

# SAN Topologies

## 2

This chapter describes the HP standard SAN topologies. You should review the SAN design considerations listed in the first chapter before starting the topology selection process. The design considerations enable you to generate a list of prioritized requirements for your SAN design. This list of requirements provides a basis for selecting the optimum fabric topology.

There are three approaches that you can choose when designing your SAN. You can choose to implement an HP standard SAN topology design, a subset or variation of an HP design, or you can design a custom SAN topology. Regardless of which approach you use, the final SAN design must adhere to the SAN design rules described in Chapter 3, "[SAN Fabric Design Rules](#)" and Chapter 4, "[Heterogeneous SAN Platform and Storage System Rules](#)."

Before choosing your design, you should review the HP standard SAN topology section in this chapter to get a good understanding of the important aspects of SAN implementation. HP recommends that you first consider implementing one of the HP standard topologies or a variation of one of these designs. If your requirements cannot be met by a standard design, then you can implement a customized SAN topology design—provided you follow the design rules.

## Why Design Rules?

HP performs extensive qualification of all HP SAN components, including application software, operating systems, host bus adapters, Fibre Channel switches, storage systems, and storage system management appliances. A sophisticated process is used to verify the interoperability of SAN components across a wide spectrum of supported configurations, taking into consideration potential customer requirements for mixed applications, servers, operating systems, and storage systems in a single environment. In order to ensure that a new SAN installation will function properly, certain guidelines must be followed. These guidelines are the result of actual laboratory testing—not theoretical projections of “what should work”—and reflect the designed-in capabilities of the various software and hardware components that are used in SAN storage systems.

The guidelines are captured in rule form to make it easier to design a SAN that will work properly and will be supported by HP. Chapter 3, "[SAN Fabric Design Rules](#)" and Chapter 4, "[Heterogeneous SAN Platform and Storage System Rules](#)", define the configuration rules. For additional information on operating system HBA/driver/firmware/software support, refer to individual array Streams connectivity documents available through HP field representatives or partners via the HP internal [SPOCK website](#).

## Product Lines

HP supports three product lines of Fibre Channel switch products that may be used to build SAN fabrics. Each product line provides certain advantages that apply to specific applications. For more information on each individual switch, and additional detail on selecting between the three product families, refer to the product information at [www.hp.com](http://www.hp.com).

The switches from the three product lines may not be mixed in a single fabric or multi-fabric SAN. If you have an existing SAN that uses switches from one product line, then you must construct another SAN in order to use switches from the other product line.

The B-Series Product Line includes a wide range of Fibre Channel switches, described as "SAN switches" and "Core switches." A partial list of products in this family includes the hp StorageWorks SAN Switch 2/16 and the hp StorageWorks Core Switch 2/64. This product line includes switches with 8, 16, 32, and 64 ports, including both full-function and entry-level models. The hp StorageWorks Core Switch 2/64 includes a pair of independent 64-port switches in a single chassis with a high level internal redundancy.

The M-Series Fabric Product Line includes a wide range of Fibre Channel switches described as "Directors" and "Edge switches." A partial list of products in this family includes the hp StorageWorks Director 2/64 and the hp StorageWorks Edge Switch 2/32. This product line includes switches with 16, 24, 32, 64, and 140 ports, and all models use the same version of internal microcode. The hp StorageWorks Director 2/64 and 2/140 have a high level of internal redundancy.

The C-Series Product Line currently includes the Cisco MDS 9509 Multilayer Director and the Cisco MDS 9216 Multilayer Fabric Switch. The MDS 9509 is supported with 176 ports, over 7 modular chassis consisting of both 16 port and 32 port modules. The MDS 9216 has a basic configuration with 16 ports. It has an expansion slot that supports either a 16 or a 32 port card, for 32 or 48 ports in total. Increased levels of support are planned for the C-Series Product Line.

The switch model numbering convention is the same in B-Series and M-Series product families. The number preceding the slash indicates the highest speed at which the switch ports can operate, measured in Gigabits per second (Gbps), and the number following the slash indicates the number of ports on the switch. The hp StorageWorks SAN Switch 2/16 is a 2 Gbps switch with 16 ports, and the hp StorageWorks Edge Switch 2/32 is a 2 Gbps switch with 32 ports.

Entry level models are indicated by the suffix "-EL" in their product name.

All HP 2 Gbps Fibre Channel products implement ports that auto negotiate their signaling speeds. Each pair of ports uses the lower of the supported speeds, so if a 2 Gbps port is connected to a 1 Gbps port, they both run at 1 Gbps in each direction. This applies to pairs of ports that are directly connected together. The speed of a remote port does not affect the local port speed, because speed matching is done within the switches in the fabric.

## Definitions

In order to understand SAN design, it is important to understand the relationship between a SAN and a fabric.

The Storage Networking Industry Association (SNIA) offers the following (slightly reworded) definitions of these two terms. (Refer to [www.snia.org](http://www.snia.org))

**SAN (Storage Area Network):** A network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements. A SAN consists of a communication infrastructure, which provides physical connections, and a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust. The term SAN is usually (but not necessarily) identified with block I/O services rather than file access services. A storage system consisting of storage elements, storage devices, computer systems, and/or appliances, plus all control software, communicating over a network.

**Fabric:** A Fibre Channel switch or two or more Fibre Channel switches interconnected in such a way that data can be physically transmitted between any two N\_Ports on any of the switches. The switches that constitute a Fibre Channel fabric are capable of routing frames using only the D\_ID in a FC-2 frame header.

This document discusses Fibre Channel SANs.

The SNIA definition of "fabric" refers to Fibre Channel and expresses the concept of fabric that is used in this document. An N\_Port is an endpoint in the fabric, which is present in a server's Host Bus Adapter (HBA), a storage system, or a storage management appliance connected to the fabric. The D\_ID is the address of the destination N\_Port, and is contained in the header of every Fibre Channel packet. A Fibre Channel packet is also called a frame. FC-2 is the layer in the Fibre Channel protocol associated with packet routing.

Another way to define a fabric is to view it as a single FC-2 address space. Any valid FC-2 address that can be reached from a given N\_Port in a fabric is part of that fabric.

The SNIA definition of "SAN" does not require that a SAN be implemented with Fibre Channel technology. When the term SAN is used in connection with Fibre Channel technology, use of a qualified phrase such as "Fibre Channel SAN" is encouraged. This usage is usually not required, because in most cases today, a SAN is a Fibre Channel SAN. According to the SNIA definition, an Ethernet-based network whose primary purpose is to provide access to storage elements would be considered a SAN. SANs are also sometimes used for system interconnection in clusters.

A given SAN can contain one or more fabrics. Most small SAN configurations use a single fabric. Larger storage environments may require very high levels of availability, and the best way to obtain this is by using a pair of redundant fabrics. HP storage systems, servers, and operating systems support both design approaches.

## HP Standard SAN Topologies

The HP standard topology designs reflect the proper application of the HP SAN design rules. Each of the standard designs is tailored for a particular data access and connectivity need. Collectively, these designs provide a wide range of options for selecting the appropriate SAN design for your specific requirements. Variations of these designs, including additions or changes, can be validated by adhering to the appropriate rule set for each topology type. A subset of a standard design is always acceptable.

The different types of HP standard SAN topologies are described in detail in the following sections.

## SAN Fabric Topologies

SAN fabric topology designs include:

- Cascaded Fabrics
- Meshed Fabrics
- Ring Fabrics
- Backbone Fabrics

Each of the design types can be:

- Implemented as a separate SAN for specific departments or applications within a company, to accommodate different data access needs.
- Implemented with centralized backup capabilities, reducing the cost of backup and restore operations.
- Deployed in one or more co-located groups.
- Deployed across a wide area with inter-switch distances up to 35 km (2 Gbps) and 100 km (1 Gbps.)

---

**Note:** Refer to Chapter 8, "[SAN Extension](#)", for additional information on extending SANs over long distances.

---

- Used to begin an ongoing deployment process using SANs and Fibre Channel technology in a modular, controlled approach. Storage consolidation can be implemented on a departmental or independent SAN basis. Future capabilities will allow for more switches within a single SAN, interconnection of multiple SANs to build larger fabrics, and provide for additional consolidation, if desired, or broader server-to-storage access.
- Centrally managed.
- Implemented with a wide range of SAN availability levels. See "[Levels of Availability](#)."
- Upgraded to higher capacity topologies or topologies optimized for different data access types if needs change.


## Single-Switch Fabrics

The smallest SAN consists of a single Fibre Channel switch, server, and storage system. This topology is a subset of all the other topologies, and forms the basis of the range of HP SAN solutions.

By choosing among various HP Fibre Channel switches, you can construct a wide range of single-switch SAN solutions. The smallest supported HP SAN uses a single 8-port switch, the hp StorageWorks SAN switch 2/8. If you need a large single switch SAN, the hp StorageWorks SAN Director 2/140 offers 140 ports. Furthermore, if there is a need for high SAN availability, two independent single-switch SAN fabrics may be used in a dual-fabric environment to give a total of 280 user ports.

A single-switch fabric maximizes SAN performance, because every port on the switch has full connectivity to every other port on the switch. This design is also very easy to install and configure, since there are no connections from one switch to another.

An example of a simple single-switch SAN is shown in Figure 1.



**Figure 1: Single-switch SAN**

Starting from a single-switch configuration, you can add more switches to your SAN fabric—following the support limits listed for each of the fabric topology designs—to increase the number of connections for servers and storage.


You can also view the various topology options as a way to connect existing smaller SANs or SAN islands. If you have already deployed several small SANs, you can connect them together to make a larger SAN. For example, if you have two single-switch SANs, you can connect them together into a cascaded fabric. Or, if you have deployed two four-switch meshed SANs as separate SANs, you can merge these into a larger single 8-switch meshed SAN as shown in Figure 4, "[Modified Meshed Fabric SAN](#)". If you have multiple single-switch SAN fabrics, you can connect these into a single larger SAN fabric by connecting them in a ring, or to a central backbone, using the backbone fabric topology.

## Cascaded, Meshed, and Ring SAN Fabrics

The first three fabric topologies that involve more than a single switch are organized so that all of the switches in the fabric are used for connecting servers and storage. Every switch has at least one user port. Typically, in these types of fabric arrangements, a small percentage of the total number of switch ports is used for inter-switch connectivity in the form of Inter-Switch Links (ISLs). Refer to [Chapter 1, "Connectivity"](#), for more information about trading ports used for ISLs for ports used for servers and storage.

### Cascaded Fabrics

A cascaded fabric SAN (see Figure 2) is a set of switches connected together, by one or more ISLs, in a tree-like arrangement.



SHR-2552B

**Figure 2: Cascaded Fabric SAN**

Cascaded fabric designs are well suited to environments with local data access patterns. In these cases, I/O requests from servers attached to a given switch are made most often to storage systems that are attached to the same switch. Groups of servers and their storage systems can be connected to the same switch to provide the highest level of I/O performance. Cascading provides a means to scale the SAN for additional connectivity of servers and storage, and allows for centralized management and backup, while maintaining the high I/O performance of local access.

Cascaded designs can also be used for centralized or distributed access; however, traffic patterns should be well understood and should be factored into the design to ensure that there are an adequate number of ISLs to meet performance requirements. Using more than one ISL between switches in a cascade also provides redundant paths between a given pair of switches in the fabric. HP highly recommends that cascaded designs be implemented with a minimum of two ISL connections on each switch, either as a pair of ISLs between the same two switches or by connecting every switch to at least two other switches in the fabric.

### Very Large Cascaded Director Plus Edge Switch Fabric

The largest fabric supported by HP at this time is based on a cascaded fabric using hp StorageWorks Director 2/140 and StorageWorks Edge switches. Because the cascade configuration may be used to maximize user port count—if relatively low bandwidth between user ports that are on separate Directors is acceptable—this configuration can provide a level of connectivity that cannot be obtained by other means. Because of the availability features built into the SAN Director product, this configuration also offers high availability in a single-fabric topology.


The SAN Director product family is limited to a three-hop maximum. Hop count is a measure of distance or links between switches, measured in ISLs. A cascade configuration with 24 SAN Directors and Edge switches that meets the three-hop rule has one Director at the top of the tree, with 7 Directors attached under it, and 16 Edge switches under them. This gives a total of 24 SAN switches in one fabric. This supported configuration provides 1024 user ports.

## Advantages of Cascaded Fabrics

- Accommodates diverse geographic conditions
- Scales easily for additional connectivity
- Shared backup is supported
- Shared management is supported
- Optimal local access is inherent in the fabric design
- Most efficient in cost per port

## Meshed Fabrics


In a meshed fabric design, all of the switches are interconnected so there are at least two paths or routes from any one switch to any other switch in the fabric. This type of connectivity provides fabric resiliency. If a single ISL or ISL port interface fails, the fabric can automatically re-route data through an alternate path. The new route can even pass through additional switches in the fabric. An example of a meshed fabric is shown in Figure 3.



**Figure 3: Meshed Fabric**

As switches are added to a meshed topology, the number of ISLs required to maintain full connectivity between any switch and any other switch becomes excessive. This reduces the number of user ports in comparison to the total number of ports, which is a measure of the connection efficiency of the fabric. The connection efficiency of this fabric design can be improved by implementing a slightly modified mesh design, as shown in Figure 4. In this case the connectivity between switches is reduced, but the fabric availability is maintained because there are still multiple paths between switches.





**Figure 4: Modified Meshed Fabric SAN**

In this example diagram, as switches are added, they are only connected to adjacent switches, not all other switches in the fabric. This still provides the benefits of full many-to-many connectivity without a decrease in connection efficiency.


Meshed fabrics are well suited to applications where data access is a mix of local and distributed. The full connectivity (or high connectivity, in the case of modified meshes) supports many-to-many access, while at the same time allowing localized access to individual switches, servers and storage.

## Advantages of Meshed Fabrics

- Can be configured for any to any or local data access, or a mix  
Reduces staff effort by minimizing reconfiguration and re-cabling of existing Fibre Channel switches. Adapts easily to new or different storage needs.
- Provides protection against link and switch port failures  
Fabric design allows Fibre Channel switches to automatically re-route under failure conditions, saving time and effort to manually trace the problem and re-route.
- Scales easily  
The mesh design can be extended from a four-switch fabric to six or eight switches easily, and without disruption to the existing SAN. The mesh design affords ease of adding servers to the SAN without impacting existing connections or equipment. This is especially useful for companies where there is rapid growth, or computing and storage needs are changing frequently.
- Shared backup is supported  
One or more Automated Tape Libraries can be added to the mesh fabric at various points without impacting performance or management.
- Shared management is supported  
All Storage Management Appliance tools can navigate and manage the Storage Area Network in the mesh fabric, saving time and effort.
- Optimal distributed access is inherent in the fabric design

## Ring Fabric


A ring fabric (see Figure 5, "Ring Fabric SAN") is a continuous ring of switches connected together into a single fabric. Each switch is connected to adjacent switches, with the last switch in the ring connected back to the first. This arrangement of switches provides almost the same level of fabric resiliency as the mesh design, with full fabric connectivity and at least two internal fabric paths or routes.



SHR-2154B

**Figure 5: Ring Fabric SAN**

If you use fewer than 12 switches in a ring constructed using B-Series Product Line switches, you can add additional switches to the outside of the ring. These satellite switches provide additional user ports with only a slight reduction in fabric availability. For example, 11 satellite switches can be connected to a 11-switch ring. This results in a 22-switch fabric and maintains the overall seven hop limit. Figure 6 shows a 22-switch fabric.



SHR-2544A

**Figure 6: Ring Fabric SAN with Satellite Switches**

Ring fabric designs are well suited to applications where data access is always localized. Servers and the storage that is accessed are on the same switch, and the majority of data traffic is handled within that switch. This implementation provides a way to scale the fabric in a modular fashion by adding a switch and groups of servers and storage as a cell, using a building block approach to increase the size of the SAN over time. This is particularly useful in situations where the storage capacity requirements vary over time, such as in a storage service provider environment.

A ring fabric can be pre-configured and installed before the server requirements are known. This is useful because the ability to install the fabric infrastructure beforehand can greatly simplify the installation of each incremental storage system or server. Interconnecting the switches in a ring topology provides a communication path that supports centralized SAN management and centralized backup.

The ring fabric is not recommended for applications that require many-to-many connectivity.

## Advantages of Ring Fabrics

- Easy to build  
Each Fibre Channel switch can support servers and storage, thus saving time and effort on SAN design and implementation.
- Scaling is simple and non-disruptive  
Fibre Channel switches can be added one at a time, as storage and connection needs dictate. Each Fibre Channel switch can support identical servers and storage for controlled growth, or can support a variety of heterogeneous systems for new demands of the business.
- Shared backup is supported  
One or more Automated Tape Libraries can be added to the ring fabric at various points without impacting performance or management.
- Shared management is supported  
All Storage Management Appliance tools can navigate and manage a SAN with a ring topology, saving time and effort.
- Optimal local access is inherent in the fabric design  
The majority of the data traffic is within each switch in the ring, minimizing any allocation, fabric and performance issues.
- Modular design  
Saves time and effort on design and implementation by complementing the basic modularity of all StorageWorks products, including the raid array controllers, universal packaging, and secondary storage (Automated Tape Libraries).


## Backbone Fabric

A backbone fabric has one or more Fibre Channel switches primarily dedicated to connecting to other switches within the fabric. The backbone switches provide high bandwidth and redundant connectivity to the other switches. This type of implementation offers the best "many-to-many" connectivity.

Backbone fabrics are well suited for implementations where the primary requirement is for full network “many-to-many” connectivity with high performance. They are the most conservative design approach in cases where the I/O traffic patterns are unknown or varying. They are also the best design to choose if you plan to implement SAN-wide storage pooling and sharing, and for environments that use storage virtualization.

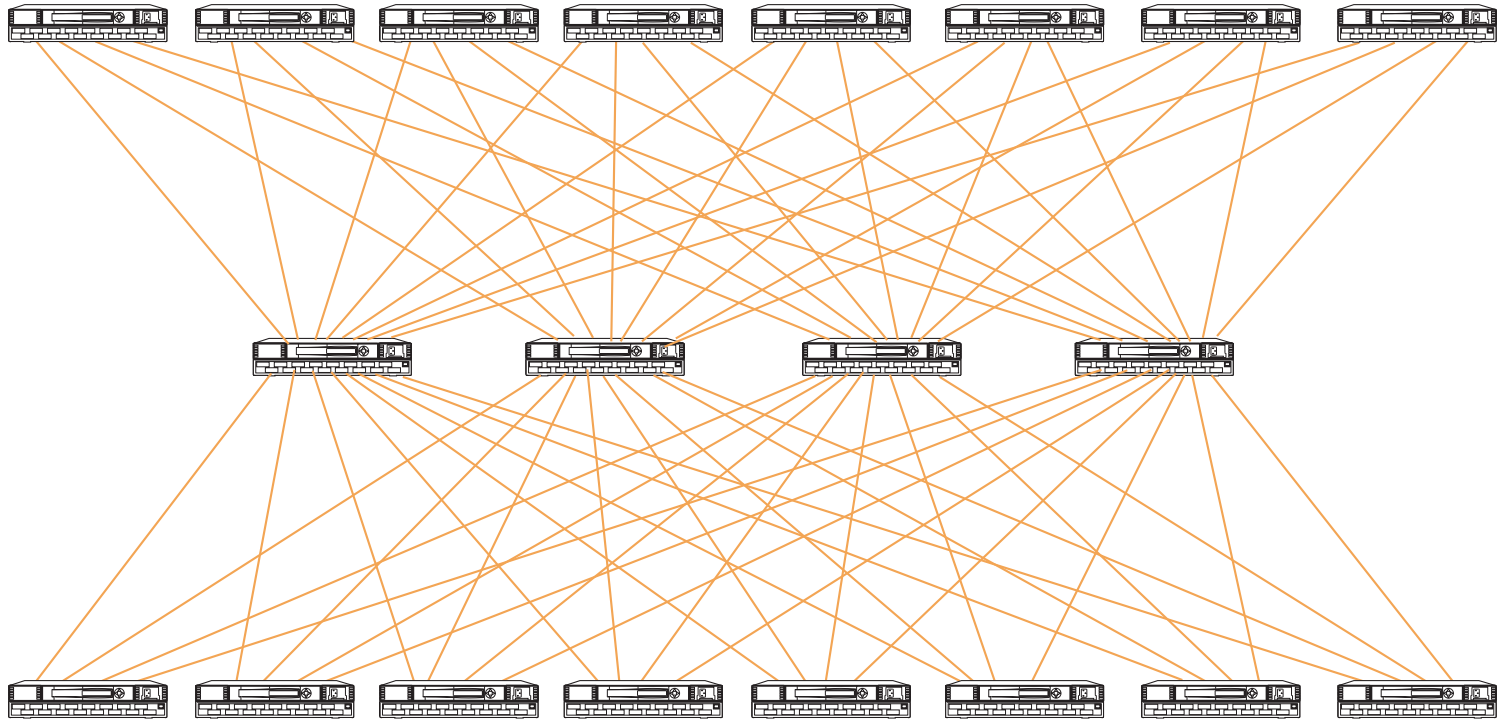
Figure 7 shows a backbone fabric where the switches in the center are dedicated to providing connections between the other switches. The switches to which servers and storage can be connected are called “edge switches”, and the switches in the center are called “backbone switches”. Servers and storage can be connected to the user ports on any of the edge switches, which maximizes the flexibility of how you use the user ports.

If required, you may choose to connect centralized primary (disk) or secondary (tape) storage directly on the backbone switches. This approach may be useful if excess ports are available on the backbone switches.



**Figure 7: Backbone Fabric SAN**

HP StorageWorks SAN fabrics currently support a large number of switches in a fabric. Figure 8 shows a large backbone fabric SAN with 20 switches. Configurations of this type can be used to support the most demanding requirements for storage system size and performance.




SHR-2164A

**Figure 8: Backbone SAN with 20 Switches**

## Fat Tree and Skinny Tree Designs

The SAN shown in Figure 7 can be drawn with the backbone switches at the top and the edge switches collected together into a row at the bottom, as shown in Figure 9. This method of illustrating the topology is helpful when evaluating the potential performance of a given backbone configuration. It shows the switches in a hierarchy, where the edge switches form a layer that provides access to the SAN and the backbone switches form a layer that distributes I/O requests between edge switches.

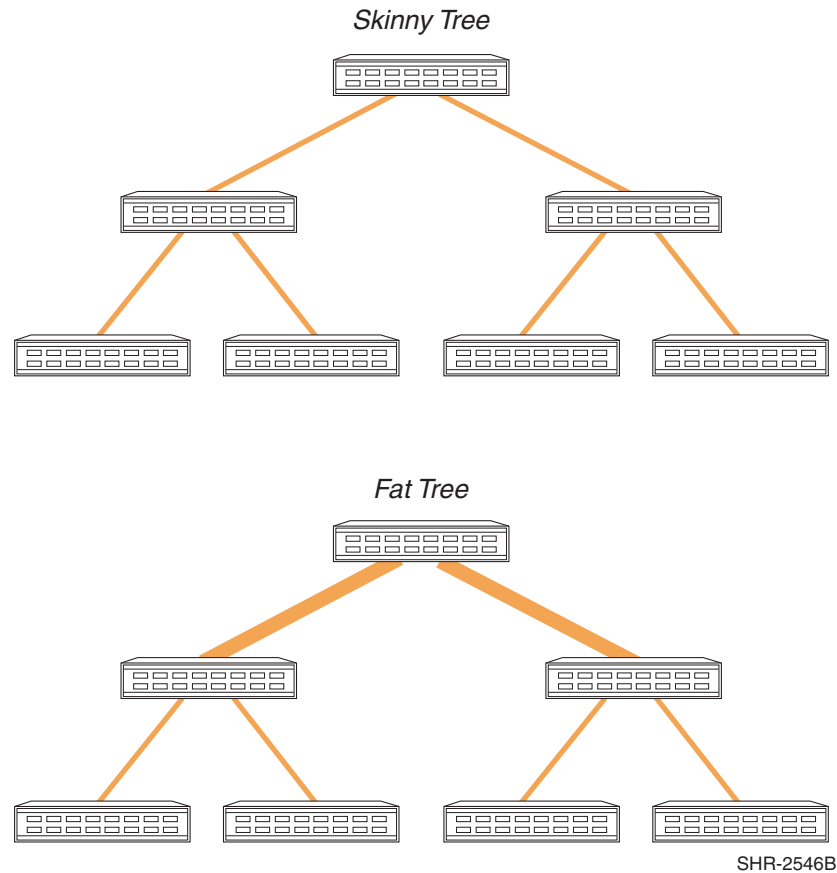


SHR-2545B

**Figure 9: Backbone SAN, Drawn Hierarchically**

Depending on how a backbone SAN is designed it can be classified as a “fat tree” or a “skinny tree”. The difference between fat and skinny trees is the number of ISLs used to connect the edge switches to the backbone switches. The number of ISLs subtracts from the number of end ports and therefore affects the total number of switches needed for a particular configuration. Fat trees use 50% of the edge switch ports as ISL connections while skinny trees use fewer than 50%.

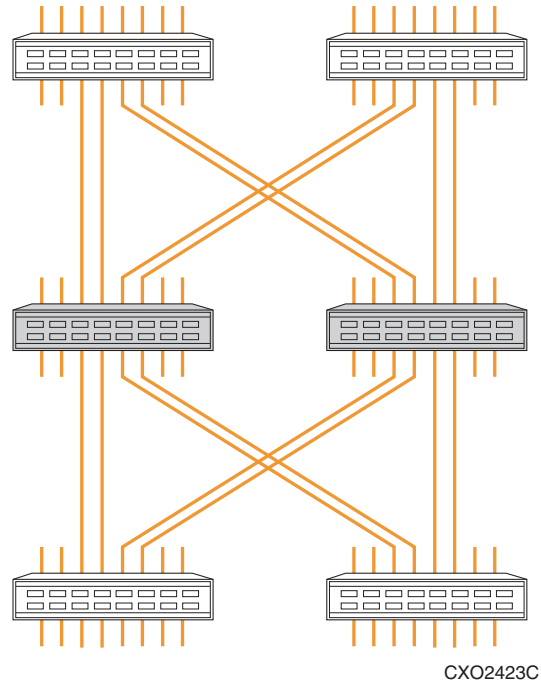
Figure 10 shows two hierarchical SAN fabrics. The skinny tree illustration shows that if all the devices connected to user ports on the left-hand side of the tree want to communicate with the devices on the right-hand side of the tree, then there are not enough connections to the switch at the top of the tree for the required traffic. There are too many port-pairs trying to use the ISLs on the switch at the root of the tree. The fat tree illustration shows that by providing additional ISLs between the switches that are further up in the hierarchy, full-performance bandwidth can be provided for *any combination of port-pairs*. This important benefit is the basis for the use of fat tree configurations in application environments where the highest level of performance and capability are required.



**Figure 10: Skinny Tree and Fat Tree**

This distinction in the number of ISL connections between fat and skinny trees results in two major differences:

1. Skinny trees require fewer switches than fat trees to supply the same number of user ports. Figure 11 shows how six 16-port switches in a skinny tree configuration yield 64 user ports, while the same switches wired in a fat tree as shown in Figure 12 yield only 32 user ports.
2. Fat trees have more ISL connections and therefore have higher cross sectional bandwidth capabilities than skinny trees. The term cross sectional bandwidth is used to refer to the maximum amount of data that can pass through the ISL connections at the midpoint of the fabric.




**Figure 11: 64-Port Skinny Tree**

The configuration shown in Figure 11 has six 16-port switches. Two are backbone switches (shaded) and four are edge switches. The edge switches each have 12 user ports available to connect to servers or storage. With the 48 ports on the edge switches and with the 16 available ports on the backbone switches the total number of user ports is 64.

There are 8 ISLs on each side of the backbone switches.

Note that a mix of 1 Gbps and 2 Gbps switches or devices makes the evaluation of this type of configuration more complicated.





**Figure 12: 32-Port Fat Tree**

The configuration shown in Figure 12 also has six 16-port switches, with two backbone switches and four edge switches. The edge switches each have eight user ports available to connect to servers or storage. No ports are available on the backbone switches for user ports, so the total number of user ports is 32.

### Pre-configured Backbone SANs

The hp StorageWorks SAN Switch Integrated 32 is a pre-configured 6-switch fat tree SAN, and the hp StorageWorks SAN Switch Integrated 64 is a pre-configured 6-switch skinny tree SAN topology. Both products house six 16-port switches in a rack mountable enclosure, and come with all the inter-switch cabling pre-connected. Available edge switch user ports are clearly marked and identified, allowing for expedient deployment of a SAN fabric.


For information about implementing and configuring fat and skinny tree SAN topologies using existing Fibre Channel switches refer to the application note “*Combining 16-Port Switches to Construct Higher Port Count Switches*”, part number AA-RPH8A-TE.

### Backbone SANs Using Core Switches and Directors

With the introduction of switches with a higher number of ports, backbone fabrics may be constructed using a mix of small and large switches. The HP standard SAN designs using this approach are as follows.

The “4 x 12” configuration, shown in Figure 13, uses a mixture of hp StorageWorks Core Switch 2/64 and SAN Switch 2/16 or SAN switch 16 devices. Four core switches are connected in a mesh configuration with two ISLs between each pair of switches in the mesh. Twelve SAN switches are connected to the mesh with four ISLs between each SAN switch and the mesh. This gives a total of 304 user ports, including 160 on the mesh and 144 on the SAN switches.

This configuration is useful for situations where a large port count is required in a single fabric, and where many high-performance systems can make use of direct connections to the mesh. For example, high performance storage systems can make good use of a 2 Gbps connection, and should be connected directly to the mesh. Since each 16-port SAN switch has four ISLs connecting it to the core, the 8 Gbps (4 x 2 Gbps) total bandwidth between the SAN switch and the mesh must be shared between the 12 user ports on that switch. This means that the I/O performance requirement for each server attached to a SAN switch in this configuration must be reviewed to make sure that the ISLs are not an I/O bottleneck. Because of the limited number of ISLs between the switches in the mesh, this configuration is not appropriate for environments where many-to-many traffic patterns are predominant.




SHR-2554A

**Figure 13: “4 x 12” Backbone SAN**

The “4 x 24” configuration also uses a mixture of HP StorageWorks core switch and SAN switch devices. By using a larger number of switches, this topology design provides adequate internal connectivity within the mesh and also provides and more user ports than the 4 x 12 configuration.

In the 4 x 24 configuration, shown in Figure 14, four core switches are connected in a mesh with four ISLs between each pair of switches in the mesh. 24 SAN switches (1 Gbps or 2 Gbps models) are connected to the mesh, with a total of four ISLs on each SAN switch (just the same as in the 4 x 12 configuration). This gives a total of 400 user ports, including 112 on the mesh and 288 on the SAN switches. This configuration may provide a better trade-off of high-performance connections—directly to the core switches—and lower-performance connections on the SAN switches.



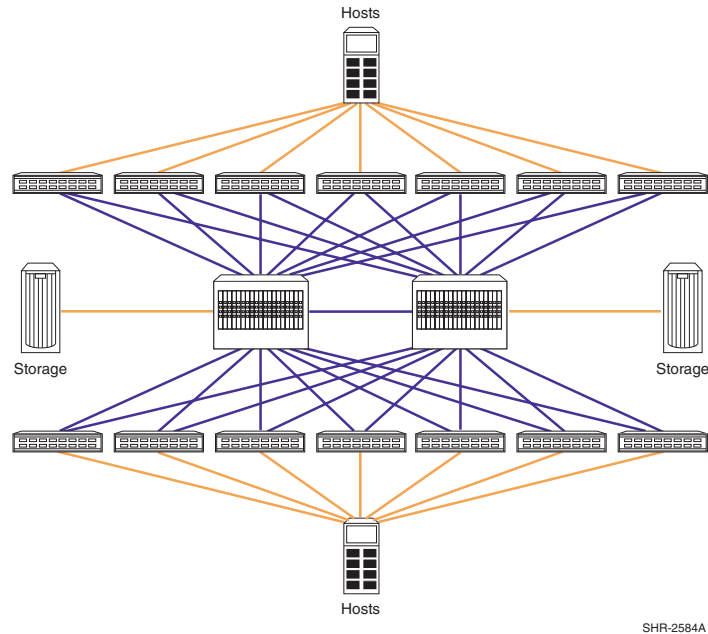
**Figure 14: "4 x 24" Backbone SAN**

## Director Fabrics

Backbone fabrics may be constructed using Director and Edge switches. These products may be used in configurations where no more than half of the ports on a given switch is used for ISL connections. This means that a fat tree configuration cannot be constructed because the distribution layer cannot be fully populated. On the other hand, there are fewer topology restrictions for these products, so you can combine a mixture of Director and Edge switches, with a single Director as the backbone, to make a skinny tree fabric.

For example, a skinny tree using the design shown in Figure 11, but using 64-port Directors, meets the ISL port count and hop count requirements for this family of products while providing 256 user ports.

Figure 15 shows a "Director plus Edge switch" tree design using two HP StorageWorks Directors and 14 HP StorageWorks Edge switches. In this topology, high performance storage systems are connected directly to the Directors, because they can make use of the full bandwidth of the 2 Gbps Fibre Channel connections. Servers are connected to the edge switches, because they require only a smaller amount of bandwidth. Depending on the bandwidth requirements of the servers, the number of ISLs between the edge switches and the Directors can be varied.



**Figure 15: Director plus edge switch SAN**

## Advantages of Backbone SANs

- Efficient port expansion: new switches need only be connected to backbone switches. Saves time and effort during the design and implementation phases by isolating the new switches from the existing SAN backbone.
- All edge switches are only two hops apart. Saves design effort for adding new servers and storage to any point on the SAN. The uniformity of access supports new usage patterns without requiring redesign and re-cabling.
- When implemented with two or more backbone switches, provides a level of switch redundancy in a single fabric. Backbone design allows Fibre Channel switches to automatically re-route under failure conditions, saving time and effort to manually trace the problem and re-route.
- Can be centrally managed. All Storage Management Appliance tools can navigate and manage the Storage Area Network in a tree backbone fabric, saving time and effort.
- Full "many-to-many" connectivity with evenly distributed bandwidth and redundant connectivity. Supports varying connection and performance demands regardless of the location within the SAN. At the same time, provides uniform routing and redundancy from a single SAN design.
- Improved bandwidth with multiple parallel ISLs. Additional ISLs ensure that all data traffic within the tree backbone SAN will be managed with less performance degradation, regardless of the location of servers and storage relative to each other.

- Offer maximum flexibility for implementing mixed access types: Local, Distributed, or Centralized.  
Saves effort planning data traffic patterns; the tree backbone supports all access patterns.
- Can be implemented with centralized backup capabilities, reducing the cost of backup and restore operations
- Can be implemented with all availability levels  
Saves effort in the design and implementation phases by offering a single design for a variety of usage requirements.
- Can be an upgrade path from other SAN designs. Backbone SAN designs offer evenly distributed bandwidth and full many-to-many connectivity; they are the best solution for flexible SAN-wide storage pooling and sharing.
- Well-suited to take full advantage of expected future technological developments such as storage virtualization  
Saves the investment made in the SAN by continuing its use as more advanced tools, products, and designs become available.

## Topology Data Access Usage

The various SAN topology options can be characterized by how well they support specific data access patterns. Refer to Chapter 1, "Data Locality". Table 2 provides a general characterization of the different topology designs as a means to compare each of the design types by optimal data access capabilities. Use the table as a basis for selecting the best-suited topology for your expected access needs.

Individual topologies can be tailored or modified to better meet specific requirements. For example, choosing a fat tree backbone design provides the best overall "many-to-many" connectivity, and allows portions of the tree implementation to be configured for local access. This can be accomplished by connecting servers and storage typically accessed on the same switch within portions of the tree backbone.

**Table 2: Topology Usage Rating**

SAN Topology	Data Locality		
	Local "One-to-One"	Centralized "Many-to-One"	Distributed "Many-to-Many"
Cascaded	Highest	Not Recommended	Not Recommended
Meshed	Medium	Medium	High
Ring	Highest	Medium	Not Recommended
Skinny Tree Backbone	Medium	High	High
Fat Tree Backbone	High	Highest	Highest
Single Switch	Highest	Highest	Highest

## Topology Maximums

Table 3, Table 4, and Table 5 indicates the maximum number of switches and ports supported for each of the HP standard SAN topologies.

**Note:** The maximums shown assume the use of the minimum number of ISLs. Depending on your specific application, you may need more ISLs. This reduces the overall number of ports available for servers and storage. Attaching the Storage Management Appliance also reduces the total number of ports available for servers and storage. See Chapter 6, “[SAN Management](#).”

**Table 3: Topology Maximums when using B-Series Product Line Switches**

SAN Topology	Maximum Number of Switches	Maximum Total Number of Ports	Maximum Number of User Ports
Single Switch	1	64	64
Cascade	28	896	840*
Mesh	28	896	734*
Ring	15	480	450*
Ring with Satellite Switches	22	704	660*
Backbone (4 by 12 with 4 core switches and 12 SAN switches)*	28	640	608
Backbone (4 by 12 standard configuration)	28	640	496
Backbone (4 by 24)	28	1024 - 1280 (Refer to Chapter 3, <a href="#">SAN Fabric Design Rules</a> , for specific configuration rules.)	728 - 1200

**Note:** \*While this is a valid configuration, it achieves a high user port count by severely limiting the connectivity within the SAN by using 1 ISL between each switch pair.

**Note:** Refer to Chapter 3, [SAN Fabric Design Rules](#), for specific configuration maximums.

**Table 4: Topology Maximums when using M-Series Product Line**

SAN Topology	Maximum Number of Switches	Maximum Total Number of Ports	Maximum Number of User Ports
Single Switch	1	140	140
Cascade	24 (maximum of 8 Directors)	1,632	1024 (cascade with 8 Directors, plus 16 edge switches*)
Mesh	Not Applicable, Exceeds Hop Count Limit	Not Applicable	Not Applicable
Ring	7	980	966*
Ring with Satellite Switches	Not Applicable	Not Applicable	Not Applicable
Backbone	24	1632	1024

**Note:** \*While this is a valid configuration, it achieves a high user port count by severely limiting the connectivity within the SAN by using 1 ISL between each switch pair.

**Note:** Refer to Chapter 3, [SAN Fabric Design Rules](#), for specific configuration maximums.

**Table 5: Topology Maximums when using C-Series Product Line Switches**

SAN Topology	Maximum Number of Switches	Maximum Total Number of Ports	Maximum Number of User Ports
Single Switch (MDS 9216)	1	48	48
Single Switch (MDS 9509)	1	176	176

The Cisco MDS 9509 and the MDS 9216 switches are currently supported in single switch fabrics only.

## Data Availability in a SAN

Data availability in a computer installation is influenced by many factors, including the application software and operating systems in the servers, the server hardware, the SAN fabric infrastructure, and the primary and secondary storage. Operational parameters including backup schedule and machine room procedures, as well as personnel issues and overall administrative practice all make important contributions to the availability of data in a computer system environment.

In some environments, adequate data availability is established by a routine backup procedure performed on a scheduled basis. In other cases, online dynamic backup of primary data to a remote site is required. Some environments use clustered servers and redundant fabrics in their SAN systems in order to achieve their data availability goals.

When considering SAN fabric topology selection, the number of Fibre Channel switches and the number of ISLs between the switches have the largest effect on the data availability. The number of connections or paths between a given server or clustered servers and the fabric, and the number of storage controller connections or paths into the fabric also affect data availability.


From the perspective of SAN architecture and fabric topology design, fabric availability can be classified into at least four categories or levels. The different categories offer a range of availability levels from the most basic interconnect scheme with no redundancy, up to fully redundant No Single Point Of Failure (NSPOF) designs.

### Levels of Availability

1. Single Fabric/Single Server and Storage Paths
2. Single Meshed Fabric/Single Server and Storage Paths
3. Single Meshed Fabric/Multiple Server and Storage Paths
4. Multiple Fabrics/Multiple Server and Storage Paths

#### **Level 1: Single Non-meshed Fabric/Single Server and Storage Paths**

These designs are implemented with single links between each switch, connected in one fabric. The Fibre Channel switches are arranged so that servers and storage connect into the fabric using single paths. This type of design does not provide any level of fabric or fabric path redundancy.




**Figure 16: Level 1: Maximum Connectivity**

#### **Level 2: Single Meshed or Cascaded Fabric/Single Server and Storage Paths**

These designs have more than one ISL between switches and/or multiple paths or routes to all switches in the fabric. Servers and storage connect into the fabric using single paths. This provides the benefit of fabric resiliency. If a single switch port or a link between two switches fails, the fabric automatically re-routes data to an alternate fabric link or route. The servers see no interruption in their I/O flow.






**Figure 17: Level 2: Fabric Resiliency**

### ***Level 3: Single Meshed or Cascaded Fabric/Multiple Server and Storage Paths***

These designs are the same as Level 2 with the addition of multiple data paths between servers and storage connecting into one fabric. Level 3 offers the benefits of both fabric resiliency and multiple server and storage paths. In the unlikely event of a switch, host bus adapter, or path failure, data is automatically re-routed to an alternate path in the servers and storage, and through the fabric. The servers see no interruption in their I/O flow. Level 3 may require (depending on the O/S) the use of fabric zoning to define a minimum of two separate paths in a single fabric. To ensure high availability, each HBA must be cabled to a different switch and be configured for access to a different storage system controller when set in multiple-bus failover mode. Each controller must be cabled to a different switch, as shown in Figure 18.



**Figure 18: Level 3: Single Fabric High Availability Multi-Pathing**


### ***Level 4: Multiple Fabrics/Multiple Server and Storage Paths***

Like Level 3, Level 4 provides for multiple data paths between servers and storage, but in the Level 4 designs these paths are connected to physically separate fabrics. This type of design provides the highest level of availability and offers no single point of failure protection (NSPOF). Any event that may affect the fabric performance or usability will be overcome by routing data to another alternate fabric. The servers see no interruption in their I/O flow.

The Level 4 design eliminates any vulnerability to fabric failures, for example, human error such as improper switch replacement procedure, inadvertent erroneous fabric configuration settings, or a fabric service failure. This type of design also provides the highest level of performance and a higher number of available ports, since all fabrics can be accessed and utilized simultaneously during normal operations. This also allows for nondisruptive upgrades.

This level of protection is available for all HP standard SAN topologies by replicating the chosen design in two separate fabrics. HP recommends that the two fabrics have similar or identical topologies. Although this may increase the overall cost of the implementation, the added benefit beyond the increase in data availability is an increase in total available ports. For example, choosing to implement a single meshed fabric design using four switches provides


up to 52 ports for server and storage connectivity. Implementing the same topology using two fabrics provides up to 104 ports for server and storage connectivity. This is shown in Figure 19.



SHR-2547A

**Figure 19: Single Fabric and Dual Fabric SANs**

Using two fabrics allows for non-disruptive software and firmware code updates. For example, given the two fabrics shown in Figure 20, you can failover operations to Fabric B, upgrade Fabric A, then failback operations to Fabric A. The procedure can then be repeated in reverse to upgrade Fabric B.



SHR-2168B

**Figure 20: Level 4: Dual Fabric High Availability Multi-Pathing Fault Tolerant**

Table 6 characterizes data availability and indicates the supported topologies for each level.

**Table 6: Fabric Design Data Availability**

Fabric Design		Availability Level	SAN Topologies
Single Fabric (Non-Meshed)	1	No Redundancy	Single Switch or Multiple Switches with Single ISL Cascade
Single Meshed Fabric Multiple Fabric Paths	2	Medium	Two ISL Cascade, Meshed, Ring, Tree
Single Meshed Fabric Multiple Fabric Paths Multiple Server and Storage Paths <sup>1</sup>	3	High	All
Two (or more) Fabrics Multiple Server and Storage Paths	4	Highest (NSPOF)	All

1. May require the use of zoning to define a minimum of two separate data paths within the single fabric. This is platform dependent.

## Availability Design Considerations

Two major considerations in choosing an availability level are the criticality of data access and cost. For mission critical applications, first consider full redundant fabric designs. The additional cost can usually be justified when you consider the cost associated with the loss of access to critical data.

You should also remember that the additional cost of more than one fabric provides more than redundancy since the number of available ports will typically double. If this increased connectivity can be utilized by adding more servers and storage to the SAN, the cost factor is minimized. Figure 7 characterizes data availability levels relative to cost and total number of available ports.

**Table 7: Availability Cost Factors**

Fabric Design	Level	Hardware Cost Factor <sup>1</sup>	Available Ports <sup>2</sup>
Single Fabric (Non-Meshed)	1	x	n - #ISL Ports
Single Meshed Fabric Multiple Fabric Paths	2	x + Additional ISLs	n - #ISL Ports
Single Meshed Fabric Multiple Fabric Paths Multiple Server and Storage Paths <sup>3</sup>	3	x + Additional ISLs + Additional HBAs	n - #ISL Ports - Additional HBA Ports
Two (or more) Fabrics Multiple Server and Storage Paths	4	x + Additional ISLs + Additional HBAs + Additional Switches	2n - #ISL Ports - Additional HBA Ports

1. The variable x is the cost of a single non-meshed fabric. It is used as a reference for comparison.
2. The variable n is the total number of ports available for devices in a SAN fabric.
3. May require the use of zoning to define a minimum of two separate data paths within the single fabric. This is platform dependent.

## Scalability and Migration

Each of the HP standard SAN topologies can be scaled incrementally to increase connectivity and overall capacity. You should always plan for expected future growth in your initial SAN design to minimize disruption when expanding capabilities and capacity over time. If you do exceed the capacity of a given topology, or find that data access needs have changed, it is possible to migrate one topology to another. Refer to Chapter 10, "Best Practices" for information about migrating topologies.

Table 8 lists the migration paths and the options for scalability for all topologies.

**Table 8: Topology Migration & Scaling**

SAN Topology	Migration	Scalability (For All Topologies)
Cascaded	Convert to Meshed, Ring or Tree	<ul style="list-style-type: none"> <li>• Increase the number of switches</li> <li>• Use higher port count switches</li> <li>• Transition to a different topology</li> <li>• Deploy multiple fabrics</li> </ul>
Meshed	Convert to Ring, or Tree	
Ring	Convert to Meshed or Tree	
Tree	Add additional backbone switches	

## Custom-Designed SAN Topologies

The HP standard SAN topologies, or subsets of these topologies, as discussed in this chapter, can meet most SAN implementation requirements. There may be specific cases where HP standard topologies (or variants) do not meet your specific needs or requirements. In these cases, a custom SAN design can be created if the SAN design rules described in this document are strictly followed. Refer to Chapter 3, "SAN Fabric Design Rules", Chapter 4, "Heterogeneous SAN Platform and Storage System Rules", and Chapter 10, "Best Practices".