

**STORAGE AREA
NETWORK**

Designing Next-Generation SANs with Brocade 4 Gbit/sec Fibre Channel Solutions

Brocade 4 Gbit/sec Fibre Channel solutions provide a strategic migration path with proven investment protection.

BROCADE

The Fibre Channel Industry Association (FCIA) has officially supported the 4 Gbit/sec Fibre Channel initiative and major switch, router, Host Bus Adapter (HBA), and storage vendors are now offering 4 Gbit/sec products in all key market segments. Brocade® was the first vendor to release 4 Gbit/sec Storage Area Network (SAN) solutions in 2004, and now provides the industry's most comprehensive 4 Gbit/sec product family:

- **Brocade 200E, Brocade 4100, Brocade 4900 Switches, and Brocade 7500 SAN Router**
- **Brocade 48000 Director**
- **Brocade FR4-18i Director Blade**
- **Embedded 4 Gbit/sec Fibre Channel SAN switches for blade servers**

4 GBIT/SEC: THE INDUSTRY'S STRATEGIC DIRECTION

As 4 Gbit/sec Fibre Channel solutions begin making their way to market, today's IT organizations will need to determine how they might benefit from deploying this new technology. In addition to the significant increase in performance, there are both technical and economical reasons for deploying new 4 Gbit/sec Fibre Channel devices.

In many IT environments, higher speeds are already needed to keep up with servers and storage devices that have been saturating 2 Gbit/sec interfaces. Newer 4 Gbit/sec technology can satisfy the needs of the highest-performing applications and growing SAN infrastructures. For instance, 4 Gbit/sec devices can reduce the amount of time required for data backups and help organizations avoid purchasing costly additional servers or storage devices. Even if server data rates are not running at the full 2 Gbit/sec, for example, a 4 Gbit/sec storage device can provide the ability to fan out to a greater number of servers. 4 Gbit/sec technology can also help organizations keep up with the increased amount of data that must be managed and stored in order to comply with new business regulations.

Moreover, 4 Gbit/sec technology provides a high Return On Investment (ROI) from a business perspective. Utilizing 4 Gbit/sec Inter-Switch Links (ISLs) can reduce the overall number of ISLs needed to support the same amount of data traffic. Accordingly, the SAN ports no longer needed for ISLs can be used to support the connectivity of more servers and storage devices, improving the overall value of technology investments.

Because 4 Gbit/sec technology utilizes the same standards as 1 and 2 Gbit/sec Fibre Channel technology, vendors are able to market new 4 Gbit/sec products at approximately the same price as comparable 2 Gbit/sec products. This means that organizations can continue to utilize existing 1 and 2 Gbit/sec devices and management tools that are fully capable of solving many business requirements today and for years to come.

BROCADE 4 GBIT/SEC ARCHITECTURE ENHANCEMENTS

As the industry's most comprehensive 4 Gbit/sec product family, Brocade products include several new features and capabilities that influence how organizations design their SAN environments. With all ports auto-sensing for 1, 2, and 4 Gbit/sec operations, Brocade 4 Gbit/sec products were designed to integrate seamlessly with existing 1 and 2 Gbit/sec systems—enabling deployment in existing SAN infrastructures.

Modular products such as the 4 Gbit/sec Brocade 48000 Director support flexible scaling through additional port blades such as the Brocade FR4-18i. This scaling advantage has also been extended to the non-modular Brocade 4 Gbit/sec switches (the Brocade 200E, Brocade 4100, and Brocade 4900) through the “Ports On Demand” feature. This software license key enables non-disruptive scaling from 32 to 48 or 64 ports in the Brocade 4900, from 16 to 24 or 32 ports in the Brocade 4100, and from 8 to 12 or 16 ports in the Brocade 200E.

To ensure higher performance, the Brocade 4 Gbit/sec ASICs provide non-blocking Fibre Channel operations—meaning that all ports can run at 4 Gbit/sec full duplex in any configuration. Another key performance enhancement is the advanced buffer-to-buffer credits shared among all ports in the switch or on a blade—supporting SAN configurations up to 500 kilometers. Additional advancements include enhanced ISL Trunking capabilities, such as 32 Gbit/sec frame-level trunks (with up to eight 4 Gbit/sec links in a single trunk) for the Brocade 48000, Brocade 4900, and Brocade 4100, and 16 Gbit/sec frame-level trunks on the Brocade 200E (up to four 4 Gbit/sec links). In addition, Dynamic Path Selection (DPS) provides exchange-based load balancing between trunk groups in all Brocade 4 Gbit/sec products.

The following sections describe how organizations can utilize these capabilities as they design new SAN environments or expand existing SANs with 4 Gbit/sec devices.

DESIGNING SAN ENVIRONMENTS WITH 4 GBIT/SEC DEVICES

When designing SANs with 4 Gbit/sec devices, organizations should consider a few key items. One of the first design considerations might be where and how to deploy 4 Gbit/sec devices in an existing SAN environment. One of the most common uses would be to utilize 4 Gbit/sec technology for ISLs connecting switches in an extremely high-performance core fabric. The higher per-port bandwidth would reduce the number of ports used as ISLs to achieve the same performance level, thereby freeing up ports for servers and storage devices.

Many organizations have already deployed 4 Gbit/sec solutions in new SAN configurations, as either the cores and/or edges of smaller core-to-edge networks. When deploying Brocade 4 Gbit/sec products in an existing 2 Gbit/sec SAN, organizations would likely do so at the edge of a core-to-edge fabric. This means that switches would be connected to the existing core switches that have 2 Gbit/sec interfaces. In this case, organizations would follow their existing ISL over-subscription ratio, because the ISLs would run at 2 Gbit/sec and all the other devices in the fabric would be 2 Gbit/sec. To help ensure investment protection, Brocade 4 Gbit/sec products are fully backward compatible with Brocade 2 Gbit/sec ISL Trunking.

As soon as organizations begin deploying 4 Gbit/sec devices, they might consider adding ISLs or upgrading the core switches to 4 Gbit/sec (using tools such as Brocade Advanced Performance Monitoring or the Brocade SAN Health utility to collect ISL usage statistics and determine whether this step is necessary).

In addition to utilizing the native capabilities of the 4 Gbit/sec Brocade products, organizations can non-disruptively install Brocade 4 Gbit/sec blades into existing Brocade 24000 Directors. By adding 4 Gbit/sec blades to a Brocade 24000, any 2 Gbit/sec directors at the edge of the fabric can leverage 4 Gbit/sec capabilities for high-performance devices connected to the blade, ISL Trunking to 4 Gbit/sec switches/directors, and extended distance connectivity. Likewise, Brocade 24000 Directors at the core of large SANs can take new 4 Gbit/sec blades in available slots for the same beneficial reasons. Or, in some cases, the entire core system might be migrated to the edge of the fabric.

With 4 Gbit/sec fabrics, there are additional design points to consider:

- 1. Long-Distance Trunking:** If a SAN is extended over dark fiber, CWDM, DWDM, or SONET, the 4 Gbit/sec Brocade products provide new capabilities that can significantly enhance performance over longer distances. In fact, because these products support up to 255 buffer credits on a single port, they can enable full-speed 1 Gbit/sec operations at approximately 500 kilometers, 2 Gbit/sec operations at approximately 250 kilometers, and 4 Gbit/sec operations at approximately 125 kilometers. Moreover, trunked links can provide superior throughput: For example, 12 Gbit/sec with three 4 Gbit/sec links at up to 100 kilometers and 6 Gbit/sec with three 2 Gbit/sec links at up to 250 kilometers.
- 2. Link Balancing:** If Brocade 4 Gbit/sec products are part of an exclusively 4 Gbit/sec network or if a certain part of the network has only 4 Gbit/sec ISLs or trunks, organizations should consider reducing the total amount of ISLs by half compared to what they would typically use in 2 Gbit/sec fabrics. (As always, a minimum of two ISLs between any edge switch and the core would be required to ensure high availability.) If the fabric has mixed speeds (both 2 and 4 Gbit/sec), the number of ISLs would likely fall somewhere in between, depending on traffic patterns. Note that Brocade 4 Gbit/sec products support DPS between trunks. In many cases, this feature can improve performance as much as 4 Gbit/sec interfaces do, and DPS can operate even when Brocade 4 Gbit/sec products are connected to 2 Gbit/sec switches.

LONG-DISTANCE TRUNKING CONSIDERATIONS

To increase trunking distance for extended SAN environments, the Brocade 4100, Brocade 4900, Brocade 7500, and 4 Gbit/sec director blades such as the Brocade FR4-18i have a total of 1024 buffer credits shared among as many as 16 ports. Twenty-four of these buffer credits are used for the “back-end” embedded ports, and the rest are available for user consumption.

F_ and FL_Ports receive eight buffer credits by default, and local E_Ports (LO mode) receive 26 buffer credits (the same amount of credits as 2 Gbit/sec ports). A minimum of eight buffer credits is reserved for each port to make sure that no ports are starved of credits. The remaining credits are available in a buffer pool, which can be configured for use by any of the 16 ports.

In contrast to previous 2 Gbit/sec Brocade switches, buffer credits are no longer automatically assigned from the pool, because line speed can easily be achieved with eight buffer credits for a local device. Organizations could assign more buffer credits to a specific port with the “portCfgLongDistance” command as long as there is a valid Brocade Extended Fabrics software license.

As a result, a single port on the Brocade 4100, Brocade 4900, Brocade 7500, and 4 Gbit/sec director blades such as the Brocade FR4-18i could have up to 255 buffer credits, providing the enhanced distance capabilities shown in Table 1. To achieve these longer distances, both ends of the long-distance link would need to have 4 Gbit/sec capabilities.

Ports Trunked	4 Gbit/sec		2 Gbit/sec		1 Gbit/sec	
	Throughput	Distance	Throughput	Distance	Throughput	Distance
8	32 Gbit/sec	30 km	16 Gbit/sec	60 km	NA	NA
4	16 Gbit/sec	60 km	8 Gbit/sec	125 km	NA	NA
2	8 Gbit/sec	125 km	4 Gbit/sec	175 km	NA	NA
1	4 Gbit/sec	125 km	2 Gbit/sec	250 km	1 Gbit/sec	500 km

Table 1.
Extended trunking distances and data transfer speeds.

LINK BALANCING CONSIDERATIONS

Even in over-provisioned networks, organizations might detect “hot spots” of congestion, with some data paths running at their limit while others remain unused. In other words, the network might have a performance bottleneck even if it has sufficient bandwidth to deliver all data flows without constraint. Brocade has three options to balance data traffic while increasing both availability and performance:

- Source-port route balancing via Fabric Shortest Path First (FSPF), also known as Dynamic Load Sharing (DLS)
- Frame-level trunking between ASICs, also known as Advanced ISL Trunking
- Exchange-level trunking, also known as Dynamic Path Selection (DPS)

Unique to Brocade 4 Gbit/sec products, exchange-based DPS can optimize fabric-wide performance and load balancing by automatically routing data to the most efficient available path in the fabric. It augments ISL Trunking to provide more effective load balancing in certain configurations, such as routing data between multiple trunk groups. As a result, a combination of DPS and ISL Trunking provides the greatest design flexibility and the highest degree of load balancing.

How DPS Works

DPS works by striping Fibre Channel exchanges across equal-cost paths (the sender places an exchange ID into every Fibre Channel frame header). In normal operation, the exchange ID remains consistent for the duration of a SCSI operation. When a DPS-enabled platform receives a frame, it takes all equal-cost routes and calculates the egress port from that set based on a formula using the Sender PID (SID), Destination PID (DID), and Exchange ID (OXID). The formula always selects the same path for a given SID-DID-OXID set.

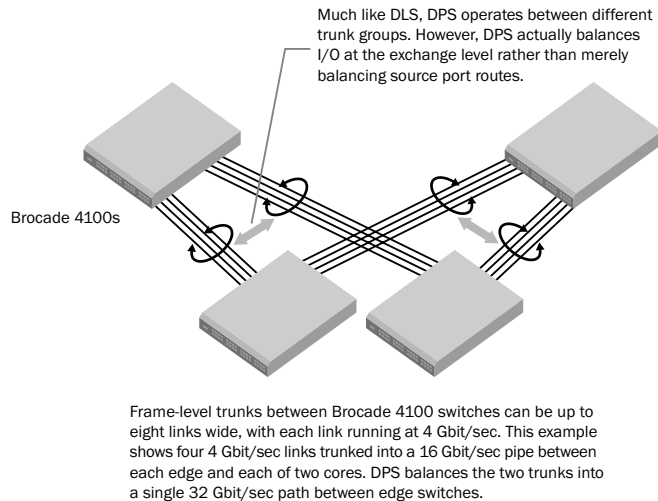
For most Fibre Channel devices, DPS stripes I/O at the SCSI level. For a given “conversation” between a server and storage port, one SCSI command would travel the first path, and the next command would travel a different path. All frames within a given exchange would be delivered in-order by virtue of traveling the same network path. Although the potential exists for out-of-order delivery between different SCSI operations, all devices tested to date are capable of handling this gracefully. Consider the following: If two servers are writing to two different storage ports across the same network, in-order delivery between the different servers is not important. It only matters that in-order delivery occurs within the data stream sent by each server, not between the two different and unrelated data streams.

The Technical Advantages of DPS

Based on exchange-level trunking, DPS matches or outperforms all similar features from any vendor *except* for Brocade frame-level trunking. However, because DPS can be combined with frame-level trunking, organizations can achieve both maximum performance and availability as well as a variety of other benefits.

Firstly, DPS does not need to occur within ASIC port groups the way frame-level trunking does. This enables load balancing across different core switches in a core-to-edge network, or different blades in a director, rather than mere DLS-based route balancing (see Figure 1).

Figure 1.
Frame-level trunking plus DPS.



Organizations could balance several groups of ports by using frame-level trunking, and then balance the resulting trunk groups by using exchange-based DPS. This approach provides the optimal balance of performance (frame-level trunking is faster) and availability (exchange-level trunking balances high-availability network topologies).

In addition, DPS can balance I/O sent from an enabled platform to any other platform even if the destination does not support the DPS feature. The transmitting switch selects the path, and the receiver does not need to do anything to ensure in-order delivery. This approach enables full backward compatibility with existing switches as well as performance benefits even if not all the switches in a fabric are using the latest technology (see Figure 2).

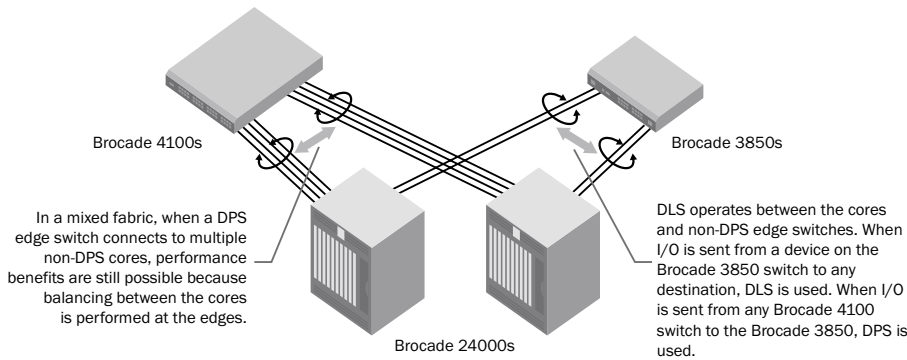


Figure 2.
DPS in mixed-speed fabrics.

DPS can also balance I/O across long-distance configurations not supported by frame-level trunking. For example, if two SAN sites have two links that take substantially different paths, there might be too much skew to form a frame-level trunk. In contrast, DPS would still be able to balance these links since it does not rely on de-skew timers for in-order delivery (see Figure 3).

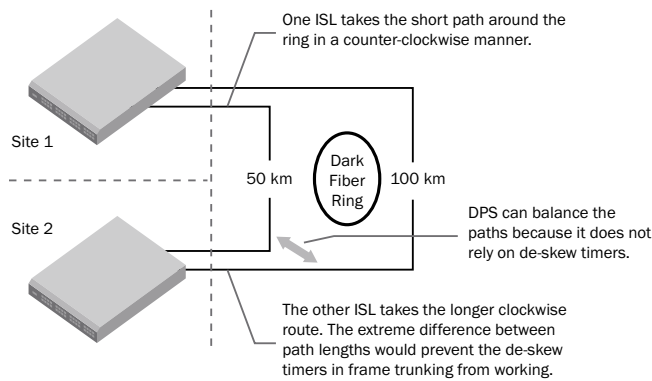


Figure 3.
DPS in large fiber ring configurations.

When considering how to utilize these new options, organizations should note that any SAN with more than one ISL can benefit from link balancing to ensure the highest level of performance. The most effective designs combine frame-level trunking for the best performance, and DLS or DPS to balance multiple trunk groups for the highest level of availability.

4 GBIT/SEC FIBRE CHANNEL ROUTING

The Brocade 7500 and the Brocade FR4-18i director blade combine the industry's first 4 Gbit/sec Fibre Channel routing capability with powerful hardware-assisted traffic forwarding for Fibre Channel over IP (FCIP). As a result, organizations can utilize these new solutions to interconnect their SAN islands for greater resource utilization and long-distance extension.

By providing this unique level of connectivity without the associated risk and complexity of physically merging SAN islands into a single large fabric, the Brocade 7500 and Brocade FR4-18i support strategic business initiatives such as disaster recovery, data migration, and ongoing technology upgrades. Although the SANs islands are physically connected, organizations can still control which devices are shared to ensure the appropriate level of SAN fabric isolation.

By providing such a highly scalable approach for extending SAN infrastructures, the Brocade 7500 and Brocade FR4-18i support key business objectives such as:

- Migrating from old to new SANs
- Consolidating data centers and rebalancing storage resources
- Migrating from test to production networks
- Moving equipment on and off lease

FOR MORE INFORMATION

With the introduction of the comprehensive 4 Gbit/sec Brocade family of products, IT organizations can begin implementing higher-performance SANs featuring next-generation technology. In addition to providing the performance required for the most data-intensive applications, these 4 Gbit/sec devices can help reduce deployment costs since SAN designs would require fewer ISLs (and less management overhead) to achieve the same level of performance. Moreover, backward compatibility with existing 1 and 2 Gbit/sec SAN environments can help protect existing technology investments and provide excellent ROI for years to come.

To learn more about Brocade 4 Gbit/sec SAN infrastructure solutions, visit www.brocade.com.

For additional Brocade technical information and resources, visit Brocade Connect, the Web portal designed exclusively for Brocade customers at www.brocade.com/jointheclub.

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

© 2007 Brocade Communications Systems, Inc. All Rights Reserved. 01/07 GA-WP-739-03

Brocade, the Brocade B-weave logo, Fabric OS, File Lifecycle Manager, MyView, Secure Fabric OS, SilkWorm, and StorageX are registered trademarks and the Brocade B-wing symbol and Tapestry are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. FICON is a registered trademark of IBM Corporation in the U.S. and 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