC10-6 blade has buffering to drive 10Gb connectivity up to 120 km per port and exceed the capabilities of 10Gb XFPs available in short-wave, 10 km, 40 km, and 80 km long-wave versions. While potential oversubscription of a fully populated blade is small (1.125:1), Local Switching is supported in groups consisting of ports 0 to 2 and ports 3 to 5, enabling maximum port speeds ranging from 8.9 to 10 Gbit/sec full duplex.

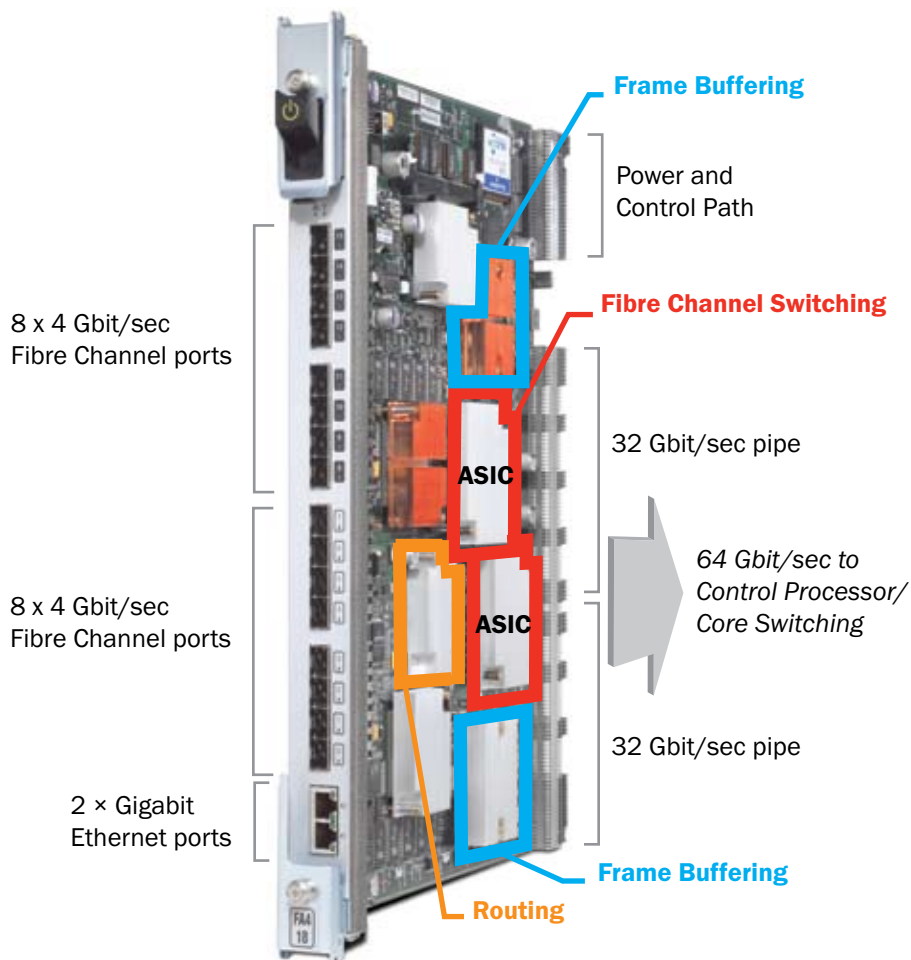
### Storage Application Blade

The Brocade FA4-18 Application Blade has sixteen 4Gb Fibre Channel ports and two auto-sensing 10/100/1000 Mbit/sec Ethernet ports for LAN-based management. It is tightly integrated with several enterprise storage applications that leverage the Brocade Storage Application Services (SAS) API—an implementation of the T11 FAIS standard—to provide wire-speed data movement and offload server resources. These fabric-based applications provide online data migration, storage virtualization, and continuous data replication and protection, and include Brocade Data Migration Manager (DMM) and other partner applications.

Figure 9 shows a photograph and functional diagram of this blade.

**Figure 9.**

FA4-18 blade design.



**16-port 4Gb Fibre Channel Blade**

On the FC4-16 blade, there are 16 user-facing ASIC ports and 16 ASIC ports facing the backplane, so the blade has a 1:1 subscription ratio. It is useful for extremely high-performance servers, supercomputing environments, high-performance shared storage subsystems, and FICON and SAN environments with unpredictable traffic patterns.

Figure 10 shows a photograph and functional diagram of the FC4-16 blade.

**Figure 10.**  
FC4-16 blade design.


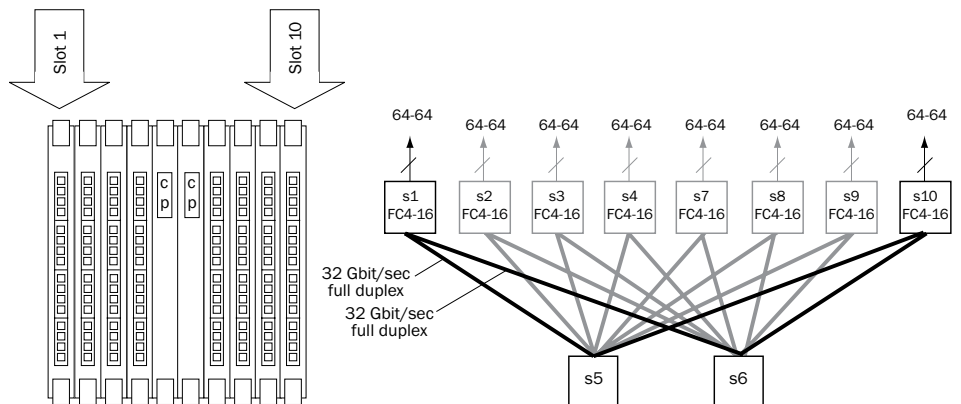


Figure 11 shows how the blade positions in the Brocade 48000 are connected to each other using FC4-16 blades in a 128-port configuration.

**Figure 11.**  
Overview of a Brocade 48000  
128-port configuration  
using FC4-16 blades.



On the left is an abstract cable-side view of the director, showing the eight slots populated with FC4-16 blades. On the right is a high-level diagram of how the slots interact with each other over the backplane.

Each thick line represents 32 Gbit/sec full duplex of internal links (8 links each at 4 Gbit/sec full duplex) connecting the port blades with the Control Processor (CP4) blades. The CP4 blades contain the ASICs that switch between the ASICs on the port blades. As every port blade is connected to both control processors, the aggregate bandwidth of these internal links is equal to the aggregate bandwidth available on external ports (64 Gbit/sec per blade x 8 blades).

With a 1:1 backplane subscription ratio, it is not necessary to use Local Switching to achieve maximum performance. While Local Switching on the FC4-16 blade reduces port-to-port transfer speed—frames cross the backplane in 2.4  $\mu$ s, whereas locally switched frames cross the blade in only 800 ns—the latency from crossing the backplane is still 50 times faster than disk access times and is much faster than any competing product.

### 32-port 4Gb Fibre Channel Blade

The FC4-32 blade operates at full 4 Gbit/sec speed per port for Local Switching and up to 2:1 oversubscribed for non-local switching. Even with no effort to utilize Local Switching, 2Gb devices and “bursty” I/O profiles can allow non-congested operation on all 32 ports simultaneously.

Figure 12 shows a photograph and functional diagram of the FC4-32 blade.

**Figure 12.**  
FC4-32 blade design.


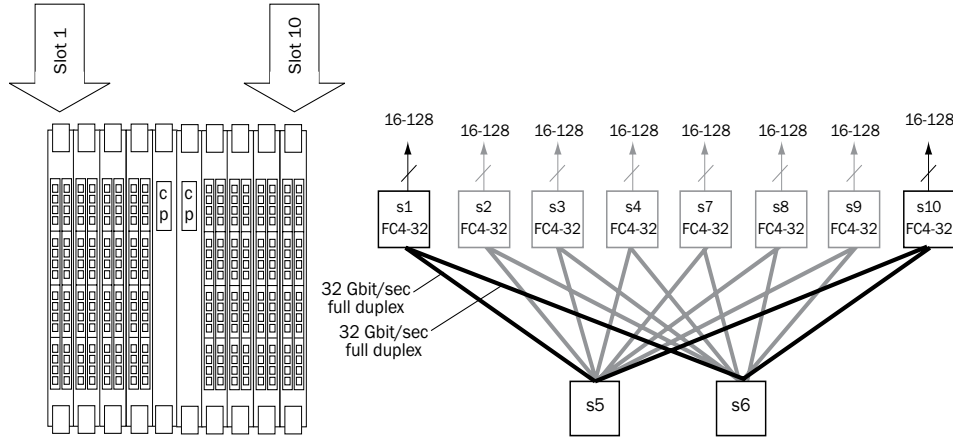


Figure 13 shows how the blade positions in a Brocade 48000 Director are connected to each other using FC4-32 blades in a 256-port configuration.

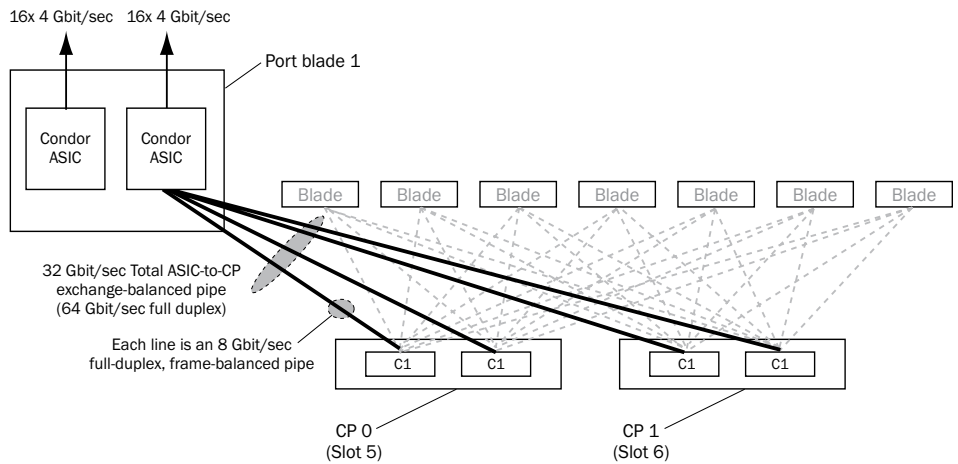
**Figure 13.**  
Overview of a Brocade 48000  
128-port configuration  
using FC4-32 blades.



When connecting a large number of devices that need sustained 4 Gbit/sec transmission line rates, IT organizations can design for Local Switching to avoid congestion. The blade is divided into two 16-port groups for Local Switching. The physically lower 16 ports (ports 0 to 7 and ports 16 to 23) form one group and the upper ports (ports 8 to 15 and ports 24 to 31) form the other group.

Figure 14 illustrates the internal connectivity between 32-port blades and the control processors.

**Figure 14.**  
FC4-32 blade  
internal connectivity.



There are two ASICs on each port blade and each ASIC has a group of 16 user-facing ports. Each thick line represents an 8 Gbit/sec internal link per group (each consisting of two 4 Gbit/sec links) between an ASIC on the port blade and the ASICs on the Control Processor (CP4) blades, providing a total of 32 Gbit/sec switching capacity across the backplane per port group (64 Gbit/sec per blade). Traffic is balanced with DPS across the four 8 Gbit/sec full-duplex links. This internal workload balancing and the resulting optimized performance represent the automatic behavior of the architecture and require no administration.

If more than 32 Gbit/sec of total throughput is needed for each 16-port group, high-priority connections can be localized within the group—ensuring that up to 16 devices or ISLs have ample bandwidth to connect to devices on other blades. These Local Switching connections do not use the backplane bandwidth. Regardless of the number of devices communicating over the backplane, locally switched devices are guaranteed the full bandwidth capacity of the port on a blade. This Brocade-unique technology for Local Switching helps preserve bandwidth to reduce the possibility of congestion in higher-density configurations.

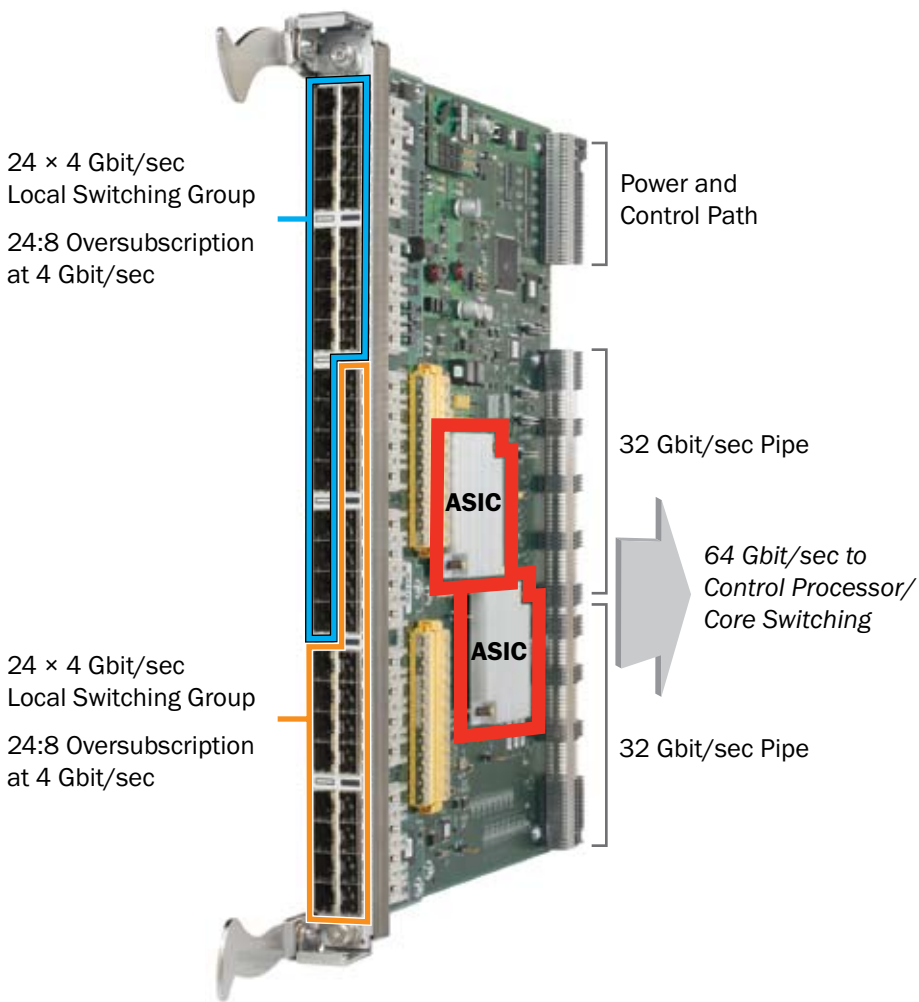
### 48-port 4Gb Fibre Channel Blade

The FC4-48 blade has a higher backplane oversubscription ratio but larger port groups to take advantage of Local Switching. While the backplane connectivity of this blade is identical to the FC4-32 blade, the FC4-48 blade exposes 24 user-facing ports per ASIC rather than 16.

This blade is especially useful for high-density SAN deployments where:

- Large numbers of servers need to be connected to the director.
- Some or all hosts are regularly running below full connection speed.
- Localization of traffic flows is easily achievable.

Figure 15 shows a photograph and functional diagram of the FC4-48 blade.



**Figure 15.**  
FC4-48 blade design.

## THE BENEFITS OF A CORE EDGE NETWORK DESIGN

The core/edge network topology has emerged as the design of choice for large-scale, highly available, high-performance SANs constructed with multiple switches of any size. The Brocade 48000 uses an internal architecture analogous to a core/edge “fat-tree” topology, which is widely recognized as being the highest-performance arrangement of switches. Note that the Brocade 48000 is not literally a fat-tree network of discrete switches, but thinking of it in this way provides a useful visualization.

While IT organizations could build a network of 40-port switches with similar performance characteristics to the Brocade 48000, it would require ten 40-port switches connected in a “fat-tree” fashion. This network would require complex cabling, management of ten discrete switching elements, support for higher power and cooling, and three times the number of SFPs to support ISLs. In contrast, the Brocade 48000 delivers the same high level of performance without the associated disadvantages of a large multi-switch network, bringing fat-tree performance to IT organizations that could previously not justify the investment or overhead costs.

It is important to understand, however, that the internal ASIC connections in a Brocade 48000 are not E\_Ports connecting a network of switches. The Fabric OS and ASIC architecture enables the entire director to be a single domain and a single hop in a Fibre Channel network. Unlike a situation in which a switch is removed from a fabric, a fabric reconfiguration is not sent across the network when a port blade is removed, further simplifying operations.

In comparison to a multi-switch, fat-tree network, the Brocade 48000:

The Brocade 48000 architecture enables the entire director to be a single domain and a single hop in a Fibre Channel network.

- Is easier to deploy and manage
- Simplifies the cable plant by eliminating ISLs and additional SFP media
- Is far more scalable than a large network of independent domains
- Is lower in both initial and operating cost
- Has fewer active components and more component redundancy for higher reliability
- Provides multiprotocol support and routing within a single chassis


## PERFORMANCE IMPACT OF CONTROL PROCESSOR FAILURE MODES

Any type of failure on the Brocade 48000—whether a control processor or core ASIC—is extremely rare. According to reliability statistics from Brocade OEM Partners, Brocade 48000 control processors have a calculated Mean Time Between Replacement (MTBR) rate of 337,000 hours (more than 38 years). However, in the event of a failure, the Brocade 48000 is designed for fast and easy control processor replacement. This section describes potential (albeit unlikely) failure scenarios and how the Brocade 48000 is designed to minimize the impact on performance and provide the highest level of system availability.

The Brocade 48000 has two control processor blades, each of which contains a processor complex and a group of ASICs that provide the core switching capacity between port groups. The control processor functions are redundant active-passive (hot-standby) while the switching functions are redundant active-active. The blade with the active control processor is known as the “active control processor blade,” but both active and standby control processor blades have active core ASIC elements. In some failure scenarios, it is also necessary for Brocade Fabric OS to automatically move routes from one control processor to another.

The CP4 ASICs and processor subsystems have separate hardware and software, with the exception of a common DC power source and printed circuit board.

Figure 16 shows a photograph and functional diagram of the control processor blade, illustrating the efficiency of the design and the separation between the ASICs and processor.



**Figure 16.**  
CP-4 blade design.

### **Control Processor Failure in a CP4 Blade**

If the processor section of the active control processor blade fails, it affects only the management plane—the core ASICs are functionally separate and continue switching frames without interruption. It is possible for a control processor block to fail completely, while the core ASICs continue to operate without switching degradation, or vice versa.

A control processor failure has no effect on the data plane: the standby control processor automatically takes over and the director continues to operate without dropping any data frames. Only during a short service procedure, during which the control processor is physically replaced, will a temporary degradation of 50 percent of available bandwidth be experienced between port card ASICs.

### **Core Element Failure in a CP4 Blade**

The potential impact of a core element failure to overall system performance is straightforward. If half of the core elements went offline due to a hardware failure, half of the aggregate switching capacity over the backplane would be offline until the condition is corrected. A Brocade 48000 with just one CP4 can still provide 256 Gbit/sec aggregate bandwidth, or 32 Gbit/sec to every director slot.

The impact of a core element failure depends on many factors. For example, the possibility of Out-of-Order Delivery (OOD) of frames depends on the fabric-wide In-Order Delivery (IOD) flag: if the flag is set, no OOD occurs. If it is not set, the application impact of OOD depends on the HBA, target, SCSI layer, file system, and application characteristics. Generally, this flag is set during installation by the OEM Partner or reseller responsible for supporting the SAN fabric and is optimized for the application environment. Most known currently shipping applications can withstand these OOD behaviors.

Data flows would not necessarily become congested in the Brocade 48000 with one core element failure. A worst case scenario would require the director to be running at or near 50 percent of bandwidth capacity on a sustained basis. With typical I/O patterns and some Local Switching, however, aggregate bandwidth demand is often below 50 percent maximum capacity. In such environments there would be no impact, even if a failure persisted for an extended period of time. For environments with higher bandwidth usage, performance degradation would last only until the failed core blade is replaced, a simple 5-minute procedure.

## **SUMMARY**

With an aggregate chassis bandwidth far greater than competitive offerings, the Brocade 48000 director is architected to deliver congestion-free performance, broad scalability, and high reliability for real-world enterprise SANs. As demonstrated by Brocade testing, the Brocade 48000:

- Delivers 8 Gbit/sec and 4 Gbit/sec Fibre Channel and FICON line-rate connectivity on all ports simultaneously
- Provides Local Switching to maximize bandwidth for high-demand applications
- Offers port blade flexibility to meet specific connectivity, performance, and budget needs,
- Provides routing, extension, and iSCSI connections using a single domain.
- Performs fabric-based data migration, protection, and storage virtualization
- Delivers "five-nines" availability

For more information about the Brocade 48000 Director, visit [www.brocade.com](http://www.brocade.com).

**Corporate Headquarters**

San Jose, CA USA  
T: (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

© 2008 Brocade Communications Systems, Inc. All Rights Reserved. 07/08 GA-WP-879-02

Brocade, Fabric OS, File Lifecycle Manager, MyView, and StorageX are registered trademarks and the Brocade B-wing symbol, DCX, and SAN Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services 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**