

Understanding SANs

1

Storage Area Networks (SANs) provide the data communication infrastructure needed for the most advanced and most cost-efficient computer mass storage systems. SANs minimize total cost of ownership for both large and small storage systems. SAN technology supports the management features and I/O price/performance demanded in today's competitive IT environment, and offers the storage component investment protection needed to minimize capital expense.

This chapter summarizes the benefits that can be obtained by using storage area networks

Why a SAN?

SANs provide unprecedented levels of flexibility in system management and configuration. Servers can be added and removed from a SAN while their data remains in the SAN. Multiple servers can access the same storage for more consistent and rapid processing. The storage itself can be easily increased, changed, or re-assigned.

In a SAN, multiple compute servers and backup servers can access a common storage pool. The SAN offers configuration choices that emphasize connectivity, performance, resilience to outage, or all three.

SANs bring enterprise-level availability to open systems servers. Properly designed SAN storage is always available. This allows many open servers to access a common storage pool with the same degree of availability previously reserved for mainframes.

SANs improve staff efficiency by supporting a variety of operating systems, servers, and operational needs. A SAN is a robust storage infrastructure that can respond quickly to new business models, unexpected growth surges, and corporate mergers.

SANs can reduce application response time, improve processing throughput, and support high-performance backup and rapid restores. SANs enable new functionality concepts such as zero backup time, they support remote data copies at nearly unlimited distances, and they support improved business continuance scenarios involving disaster recovery planning and disaster tolerant configurations. SANs also support the latest storage security measures, and they can be managed by Web-based tools from any location.

In a well-designed SAN, these features are complementary and cumulative; that is, a SAN can incorporate all of these features, or you can start with a SAN designed for any one of them and add other features later. SANs enable economy of scale that was previously unavailable to open systems in the areas of backup, management, growth, and performance. Because of this flexibility, a SAN can grow and adapt to the changing computer storage system needs of the most challenging business environment.

HP SAN Design Philosophy

The HP SAN design philosophy is to use a mix of moderately sized components to meet a wide range of storage system requirements. This approach maximizes customer value by optimizing the use of the features and functionality provided by HP SAN products. The same set of products can be used in small office environments, large business systems, and in the most demanding enterprise-class installations. HP StorageWorks products support SAN designs that work with a heterogeneous mixture of applications, operating systems, servers, storage systems, and SAN infrastructure components.

This approach to storage system design provides the following advantages.

- Greater flexibility in SAN design, to meet the widest possible range of requirements.
- Incremental scaling over time, by the addition of capacity and features as they become required.
- Support for diverse geographic and data locality requirements.

The design of HP SAN systems is facilitated by using multiple port functionality and a simple, standardized approach to storage configuration.

Multiple Port Functionality

- **Product line components enable large SANs.**

HP Fibre Channel switches have multiple ports, and can be interconnected across long distances to achieve large network configurations. A set of interconnected switches is called a Fibre Channel fabric. Each fabric has ports into which several computer servers, storage systems, and related components can be integrated. Multiple fabrics may be included in a single SAN if needed to meet connectivity or availability requirements.

- **HP StorageWorks RAID Array controllers enable heterogeneous SANs.**

Each HP controller has multiple ports for connection to the SAN. In addition, each storage controller can support a heterogeneous mixture of servers. Storage controllers and servers can be added to an existing installation in an incremental fashion. This maximizes the flexibility of the configuration and supports a number of scalable growth paths from any given SAN configuration starting point.

Approaches to Simplified Design

HP provides three approaches to SAN design and implementation. You can design and implement your SAN using an HP standard design, create a variation of an HP design, or create a custom design by following the HP StorageWorks SAN design rules. These approaches are listed in order of increasing effort and experience required to implement a particular approach.

1. **HP standard design.**

HP standard designs specify the arrangement of Fibre Channel switches within a SAN fabric and are optimized for specific [Data Locality](#) needs and typical workloads. Copying a standard design is the simplest approach to SAN design, and is recommended for anyone just starting with SAN technology.

2. **Variation of an HP design.**

Each of the HP standard topologies is optimized for a particular data locality type and offers different levels of connectivity. Often, you can select a standard SAN design that is close to your specific needs, and then modify it to meet the data locality and connectivity

requirements. By starting from a standard design, you use the SAN experience of HP to leverage your own design efforts. This approach is recommended for anyone with an intermediate level of SAN experience.

3. **Custom design using the HP StorageWorks SAN design rules.**

The SAN design rules specify the maximum limits and guidelines for custom-designed topologies and also allow SAN designs that can be tailored to meet unique or specific storage and access requirements. The design rules contain the essential information about HP StorageWorks SANs. This information is accessible and useful to anyone with an intermediate or advanced level of experience with SANs. For further information, refer to Chapter 3, "[SAN Fabric Design Rules](#)" and Chapter 4, "[Heterogeneous SAN Platform and Storage System Rules](#)."

Note: In this document "SAN topology" refers to the arrangement of Fibre Channel switches within a fabric.

Design Considerations

When a new SAN design is under development, or an existing SAN is to be modified, a number of design considerations must be evaluated. Such considerations include:

- Geographic Layout
- Data Locality
- Connectivity
- Storage Capacity
- Heterogeneous Platforms and Operating Systems
- Scalability and Migration
- Backup and Restore
- Data Availability
- Disaster Tolerance
- Switch and Hop Counts
- Oversubscription
- Performance and Application Workloads
- Manageability
- Fabric Zoning
- Selective Storage Presentation
- SAN Security

Each of these is discussed in additional detail below.

Geographic Layout

The geographical location of building sites, campuses, and facilities, and the location of servers and storage within individual buildings can be a major factor in determining the appropriate SAN design. Various SAN options are available to support long distance connections within a SAN, so a wide range of varying geographical needs can be met. An HP StorageWorks SAN can be implemented with multiple inter-switch cable segments. For more information on supported distances, refer to Chapter 3, “[SAN Fabric Design Rules](#).”

Support for these long distances also provides for interconnection of existing independent SAN islands into a single geographically distributed SAN. Refer to Chapter 8, “[SAN Extension](#)” and Chapter 11, “[Merging SAN Fabrics](#)” for further information.

Data Locality

A major factor in determining the optimal SAN topology design is the set of requirements associated with the data access patterns between storage and the associated servers. Storage and server deployment should be based on the specific application requirements for data locality. A high frequency of data reference and a short response time requirement implies a requirement for greater data locality—the SAN must be designed so that the servers and storage have a high capacity path between them.

In the context of SAN topology design, locality refers to the placement of storage systems in the SAN relative to the placement of the servers accessing the storage. Possible placements include:

- Local, “one-to-one”
- Centralized in a single storage pool or centralized pools, “many-to-one”
- Distributed among storage pools throughout the fabric, “many-to-many”

Local or "one-to-one" is where the primary data access is between individual servers and individual storage systems. In many cases this implies that they should both be connected to the same Fibre Channel switch. Centralized or "many-to-one" data access is where the primary access pattern is between many servers and a single centrally located storage system. Distributed or "many-to-many" is where data access occurs between many different servers and many different storage systems. Many-to-many data access is encountered in environments that use SAN-wide storage pooling and sharing.

Selection of the appropriate SAN topology design should primarily be based on the expected primary data locality need, however consideration should also be given to corporate, departmental, and organizational requirements relative to data grouping and accessibility.

Connectivity

Connectivity is the total number of Fibre Channel ports needed to connect servers and storage to the fabric. Ports available for server or storage connections are called "user ports". In multiple switch fabrics the number of user ports in a SAN will be less than the total number of ports in the SAN because of the need for inter-switch links (ISLs) to connect the multiple switches together. One ISL port is the minimum needed for connectivity, however more ISL connections may be required in order to provide the required performance.

The need for ports for Inter-Switch Link connectivity at the required performance level directly affects the total number of Fibre Channel switches required in a SAN. Data locality and geographic requirements must be taken into consideration when deciding on the number of ISL connections for the fabric. You should also consider future connectivity requirements and develop a design that can scale or that can migrate to a topology design with more capacity.

If the total number of ports required exceeds what is supported in a given topology, you must consider higher capacity topologies, or perhaps deploy multiple independent SANs.

Storage Capacity

The total storage capacity requirement, including expected future growth, should be calculated to ensure that the design is adequate to meet your needs. There are two aspects to storage capacity. The first is the total required capacity measured in GigaBytes or TeraBytes. Storage capacity can be increased by adding larger capacity disks, adding additional disks, or by deploying additional storage systems in the SAN.

A second aspect of capacity is performance. As disk drive sizes increase it is possible you will use fewer disks. This approach makes it easier to design a storage system that will have the required size but does not aid in designing for required performance. There are workloads whose performance is limited by a lower number of disk drives. Consider the performance impact of using a lower number of higher capacity disk drives versus a higher number of lower capacity disks if high application performance is a critical requirement.

Heterogeneous Platforms and Operating Systems

HP heterogeneous Open SANs support a wide range of multi-vendor hardware platforms and operating systems in a mixed environment. You can tailor your SAN for the specific platforms and operating systems you require. HP storage controllers can be shared across many different platforms and operating systems, all managed within the same SAN. Specific support limits of individual platforms and operating systems may vary and need to be understood and considered when evaluating SAN designs.

Scalability and Migration

A major benefit provided by HP standard SAN designs is the capability to grow or scale incrementally over time as storage and connectivity needs increase. For all designs, consideration should be given to choosing a design that will accommodate expected future growth and usage requirements.

HP-designed SAN topologies can address immediate needs and requirements, and accommodate future changes. There are migration paths for each of the topologies to provide for configuration flexibility, expansion, and increased capabilities. Refer to Chapter 11, "[Best Practices](#)" for information about scaling and migrating different SAN topology designs, as some transitions are easier to perform than others. All aspects of scaling and migration should be understood when choosing a topology design.

For further information, refer to Chapter 2, "[SAN Topologies](#)."

Backup and Restore

SAN-based backup provides high bandwidth and centralized control for your backup and restore operations. This can provide significant savings in time and management complexity over individual server or network based backup and restore implementations. SAN designs should provide adequate connectivity and bandwidth for backup, to maximize the benefits of SAN based backup. If your SAN design does not consider or accommodate backup bandwidth requirements, then you may affect backup performance. Centralized backup implies lower data locality within the SAN. Backup is an operation where data is accessed infrequently and where latency is not a concern. Refer to Chapter 5, "[Enterprise Backup Solution](#)."

Data Availability

Data availability is a broad measure of how reliable a storage system is in routine operation. Depending on the specific requirements of a given application, you can choose from a wide range of methods. In some cases, a routine tape backup on a periodic basis provides enough availability. In other cases, a SAN with multiple paths between servers and storage within a single fabric is adequate. In the most demanding environments, you can configure an HP SAN that provides No Single Point Of Failure (NSPOF) in the data access paths and in the storage systems. A mixture of different availability levels can be implemented within the same SAN, depending on the level of protection required for specific applications or data. For further information, refer to Chapter 2, "[Data Availability in a SAN](#)."

Disaster Tolerance

Disaster tolerance is a measure of how reliably data can be accessed and restored in the event of complete failure of a facility or site.

Consideration must be given to the criticality of data in the event of unforeseen catastrophic site failures. Remote data replication requirements should be considered in the SAN design to ensure protection against site failures and full recovery of critical data. Selected data can be copied to remote storage arrays, automatically providing recovery capabilities in the case of a primary site interruption or possible loss. Using multiple storage arrays, portions of the SAN can be configured for disaster tolerance, providing a common SAN with mixed data protection levels.

Switch and Hop Counts

Data routing through the fabric is described in terms of hops, where a single hop is one or more ISLs between any two switches. The general rule is that you should minimize the number of hops between devices that will communicate regularly in a SAN.

Oversubscription

Oversubscription is a normal part of any SAN design and is essentially required to help reduce the cost of the SAN infrastructure. Oversubscription refers to the fan-out ratio of available resources such as ISL bandwidth or storage system I/O capacity, to the consumers of the resource. A general rule-of-thumb relates oversubscription to the cost of the solution such that the higher the oversubscription, the less costly the solution.

Oversubscription occurs when multiple data streams on multiple ports are funneled into a single data stream on a single port. Since all ports have equal bandwidth, there is a bandwidth mismatch when the multiple parallel data streams are directed into a single port.

Oversubscription or congestion can occur in a fabric with multiple switches when data from multiple sources must be sent to a single destination port, or when data is required to be sent across an ISL from multiple input ports. In situations where this occurs, the Fibre Channel switches utilize fairness algorithms to ensure that all devices are serviced. The switches use the fairness algorithm to interleave frames from multiple devices, thus giving fractional bandwidth to all devices. If this occurs often, then overall performance in the fabric will be reduced. Oversubscription can be minimized by ensuring that your fabric design provides for an adequate number of ISLs between all switches, and by minimizing the cases where many devices or ports are attempting to share a single switch port.

Performance and Application Workloads

Performance requirements need to be considered in any SAN fabric design. This can be difficult, because data traffic in a SAN is not always predictable. Applications can usually be classified as high bandwidth or high throughput. What is important is that the SAN provides an adequate level of performance based on the workload presented by the applications.

Other factors to consider are the locality of data in relation to the servers most likely to access the data, and the number and placement of ISLs between switches in the fabric. In general, SAN topology designs with fewer switch hops between devices provide better performance due to a lower probability of oversubscription or congestion.

Manageability

SAN management can be centralized using a dedicated Storage Management Appliance, regardless of the arrangement or location of the storage components. The Storage Management Appliance connects directly to the SAN through a Fibre Channel switch, providing it with connectivity to all devices connected to the SAN. Refer to Chapter 6, "[SAN Management](#)", for more information on SAN management.

Fabric Zoning

Zoning is a fabric management service used to define logical subsets within a SAN. Zoning enables resource partitioning for management and access control by dividing a physical SAN into multiple overlapping logical zones, where each zone is defined by the set of ports or devices that are addressable within the zone. The HP Fibre Channel switch zoning feature provides a way to control SAN access at the device or port level. This capability allows you to set up barriers between different operating environments, to deploy logical fabric subsets by creating defined server and/or storage groups, or to define temporary server and storage zones for tape backup. Zones can be configured dynamically.

Selective Storage Presentation

HP storage systems implement a LUN masking feature called Selective Storage Presentation (SSP). This feature allows you to assign or selectively present logical units on a given storage system to one or more servers in the SAN. This provides protection against unintended access of a given storage set by a given server, while allowing accesses that are needed for proper operation. This provides a level of data access security, and allows multiple operating systems to be used in a single SAN.

Utilization of both SSP and Fabric Zoning provides for the most flexible SAN node and device access management. These features should be viewed as complementary in that usage of both provides the greatest range of SAN storage access management capabilities. For more information, refer to Chapter 6, "[Selective Storage Presentation](#)."

SAN Security

Security is provided in HP SANs and storage systems by a combination of product features and system management practices. HP StorageWorks SAN hardware and SAN management tools provide reliable access to data, robust data storage, and enforcement of data access restrictions, and careful system management ensures that proper security practices are followed. For more information, refer to Chapter 9, "[SAN Security](#)."

Summary

SAN design requires the consideration of many factors. To successfully complete a SAN design and implementation, you provide the requirements for your SAN and HP provides the product features and capabilities needed to meet the requirements. HP offers a design philosophy and standard SAN topologies that can be used as guidance. In addition, HP provides a comprehensive set of design rules. By following these rules, you will have a SAN configuration that is supported by the HP storage engineering organization.

This guide provides the detailed information you need to design a SAN that meets your unique storage system requirements.