WHITE PAPER

BUILDING ENTERPRISE SANS THROUGH INTELLIGENT NETWORKING

Networked Brocade Fibre Channel fabrics provide the flexibility to build highly manageable and cost-effective enterprise SANs

BROCADE
The intelligent platform for networking storage

TABLE OF CONTENTS

Executive Overview **2**

Intelligent Networking in the Fabric **3**

Key Requirements for an Enterprise SAN **4**
Superior Scalability 4
Higher Availability 6
Improved Performance 6
Simplified Manageability 8
Increased Security 8
Long-Term Investment Protection 9

Design Criteria: Network Intelligence **10**

Overall Network Design **11**
Basic Dual-Switch SAN Design 12
Departmental SAN Design 13
Two-Tier Layered Design 13
Backbone Design 14
Redundant Star Design 15
Fan-Out Design 16

Real-World SAN Implementations **16**
Banking 16
Telecommunications 17
Software Development 18
Storage Services Provider 19

A Better Way to Deploy Enterprise SANS **20**

As organizations increasingly realize the advantages of their Storage Area Network (SAN) infrastructures, they are looking for new ways to extend those benefits throughout their enterprises. To broaden the scope of their SAN deployments, these organizations must identify the full range of their business requirements and consider potential design configurations that provide the necessary levels of scalability, performance, and availability at the enterprise level. One of the key factors in implementing this type of SAN is network intelligence.

Only a networked SAN fabric design can provide the high level of intelligence and scalability that an enterprise SAN deployment requires. Single director switch designs do not provide the necessary configuration flexibility for enterprise SAN fabrics. Today, organizations around the world are turning to Brocade® Fibre Channel switch-based SAN fabrics to provide a flexible and reliable infrastructure for their enterprise SANs. As a result, they are realizing the competitive advantages a networked SAN infrastructure can offer.

EXECUTIVE OVERVIEW

As Storage Area Networks (SANs) grow in size and scope across the enterprise, they are providing a wide range of benefits, such as increased efficiency, reduced costs, greater availability, and long-term investment protection. However, some SAN fabric configurations provide greater advantages than others.

A networked SAN approach addresses business requirements much better than other design alternatives, such as single-switch director models where all hosts and storage are connected to one or sometimes multiple directors. Although this ESCON-attached director architecture has functioned well in mainframe environments, it is simply not designed for today's open systems UNIX and Windows NT environments that require exceptional scalability, flexibility, and interoperability. After all, high-port-count switches do not necessarily provide a more robust enterprise SAN solution simply because of their size.

To be truly effective in enterprise SAN deployments, today's high-port-count switches—such as Brocade 64- and 128-port Core Fabric Switches—must combine greater port density with the ability to network efficiently with other switches. The key distinction between Brocade high-port-count switches and other director models is that Brocade switches are designed to form the core of the network, primarily connecting to other switches—not to host servers and storage devices. For these core fabric switches to be effective in an enterprise SAN, they must provide advanced networking services, such as Inter-Switch Link (ISL) Trunking. ISL Trunking enables highly concentrated bandwidths of up to 8 Gbit/sec, which is essential for a high-performance core or concentration point in a network.

Today, a wide variety of organizations are taking advantage of this innovative network intelligence to build some of the world's largest enterprise SANs that satisfy the most demanding requirements for performance, availability, scalability, and investment protection. This paper describes how an intelligent, networked Brocade SAN fabric can help ensure a flexible, scalable, and highly available storage environment. Included is a discussion of the primary requirements of building and deploying an enterprise fabric. Also included are design guidelines, possible configuration models, and real-world examples of Brocade-based enterprise SAN fabrics in operation today.

INTELLIGENT NETWORKING IN THE FABRIC

When it comes to building enterprise SANs, a networked SAN fabric provides the flexibility that enables organizations to easily deploy and scale their SAN environment. Using the intelligence built directly into the network, organizations can quickly react to changing business requirements and introduce new technologies as they become available. In this way they can continue to protect their existing SAN investment while adding new components to scale a relatively small SAN environment to the enterprise level. At the same time, these organizations can easily design an enterprise SAN from the outset if that is what their requirements dictate.

To facilitate the deployment of these types of enterprise SAN fabrics, Brocade has focused on developing the most intelligent Fibre Channel fabric in the industry. This intelligence includes industry-leading Application Specific Integrated Circuit (ASIC) chip technology and the only distributed fabric operating system (Brocade Fabric OS), which features publicly available APIs. A real-time Fabric OS simplifies the development of value-added software and applications to address a variety of business and technical requirements. The Fabric Access API enables Brocade software partners to fully integrate their high-level applications—such as SAN management—with a Brocade fabric.

KEY REQUIREMENTS FOR AN ENTERPRISE SAN

A SAN fabric is essentially a well-defined network of highly intelligent switches that provides enterprise-class scalability, availability, performance, manageability, security, and investment protection. At a more detailed level, these advantages include:

- Scalability through networking in an incremental, non-disruptive manner
- Availability through a resilient network, or continued access to data no matter what types of outages occur
- High performance and dynamic tuning capabilities
- Manageability through automatic monitoring, integration with higher level tools, and proactive management
- Highly scalable, industry-leading security
- Investment protection with the ability to support legacy equipment, upgrade existing switches with a simple license key, and dynamically add switches to existing fabrics

Superior Scalability

Scalability is more than just determining how large an enterprise can grow or how many hosts and storage devices can be supported. It isn't important how many hosts and storage devices reside on a single switch, but rather how many hosts and storage devices can be supported in a fabric. Current Brocade fabrics already support hundreds of ports, and many organizations are beginning to implement fabrics of 1,000 ports and larger.

Although support for very large fabrics is important, the ability to incrementally grow into larger SANs as part of a "pay-as-you-grow" strategy is even more critical. That's because most organizations do not implement a new technology—such as a SAN—to completely replace their IT infrastructures. Instead, they tend to select a defined business problem to address. After they are content with the results, organizations then expand the new technology to solve other business issues.

For example, an organization might want to share resources in its UNIX environment. As a result, the organization might consolidate UNIX-based servers and storage, perhaps over a few floors in a particular building, by deploying a 4-switch Brocade SAN. The organization might then decide to address another business concern—its growing Windows NT server environment.

The organization might consolidate servers and storage for Windows NT in the same building by adding four more switches to the SAN to create an 8-switch SAN.

Next, the organization might choose to take advantage of its new Fibre Channel infrastructure to offload data backup traffic from its overworked production TCP/IP-based network. The organization could implement a LAN-free backup solution to centralize backup for both UNIX and Windows NT environments. The Fibre Channel network enables the organization to connect over longer distances to perform the centralized LAN-free backup by using tape resources in a different building. To accomplish this, the organization simply needs to add two more Brocade switches, for example, to the existing 8-switch fabric (see Figure 1).

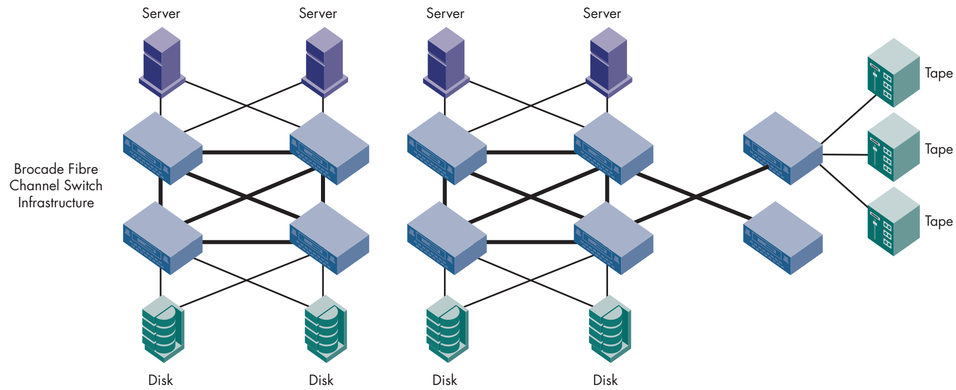

Figure 1.
Pay-as-you grow scalability for a SAN fabric

The key advantage of this incremental deployment is that one environment is not necessarily disrupted when new switches are added. This flexible approach to expansion enables an organization to grow into a larger SAN fabric at a pace based on specific business requirements. In addition to reducing the overall cost of deployment, this strategy supports incremental growth whenever an organization wants to deploy a new technology or architecture.

Simplified scalability also helps organizations that need to grow without impacting current business operations. Brocade switches can be non-disruptively added to an existing fabric without necessarily causing downtime for existing servers and storage devices. Redesign is not necessarily required since existing switches automatically recognize the new switches and even pass along certain fabric-wide information, such as addressing and zoning information. The ability to incrementally add switches—instead of completely replacing existing switches with a single bigger switch—avoids the necessity of a forklift upgrade, which is inherently disruptive to business operations.

Another scalability advantage of a networked fabric is longer connectivity distances. With single-switch Fibre Channel SANs, organizations usually deploy a 10 km maximum distance between host servers and storage. With standard Brocade Inter-Switch Link (ISL) connections, each switch can reside up to 10 km away from another switch. This enables organizations to separate host servers and storage by up to 70 km (Brocade supports up to seven hops between host servers and storage).

Although this type of network is technically feasible, there are much better alternatives to increase long-distance options, such as Brocade Extended Fabrics software used in conjunction with Dense Wave Division Multiplexing (DWDM) or extended Gigabit Interface Converters (GBICs). However, some switch cascading (networking) to achieve longer distances might be the best solution in certain situations. The key considerations are whether organizations have facilities to place each switch 10 km apart and whether the Fibre Channel cables can be laid across that distance. For some situations this might be possible on a limited basis. However, with a single, monolithic switch approach, it is not even a possibility: organizations are limited to approximately 10 km maximum distance between host servers and storage.

Higher Availability

High availability is all about designing a network to avoid single points of failure, such that applications can access data no matter what type of failure occurs in the fabric. As a result, high availability is primarily achieved through network resilience and appropriate design methods.

Although many types of failures can affect data access, true high availability means protecting data against all types of failures. For example, organizations need to guard against hardware, software, and localized physical failures/disasters, as well as the most common type of failure: user error. To protect against user error, organizations need to design a SAN that requires as little human interaction as possible. The SAN should be auto-configuring and auto-reconfiguring. For example, the SAN should automatically recognize when a new switch is added to the fabric and then educate the new switch about the rest of the SAN characteristics (such as zoning information). Brocade networks are auto-configuring and auto-reconfiguring, learning the network topology non-disruptively. Automating as much as possible—such as automatically recognizing a new switch in a fabric or automatically sensing the type of connection to a port (universal ports)—can help minimize the possibility of a user error impacting data access.

High-availability SANs can also heal themselves without user intervention, or even host server intervention. If a route between a host server and its storage is unavailable, Brocade switches simply reroute data to the next fastest route. The unavailability of the original route does not affect application access to data. In fact, with many types of outages Brocade switches automatically perform rerouting without the host servers ever knowing. In contrast, monolithic single-switch SANs represent a potential single point of failure—which violates the basic rule of a high-availability environment.

Most organizations are currently striving to achieve at least 99.999 percent (the five "nines") availability in their computing systems—a figure equivalent to less than 5.3 minutes of downtime a year. Additional downtime can severely impact business operations and cost valuable time, money, and resources. To help these organizations avoid or withstand a variety of failures, resilient Brocade SANs incorporate a wide range of capabilities, including:

- Highly available components with built-in redundancy and hot-plugging capabilities
- No single points of failure
- Intelligent routing and rerouting
- Dynamic failover protection
- Non-disruptive server and storage maintenance
- Hardware zoning for creating safe and secure environments
- Predictive fabric management
- High-availability networked designs

Improved Performance

High performance (throughput) is a key requirement for enterprise SANs: Any organization installing a network between its servers and storage needs assurance that application performance can reach very high rates. The most basic requirement of this type of SAN is the ability to select a high-throughput route to access data. To simplify this process, Brocade switches use extremely intelligent routing algorithms based on both hop count as well as cost analysis. The Brocade routing algorithm—Fabric Shortest Path First (FSPF)—has been approved by the ANSI T11 committee as the industry's new interoperability routing standard. In fact, Brocade technology often becomes the industry standard.

In current SAN environments, each Fibre Channel link is 1or 2 Gbit/sec. Thus, the primary determining factor for the best path is actually the fewest number of links to traverse, also referred to as "hop count." Data sent from one switch to another constitutes one hop. If a server and its storage device are connected to the same switch, that would be zero hops. With more recent Brocade Fabric OS versions, organizations have the flexibility to force routing by changing the link values to indicate that one link is slower than another link even though they are technically the same speed.

Although organizations might want to change the value to force a certain data path, this is not usually a recommended approach. However, with 2 Gbit/sec links and switches, organizations will likely deploy a mixed 1 Gbit/sec and 2 Gbit/sec environment. The routing algorithms must then account for the different link speeds when determining the fastest route. With a mixed link-speed environment, both hop count and cost basis determine the shortest, or fastest, path.

However, the path with the fewest number of links might not necessarily be the fastest path. Additional hops with higher bandwidth links might actually be faster. Whatever the case might be, the network intelligence—which is already built into the Brocade FSPF routing algorithms—is critical for meeting high-performance expectations.

Another key performance aspect is the ability to dynamically tune the environment to meet specific requirements. In Figure 2, Server A has the shortest path to its storage. But what happens when Server B needs to access disk volumes in the same storage array and the fastest path chosen is through the same two cascaded switches? If Server A is using the entire throughput of the single ISL between the two switches, an organization could dynamically add another ISL between the two switches so Server B would have its own shortest path to the disk volumes. This dynamic change would not necessarily interrupt service to Server A.

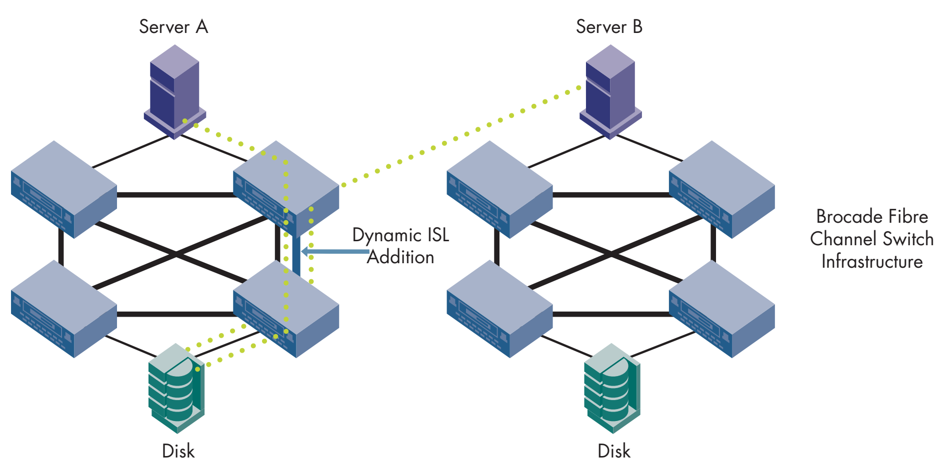

Figure 2.
Dynamic performance tuning

Simplified Manageability

With data management costs steadily rising, many organizations want to simplify storage management through centralization. To simplify fabric management, Brocade has created a variety of tools, such as Brocade WEB TOOLS, which provides a view of the entire fabric on a single, Web-based screen. Organizations might also want to manage their servers and HBAs, view a topology map of their networks, and manage their storage devices and adapters. To handle these tasks, they can use higher level management tools that work in conjunction with the industry-standard open Brocade Fabric OS API. Leading management products—such as VERITAS SANPoint Control and Tivoli Storage Network Manager—can provide excellent management value, especially in heterogeneous environments.

Manageability also means more than simply monitoring a SAN fabric environment. It also means proactively managing the environment by defining policies and thresholds to prevent an interruption of service. Brocade enables proactive management with Fabric Watch, a software offering that enables each switch to constantly monitor the SAN fabric for potential faults and automatically alert administrators about potential problems before they become costly failures.

Using Fabric Watch, organizations can implement policies to monitor ISL utilization. Consider the previous dynamic performance-tuning example. How would an organization know when to add another ISL for Server B before performance slowed down, or even worse, stopped? The organization could set a threshold monitor that would alert an administrator or automatic routine to dynamically add an ISL whenever the existing ISL reached 70 percent utilization, for example. As a result, Server B would experience no disruption of service. This type of proactive management and dynamic tuning helps enable a true high-availability environment capable of providing the five "nines" (99.999 percent) of availability.

Increased Security

Security is another key requirement to building a successful enterprise-level SAN. Switch security typically involves zoning, the process of restricting certain devices from specified parts of the fabric. Most organizations want to control access to resources, usually giving each server access to just a portion of the fabric's disk or tape resources. Some organizations implement a time-based resource allocation model where, for example, UNIX servers might have access to the tape resources from 5:00 P.M. to midnight to complete their backups and Windows NT servers might have access from midnight to 5:00 A.M.

Regardless of their needs, organizations can utilize industry-leading Brocade security solutions that provide flexible hardware- and software-enforced zoning. Hardware zoning means that the access is controlled by the ASIC such that the switch acts like a firewall to restrict data flow. Software zoning means that the name server, or database, in the switch controls data access. Although some vendors claim that port-based zoning is the same as hardware-based zoning, it isn't. True hardware-based zoning is done in the ASIC and is available in all Brocade SilkWorm 2000 and 3000 family switches.

The highest level of security usually stems from a combination of hardware and software zoning. Both hardware and software zoning should enable functions such as overlapped zones. In other words, a component such as a storage device should be able to reside in more than one zone if it is shared among many different servers. Brocade enables this capability along with the ability to distribute zoning information to new switches as they join an existing SAN fabric.

Another important function of zoning is scalability. Brocade switches provide an unlimited number of zones, while some vendor products limit the number of zones to as few as 16—a maximum that would not enable a very scalable enterprise SAN.

Brocade has also recently developed Secure Fabric OS, a comprehensive security solution for Brocade-based SAN fabrics. With its flexible design, Secure Fabric OS enables organizations to customize SAN security to address specific policy requirements. The most complete solution available for securing SAN fabric infrastructures, Secure Fabric OS includes the following sets of features:

- Fabric Configuration Servers
- Management Access Controls
- Device Connection Controls
- Switch Connection Controls
- Secure Management Communications

Long-Term Investment Protection

One of the most important characteristics of an enterprise-level SAN is investment protection. SANs should support the current technology infrastructure by accounting for legacy servers and devices, such as those that support only private loop protocols. Brocade supports legacy servers with QuickLoop and legacy devices with translative mode operation. Brocade also provides some full fabric benefits for these servers and devices. For instance, QuickLoop enables multiple simultaneous 1 and 2 Gbit/sec connections connections between private loop servers and devices, instead of the typical shared bandwidth. This results in a significant throughput advantage. In addition, public servers can access private devices.

Another aspect of investment protection is being able to start with a smaller, entry-level SAN environment and upgrade it to full fabric functionality with a simple license key. Brocade license key upgrades enable organizations to maintain their existing switch investments while increasing their value with new functions. No hardware or software changes are necessary.

Being able to scale the SAN environment without replacing existing switches is also important. Brocade switches feature forward- and backward-compatibility to enable interoperation with earlier models in a single SAN fabric. Moreover, new switches can be added non-disruptively to a fabric. If organizations need to increase the number of host server and device attachments or want to implement new technology, they can do so without making a forklift upgrade. In contrast, organizations that have designed SANs as a single large switch—a monolithic director approach—are very limited in their options. The lack of networkability seriously constrains incremental growth.

DESIGN CRITERIA: NETWORK INTELLIGENCE

Proper upfront design of an enterprise SAN can significantly improve scalability and manageability over time, and the most important design factor is the intelligence of the network. The only real way to build an intelligent networked fabric is to avoid single-switch configurations, because they are extremely limiting. No matter what size the switch is (the number of ports it has), it is usually the wrong size—either too big or too small. For example, an organization might start by calculating how many host server and storage connections it needs for a SAN. In this case, the total might be 33 ports. If a vendor sells only a 32-port switch, the organization would have to buy 64 ports, or two 32-port switches. Until at least some point in the future, the organization would be paying for 31 ports it doesn't really need.

If an organization realizes it needs 300, 500, or even 1,000 ports, no vendors are likely to support a single switch with that many ports (nor should they). If the organization needs more ports, it would eventually have to forklift upgrade its existing big switches to an even larger big switch, or at least add more big switches. But adding more big switches is essentially the same as building a fabric through networking—just more expensive.

Any enterprise-level organization realizes that high availability and business continuance require duplication for servers, applications, user access to applications, storage, and so on. And the storage network is no exception. The only way to avoid single points of failure is with redundant hardware, software, and people. A single large switch is inherently not highly available because it is only one device. Unless there is networking in the design, the system cannot be considered truly resilient.

The other primary limiting factor with a single large switch SAN design is the general lack of scalability, because growth typically requires redesign and other disruptive impact to the SAN environment. For instance, an organization might need to grow its environment simply because it needs to support more host servers or storage devices. Or it might need to adjust to changing business requirements or technology enhancements—such as 2 Gbit/sec links. With a networked switch fabric, the organization would be able to implement new technology non-disruptively. With a single large-switch design, the only way to add new technologies is with forklift upgrades or other disruptive methods.

One way to consider network intelligence is as a hierarchy of needs with the fundamental goal being continuous application access to data and the least interruption of service. In other words, organizations should prioritize and optimize SAN design for factors that are most likely to affect this goal. Figure 3 shows one possible hierarchy.

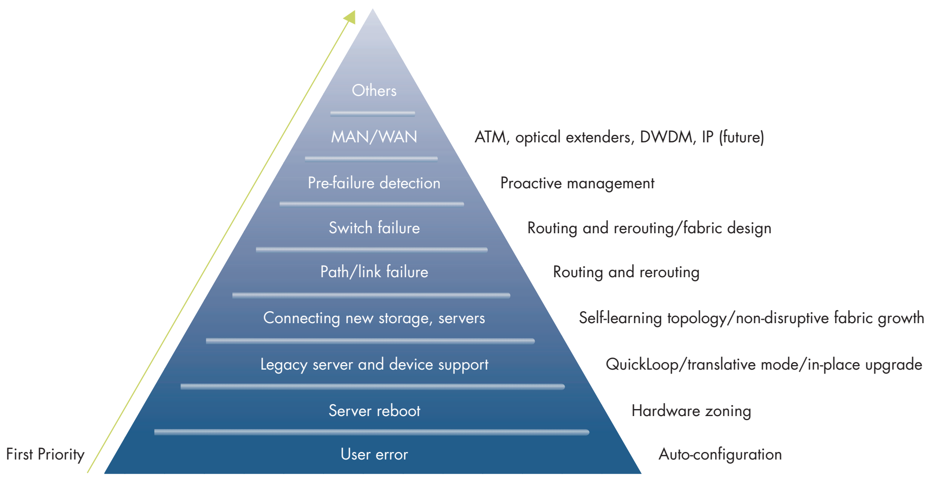

Figure 3.
A hierarchy of needs for SAN design

Because the most likely failure is user error, the highest priority is for network intelligence to provide data access even after a user error. Brocade functions such as auto-configuration and dual-fabric designs help minimize the impact of such errors.

The next most-likely failure might be a Windows NT server reboot, for example. If an organization uses only software zoning by World Wide Name (WWN), it runs the risk of a Windows NT server (after reboot) accessing storage it should not. This is because Windows NT on reboot might try to discover every device in the environment. By acting like a firewall controlled by the ASIC, Brocade hardware zoning prevents Windows NT from accessing storage it should not. Instead, Windows NT sees storage only through the ports for which it is zoned.

This hierarchy helps to demonstrate why it doesn't make sense to worry about CPU or switch failures unless an organization has first protected itself against user errors or server reboots, which are much more likely to occur than an extremely rare CPU or switch failure. In fact, Brocade has determined that its average Mean Time Between Failure (MTBF) rate for SilkWorm switches is 18 years. The actual measured MTBF of the thousands of installed switches is more than 50 years, so any switch failure is highly unlikely. Even if a switch does happen to fail, the resilient network design with alternate paths enables applications to maintain access to data.

OVERALL NETWORK DESIGN

Organizations can design SANs for a specific set of nodes or applications and then grow the design to an enterprise scale, or they can design an enterprise SAN from the beginning. First they need to determine how large the SAN might need to be. Size or scalability in terms of number of ports and the predicted growth rate are key factors. However, the total number of ports required for servers should also include the number of HBA connections, and for storage, the number of storage connection points.

Organizations also need to design the SAN environment for a desired level of availability, using duplicate hosts, storage, and networks, as well as long-distance connectivity options. In addition, organizations must design for the level of performance they need, ensuring that oversubscribed ISLs do not affect throughput. They might also have specific requirements for certain applications—such as the ability to segment the SAN fabric so that a long-running application job has no impact on the remaining SAN, for example. With a typical, single large-switch SAN design, each host server and storage device is connected to the switch—an approach that severely restricts scalability and growth. Eventually organizations run out of ports to connect hosts and storage and face a disruptive upgrade or redesign.

Some of today's most common SAN fabric topologies include the following approaches, which are described in more detail in the subsequent sections:

- Basic dual-switch SAN design
- Departmental SAN design
- Layered 2-tier design
- Backbone design
- Redundant star design
- Fan-out design

Basic Dual-Switch SAN Design

One of the most basic types of SANs is a dual-switch design, which generally includes a small number of hosts and storage devices (see Figure 4). It is not highly scalable, because all servers and storage must be connected to both switches to achieve high availability. Dual fabrics (fabrics in which switches are not connected to each other) protect organizations against errors such as a user erasing or changing the zoning information. Because the zoning information is separate for each fabric, altering the zoning in one fabric does not affect the second.

Figure 4.
A basic dual-switch design

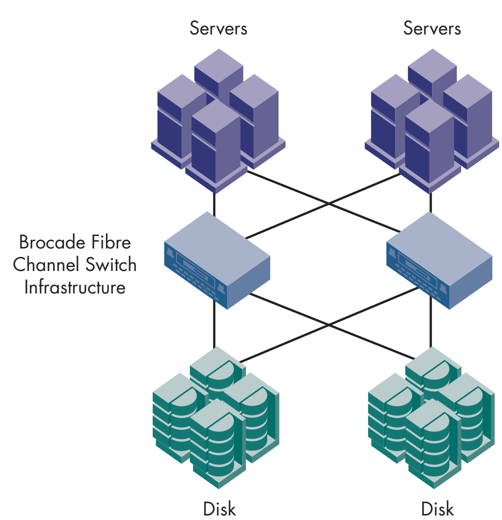

A dual-switch, dual-fabric design provides very high availability. Dual HBAs in each server and storage device must have at least two ports. Failover for a failed path or even a failed switch depends on host failover software, such as VERITAS Dynamic MultiPath (DMP), EMC Powerpath, or HP StorageWorks. The switches do not reroute traffic for a failed link, because there is no fabric or network with this design. Each switch is a single fabric.

Departmental SAN Design

A more scalable design is the departmental SAN, which is typically a mesh or a full-mesh network with an optimal design of four switches (see Figure 5). In a full-mesh design each switch connects to every other switch. This departmental SAN design has limited scalability: As the number of switches in the mesh increases, the number of ports available for server and storage nodes decreases.

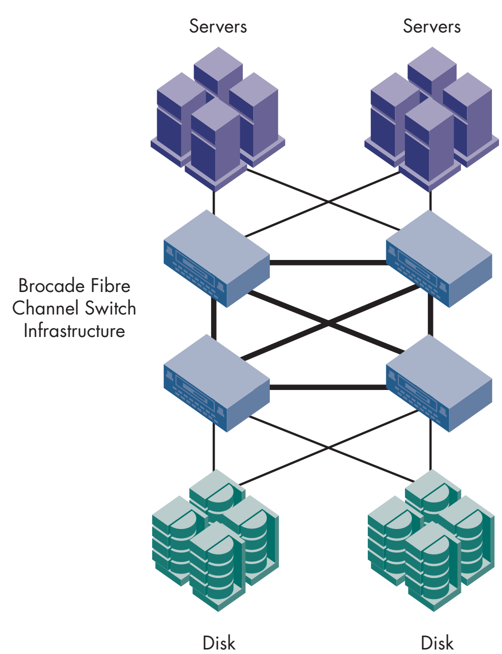

Figure 5.
A departmental SAN design

In contrast, a benefit of the full-mesh design is the ability to handle failures by the switch FSPF rerouting functions instead of relying on only the host failover software. For example, GBIC, link, path, and entire switch failures can often be handled by a simple rerouting through the network. The host server never sees the failure. However, the host failover software is still required for certain failures (such as HBA or certain types of link failures). As a result, the highest level of availability stems from a combination of host failover software and the intelligence in the Brocade fabric.

Two-Tier Layered Design

The next more scalable design is a 2-tier layered design. Organizations can implement this design with dual fabrics, which essentially look like two versions of Figure 6. A 2-tier design typically has a host tier and a storage tier of switches. All host servers are connected to the host tier, and all storage is connected to the storage tier. A path from a host server to its storage is always a single hop away. This design is often used for connecting a few large servers to a few large storage devices.

Figure 6.
A 2-tier layered design

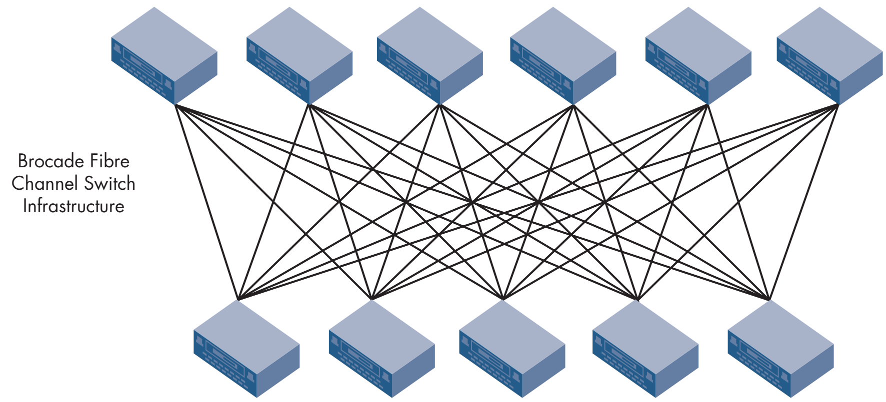

Although this design is more scalable than some others, it has limited scalability overall. There is a limit to how many switches can reside in either the host or tier layers, because all host tier switches are connected to all storage tier switches and vice versa. Organizations would eventually run out of ports if they added too many switches to either layer. However, this design can support hundreds of end-node connections with high throughput and high availability.

Backbone Design

A much more scalable approach is a backbone design borrowed from traditional networking methodologies that include core switches and edge switches. The edge switches are networked in any topology (such as dual switch, mesh, full mesh, 2-tier, or redundant star) to form SAN islands. The SAN islands then connect to each other to form a single fabric by using the core switches. Servers and storage devices are usually connected to the switches in the SAN islands, not directly to the switches in the core, which is typically reserved for connecting the islands. One exception to this rule might be a centralized tape backup solution, where all SAN islands can access a large tape library connected to the core switches. However, attaching a tape library to the core reduces the number of core ports available to connect the islands. Figure 7 shows two sets of core switches in a 4-switch mesh design. However, many other core switch topologies are possible.

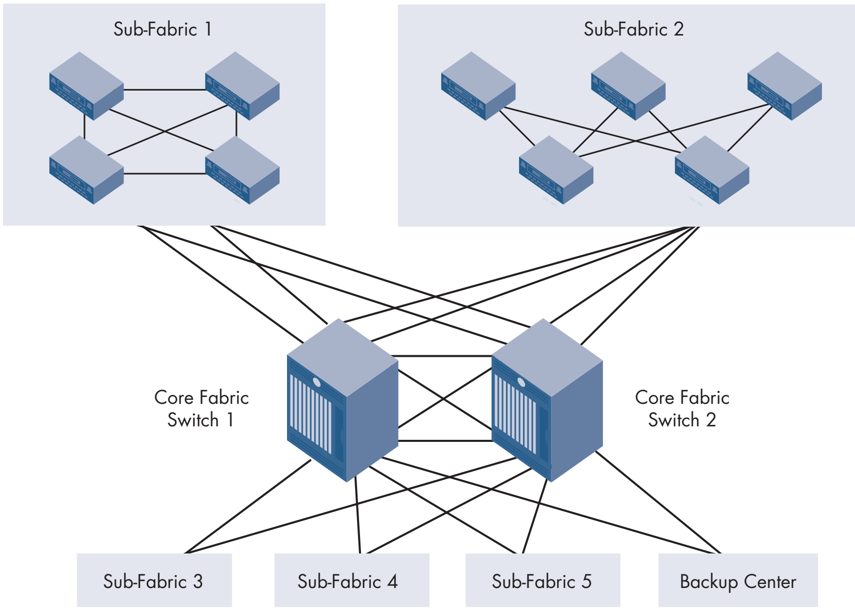

Figure 7.
A backbone design with edge and core switches

SAN backbone designs are an easy way to deploy enterprise fabrics, even if organizations begin with several smaller fabrics. Backbone designs provide a simple migration path to a single fabric without requiring complete SAN redesign, because the SAN islands are typically dispersed either geographically or logically. Geographic dispersion might be for departments or applications that are many miles apart. Logical dispersion might be a way to separate platforms or applications from each other even if they are in the same physical location. The most logical use for a large, high-port-count switch might be as a core switch in a SAN backbone. For the highest availability requirements, however, multiple core switches might be necessary.

Redundant Star Design

The most scalable, enterprise-level design is a topology borrowed from the traditional networking world: the redundant star topology consisting of a 3-layer switch design. The left central switch is connected to all switches above and below it as shown in Figure 8, forming a star. Higher availability requires a redundant star (the right central switch). The central switches are not connected to each other. This is a highly scalable design because it can grow both vertically and horizontally, and organizations can connect servers to one tier and storage to another tier.

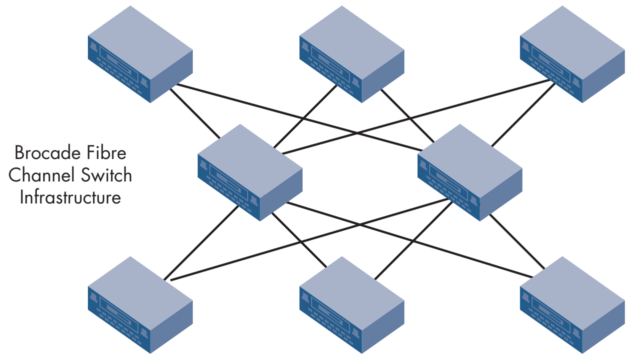

Figure 8.
A redundant star design

WHITE PAPER

Fan-out Design

The fan-out design typically connects many servers to a large storage array with very high availability. As shown in Figure 9, this redundant fabric design with a 16-port storage array enables 160 host server ports per fabric (using 16-port Brocade switches) and two ports per switch for tape backup connections.

Figure 9.
A redundant fan-out design

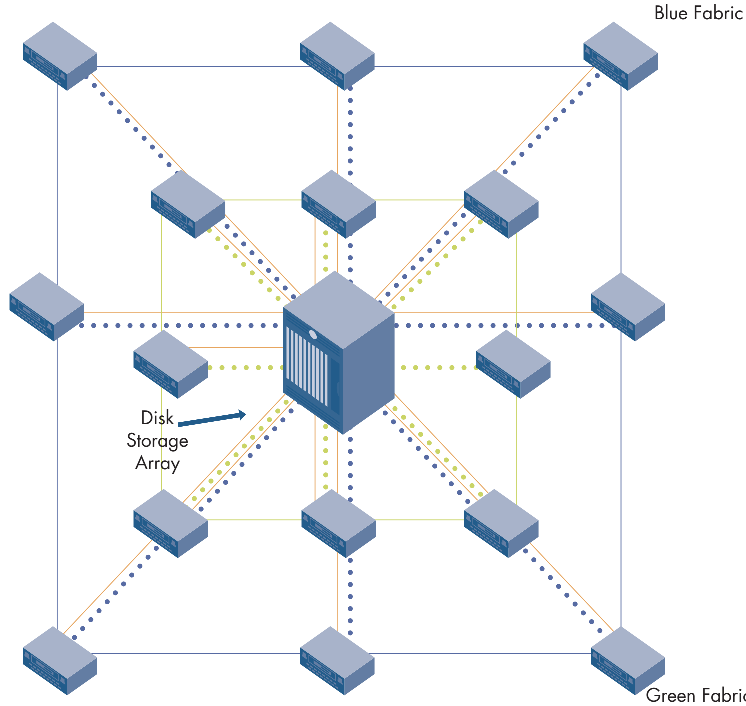

REAL-WORLD SAN IMPLEMENTATIONS

Many leading companies around the world are already enjoying the competitive advantages of a Brocade SAN implementation. Following are design examples for four Brocade customers: a very large bank, a telecommunications company, a software company, and an industry-leading Storage Service Provider (SSP). These examples often represent only a particular division within a very large company. Many companies start with smaller pilot projects and then expand the infrastructure to encompass more environments or to take advantage of new SAN technologies or applications.

Banking

One of the largest private sector banks in Europe wanted to implement a highly available, high-performance SAN environment. The bank had experienced a situation where an administrative user error took down a production server running a large banking transaction. Because strict European regulations require banks to pay a certain percentage of the transaction for every minute of downtime, the error was extremely costly.

To prevent this type of problem from recurring, the bank implemented a SAN infrastructure designed to address all types of outages: hardware, software, and physical disasters, as well as user errors. The bank implemented two SAN fabrics separated by 35 km (see Figure 10). The primary site has several large Sun E10000 servers and EMC Symmetrix storage, and the disaster-ready site duplicates the other. EMC remote mirroring through Symmetrix Remote Data Facility (SRDF) is used between the two sites.

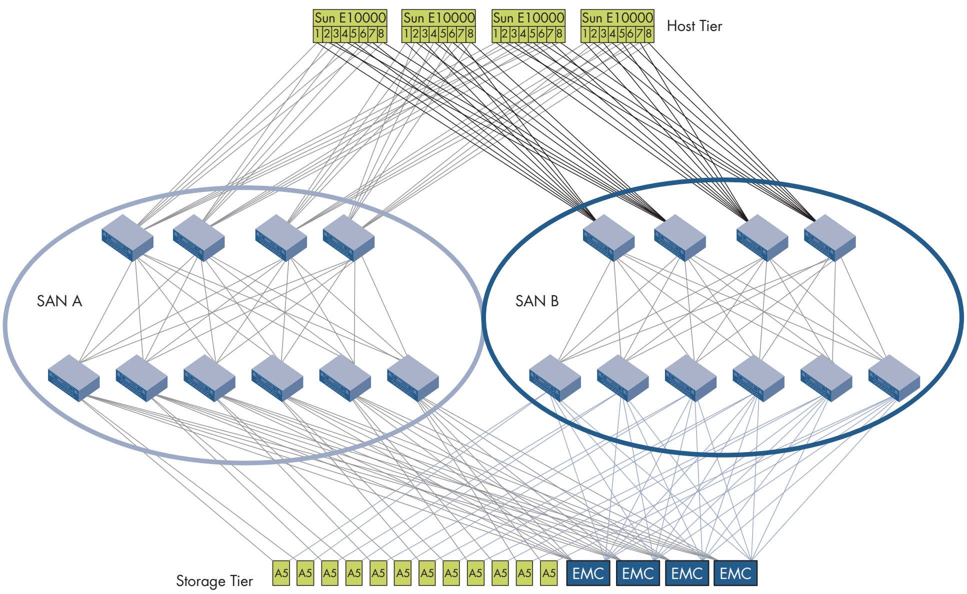

Figure 10.
Fault-resilient any-to-any connectivity enabled by two 10-switch Brocade SANs

The SAN infrastructure in the disaster-ready site is designed to protect against all types of outages and provide a very high-performance environment. The environment consists of four Sun E10000 servers, each running eight domains (similar to having 32 host servers). Four EMC Symmetrix arrays house most of the data, and several Sun StorEdge A5200 storage devices are used for the boot images. Any of the 32 domain images can boot from any of the boot images.

The SAN fabric is a 2-tier layered design that provides only a single hop between host servers and storage. Multiple HBAs from each domain, multiple storage connection points, and multiple ISLs enable a high-availability, high-throughput design. A redundant fabric provides an additional level of availability. No switches in the first SAN are connected to any switches in the second SAN, even though all host servers and storage are connected to both SANs. This dual-fabric design protects against a fabric-wide outage.

Telecommunications

A large telecommunications company had some very interesting SAN design goals. First, it wanted to localize data traffic wherever possible, connecting servers and disk storage to the same switch or same small group of switches. However, it also wanted all data to be available to the entire fabric: The servers and storage needed access to centrally located enterprise-class tape systems.

Another design goal was to increase flexibility by disabling or dropping ISLs to segment the fabric. Creating multiple fabrics would enable the company to isolate long-running batch applications to certain hosts and storage without impacting the rest of the I/O activity in their environment. To address these requirements, the company deployed a SAN backbone design with SAN islands that enable it to localize data traffic while making data available to the entire fabric (see Figure 11). The flexible SAN backbone design makes it easy to segment the fabric by disabling ISL connections between the SAN islands and the core backbone switches.

Figure 11.
A traditional backbone design

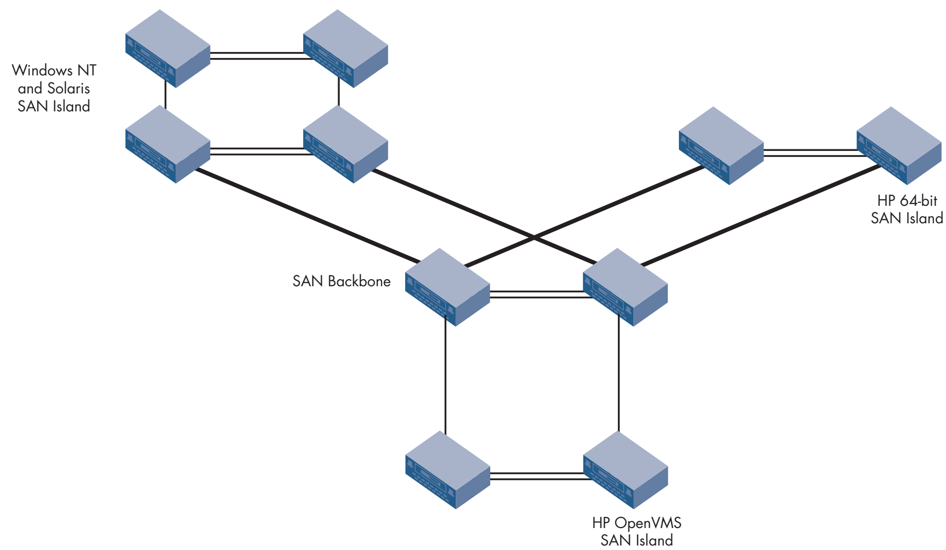

Future plans include implementing longer distance SANs to provide better disaster recovery support, as well as some dual-fabric designs for higher availability. The company has decided not to duplicate the entire SAN infrastructure, but instead, duplicate only key applications that require a higher level of availability.

Software Development

A large software company is currently implementing a Brocade-based 3-tier SAN design with plans for ongoing expansion (see Figure 12). The initial design has a host tier, core tier, and storage tier composed of 33 switches. The host switches are connected to only one core switch, but all the storage switches are connected to all three core switches.

Figure 12.
A 3-tier SAN fabric design

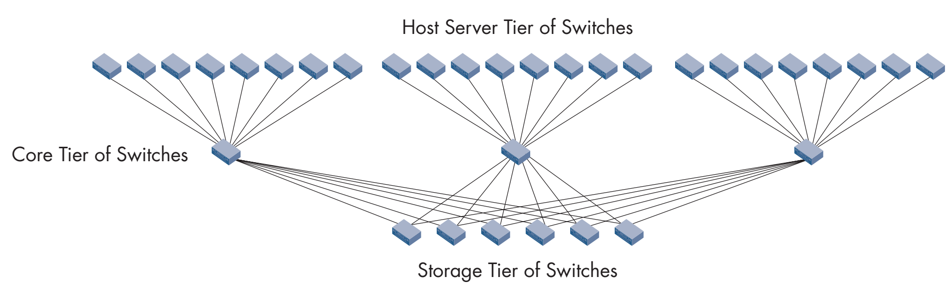

Although acceptable, this design is not ideal, because the single core switches represent a single point of failure. To address this issue, the company plans to implement dual fabrics (so the total number of switches will be 66 switches). Hardware- and software-based mirroring would provide failover capabilities. The dual fabric with mirroring guarantees no single point of failure. However, having at least dual-switch core switches would enable certain outages to be handled by Brocade rerouting instead of host failover software or a dual-fabric design. Dual-switch core switches would enable Brocade network intelligence to provide additional resiliency. In the next two years, the company plans to horizontally scale this design with more host, core, and storage switches to 62 switches per SAN fabric (or 124 switches for the dual fabric).

Storage Services Provider

A leading SSP wanted to implement one of the world's largest SAN fabrics to support multiple customers' storage requirements. The SSP has a single long-distance fabric for four customers in the metropolitan New York area, with a remote disaster recovery facility in New Jersey.

To support multiple customers over long distances, the SSP has implemented a SAN backbone design in which each SAN island isolates a customer environment. The SAN backbone (core switches) connects the centralized storage and Optical Networks, Inc. (ONI) DWDM equipment, which enables native Fibre Channel distances of up to 100 km. The ONI equipment is connected to a SAN island in New Jersey for long-distance disk mirroring.

One of the main applications in this SAN fabric is an integrated Network Attached Storage (NAS) solution with a back end of Fibre Channel-attached storage. In addition to the NAS workload, there is Oracle online transaction processing and data warehousing. The heterogeneous server environment includes Solaris, Windows NT, AIX, and HP-UX. Approximately 100 TB of storage is managed by the SSP on a variety of disk storage arrays. SAN management is accomplished by sending Brocade SNMP MIBs to CA Unicenter. Traditional enterprise LAN-based backup is performed by each customer using VERITAS NetBackup, but future plans include the SSP offering a LAN-free backup service.

The single SAN fabric includes 44 Brocade Silkworm 2800 switches and integrates some of the latest technology for long-distance SANs. Figure 13 shows the topology of this large fabric, with four SAN islands in New York—partial meshes of eight switches each (for a total of 32 switches). The four islands are connected to a backbone of six core switches in a partial mesh topology. The SAN island in New Jersey connects mirrored EMC and Compaq disks used for disaster tolerance. The SSP uses Brocade Extended Fabrics software to obtain near-native Fibre Channel performance over long distances.

Figure 13.
A large SSP SAN fabric

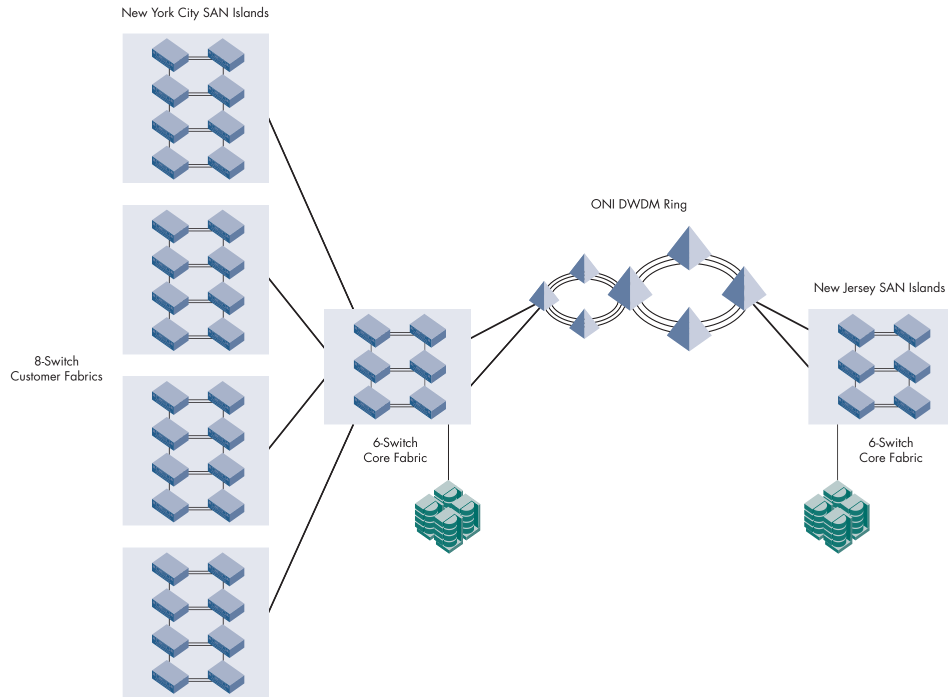

A BETTER WAY TO DEPLOY ENTERPRISE SANS

Organizations that face some of the most demanding requirements for high availability and performance have chosen to implement their SAN infrastructures with Brocade switches. Brocade provides enterprise-level SANs with best-in-class performance, availability, scalability, and investment protection. Moreover, Brocade switches strictly adhere to standards, and provide much of the network intelligence that has become the industry standard.

Today, organizations can start with a small SAN fabric and scale up to larger designs, or start right away by implementing enterprise and long-distance SANs in easy-to-manage fabric configurations. However, flexible enterprise SAN fabric design is not about implementing larger port count switches (such as a monolithic director-class switch) in standalone configurations. Implementing a successful enterprise SAN is all about providing the right network intelligence in conjunction with a proper SAN design to enable highly scalable and available environments.

For more information about Brocade enterprise SAN infrastructures, visit www.brocade.com.

BROCADE
The intelligent platform for networking storage

Corporate Headquarters
1745 Technology Drive
San Jose, CA 95110
T: (408) 487-8000
F: (408) 487-8101
info@brocade.com

European Headquarters
29, route de l'Aéroport
Case Postale 105
Geneva 15, Switzerland 1215
T: +41 22 799 56 40
F: +41 22 799 56 41
europe-info@brocade.com

Asia Pacific Headquarters
Shiroyama JT Trust Tower 36th Fl.
4-3-1 Toranomon
Minato-ku, Tokyo 105-6036
Japan
T: +81 3 5402 5300
F: +81 3 5402 5399
apac-info@brocade.com